The Minimization of Random Hypergraphs[*]

Thomas Bläsius[1], Tobias Friedrich[2], and Martin Schirneck[2]

[1]*Karlsruhe Institute of Technology, Karlsruhe, Germany*
[2]*Hasso Plattner Institute, University of Potsdam, Germany*
`firstname.lastname@{kit.edu|hpi.de}`

Abstract

We investigate the maximum-entropy model $\mathcal{B}_{n,m,p}$ for random n-vertex, m-edge multi-hypergraphs with expected edge size pn. We show that the expected size of the minimization of $\mathcal{B}_{n,m,p}$, i.e., the number of its inclusion-wise minimal edges, undergoes a phase transition with respect to m. If m is at most $1/(1-p)^{(1-p)n}$, then the minimization is of size $\Theta(m)$. Beyond that point, for α such that $m = 1/(1-p)^{\alpha n}$ and H being the entropy function, it is

$$\Theta(1) \cdot \min\left(1, \frac{1}{(\alpha - (1-p))\sqrt{(1-\alpha)n}}\right) \cdot 2^{(\mathrm{H}(\alpha)+(1-\alpha)\log_2 p)n}.$$

This implies that the maximum expected size over all m is $\Theta((1+p)^n/\sqrt{n})$.

Our structural findings have algorithmic implications for minimizing an input hypergraph, which in turn has applications in the profiling of relational databases as well as for the Orthogonal Vectors problem studied in fine-grained complexity.

The main technical tool is an improvement of the Chernoff–Hoeffding inequality, which we make tight up to constant factors. We show that for a binomial variable $X \sim \mathrm{Bin}(n,p)$ and real number $0 < x \leq p$, it holds that

$$\mathrm{P}[X \leq xn] = \Theta(1) \cdot \min\left(1, \frac{1}{(p-x)\sqrt{xn}}\right) \cdot 2^{-\mathrm{D}(x \,\|\, p)\,n},$$

where D denotes the Kullback–Leibler divergence between Bernoulli distributions. The result remains true if x depends on n as long as it is bounded away from 0.

1 Introduction

A plethora of work has been dedicated to the analysis of random graphs. Random hypergraphs, however, received much less attention. For many types of data, hypergraphs provide a much more natural model. This is especially true if the data has a hierarchical structure or reflects interactions between groups of entities. In non-uniform hypergraphs, where edges can have different numbers of vertices, a phenomenon occurs that is unknown to graphs: an edge may be contained in another, with multiple edges even forming chains of inclusion. We are often only interested in the endpoints of those chains, namely, the collections of inclusion-wise minimal or maximal edges. This is the *minimization* or *maximization* of the hypergraph, respectively.

We investigate the maximum-entropy model $\mathcal{B}_{n,m,p}$ for random multi-hypergraphs with n vertices and m edges and expected edge size pn for a constant sampling probability p. In other words, out of all probability distributions on hypergraphs that result in an expected edge size of pn, $\mathcal{B}_{n,m,p}$ is the one of maximum entropy. We are interested in the expected size of the minimization/maximization of $\mathcal{B}_{n,m,p}$, that is, the expected number of its minimal/maximal edges. Our results are phrased in terms of the minimization, but replacing the probability p with

---

$1-p$ immediately transfers them to the maximization. We show that the size of the minimization undergoes a phase transition with respect to the number of edges m with the point of transition at $m^* = 1/(1-p)^{(1-p)n}$. While the number of edges is still small, a constant fraction of them is minimal and the minimization grows linearly in the total sample sizes. For $m > m^*$, the size of the minimization is instead governed by the entropy function of the exponent α such that $m = 1/(1-p)^{\alpha n}$ (see Theorem 2 for a precise statement). We show that the ratio of minimal edges goes down exponentially as m increases. This seems counter-intuitive at first as it decouples the behavior of the minimization from the growth of the underlying hypergraph. We further determine that the maximum expected number of minimal edges over all m is of order $\Theta((1+p)^n/\sqrt{n})$, the maximum is attained at $m = 1/(1-p)^{\frac{n}{1+p}}$.

Our results establish a close connection between the size of the minimization and the binomial distribution.[1] Theorem 1 (see below) characterizes the number of minimal edges in terms of the total number of edges m and the likelihood that a binomial variable deviates from its expectation. The main tool in our analysis is the Chernoff–Hoeffding theorem bounding the tail of the distribution function via the Kullback–Leibler divergence from information theory. However, the existing inequalities are not sharp enough to derive tight statements on the expected size of the minimization. So far, there is an $O(\sqrt{n})$ gap between the best-known upper and lower estimates. In this work, we improve both sides such that they match up to constant factors. The improvement transfers directly to the weaker, but more practical, versions of Chernoff bounds widely used in the literature. While our findings concentrates on the asymptotics, we also give an explicit interval for the leading constants involved and our proofs indicate how large n must be to be applicable. This makes the result useful also in a non-asymptotic setting and in areas beyond the treatment of random hypergraphs. We illustrate this by giving an alternative proof that the binomial distribution does not vanish outside of a constant number of standard deviations from its mean. Our proof avoids the Berry–Esseen inequality, that is, normal approximation. We also improve the rate of convergence in Cramér's theorem.

Regarding our main topic of hypergraphs, our structural insights have algorithmic implications for the task of actually computing the minimization $\min(\mathcal{H})$ of an input hypergraph $\mathcal{H}$. We give two examples from fine-grained complexity as well as data profiling. There is reason to believe that there exists no minimization algorithm running in time $O(m^{2-\varepsilon}) \cdot \mathrm{poly}(n)$ for any $\varepsilon > 0$ on m-edge, n-vertex hypergraphs. The argument uses the Sperner Family problem, which is to decide whether $\mathcal{H}$ contains two edges such that one is contained in the other, that is, whether $|\min(\mathcal{H})| < |\mathcal{H}|$. The latter is equivalent to the, arguably more prominent, Orthogonal Vectors problem studied in fine-grained complexity. A truly subquadratic algorithm for the minimization would thus falsify the Orthogonal Vectors Conjecture and in turn the Strong Exponential Time Hypothesis (SETH), resulting in a major breakthrough in Boolean satisfiability. On the other hand, partitioning the edges of the hypergraph $\mathcal{H}$ by their cardinality and processing them in order of increasing cardinality gives an $O(mn|\min(\mathcal{H})|)$-time algorithm, which is $O(m^2 n)$ in the worst case. When looking at the average-case complexity for the $\mathcal{B}_{n,m,p}$ distribution, we get a run time of $O(mn\,\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|])$. Our results therefore show that the algorithm is subquadratic on average for all m beyond the phase transition, it is even linear for m larger than $1/(1-p)^n$.

There is also a connection to the profiling of relational databases. Data scientists regularly need to compile and output a comprehensive list of metadata, like unique column combinations, functional dependencies, or, most general, denial constraints. These multi-column dependencies can all be described as the minimal hitting sets of certain hypergraphs created from comparing pairs of rows in the database and recording the sets of attributes in which they differ. Computing these difference sets one by one generates an incoming stream of seemingly random subsets. Filtering the inclusion-wise minimal ones from the stream does not affect the solution, but can greatly reduce the number of sets and the complexity of the resulting hitting set instance. Minimizing the input is therefore a standard preprocessing technique in data profiling. In real-world databases, there are often fewer minimal difference sets than rows in the database, let alone pairs thereof. Therefore, the space needed to store the sets usually makes up only a small fraction of the original input size. The upper bounds given in Theorems 2 provide a theoretical explanation for this observation. We show that only a few difference sets can be expected to be

---

[1]The notation $\mathcal{B}_{n,m,p}$ is mnemonic of the binomial distribution emerging in the sampling process.

minimal and their number even shrinks as the database grows larger. Conversely, the difference sets and the corresponding multi-column dependencies are mutually dual, which allows to recover the minimized input from the the collection of all solutions. In this sense, the matching lower bounds in Theorems 2 can be seen as the smallest amount of data any enumeration algorithm needs to process in order to correctly cover all dependencies.

Related Work. Erdős–Rényi graphs $\mathcal{G}_{n,m}$ [30] and Gilbert graphs $\mathcal{G}_{n,p}$ [35] are arguably the most discussed random graph models in the literature. We refer the reader to the monograph by Bollobás [16] for an overview. A majority of the work on these models concentrates on various phase transitions with respect to the number of edges m or the sample probability p, respectively. This intensive treatment is fueled by the appealing property that Erdős–Rényi graphs are "maximally random" in that they do not assume anything but the number of vertices and edges. More formally, among all probability distributions on graphs with n vertices and m edges, $\mathcal{G}_{n,m}$ is the unique distribution of maximum entropy. The same holds for $\mathcal{G}_{n,p}$ under the constraint that the expected number of edges is $p\binom{n}{2}$, see [2].

The intuition of being maximally random is captured by the Shannon entropy, which is the central concept in information theory [22, 63]. A discrete stochastic system described by the probability distribution $(p_i)_i$ has a (binary) entropy of $H((p_i)_i) = -\sum_i p_i \log_2 p_i$. The self-information of a single state with probability p is $-\log_2 p$, the entropy is the expected information of the whole system. It is a measure of surprisal or how "spread out" the distribution is. Originally stemming from thermodynamics [49], the versatility of this definition is key to the successful application of information theory to fields as diverse as cryptography [19], machine learning [36], quantum computing [55], and of course network analysis [54]. Jaynes' *principle of maximum entropy* states that out of an ensemble of probability distributions that all describe the observed phenomena equally well, the one of maximum entropy is to be preferred in order to minimize any outside bias [40, 41, 47]. Properties of the maximum-entropy model then pertain to the average system matching the observations. In the context of random graphs, it is used to define so-called null models [66]. After certain graph statistics observed in real-world networks are fixed, one chooses the maximum-entropy distribution that meets these constraints. By comparing the original network with a "typical" graph drawn from the null model, one can infer whether other properties are correlated with the constraints. This method was made rigorous by Park and Newman [59] building on earlier work in general statistics. Prescribing the exact or expected number of edges leads to the $\mathcal{G}_{n,m}$ or $\mathcal{G}_{n,p}$ distributions, respectively. The configuration model fixes the whole degree sequence of the graph [17] and the soft configuration model relaxes these constraints to hold in expectation [11, 34].

Many early attempts to transfer the concept of null models to hypergraphs have been only indirect in that they have studied hypergraphs via their clique expansion [53] or as bipartite graphs [61]. This is unsatisfactory since the projections alter relevant observables, like node degrees or the number of triangles. Only recently, Chodrow generalized the configuration model directly to multi-hypergraphs [20], which has subsequently been refined by Arafat et al. [3]. There also seems to be not much work on hypergraph models that can be cast into the maximum-entropy framework without being intentionally designed as such, a notable exception is the work by Schmidt-Pruzan and Shamir [62]. They fix the exact (respectively, expected) edge sequence such that the largest edge has cardinality $O(\log n)$ and show a "double jump" phase transition in the size of the largest connected component. Most of the recent literature on random hypergraphs concentrates on the k-uniform model where every edge has exactly k vertices [7, 8, 43] or, equivalently, on random binary matrices with k 1s per column [21]. Our model neither prescribes the exact cardinalities of the edges nor a bound on their maximum size, instead it only requires that the *expected* edge size is pn.

Probably closest to our work is a string of articles by Demetrovics et al. [26] as well as Katona [44, 45]. They investigate random databases and connect the Rényi entropy of *order 2* of the logarithmic number of rows with the probability that certain unique column combinations or functional dependencies hold. In contrast, we connect the Shannon entropy of the logarithmic number of *pairs* of rows (quantity α, see Section 2) with the expected number of minimal difference sets. Unique column combinations, functional dependencies, and their generalization to denial constraints are dual to the difference sets, one are the hitting sets of the

other [1, 10, 12, 32]. Also, the Shannon entropy is equal to the Rényi entropy of *order 1* [22]. In this sense, we complement the result by Demetrovics et al. and show that the duality also pertains to the order of entropy. Furthermore, it has often been observed in practice that the collection of minimal difference sets of real-world databases is much smaller than the original instance size [14, 15, 50, 58], our findings provide a theoretical explanation for this phenomenon.

The minimization of hypergraphs also occurs in fine-grained complexity in form of the Sperner Family problem. It is subquadratically equivalent to the Orthogonal Vectors problem [18, 33], which in turn admits a fine-grained reduction from CNF-Satisfiability [65]. Any truly sub-quadratic algorithm for computing the minimization of a hypergraph *in the worst case* would mark a major breakthrough in satisfiability. Very recently, ideas from fine-grained complexity have been extended to to the average case [6, 24, 42]. We show that a simple algorithm for Sperner Family is subquadratic *on average* on hypergraphs with expected edge size pn.

The analysis of random (hyper-)graphs naturally builds on tools from combinatorics and probability theory. Conversely, it has always helped to advance those fields by improving on known techniques [9, 16, 39]. The binomial distribution is used all throughout science to describe complex systems emerging from the overlapping effects of many independent choices. Its typical behavior is well understood, as described in the central limit theorem and the strong law of large numbers. Bounding its tails, however, remains the subject of ongoing research. The resulting concentration inequalities play a significant role in the theory of large deviations [25, 27], the analysis of randomized algorithms and search heuristics [28, 29, 51], as well as computational learning [46], to name a few examples. In this work, we sharpen the tail inequalities of the Chernoff–Hoeffding theorem [38]. We prominently use a result by Klar [48] on the relation between the distribution function and the probability mass function. Some refined inequalities have been known before. By Cramér's theorem [23], Chernoff–Hoeffding is asymptotically tight up to subexponential factors and the gap has subsequently been reduced to $O(\sqrt{n})$, see [4]. We close it down to a constant. There also exist comparatively tight bounds based on the normal limit of the binomial distribution. Contributions by Prokhorov [60] and later Slud [64] founded major lines of research in that direction. However, we avoid this approach and give a purely combinatorial argument since the normal approximation cannot be expressed in terms of elementary functions. Also, it tends to place unnecessary restrictions on the sampling probability p for non-asymptotic results.

Outline. In the next section, we introduce the hypergraph model and state our results in full detail. We review some notation and general concepts in Section 3. Section 4 is dedicated to the Chernoff–Hoeffding theorem. Section 5 adds further technical contributions as the foundation of the subsequent proofs. The main theorems on the size of the minimization are then proven in Section 6. Section 7 concludes the work.

2   Model and Main Results

Fix a probability p and positive integers n and m. The random multi-hypergraph $\mathcal{B}_{n,m,p}$ is generated by independently sampling m (not necessarily distinct) subsets of $[n]$. Each set contains any vertex $v \in [n]$ with probability p independently of all other choices.[2]

We quickly argue that it is indeed the maximum-entropy model. Besides the size of the universe n and the number of edges m, the only constraint is the expected edge size pn. The independence bound on the entropy reads as follows. Let X_1 to X_m be random variables with joint distribution $P_{X_1,\dots,X_m}$ and marginals P_{X_j}. Then, their entropies observe the inequality $H(P_{X_1,\dots,X_m}) \leq \sum_{j=1}^m H(P_{X_j})$, with equality holding if and only if the X_j are independent [22]. This implies that the edges need to be sampled independently in order to maximize the entropy, the same is true for the vertices inside one edge. Finally, the fact that setting the sampling probability of the vertices to be all equal indeed gives the maximum entropy under a given mean set size was proven by Harremoës [37].

We are interested in the expected number of inclusion-wise minimal edges in $\mathcal{B}_{n,m,p}$, denoted by $E[|\min(\mathcal{B}_{n,m,p})|]$. We describe the asymptotic behavior of this expectation with respect to

---

[2]In the context of Boolean functions, this is called the *p-biased distribution* [56].
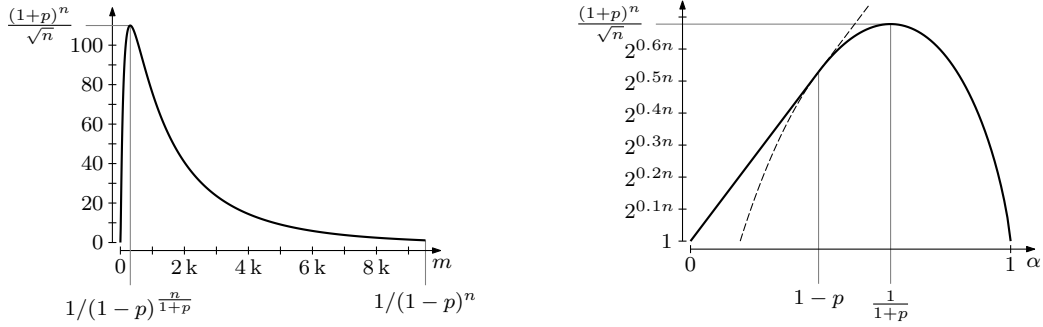
(a) The expected size of the minimization as a function of m for $n = 10$ and $p = 0.6$ in the information-theoretic regime. The vertical line at $m = 1/(1-p)^{\frac{n}{1+p}}$ indicates the position of the maximum (Lemma 27). For $m > 1/(1-p)^n$, the size goes to 1. The linear bound for $m \leq 1/(1-p)^{(1-p)n}$ is not shown as it is too close to 0.

(b) The expected size of the minimization as a function of α for $p = 0.6$ (the plot is independent of n). The vertical line at $\alpha = 1 - p$ indicates the phase transition between the linear and the information-theoretic regime. The respective bounds are continued as dashed lines into the other regime. The vertical line at $\alpha = 1/(1+p)$ indicates the position of the maximum (Lemma 27).

Figure 1: Illustration of Theorem 2 showing the expected size of the minimization of a random hypergraph depending on the number of edges m (a) and on α (b). As α grows logarithmically in m, (b) shows the same plot as (a) but with both axes being logarithmic.

n. In more detail, we view the number of edges $m = m(n)$ as a function of n and bound the univariate asymptotics of $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|]$ with respect to n for any choice of m. The sampling probability p is considered to be a constant throughout.

We show first that the expected size of the minimization can be described precisely in terms of m and the binomial distribution $\mathrm{Bin}(n,p)$. To state our result in full detail, we define

$$\alpha = \log_{\frac{1}{(1-p)^n}} m = -\frac{\log_{1-p} m}{n}.$$

The quantity α is a non-negative function of p, n, and m, it is well-defined for all $0 < p < 1$ and $n, m \geq 1$. Asymptotically in n, it is of order $\Theta((\log m)/n)$. If p and n are fixed, choosing a value for α determines m since we can rewrite m as $1/(1-p)^{\alpha n}$.

Theorem 1. *Let n, m be positive integers and p a probability. If $p = 0$ or $p = 1$, then $|\min(\mathcal{B}_{n,m,p})| = 1$ holds deterministically. For $0 < p < 1$, let $X \sim \mathrm{Bin}(n,p)$ be a binomially distributed random variable.*

(i) *For any $\varepsilon > 0$ and all $m \leq 1/(1-p)^{(1-\varepsilon)n}$, i.e., all $0 \leq \alpha \leq 1 - \varepsilon$,*

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(m) \cdot \mathrm{P}[X \leq (1-\alpha)n].$$

(ii) *There exists a $c > 0$ such that for all $m \geq 1/(1-p)^{n+c \ln n}$, i.e., all $\alpha \geq 1 + \frac{c \ln n}{n}$,*

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = 1 + \mathrm{o}(1).$$

The asymptotic estimate in the first statement is tight up to constants, the second statement is even tight up to lower-order terms. The constants hidden in the big-O notation are universal in the sense that they, of course, do not depend on m or n, but also not on α, which describes the relation between the former two. We note that the constants may depend on p and ε. There is a gap in the theorem at $m = 1/(1-p)^n$, which can, however, be made arbitrarily small. Let $C = 1/(1-p)$, then Statement (i) holds if $m \leq (C-\delta)^n$ for any constant $\delta > 0$ and Statement (ii) takes over at $m \geq (C + \mathrm{o}(1))^n$. Unlike what one might expect, we show (in Lemma 26) that Theorem 1 (i) cannot be extended to the case where α converges to 1.

The characterization of the size of the minimization via the binomial distribution allows us to give estimates in information-theoretic terms, namely, using the Shannon entropy of α. This reveals a phase transition at $m^* = 1/(1-p)^{(1-p)n}$, i.e., for $\alpha = 1 - p$. Let $\mathrm{H}(x) = -x \log_2 x - (1-x) \log_2(1-x)$ denote the binary entropy function.

Theorem 2. *Let n, m be positive integers and $0 < p < 1$ a non-trivial probability.*

(i) If $m \leq 1/(1-p)^{(1-p)n}$, then $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(m)$.

(ii) For any $\varepsilon > 0$ and all $1/(1-p)^{(1-p)n} \leq m \leq 1/(1-p)^{(1-\varepsilon)n}$, i.e., all $1 - p \leq \alpha \leq 1 - \varepsilon$,

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(1) \cdot \min\left(1, \frac{1}{(\alpha - (1-p))\sqrt{(1-\alpha)n}}\right) \cdot 2^{(\mathrm{H}(\alpha) + (1-\alpha)\log_2 p)\,n}$$

$$= \Theta(1) \cdot \min\left(1, \frac{1}{(\alpha - (1-p))\sqrt{(1-\alpha)n}}\right) \cdot \left(\frac{p^{1-\alpha}}{(1-\alpha)^{1-\alpha}\,\alpha^\alpha}\right)^n.$$

The bounds in the two cases are very different in nature. They are visualized in Figure 1 showing the expected size of the minimization as a function of the number of edges m and of α. To distinguish the cases also in writing, we use the term *linear regime* if m is between 1 and $1/(1-p)^{(1-p)n}$, corresponding to $0 \leq \alpha \leq 1 - p$, likewise, we refer to m being between $1/(1-p)^{(1-p)n}$ and $1/(1-p)^n$, i.e., $1 - p \leq \alpha \leq 1$, as the *information-theoretic regime*. In Statement *(ii)* of the above theorem, the leading constant may depend on the choice of ε.

In order to derive Theorem 2 from Theorem 1, we tighten the Chernoff–Hoeffding inequality on the binomial distribution function. Let $\mathrm{D}(x \parallel y) = -x\log_2\left(\frac{y}{x}\right) - (1-x)\log_2\left(\frac{1-y}{1-x}\right)$ denote the binary Kullback–Leibler divergence.

Theorem 3. *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n,p)$ a binomial variable. Suppose the function $x = x(n)$ takes real values in the interval $[\varepsilon, 1 - \varepsilon]$ for some $\varepsilon > 0$. Let φ and ψ denote the functions*

$$\varphi(n,p,x) = \min\left(1, \frac{1}{(p-x)\sqrt{xn}}\right) \quad and \quad \psi(n,p,x) = \min\left(1, \frac{1}{(x-p)\sqrt{(1-x)n}}\right),$$

with additionally $\varphi(n,p,p) = 1$ and $\psi(n,p,p) = 1$.

There exist constants $C_1, C_2, C_3, C_4 > 0$, independent of n and x but possibly dependent on p and ε, such that the following statements hold for all n sufficiently large.

(i) If $x \leq p$, then $C_1 \varphi \cdot 2^{-\mathrm{D}(x\parallel p)\,n} \leq \mathrm{P}[X \leq xn] \leq C_2 \varphi \cdot 2^{-\mathrm{D}(x\parallel p)\,n}$.

(ii) If $x \geq p$, then $C_3 \psi \cdot 2^{-\mathrm{D}(x\parallel p)\,n} \leq \mathrm{P}[X \geq xn] \leq C_4 \psi \cdot 2^{-\mathrm{D}(x\parallel p)\,n}$.

3 Preliminaries and Notation

Multi-Hypergraphs. A *hypergraph* on the vertex set $[n] = \{1, \ldots, n\}$ is a set of subsets $\mathcal{H} \subseteq \mathcal{P}([n])$, called the *(hyper-)edges*. If $\mathcal{H}$ is a multiset instead, we have a *multi-hypergraph*. We do not allow multiple copies of the same vertex in one edge. The *minimization* of a hypergraph $\mathcal{H}$ is the collection of its inclusion-wise minimal edges, $\min(\mathcal{H}) = \{E \in \mathcal{H} \mid \forall E' \in \mathcal{H} \colon E' \subseteq E \Rightarrow E' = E\}$. We extend this notion to multi-hypergraphs by requiring that, whenever a minimal edge has multiple copies, only one of them is included in the minimization. This way $\min(\mathcal{H})$ is always a mere hypergraph (a set). For a multi-hypergraph $\mathcal{H}$, we use $|\mathcal{H}|$ to denote the total number of edges counting multiplicities, and $\|\mathcal{H}\|$ for the number of distinct edges, that is, the cardinality of the support of $\mathcal{H}$. Evidently, we have $|\min(\mathcal{H})| \leq \|\mathcal{H}\| \leq |\mathcal{H}|$.

Information Theory. We intend the expressions $0\log_a 0$ and $0\log_a\left(\frac{0}{0}\right)$ to both mean 0 for any positive real base a. Note that this convention implies $0^0 = a^{0\log_a 0} = 1$ and $\left(\frac{0}{0}\right)^0 = 1$. We use $\mathrm{ld}\,x$ for the binary (base-2) logarithm of x. The (binary) *entropy function* H is defined for all probabilities x as

$$\mathrm{H}(x) = -x\,\mathrm{ld}\,x - (1-x)\,\mathrm{ld}(1-x).$$

The entropy function is the Shannon entropy (equivalently, the Rényi entropy of order 1) of the Bernoulli distribution with parameter x. In the notation of the previous sections, we have $H(x) = H((x, 1-x))$. The entropy function is symmetric around $1/2$ with $H(x) = H(1-x)$. On the open unit interval, H is positive and strictly concave, it has its maximum at $1/2$ with value $H(1/2) = 1$. The entropy power

$$2^{H(x)} = \frac{1}{x^x (1-x)^{1-x}}$$

is called the *perplexity* of x. We use it to estimate binomial coefficients, see [22].

Lemma 4. *Let n be a positive integer and $0 < x < 1$ a rational such that xn is an integer, then*

$$\frac{2^{H(x)n}}{\sqrt{8\,n\,x(1-x)}} \leq \binom{n}{xn} \leq \frac{2^{H(x)n}}{\sqrt{\pi\,n\,x(1-x)}}.$$

For any two probabilities x and y, the (binary) *Kullback–Leibler divergence*[3] between Bernoulli distributions with respective parameters x and y is

$$D(x \,\|\, y) = -x \operatorname{ld}\!\left(\frac{y}{x}\right) - (1-x) \operatorname{ld}\!\left(\frac{1-y}{1-x}\right).$$

The function $D(x \,\|\, y)$ is convex in both x and y, and attains its minimum 0 for $x = y$. The divergence observes $D(x \,\|\, y) = D(1-x \,\|\, 1-y)$. Its partial derivative with respect to x is

$$\frac{\partial}{\partial x} D(x \,\|\, y) = \operatorname{ld}\!\left(\frac{x}{1-x} \frac{1-y}{y}\right).$$

Next, we establish the fact that the divergence scales quadratically in the difference $y - x$.

Lemma 5. *Let $0 < x \leq y < 1$ be two non-trivial probabilities. Define t^+ to be the maximizer of $t(1-t)$ over the interval $[x, y]$, and t^- the minimizer. Then, it holds that*

$$\frac{(y-x)^2}{t^+(1-t^+)} \leq 2\ln(2) \cdot D(x \,\|\, y) \leq \frac{(y-x)^2}{t^-(1-t^-)}.$$

Proof. Let $\varepsilon = y - x$. $D(y - \varepsilon \,\|\, y)$ is two-times differentiable with respect to ε with derivatives

$$\frac{\partial}{\partial \varepsilon} D(y - \varepsilon \,\|\, y) = \operatorname{ld}\!\left(\frac{y}{1-y} \frac{1-y+\varepsilon}{y-\varepsilon}\right) \quad \text{and} \quad \frac{\partial^2}{\partial \varepsilon^2} D(y - \varepsilon \,\|\, y) = \frac{1}{\ln 2} \frac{1}{(y-\varepsilon)(1-y+\varepsilon)}.$$

Both the divergence itself as well as its first derivative vanish at $\varepsilon = 0$. By Taylor's theorem (with the Lagrange form of the remainder), there exists an ξ with $0 \leq \xi \leq \varepsilon$ such that

$$D(y - \varepsilon \,\|\, y) = \frac{\varepsilon^2}{2!} \cdot \frac{\partial^2}{\partial \varepsilon^2} D(y - \varepsilon \,\|\, y)\bigg|_{\varepsilon = \xi} = \frac{\varepsilon^2}{2\ln 2\,(y-\xi)(1-y+\xi)}.$$

The lemma follows from $y - \xi$ ranging over $[x, y]$. $\qquad\square$

We often use the following quantity derived from the divergence, resembling the perplexity.

$$2^{-D(x \,\|\, y)} = 2^{H(x)} \cdot y^x (1-y)^{1-x} = \left(\frac{y}{x}\right)^x \left(\frac{1-y}{1-x}\right)^{1-x}.$$

The next lemma is useful when relating quantities of this kind for different parameters.

Lemma 6. *Let $0 \leq x \leq y \leq z \leq 1$ be three probabilities, then*

$$2^{-D(x \,\|\, z)} = \left(\frac{y}{1-y} \frac{1-z}{z}\right)^{y-x} 2^{-D(x \,\|\, y)} \cdot 2^{-D(y \,\|\, z)}.$$

In particular, for any fixed z, $2^{-D(x \,\|\, z)}$ is non-decreasing in x as long as $x \leq z$.

---

[3]The divergence is sometimes called *relative entropy*, we avoid this term due to ambiguities, see [22].

Proof. The convention $\left(\frac{0}{0}\right)^0 = 1$ ensures that $\left(\frac{y}{1-y}\frac{1-z}{z}\right)^{y-x}$ is well-defined even for $z = 0$.

$$\frac{2^{-\mathrm{D}(x\,\|\,z)}}{2^{-\mathrm{D}(y\,\|\,z)}} = \frac{\left(\frac{z}{x}\right)^x \left(\frac{1-z}{1-x}\right)^{1-x}}{\left(\frac{z}{y}\right)^y \left(\frac{1-z}{1-y}\right)^{1-y}} = \frac{\left(\frac{z}{x}\right)^x \left(\frac{1-z}{1-x}\right)^{1-x}}{\left(\frac{z}{y}\right)^{y-x} \left(\frac{z}{y}\right)^x \left(\frac{1-z}{1-y}\right)^{1-x} \left(\frac{1-z}{1-y}\right)^{x-y}}$$

$$= \frac{1}{\left(\frac{z}{y}\right)^{y-x} \left(\frac{1-z}{1-y}\right)^{x-y}} \left(\frac{y}{x}\right)^x \left(\frac{1-y}{1-x}\right)^{1-x} = \left(\frac{y}{1-y}\frac{1-z}{z}\right)^{y-x} 2^{-\mathrm{D}(x\,\|\,y)}.$$

The monotonicity follows from the last two factors being at most 1 due to $x \le y \le z$. $\qquad\square$

Polynomials of Probabilities. We regularly estimate expressions of the form $(1-x)^n$ where x is a probability. The first inequality is taken from [52].

Lemma 7. *Let n be a positive integer and x a real number such that $|x| \le n$, then*

$$e^x \left(1 - \frac{x^2}{n}\right) \le \left(1 + \frac{x}{n}\right)^n.$$

We reach rather tight bounds on $(1-x)^n$ by substituting x for $-nx$ above, and combining this with the simple fact that $(1+x) \le e^x$ holds for all x.

Corollary 8. *Let n be a non-negative integer and x a probability, then*

$$e^{-nx} \left(1 - n x^2\right) \le (1-x)^n \le e^{-nx}.$$

The next inequalitiy was given by Badkobeh, Lehre, and Sudholt [5].

Lemma 9 (Lemma 10 in [5]). *Let n be a non-negative integer and x a probability, then*

$$\frac{nx}{1 + nx} \le 1 - (1-x)^n \le nx.$$

We prepare the following lemma on conditional probabilities of series of events for later.

Lemma 10. *Consider a random experiment with three outcomes A, B, and C with $\mathrm{P}[B] > 0$. In a series of m i.i.d. trials, let A_j denote the event that the outcome of the j-th trial is A, same with B. Then, it holds that $\mathrm{P}[\forall j \le m\colon \neg A_j \mid \exists k \le m\colon B_k] \le \mathrm{P}[\forall j \le m\colon \neg A_j \mid B_m]$.*

Proof. The case $\mathrm{P}[B] = 1$ is trivial, thus assume $0 < \mathrm{P}[B] < 1$. The assertion in the lemma is equivalent to

$$\mathrm{P}[\forall j \le m\colon \neg A_j \mid \neg B_m \wedge (\exists k < m\colon B_k)] \le \mathrm{P}[\forall j \le m\colon \neg A_j \mid B_m]. \tag{1}$$

To see this, observe that for any four reals x, y, z, w such that y, w, and $y + w$ are all non-zero, $\frac{x+z}{y+w} \le \frac{z}{w}$ holds if and only if $\frac{x}{y} \le \frac{z}{w}$ does. The event $[\exists k \le m\colon B_k]$ can be partitioned into $[\neg B_m \wedge (\exists k < m\colon B_k)]$ and $[B_m]$, giving

$$\mathrm{P}[\forall j \le m\colon \neg A_j \mid \exists k \le m\colon B_k] = \frac{\mathrm{P}[(\forall j \le m\colon \neg A_j) \wedge (\exists k \le m\colon B_k)]}{\mathrm{P}[\exists k \le m\colon B_k]}$$

$$= \frac{\mathrm{P}[(\forall j \le m\colon \neg A_j) \wedge \neg B_m \wedge (\exists k < m\colon B_k)] + \mathrm{P}[(\forall j \le m\colon \neg A_j) \wedge B_m]}{\mathrm{P}[\neg B_m \wedge (\exists k < m\colon B_k)] + \mathrm{P}[B_m]}.$$

Applying the observation to the real numbers $x = \mathrm{P}[(\forall j \le m\colon \neg A_j) \wedge \neg B_m \wedge (\exists k < m\colon B_k)]$, $y = \mathrm{P}[\neg B_m \wedge (\exists k < m\colon B_k)]$, $z = \mathrm{P}[(\forall j \le m\colon \neg A_j) \wedge B_m]$, and $w = \mathrm{P}[B_m]$ gives the equivalence.

The actual lemma is proven by induction over m. The case $m = 1$ is trivial as both sides simplify to $P[\neg A_1 \mid B_1]$. Suppose that $P[\forall j < m : \neg A_j \mid \exists k < m : B_k] \leq P[\forall j < m : \neg A_j \mid B_{m-1}]$ holds, it is sufficient to conclude (1). The independence of the trials imply

$$
\begin{aligned}
P[\forall j \leq m : \neg A_j \mid \neg B_m \wedge (\exists k < m : B_k)] &= \frac{P[(\forall j \leq m : \neg A_j) \wedge \neg B_m \wedge (\exists k < m : B_k)]}{P[\neg B_m \wedge (\exists k < m : B_k)]} \\
&= \frac{P[\neg A_m \wedge \neg B_m] \cdot P[(\forall j < m : \neg A_j) \wedge (\exists k < m : B_k)]}{P[\neg B_m] \cdot P[\exists k < m : B_k]} \\
&= P[\neg A_m \mid \neg B_m] \cdot P[\forall j < m : \neg A_j \mid \exists k < m : B_k].
\end{aligned}
$$

By induction, this is at most $P[\neg A_m \mid \neg B_m] \cdot P[\forall j < m : \neg A_j \mid B_{m-1}]$. The probabilities of the outcomes do not change over the trials, and also event B_m implies $\neg A_m$. Therefore,

$$
\begin{aligned}
P[\neg A_m \mid \neg B_m] \cdot P[\forall j < m : \neg A_j | B_{m-1}] &= \frac{1 - P[A_m] - P[B_m]}{1 - P[B_m]} \frac{(1 - P[A_m])^{m-2} \cdot P[B_{m-1}]}{P[B_{m-1}]} \\
&= \left(1 - \frac{P[A_m]}{1 - P[B_m]}\right) \cdot \left(1 - P[A_m]\right)^{m-2} \leq \left(1 - P[A_m]\right)^{m-1} = P[\forall j \leq m : \neg A_j \mid B_m]. \quad \square
\end{aligned}
$$

4   The Chernoff–Hoeffding Theorem

Most of the concentration bounds on the binomial distribution that are used in many different fields of science can be traced back to the Chernoff–Hoeffding theorem [29, 38]. It employs the Kullback–Leibler divergence to bound the probability that a random variable $X \sim \text{Bin}(n, p)$ deviates from its expectation $E[X] = pn$. The theorem states for all x with $0 \leq x \leq p$ that

$$
P[X \leq xn] \leq 2^{-D(x \,\|\, p)\, n} = \left(\frac{p}{x}\right)^{xn} \left(\frac{1-p}{1-x}\right)^{(1-x)n}.
$$

It follows easily that also $P[X \geq xn] \leq 2^{-D(x \,\|\, p)\, n}$ is true for all $p \leq x \leq 1$. Several weaker but more practical inequalities have been derived from the Chernoff–Hoeffding theorem, colloquially summarized as Chernoff bounds [28, 51]. Cramér's theorem asserts that the exponent $D(x \,\|\, p)$ is asymptotically tight [23, 27]. Any improvement of this inequality can therefore be at most subexponential in n. Stirling's approximation, see [4], gives the following lower bound assuming that the product xn is an integer,

$$
P[X \leq xn] \geq \frac{1}{\sqrt{8\, n\, x(1-x)}} \cdot 2^{-D(x \,\|\, p)\, n}.
$$

There is an obvious gap of order $\sqrt{n}$ between the two estimates. This section is dedicated to closing this gap by proving Theorem 3. It sharpens both inequalities, making them tight up to constant factors. While this improvement will be a valuable tool in our study of hypergraphs, it has applications also in many other areas. The proof is split into two major parts. The first one treats those x for which the product xn is an integer, the second one extends this to the general case. The parts are further subdivided depending on the limiting behavior of $x = x(n)$, namely, whether it converges to p or not.

4.1   Integral Case

For the integral case, we first use a result by Klar [48] about the connection between the distribution function and the probability mass function (PMF) of the binomial distribution.

Lemma 11 (Proposition 1 (c) in [48])**. *Let n be a positive integers, $p \neq 0$ a probability, and $X \sim \text{Bin}(n, p)$ a binomial variable. For all non-negative integers $k \leq pn$, it holds that*

$$
1 \leq \frac{P[X \leq k]}{P[X = k]} \leq \frac{p(n+1-k)}{n+1-k-(n+1)(1-p)}.
$$

We combine this with the perplexity bound on the binomial coefficient in Lemma 4. The respective lower bounds in the next lemma where known before, see for example the textbook by Ash [4, Lemma 4.7.2], we reprove them here en passant. The upper bounds are novel and we will later use them in our proof of the integral case of Theorem 3.

Lemma 12. *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n,p)$ a binomial variable. Suppose $0 < x < 1$ is a rational such that xn is an integer.*

(i) If $x < p$, then

$$\frac{1}{\sqrt{8\,n\,x(1-x)}} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \mathrm{P}[X \leq xn] \leq \frac{p\sqrt{1-x}}{(p-x)\sqrt{\pi\,xn}} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}.$$

(ii) If $p < x$, then

$$\frac{1}{\sqrt{8\,n\,x(1-x)}} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \mathrm{P}[X \geq xn] \leq \frac{(1-p)\sqrt{x}}{(x-p)\sqrt{\pi\,(1-x)n}} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}.$$

Proof. Applying the first statement to the complementary variable $\overline{X} \sim \mathrm{Bin}(n, 1-p)$ implies the second statement since $\mathrm{P}[X \geq xn] = \mathrm{P}[\overline{X} \leq (1-x)n]$. Hereby, we use that the Kullback–Leibler divergence observes $\mathrm{D}(1-x\,\|\,1-p) = \mathrm{D}(x\,\|\,p)$. We are left to prove the first statement.

Lemma 4 gives the following error bounds on the PMF $\mathrm{P}[X = xn] = \binom{n}{xn} \cdot p^{xn}(1-p)^{(1-x)n}$.

$$\frac{1}{\sqrt{8\,n\,x(1-x)}} \leq \frac{\mathrm{P}[X = xn]}{2^{\mathrm{H}(x)n} \cdot p^{xn}(1-p)^{(1-x)n}} = \frac{\mathrm{P}[X = xn]}{2^{-\mathrm{D}(x\,\|\,p)\,n}} \leq \frac{1}{\sqrt{\pi\,n\,x(1-x)}}.$$

The lower bounds follows immediately from $\mathrm{P}[X \leq xn] \geq \mathrm{P}[X = xn]$.

For the upper bound, we use Lemma 11 at the integer position $k = xn$. Let $f_{n,xn}(p)$ denote the resulting bound on the ratio $\mathrm{P}[X \leq xn]/\mathrm{P}[X = xn]$, that is,

$$f_{n,xn}(p) = \frac{p(n+1-xn)}{n+1-xn-(n+1)(1-p)} = \frac{p(n+1-xn)}{p(n+1)-xn} = \frac{p\left(1 - \frac{xn}{n+1}\right)}{p - \frac{xn}{n+1}}.$$

We claim that for all x and p with $x < p$, the function $f_{n,xn}(p)$ is increasing in n. We show this by verifying that the (partial) discrete derivative $\Delta_n(f_{n,xn})$ with respect to n is positive.

$$\Delta_n(f_{n,xn})(p) = f_{n+1,x(n+1)}(p) - f_{n,xn}(p) = \frac{p(n+2-x(n+1))}{p(n+2)-x(n+1)} - \frac{p(n+1-xn)}{p(n+1)-xn}$$

$$= \frac{p(1-p)x}{((p-x)n+p)\cdot((p-x)n+2p-x)} > 0.$$

The function $f_{n,xn}(p)$ thus converges from below to $p(1-x)/(p-x)$ as n increases, giving an upper bound on $\mathrm{P}[X \leq xn]/\mathrm{P}[X = xn]$ for all n. Multiplying with the error bounds and the divergence completes the proof. $\square$

The upper bounds above are already very close to the desired ones of Theorem 3. In fact, we will see that the lemma is enough to conclude $\mathrm{P}[X \leq xn] \leq \varphi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$ if xn is an integer and $\varphi = \min(1, 1/((p-x)\sqrt{xn}))$. The lower bound in Statement (i), however, matches the upper one only if the $x = x(n)$ is bounded away from p for all n. More work is needed for the case $x \to p$. It has already been useful to have a good estimate for the ratio $\mathrm{P}[X \leq k]/\mathrm{P}[X = k]$. Unfortunately, Lemma 11 gives only a trivial lower bound. We strengthen this in the next lemma, Lemma 14 then shows how this translates into a stronger lower bound on the binomial distribution function. Finally, Lemma 15 combines all results of this section into a version of the Chernoff–Hoeffding theorem, which is tight whenever the product xn is an integer.

Lemma 13. *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n,p)$ a binomial variable. Then, for all non-negative integers i and k with $i \leq k \leq pn$, it holds that*

$$\frac{\mathrm{P}[X \leq k]}{\mathrm{P}[X = k]} \geq (k-i+1)\left(1 - \frac{pn-i}{pn(1-\frac{k}{n})}\right)^{k-i}.$$

Proof. The PMF of X is increasing for arguments smaller than pn, therefore

$$\frac{\mathrm{P}[X \le k]}{\mathrm{P}[X = k]} = \sum_{j=0}^{k} \frac{\mathrm{P}[X = j]}{\mathrm{P}[X = k]} \ge \sum_{j=i}^{k} \frac{\mathrm{P}[X = j]}{\mathrm{P}[X = k]} \ge (k - i + 1) \frac{\mathrm{P}[X = i]}{\mathrm{P}[X = k]}.$$

The last ratio is lower-bounded by

$$\frac{\mathrm{P}[X = i]}{\mathrm{P}[X = k]} = \frac{k!\,(n-k)!}{i!\,(n-i)!} \left(\frac{1-p}{p}\right)^{k-i} = \prod_{\ell=1}^{k-i} \frac{i + \ell}{n - i - \ell + 1} \cdot \left(\frac{1-p}{p}\right)^{k-i} \ge \left(\frac{i}{n-i} \frac{1-p}{p}\right)^{k-i}.$$

For the base of the last expression we get

$$\frac{i}{n-i} \frac{1-p}{p} = \frac{pn - pi - pn + i}{pn - pi} = 1 - \frac{pn - i}{pn - pi} = 1 - \frac{pn - i}{np(1 - \frac{i}{n})} \ge 1 - \frac{pn - i}{np(1 - \frac{k}{n})}. \qquad \square$$

The first factor $k - i + 1$ increases as i gets smaller while at the same time the second factor decreases. Therefore, in order to apply Lemma 13, one has to choose a balancing cut-off point. We do so in the proof of the following lemma.

Lemma 14. *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n, p)$ a binomial variable. Suppose $0 < x < 1$ is a rational such that xn is an integer.*

(i) If $x < p$, then

$$\mathrm{P}[X \le xn] \ge \frac{p\sqrt{1-x}}{16\sqrt{2}} \cdot \min\left(1, \frac{1}{(p-x)\sqrt{xn}}\right) \cdot 2^{-\mathrm{D}(x \,\|\, p)\, n}.$$

(ii) If $p < x$, then

$$\mathrm{P}[X \ge xn] \ge \frac{(1-p)\sqrt{x}}{16\sqrt{2}} \cdot \min\left(1, \frac{1}{(x-p)\sqrt{(1-x)n}}\right) \cdot 2^{-\mathrm{D}(x \,\|\, p)\, n}.$$

Proof. The second statement follows from the first as in Lemma 12. Define an auxiliary integer function $g = g(n, p, x)$ as

$$g(n, p, x) = \left\lfloor \frac{p(1-x)}{2} \cdot \min\left(\sqrt{xn}, \frac{1}{p-x}\right) \right\rfloor.$$

Applying Lemma 13 at position $k = xn$ with the cut-off point $i = xn - g$ gives

$$\mathrm{P}[X \le xn] \ge (g+1)\left(1 - \frac{pn - xn + g}{pn(1-x)}\right)^{g} \cdot \mathrm{P}[X = xn]. \tag{2}$$

We want to lower-bound the middle factor in (2) by a constant. Bernoulli's inequality gives

$$\left(1 - \frac{pn - xn + g}{pn(1-x)}\right)^{g} = \left(1 - \frac{p - x + \frac{g}{n}}{p(1-x)}\right)^{g} \ge 1 - \frac{g(p-x) + \frac{g^2}{n}}{p(1-x)}.$$

We claim that the numerator $g(p - x) + g^2/n$ is at most $3p(1 - x)/4$. We split the argument depending on the relative size of $\sqrt{xn}$ and $1/(p - x)$. If $\sqrt{xn} \ge 1/(p - x)$, then we have $g = \lfloor p(1-x)/2(p-x) \rfloor$ and thus

$$g(p-x) + \frac{g^2}{n} \le \frac{p(1-x)}{2} + \frac{p^2(1-x)^2}{4} \cdot \frac{1}{(p-x)^2 n} \le \frac{p(1-x)}{2} + \frac{p^2(1-x)^2}{4} \cdot x.$$

Conversely, if $\sqrt{xn} \leq 1/(p-x)$, then $g = \lfloor p(1-x)\sqrt{xn}/2 \rfloor$ and

$$g(p-x) + \frac{g^2}{n} \leq \frac{p(1-x)}{2} \cdot (p-x)\sqrt{xn} + \frac{p^2(1-x)^2}{4} \cdot \frac{xn}{n} \leq \frac{p(1-x)}{2} + \frac{p^2(1-x)^2}{4} \cdot x.$$

The last expressions of both inequalities are the same and can be bounded by $3p(1-x)/4$. The middle factor is therefore at least a constant since

$$1 - \frac{g(p-x) + \frac{g^2}{n}}{p(1-x)} \geq 1 - \frac{3p(1-x)}{4} \frac{1}{p(1-x)} = \frac{1}{4}.$$

Reinserting this into Inequality (2) and applying the definition of g and Lemma 4 gives the result.

$$\mathrm{P}[X \leq xn] \geq \frac{g+1}{4} \cdot \mathrm{P}[X = xn] \geq \frac{p(1-x)}{8} \cdot \min\left(\sqrt{xn}, \frac{1}{p-x}\right) \cdot \mathrm{P}[X = xn]$$

$$\geq \frac{p(1-x)}{8} \cdot \min\left(\sqrt{xn}, \frac{1}{p-x}\right) \cdot \frac{1}{\sqrt{8nx(1-x)}} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$$

$$= \frac{p\sqrt{1-x}}{16\sqrt{2}} \cdot \min\left(1, \frac{1}{(p-x)\sqrt{xn}}\right) \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}. \qquad \square$$

Next, we prove Theorem 3 for the case that xn is an integer by combining the results above. We emphasize the facts that Lemma 15 holds for all positive integers n, not only asymptotically, and x may range over the whole interval $[0,1]$.

Lemma 15 (integral case of Theorem 3). *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n,p)$ a binomial variable. Suppose $x = x(n)$ takes rational values in the unit interval such that xn is an integer. Let φ and ψ denote the functions*

$$\varphi(n,p,x) = \min\left(1, \frac{1}{(p-x)\sqrt{xn}}\right) \quad and \quad \psi(n,p,x) = \min\left(1, \frac{1}{(x-p)\sqrt{(1-x)n}}\right),$$

with additionally $\varphi(n,p,0) = \varphi(n,p,p) = 1$ and $\psi(n,p,1) = \psi(n,p,p) = 1$.

(i) If $x \leq p$, then $\frac{p\sqrt{1-p}}{16\sqrt{2}} \cdot \varphi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \mathrm{P}[X \leq xn] \leq \varphi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$.

(ii) If $x \geq p$, then $\frac{(1-p)\sqrt{p}}{16\sqrt{2}} \cdot \psi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \mathrm{P}[X \geq xn] \leq \psi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$.

Proof. Statement (ii) follows from (i) in the usual way since $\psi(n,p,x) = \varphi(n,1-p,1-x)$. Let $C = p\sqrt{1-p}/16\sqrt{2}$. Note that C is at most 0.045 for any p. We first discuss the corner cases $x = 0$ and $x = p$ (assuming that pn is an integer). If $x = 0$, then we have $\mathrm{P}[X \leq 0 \cdot n] = (1-p)^n = \varphi(n,p,0) \cdot 2^{-\mathrm{D}(0\,\|\,p)\,n}$. If $x = p$, the upper bound $\mathrm{P}[X \leq pn] \leq 1 = \varphi(n,p,p) \cdot 2^{-\mathrm{D}(p\,\|\,p)\,n}$ holds vacuously. The lower bound follows from pn being the median of the binomial distribution, which implies $\mathrm{P}[X \leq pn] \geq 1/2 \geq C = C \cdot \varphi(n,p,p) \cdot 2^{-\mathrm{D}(p\,\|\,p)\,n}$.

Assume $0 < x < p$. The original Chernoff–Hoeffding theorem and Lemma 12 together give

$$\mathrm{P}[X \leq xn] \leq \min\left(1, \frac{p\sqrt{1-x}}{(p-x)\sqrt{\pi xn}}\right) \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \min\left(1, \frac{p}{\sqrt{\pi}} \frac{1}{(p-x)\sqrt{xn}}\right) \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}.$$

The latter is at most $\varphi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$. Finally, the lower bound in this case is an easy consequence of Lemma 14 and $p\sqrt{1-x}/16\sqrt{2}$ being larger than $C = p\sqrt{1-p}/16\sqrt{2}$. $\qquad \square$

4.2 General Case

The second major step of the argument is to extend the result above from integral products xn to arbitrary real x. The equality $\mathrm{P}[X \leq xn] = \mathrm{P}[X \leq \lfloor xn \rfloor]$ holds universally as X assumes only integer values. In Section 4.1, we have given bounds on the second probability $\mathrm{P}[X \leq \lfloor xn \rfloor]$ in terms of the ratio $\lfloor xn \rfloor/n$. To reach the generality of the Chernoff–Hoeffding theorem, we need

to infer bounds on $P[X \leq xn]$ in terms of x. In what follows, let x' abbreviate $\lfloor xn \rfloor / n$. Consider the upper bound in Lemma 15 (i) as an illustrating example. It states that

$$P[X \leq xn] = P[X \leq x'n] \leq \min\left(1, \frac{1}{(p-x')\sqrt{x'n}}\right) \cdot 2^{-D(x' \| p) n}.$$

If there exists some constant C, possibly dependent on p but independent of x and n, such that

$$\frac{1}{(p-x')\sqrt{x'n}} \cdot 2^{-D(x' \| p) n} \leq \frac{C}{(p-x)\sqrt{xn}} \cdot 2^{-D(x \| p) n},$$

then our estimate transfers to the general case. A similar reasoning applies to the other cases.

The next two lemmas prepare the necessary technical machinery to show the existence of those constants. Lemma 16 clarifies the monotonicity of the functions in question. It shows that transitioning from x' to x can only increase the upper bound, meaning that we can actually choose $C = 1$ in the above illustration. For the opposite direction, Lemma 17 asserts that this transitions incurs a multiplicative loss that is at most linear in x. Below, we often conclude the monotonic behavior of a product of functions from that of its factors. While in general the product of non-decreasing functions is not itself non-decreasing, the monotonicity transfers if all factors are additionally non-negative.

Lemma 16. *Let n be a positive integer, and p and x two probabilities. The function*

$$g_{n,p}(x) = \frac{2^{-D(x \| p) n}}{(p-x)\sqrt{xn}}$$

is non-decreasing for all x such that $1/n \leq x < p$, provided that n is sufficiently large.

Proof. Quantity $2^{-D(x \| p) n}$ is non-decreasing for $x \leq p$ (Lemma 6) and it is not hard to prove this also for $1/(p-x)\sqrt{xn}$ given that $x \geq p/3$. The main focus of this proof is to show that the divergence power dominates the monotonicity of $g_{n,p}$ also for $1/n \leq x \leq p/3$.

Taking derivatives gives

$$\frac{d}{dx} g_{n,p}(x) = \frac{1}{(p-x)^2 x \sqrt{n}} \left(\left(\frac{\partial}{\partial x} 2^{-D(x \| p) n}\right)(p-x)\sqrt{x} - 2^{-D(x \| p) n}\left(\frac{\partial}{\partial x}(p-x)\sqrt{x}\right)\right)$$

$$= \frac{1}{(p-x)^2 x \sqrt{n}} \left(n \ln\left(\frac{1-x}{x}\frac{p}{1-p}\right) 2^{-D(x \| p) n}(p-x)\sqrt{x} - 2^{-D(x \| p) n} \frac{p-3x}{2\sqrt{x}}\right)$$

$$= \frac{2^{-D(x \| p) n}}{(p-x) x^{3/2} \sqrt{n}} \left(n \ln\left(\frac{1-x}{x}\frac{p}{1-p}\right) x - \frac{p-3x}{2(p-x)}\right).$$

The first factor is positive for all n and $x < p$ and the same is true for the second one if $p/3 < x$. Assume $x \leq p/3$ in the remainder. Then, the last term of the second factor, $-(p-3x)/2(p-x)$, is at least $-1/2$. It is thus sufficient to prove the non-negativity of

$$h_{n,p}(x) = n \ln\left(\frac{1-x}{x}\frac{p}{1-p}\right) x - \frac{1}{2}$$

on the subinterval $[1/n, p/3]$ for all n large enough. We do this in two claims. First, $h_{n,p}$ is concave there and, secondly, its values at the endpoints of the interval are non-negative. Regarding the concavity, observe that the derivative

$$\frac{d}{dx} h_{n,p}(x) = n \ln\left(\frac{1-x}{x}\frac{p}{1-p}\right) - \frac{n}{1-x}$$

is the sum of two non-increasing functions of x; $n \ln(\frac{1-x}{x}\frac{p}{1-p})$ is non-increasing as it is the derivative of the concave mapping $x \mapsto -D(x \| p) n$. At the endpoint $1/n$, we have

$$h_{n,p}\left(\frac{1}{n}\right) = \ln\left((n-1)\frac{p}{1-p}\right) - \frac{1}{2},$$

which is non-negative for all $n \geq (\sqrt{e}\,(1-p)/p) + 1$. Similarly at endpoint $p/3$,

$$h_{n,p}\left(\frac{p}{3}\right) = n \ln\left(\frac{3-p}{1-p}\right)\frac{p}{3} - \frac{1}{2}$$

is non-negative for $n \geq 3/2p \ln(\frac{3-p}{1-p})$. $\qquad\square$

Lemma 17. *Let n be a positive integer, p and x two probabilities, $x' = \lfloor xn \rfloor/n$, and again*

$$g_{n,p}(x) = \frac{2^{-\mathrm{D}(x\,\|\,p)\,n}}{(p-x)\sqrt{xn}}.$$

The following inequalities hold for all x such that $1/n \leq x < p$ and n sufficiently large.

(i) $g_{n,p}(x') \geq \frac{1-p}{p\sqrt{2}}\,e^{-\frac{1}{2-2p}} \cdot x \cdot g_{n,p}(x)$;

(ii) $2^{-\mathrm{D}(x'\,\|\,p)\,n} \geq \frac{1-p}{p}\,e^{-\frac{1}{2-2p}} \cdot x \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$.

Proof. We make heavy use of the facts that $x - 1/n < x' \leq x$, and that $x \geq 1/n$ implies $x' \geq 1/n$. The relative difference between $g_{n,p}(x')$ and $g_{n,p}(x)$ is

$$\frac{g_{n,p}(x')}{g_{n,p}(x)} = \frac{2^{-\mathrm{D}(x'\,\|\,p)\,n}}{2^{-\mathrm{D}(x\,\|\,p)\,n}} \cdot \frac{p-x}{p-x'} \cdot \sqrt{\frac{x}{x'}}.$$

The last factor is at least 1 and the middle one is $\frac{p-x}{p-x'} = 1 - \frac{x-x'}{p-x'} \geq 1 - \frac{1}{pn} \geq \frac{1}{\sqrt{2}}$ given that $n \geq 2/(2-\sqrt{2})p$. The first factor can be estimated using Lemma 6 and $x' - x \leq 1/n$ as

$$\frac{2^{-\mathrm{D}(x'\,\|\,p)\,n}}{2^{-\mathrm{D}(x\,\|\,p)\,n}} = \left(\frac{x}{1-x}\,\frac{1-p}{p}\right)^{(x-x')n} 2^{-\mathrm{D}(x'\,\|\,x)\,n} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$$

$$\geq \frac{x}{1-x}\,\frac{1-p}{p}\,2^{-\mathrm{D}(x'\,\|\,x)\,n} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \geq x \cdot \frac{1-p}{p}\,2^{-\mathrm{D}(x'\,\|\,x)\,n} \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}.$$

It remains to show that $2^{-\mathrm{D}(x'\,\|\,x)\,n}$ is at least a constant, namely, we claim $2^{-\mathrm{D}(x'\,\|\,x)\,n} > e^{-\frac{1}{2-2p}}$. Let $t^- = \arg\min_{t \in [x',x]} t(1-t)$, observe that $1/n \leq t^- < p$ holds. By Lemma 5, the exponent $\mathrm{D}(x'\,\|\,x)\,n$ (to the base $1/2$) is bounded.

$$\mathrm{D}(x'\,\|\,x)\,n \leq \frac{(x-x')^2}{t^-(1-t^-)\,2\ln 2} \cdot n < \frac{\frac{1}{n^2}}{\frac{1}{n}(1-t^-)\,2\ln 2} \cdot n = \frac{1}{(1-t^-)\,2\ln 2} < \frac{1}{(2-2p)\ln 2}. \qquad\square$$

We have the tools ready to prove Theorem 3 in its entirety.

Theorem 3 (restated with explicit constants). *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n,p)$ a binomial variable. Suppose the function $x = x(n)$ takes real values in the interval $[\varepsilon, 1-\varepsilon]$ for some $\varepsilon > 0$. Let φ and ψ denote the functions*

$$\varphi(n,p,x) = \min\left(1, \frac{1}{(p-x)\sqrt{xn}}\right) \quad and \quad \psi(n,p,x) = \min\left(1, \frac{1}{(x-p)\sqrt{(1-x)n}}\right),$$

with additionally $\varphi(n,p,p) = 1$ and $\psi(n,p,p) = 1$. The following statements hold for all n sufficiently large.

(i) *If $x \leq p$, then $\frac{\varepsilon\,(1-p)^{\frac{3}{2}}}{32}\,e^{-\frac{1}{2-2p}} \cdot \varphi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \mathrm{P}[X \leq xn] \leq \varphi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$.*

(ii) *If $x \geq p$, then $\frac{\varepsilon p^{\frac{3}{2}}}{32}\,e^{-\frac{1}{2p}} \cdot \psi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n} \leq \mathrm{P}[X \geq xn] \leq \psi \cdot 2^{-\mathrm{D}(x\,\|\,p)\,n}$.*

Proof. We only need to prove the first statement. Let $x' = \lfloor xn \rfloor / n$. This implies $x' \leq x$ and makes $x'n$ an integer such that $\mathrm{P}[X \leq x'n] = \mathrm{P}[X \leq xn]$. Recall the definition of $g_{n,p}$ from Lemma 16. It is chosen such that for all n, p, and x, we have

$$\varphi(n, p, x') \cdot 2^{-\mathrm{D}(x' \,\|\, p)\, n} = \begin{cases} 2^{-\mathrm{D}(x' \,\|\, p)\, n}, & \text{if } 1 \leq \frac{1}{(p-x')\sqrt{x'n}} \text{ or } x' = p; \\ g_{n,p}(x'), & \text{otherwise.} \end{cases}$$

Lemmas 6 and 16 together establish $\varphi(n, p, x') \cdot 2^{-\mathrm{D}(x' \,\|\, p)\, n} \leq \varphi(n, p, x) \cdot 2^{-\mathrm{D}(x \,\|\, p)\, n}$ in both cases, provided that n is large enough. The upper bound in Statement (i) now follows from the integral case in Lemma 15. Regarding the lower bound, Lemma 17 gives for all n large enough,

$$\varphi(n, p, x') \cdot 2^{-\mathrm{D}(x' \,\|\, p)\, n} \geq \begin{cases} \frac{1-p}{p} e^{-\frac{1}{2-2p}} \cdot x \cdot 2^{-\mathrm{D}(x \,\|\, p)\, n}, & \text{if } 1 \leq \frac{1}{(p-x')\sqrt{x'n}} \text{ or } x' = p; \\ \frac{1-p}{p\sqrt{2}} e^{-\frac{1}{2-2p}} \cdot x \cdot g_{n,p}(x), & \text{otherwise.} \end{cases}$$

In summary, using Lemma 15 and the assumption $x \geq \varepsilon$, we have

$$\mathrm{P}[X \leq xn] = \mathrm{P}[X \leq x'n] \geq \frac{p\sqrt{1-p}}{16\sqrt{2}} \cdot \varphi(n, p, x') \cdot 2^{-\mathrm{D}(x' \,\|\, p)\, n}$$

$$\geq \frac{\varepsilon(1-p)^{\frac{3}{2}}}{32} e^{-\frac{1}{2-2p}} \cdot \varphi(n, p, x) \cdot 2^{-\mathrm{D}(x \,\|\, p)\, n}. \qquad \square$$

4.3 Applications

In this excursive section, we highlight two applications of Theorem 3 that do not fall into our main objective of studying hypergraphs, we find them instructive nevertheless. First, we give an alternative proof for the intuition that the probability of a binomial variable taking values outside of a few standard variation of its expectation may be small but does not converge to 0. The result was previously obtained via the Berry–Esseen inequality using the normal approximation of the binomial distribution [57]. Secondly, we improve the rate of convergence in Cramér's theorem.

Anti-Concentration Inequalities. Chernoff bounds are usually interpreted as concentration inequalities [29]. They state that the mass of the binomial distribution is concentrated around its mean and the probability for any other value falls exponentially in the distance to the expectation. However, it is also known that the probability of a binomial variable $X \sim \mathrm{Bin}(n, p)$ taking values outside of a constant number of standard deviations from $\mathrm{E}[X]$ does not vanish, even as n grows large. Results of the latter kind are occasionally called anti-concentration inequalities. The particular bound[4] we are interested in was given by Oliveto and Witt [57]. We give an alternative proof that avoids the normal distribution.

Corollary 18 (Lemma 6.1 in [57])**.** *Let n be a positive integer, $0 < p < 1$ a non-trivial probability, and $X \sim \mathrm{Bin}(n, p)$ a binomial variable. Let $\sigma_X = \sqrt{np(1-p)}$ denote the standard deviation of X. For every non-negative real $c \geq 0$, there exists some positive $C > 0$, independent of n but possibly dependent on c and p, such that*

$$\mathrm{P}[X \leq \mathrm{E}[X] - c \cdot \sigma_X] \geq C.$$

Conversely, for any non-negative function $f = f(n)$ with $\lim_{n\to\infty} f(n) = \infty$, it holds that

$$\lim_{n\to\infty} \mathrm{P}[X \leq \mathrm{E}[X] - f \cdot \sigma_X] = 0.$$

Proof. Let $c' = c\sqrt{p(1-p)}$. In the notation of the previous sections, we have

$$x = \frac{\mathrm{E}[X] - c \cdot \sigma_X}{n} = p - \frac{c'}{\sqrt{n}}.$$

---

[4] Lemma 6.1 in [57] only states the existence of a particular pair of constants c, C. It is easy to check that their proof remains valid for any $c \geq 0$. For a generalization to unequal probabilities, see Lemma 1.10.16 in [28].

We assume $x \geq p/2$, this does not loose generality as x converges to p. Applying Theorem 3 gives

$$\mathrm{P}[X \leq \mathrm{E}[X] - c \cdot \sigma_X] \geq \frac{p(1-p)^{\frac{3}{2}}}{64} e^{-\frac{1}{2-2p}} \cdot \min\left(1, \frac{1}{c'\sqrt{x}}\right) \cdot 2^{-\mathrm{D}(x \,\|\, p)\, n}.$$

The minimum is at least $1/cp\sqrt{1-p}$. To show that also the last factor is bounded, we use an argument very similar to the one in Lemma 17. Let $t^- = \arg\min_{t \in [x,p]} t(1-t)$. We have $p/2 \leq t^- \leq p$, resulting in

$$\mathrm{D}(x \,\|\, p) \cdot n \leq \frac{\left(\frac{c'}{\sqrt{n}}\right)^2}{t^-(1-t^-)\, 2\ln 2} \cdot n = \frac{c^2\, p(1-p)}{t^-(1-t^-)\, 2\ln 2} \leq \frac{c^2\, p(1-p)}{\frac{p}{2}(1-p)\, 2\ln 2} = \frac{c^2}{\ln 2}.$$

If we move away from $\mathrm{E}[X]$ by more than a constant number of standard deviations, already the original Chernoff–Hoeffding theorem is sharp enough to prove that the distribution function vanishes. For some non-negative $f(n) = \omega(1)$, define $x = p - (f\sqrt{p(1-p)}/\sqrt{n})$. If x is negative, the statement holds vacuously; otherwise, we have $\mathrm{P}[X \leq \mathrm{E}[X] - f \cdot \sigma_X] \leq 2^{-\mathrm{D}(x \,\|\, p)\, n}$. Let $t^+ = \arg\max_{t \in [x,p]} t(1-t)$. Lemma 5 gives

$$\mathrm{D}(x \,\|\, p) \cdot n \geq \frac{f(n)^2\, p(1-p)}{t^+(1-t^+)\, 2\ln 2} = \omega(1). \qquad \square$$

Cramér's Theorem. Theorem 3 also has implications for the rate of convergence in Cramér's theorem in the theory of large deviations. Fix some non-trivial probability p and let $(X_i)_i$ be a sequence of i.i.d. Bernoulli variables[5] with success probability p. Define

$$\Lambda^*(x) = \sup_{t \in \mathbb{R}} \left\{ tx - \ln(\mathrm{E}[e^{tX_1}]) \right\}$$

to be the Legendre transform of the cumulant-generating function of X_1. Cramér's theorem [23, 27] states that the transform observes the following limiting property for all x with $p < x < 1$,

$$\lim_{n \to \infty} \frac{1}{n} \ln\left(\mathrm{P}\left[\sum_{i=1}^{n} X_i \geq xn\right] \right) = -\Lambda^*(x).$$

For notational convenience, let $\mathrm{D}_e(x \,\|\, y) = \ln(2) \cdot \mathrm{D}(x \,\|\, y) = -x \ln\left(\frac{y}{x}\right) - (1-x)\ln\left(\frac{1-y}{1-x}\right)$ denote the natural (base-e) Kullback–Leibler divergence. It is straightforward to verify from $\mathrm{E}[e^{tX_1}] = 1 - p + pe^t$ that Cramér's function $\Lambda^*(x) = \mathrm{D}_e(x \,\|\, p)$ is in fact the natural divergence. This shows that the original Chernoff–Hoeffding inequality with $2^{-\mathrm{D}(x \,\|\, p)\, n} = e^{-\mathrm{D}_e(x \,\|\, p)\, n}$ is asymptotically tight up to sublinear terms in the exponent. However, the rate of convergence of the above limit is subject of ongoing research [25, 27, 31]. Observe that under the current assumptions, we have $1 \geq 1/(x-p)\sqrt{(1-x)n}$ for all n sufficiently large. Theorem 3 thus implies the following corollary, which further clarifies the rate in the Bernoulli case.

Corollary 19. *Let $0 < p < 1$ be a non-trivial probability and $(X_i)_i$ a sequence of i.i.d. Bernoulli variables with parameter p. Then, for any x with $p < x < 1$, it holds that*

$$\frac{1}{n} \ln\left(\mathrm{P}\left[\sum_{i=1}^{n} X_i \geq xn\right] \right) = -\mathrm{D}_e(x \,\|\, p) - \frac{1}{2}\frac{\ln n}{n} - \frac{\ln(x-p)}{n} - \frac{1}{2}\frac{\ln(1-x)}{n} \pm \mathrm{O}\left(\frac{1}{n}\right).$$

---

[5]Cramér's theorem holds more generally for any i.i.d. sequence $(X_i)_i$ such that the cumulant-generating function of X_1 is finite everywhere [27].

5   Distinct Sets and Minimality

We return to the main topic of this work, which is determining the expected number of minimal edges in the maximum-entropy multi-hypergraph $\mathcal{B}_{n,m,p}$. The vertex sampling probabilities $p = 0$ or $p = 1$ result in trivial hypergraphs, we thus assume $0 < p < 1$ in the remainder of this work, unless explicitly stated otherwise. Also n, m always denote positive integers and $X \sim \mathrm{Bin}(n, p)$ a binomial variable. Under this assumptions, every subset of $[n]$ then has a non-vanishing chance to be sampled. Such a set is minimal for $\mathcal{B}_{n,m,p}$ if and only if it is generated in one of the trials and no proper subset ever occurs. Both of these aspects influence the chance of minimality, but their impact varies depending on the cardinality of the sets in question.

The number of vertices per edge is heavily concentrated around pn and the more vertices there are in an edge, the less likely it is minimal. Intuitively speaking, almost no sets with very low cardinalities are sampled, but if so, they are very likely included in the minimization. There are plenty of edges with a medium number of vertices and there is still a good chance they are minimal. Sets of high cardinality rarely occur and usually they are dominated by smaller edges. This disparity is exacerbated by a large number of trials. Boosting m increases the probability that also sets of cardinality a bit further away from pn are sampled, at the same time the process now generates more duplicate sets that do not count towards the minimization. More importantly though, the likelihood of a larger set being minimal is smaller with many trials. Eventually, the last effect outweighs all others, creating the situation in which the only minimal edge is empty.

In this section, we make this intuition rigorous. We start by giving preliminary bounds on the number of minimal edges as a first step towards the proof of Theorem 1. These bounds have the form of are binomial sums of polynomials of probabilities, depending on which factors we include in those sums, we get an upper or a lower bound. The estimates are already tight up to constants but are rather unwieldy. They will serve as the basis for our further analysis.

Let $\mathcal{D}_{n,p}$ denote the maximum-entropy distribution on the power set $\mathcal{P}([n])$ provided that $\mathrm{E}_{S \sim \mathcal{D}_{n,p}}[|S|] = pn$, each vertex is included independently with probability p. Let $S_j \sim \mathcal{D}_{n,p}$ be the outcome of the j-th trial. Define $s_{n,p}(i, m) = \mathrm{P}[\exists j \leq m \colon S_j = [i]]$ to be the probability that the set $[i]$ (in fact, any set of cardinality i) is sampled, and $w_{n,p}(i, m) = \mathrm{P}[\forall j \leq m \colon \neg(S_j \subsetneq [i])]$ as the probability that no trial ever samples a proper subset of $[i]$.[6]

Lemma 20. *We have $s_{n,p}(i, m) = 1 - (1 - p^i(1-p)^{n-i})^m$ and $w_{n,p}(i, m) = (1 - (1-p)^{n-i}(1-p^i))^m$. The following statements hold for the minimization of multi-hypergraph $\mathcal{B}_{n,m,p}$.*

  (i) $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \geq \sum_{i=0}^{n} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i, m)$.

  (ii) $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \leq \sum_{i=0}^{n} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i, m - 1)$.

  (iii) $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \leq 1 + \frac{1}{p} \sum_{i=0}^{n} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i, m)$.

Proof. To see the closed forms of $s_{n,p}(i, m)$ and $w_{n,p}(i, m)$, observe that the random set $S_j \sim \mathcal{D}_{n,p}$ differs from $[i]$ with probability $1 - p^i(1-p)^{n-i}$. The independence of the trials gives $s_{n,p} = 1 - (1 - p^i(1-p)^{n-i})^m$. Similarly, S_j is a subset of $[i]$ if it does not contain an element of $[n] \backslash [i]$, having probability $(1-p)^{n-i}$, Conditioned on being any subset, S_j is a *proper* subset if it is missing at least one element of $[i]$. The expression for $w_{n,p}(i, m)$ follows from here.

Regarding the main statements, some fixed set $S \subseteq [n]$ is in $\min(\mathcal{B}_{n,m,p})$ if and only if it is sampled in one of the m trials and no proper subset is sampled. The probability for both events depends only on $|S|$ as all sets with the same cardinality are equally likely.

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \sum_{S \subseteq [n]} \mathrm{P}[(\exists k \leq m \colon S_k = S) \wedge (\forall j \leq m \colon \neg(S_j \subsetneq S))]$$

$$= \sum_{i=0}^{n} \binom{n}{i} \mathrm{P}[\exists k \leq m \colon S_k = [i]] \cdot \mathrm{P}[\forall j \leq m \colon \neg(S_j \subsetneq [i]) \mid \exists k \leq m \colon S_k = [i]]$$

$$= \sum_{i=0}^{n} \binom{n}{i} s_{n,p}(i, m) \cdot \mathrm{P}[\forall j \leq m \colon \neg(S_j \subsetneq [i]) \mid \exists k \leq m \colon S_k = [i]].$$

---

[6]The notation $s_{n,p}$ refers to the set being *sampled*, these probabilities are then *weighted* by the factors $w_{n,p}$.

The last factor describes the likelihood that any set with i elements is minimal, conditioned on it being sampled at all. The stated bounds differ only in the way this factor is estimated. We claim it to be at least $\mathrm{P}[\forall j \leq m\colon \neg(S_j \subsetneq [i])]$ (that is, without the condition) while at the same time being at most $\mathrm{P}[\forall j < m\colon \neg(S_j \subsetneq [i])]$ (with one fewer trial). The first inequality is obvious because conditioning on some trial producing $[i]$ itself only increases the chances of never sampling a proper subset. For the second one, we apply Lemma 10 to the events $A_j = [S_j \subsetneq [i]]$ and $B_j = [S_j = [i]]$, which shows that

$$\mathrm{P}[\forall j \leq m\colon \neg(S_j \subsetneq [i]) \mid \exists k \leq m\colon S_k = [i]] \leq \mathrm{P}[\forall j \leq m\colon \neg(S_j \subsetneq [i]) \mid S_m = [i]].$$

The proof of the claim, and thereby the one of Statements (i) and (ii), is completed by the fact that $\mathrm{P}[\forall j \leq m\colon \neg(S_j \subsetneq [i]) \mid S_m = [i]]$ is equal to $\mathrm{P}[\forall j < m\colon \neg(S_j \subsetneq [i])]$.

In order to prove Statement (iii), note that the i-th term of the sums in the first two statements have a relative difference of $w_{n,p}(i,m)/w_{n,p}(i,m{-}1)$. By independence, this is equal to $w_{n,p}(i,1)$, the probability to not sample a strict subset of $[i]$ in a single trial. For $i < n$, it is easy to see that $w_{n,p}(i,1) \geq p$. If $i = n$, we have $w_{n,p}(n,1) = p^n$ which is sub-constant. However, the statement follows anyway as the contribution of the last term to the whole sum is at most 1. $\quad\square$

The part that all three bounds of Lemma 20 have in common describes the expected number of *distinct* sets in $\mathcal{B}_{n,m,p}$. Recall that we use $\|\mathcal{H}\|$ to denote the number of distinct sets of some multi-hypergraph $\mathcal{H}$. That means, we have

$$\mathrm{E}[\|\mathcal{B}_{n,m,p}\|] = \sum_{i=0}^{n} \binom{n}{i} s_{n,p}(i,m) = \sum_{i=0}^{n} \binom{n}{i} \Big(1 - (1 - p^i(1-p)^{n-i})^m \Big).$$

We weight the terms of the sum by $w_{n,p}(i,m)$ or $w_{n,p}(i,m{-}1)$ to estimate the size of the minimization. We analyze the two parts separately, starting with the *weighting factors* $w_{n,p}$.

The behavior of the $w_{n,p}$ may have applications besides our study of random hypergraphs. Consider m trials according to the maximum-entropy distribution $\mathcal{D}_{n,p}$ on subsets of $[n]$. Then, $w_{n,p}(i,m)$ is by definition the probability that any fixed subset of cardinality i survives as minimal after m trials, equivalently, any proper subset is sampled with probability $1 - w_{n,p}(i,m)$. It is easy to see that weighting factors are non-increasing in both i and m. We prove next that the weighting factors are in fact threshold functions falling abruptly from almost 1 to almost 0 as i increases from 0 to n, the position of the transition depends on n, m, and p. Recall that $\alpha = -(\log_{1-p} m)/n$. Lemma 21 below establishes a sharp threshold behavior of $w_{n,p}(i,m)$ at

$$i^* = n + \log_{1-p} m = (1-\alpha)n.$$

Note that i^* is always at most n since $\log_{1-p} m$ is non-positive. The definition is such that it ensures the equality $m = 1/(1-p)^{n-i^*} = 1/(1-p)^{\alpha n}$. For increasing m, the threshold gets smaller relative to n. Once m grows beyond $1/(1-p)^n$, i.e., $\alpha > 1$, the quantity i^* can no longer be interpreted as a cardinality as it becomes negative. Later, in Lemma 25, we will see that m being this large is in fact irrelevant for the analysis of the minimization.

Lemma 21. *It holds that $w_{n,p}(0,m) = 1$ and $w_{n,p}(n,m) = p^{nm}$. Suppose $i = i(n)$ takes integer values such that $0 < i < n$, then*

(i) $\exp(-m(1-p)^{n-i}) \cdot (1 - m(1-p)^{2(n-i)}) \leq w_{n,p}(i,m) \leq \exp(-m(1-p)^{n-i+1})$.

In particular, the following statements hold.

(ii) *If $i = n + \log_{1-p} m + \omega(1)$, then $\lim_{n \to \infty} w_{n,p}(i,m) = 0$.*

(iii) *If $i = n + \log_{1-p} m - \omega(1)$, then $\lim_{n \to \infty} w_{n,p}(i,m) = 1$.*

(iv) *If $i = n + \log_{1-p} m \pm \Theta(1)$, then $w_{n,p}(i,m) = \Theta(1)$.*

Proof. The corner cases are elementary. Assume $0 < i < n$ for the rest of the proof. We estimate $w_{n,p}(i,m)$ using mainly Corollary 8. This yields

$$w_{n,p}(i,m) = (1 - (1-p)^{n-i}(1-p^i))^m \leq (1 - (1-p)^{n-i}(1-p))^m \leq \exp\Big(-m(1-p)^{n-i} \cdot (1-p)\Big).$$

The limiting behavior is determined by the product $m(1-p)^{n-i}$. If $i = n + \log_{1-p} m + \omega(1)$, then $m(1-p)^{n-i} = m(1-p)^{-(\log_{1-p} m) - \omega(1)} = (1-p)^{-\omega(1)}$ diverges and the weighting factor $w_{n,p}(i,m)$ converges to 0. Conversely, we get from $1 - p^i \leq 1$ that

$$w_{n,p}(i,m) \geq (1 - (1-p)^{n-i})^m \geq \exp\Big(-m(1-p)^{n-i}\Big) \cdot (1 - m(1-p)^{2(n-i)}).$$

If $i = n + \log_{1-p} m - \omega(1)$, both $m(1-p)^{n-i} = (1-p)^{\omega(1)}$ and $m(1-p)^{2(n-i)} = (1-p)^{\omega(1)}/m$ tend to 0, implying $\lim_{n \to \infty} w_{n,p}(i,m) = 1$.

Finally, if the cardinality i is close to the threshold $i^* = n + \log_{1-p} m$, the limit may not exist. We show that $w_{n,p}(i,m)$ is still bounded away from 0 for all n. Suppose $i = n + \log_{1-p} m \pm \Theta(1)$, in particular, the difference $i^* - i$ is bounded. If m is constant with respect to n, so is $w_{n,p}(i,m) \geq (1 - (1-p)^{n-i})^m \geq p^m$. Here, we used the assumption $i < n$. If m diverges, then $n - i = \log_{1-p} m \mp \Theta(1) = \omega(1)$ diverges with it and thus

$$w_{n,p}(i,m) \geq \exp\Big(-m(1-p)^{n-i}\Big)(1 - m(1-p)^{2(n-i)})$$

$$= \exp\Big(-(1-p)^{i^*-i}\Big) \cdot (1 - (1-p)^{(i^*-i)+(n-i)}) = \Omega(1). \qquad \square$$

After demonstrating a sharp threshold for the weighting factors, we turn to the number of distinct sets $\|\mathcal{B}_{n,m,p}\|$. A trivial cap for distinct sets is the total number of edges m. When starting the sampling, many different sets are generated and $\|\mathcal{B}_{n,m,p}\|$ is close to m. As the number of trials increases though, duplicates occur in the sample and the two quantities grow apart. To discuss this in more detail, we introduce some notation. For a pair of integers ℓ, u with $0 \leq \ell \leq u \leq n$, let $\|\mathcal{B}_{n,m,p}(\ell,u)\|$ denote the number of distinct samples whose cardinality is between ℓ and u, including. This is also at most as large as the total number of samples in that range. It thus makes sense to expect an upper bound in terms of the binomial distribution. We confirm this below and further prove that there is also a lower bound of the same flavor.

Lemma 22. *Let ℓ, u be integers such that $0 \leq \ell \leq u \leq n$ and $\mathbf{p} = \max_{\ell \leq i \leq u} \{p^i(1-p)^{n-i}\}$, then*

$$\mathbf{p} = \begin{cases} p^\ell(1-p)^{n-\ell}, & \text{if } p < 1/2; \\ 1/2^n, & \text{if } p = 1/2; \\ p^u(1-p)^{n-u}, & \text{otherwise.} \end{cases}$$

The expected number of distinct sets in $\mathcal{B}_{n,m,p}$ with cardinality between ℓ and u observes

$$\frac{m}{1+m\mathbf{p}} \cdot \mathrm{P}[\ell \leq X \leq u] \leq \mathrm{E}[\|\mathcal{B}_{n,m,p}(\ell,u)\|] \leq m \cdot \mathrm{P}[\ell \leq X \leq u].$$

Proof. The closed form for $\mathbf{p}$ can be seen from the equality $p^i(1-p)^{n-i} = (p/(1-p))^i \cdot (1-p)^n$ since the odds $p/(1-p)$ are strictly smaller than 1 iff $p < 1/2$. For the number of distinct sets in $\mathcal{B}_{n,m,p}(\ell, u)$, Lemma 9 implies that

$$\mathrm{E}[\|\mathcal{B}_{n,m,p}(\ell, u)\|] = \sum_{i=\ell}^{u} \binom{n}{i}(1 - (1 - p^i(1-p)^{n-i})^m) \leq m \cdot \sum_{i=\ell}^{u} \binom{n}{i} p^i(1-p)^{n-i}.$$

Conversely, we have

$$\mathrm{E}[\|\mathcal{B}_{n,m,p}(\ell, u)\|] \geq \sum_{i=\ell}^{u} \binom{n}{i} \frac{m p^i(1-p)^{n-i}}{1 + m p^i(1-p)^{n-i}} \geq \frac{m}{1 + m\mathbf{p}} \cdot \sum_{i=\ell}^{u} \binom{n}{i} p^i(1-p)^{n-i}. \qquad \square$$

6 Size of the Minimization

We now prove the main theorems with the help of the tools above. The key observation is that the minimization is dominated by the sets with cardinalities around $i^* = n + \log_{1-p} m = (1-\alpha)n$.

6.1 Binomial Characterization

Theorem 1 establishes a close connection between the expected size $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|]$ of the minimization and the binomial distribution. We split its proof into three lemmas corresponding, in that order, to the lower bound in the first statement of the theorem, the tight upper bound, and the second statement.

Note that the following lemma is slightly more general than what is claimed in Theorem 1 (*i*). It pertains to all $m \leq 1/(1-p)^n$, equivalently, it does not require α to be bounded away from 1.

Lemma 23 (Lower bound of Theorem 1 (*i*)). *For all $m \leq 1/(1-p)^n$, i.e., all $0 \leq \alpha \leq 1$, it holds that $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Omega(m) \cdot \mathrm{P}[X \leq i^*]$.*

Proof. The sought expectation is at least as large as the number of distinct minimal sets up to some cardinality i, for an arbitrary choice of $i \leq n$. As an ansatz, we set it equal to the threshold $i^* = n + \log_{1-p} m = (1-\alpha)n$. Without loosing generality, i^* is an integer; otherwise, we take $\lfloor i^* \rfloor$. Note that the assumption $m \leq 1/(1-p)^n$ ensures $i^* \geq 0$. Let $\mathbf{p} = \max_{0 \leq i \leq i^*}\{p^i(1-p)^{n-i}\}$. Lemmas 20 (*i*) and 22 imply

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \geq \sum_{i=0}^{i^*} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i, m) \geq \sum_{i=0}^{i^*} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i^*, m)$$

$$= \mathrm{E}[\|\mathcal{B}_{n,m,p}(0, i^*)\|] \cdot w_{n,p}(i^*, m) \geq \frac{1}{1 + m\mathbf{p}} \cdot w_{n,p}(i^*, m) \cdot m\,\mathrm{P}[X \leq i^*].$$

To complete the proof, we verify that the first two factors do not vanish. The first factor is immediate from the closed form for $\mathbf{p}$ (Lemma 22) since $m\mathbf{p} = \max_{0 \leq i \leq i^*}\{mp^i(1-p)^{n-i}\} \leq \max\{1, p^{i^*}\} = 1$. Lemma 21 (*i*) shows that there exists a universal constant $\delta > 0$ for all $m \leq 1/(1-p)^n$ such that $w(i^*, m) \geq \delta$. $\qquad \square$

Lemma 24 (Upper bound of Theorem 1 (*i*)). *For any $\varepsilon > 0$ and all $m \leq 1/(1-p)^{(1-\varepsilon)n}$, i.e., all $0 \leq \alpha \leq 1 - \varepsilon$, $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \mathrm{O}(m) \cdot \mathrm{P}[X \leq i^*]$. The leading constant may depend on ε.*

Proof. We know that $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \leq \sum_{i=0}^{n} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i, m-1)$. We split the sum at the threshold $i^* = (1-\alpha)n$ and handle the two parts separately. The assumption on m is such that $i^* \geq \varepsilon n$. For the first part, Lemma 22 implies

$$\sum_{i=0}^{i^*} \binom{n}{i} s_{n,p}(i, m) \cdot w_{n,p}(i, m-1) \leq \mathrm{E}[\|\mathcal{B}_{n,m,p}(0, i^*)\|] \leq m \cdot \mathrm{P}[X \leq i^*].$$

For cardinalities larger than i^*, we can no longer ignore the influence of the weighting factors. We show that the whole second part of the sum is within constant factors of $m \cdot \mathrm{P}[X = i^*]$. Let

$\ell \leq n - i^*$ be a positive integer. Consider the $(i^*+\ell)$-th term. If $\ell = n - i^*$ this is the last one and contributes at most 1. Assume $\ell < n - i^*$. As in the proof of Lemma 22, we see that the term is upper-bounded by $m \cdot \mathrm{P}[X = i^* + \ell] \cdot w_{n,p}(i^*+\ell, m)$. Dividing by $m \cdot \mathrm{P}[X = i^*]$ gives

$$\frac{\binom{n}{i^*+\ell} s_{n,p}(i^* + \ell, m) \cdot w_{n,p}(i^* + \ell, m - 1)}{m \, \mathrm{P}[X = i^*]} \leq \frac{m \, \mathrm{P}[X = i^* + \ell]}{m \, \mathrm{P}[X = i^*]} \cdot w_{n,p}(i^* + \ell, m - 1)$$

$$= \frac{\binom{n}{i^*+\ell}}{\binom{n}{i^*}} \left(\frac{p}{1 - p}\right)^{\ell} w_{n,p}(i^* + \ell, m - 1).$$

The first factor is bounded for any fixed ℓ, using $\alpha \leq 1 - \varepsilon$.

$$\frac{\binom{n}{i^*+\ell}}{\binom{n}{i^*}} = \prod_{j=1}^{\ell} \frac{n - i^* - j + 1}{i^* + j} \leq \left(\frac{n - i^*}{i^*}\right)^{\ell} = \left(\frac{\alpha}{1 - \alpha}\right)^{\ell} \leq \frac{1}{\varepsilon^{\ell}}.$$

Applying Lemma 21 (i) and the same idea as in Lemma 20 (iii) to the weighting factor yields

$$w_{n,p}(i^* + \ell, m - 1) = \frac{w_{n,p}(i^* + \ell, m - 1)}{w_{n,p}(i^* + \ell, m)} \cdot w_{n,p}(i^* + \ell, m)$$

$$\leq \frac{1}{p} \exp(-m \, (1 - p)^{n - i^* - \ell + 1}) = \frac{1}{p} \exp(-(1 - p)^{-\ell + 1}).$$

Here, we used $\ell < n - i^*$, ensuring that the ratio between subsequent factors is a constant.

Define $a = p/\varepsilon(1 - p)$ and $b = 1/(1 - p)$. So far, we have established that the ratio between the term at $i^* + \ell$ and $m \cdot \mathrm{P}[X = i^*]$ is at most

$$r(\ell) = \frac{1}{p} \frac{a^{\ell}}{\exp(b^{\ell - 1})}.$$

The function r itself is free of any dependence on n, m, or α, but in order to prove our claim, we need to bound the sum $1 + \sum_{\ell=1}^{n - i^* - 1} r(\ell)$. We show that the series $\sum_{\ell=1}^{\infty} r(\ell)$ is in fact summable.

To this end, consider the sequence $t(\ell) = r(\ell) \cdot 2^{\ell}$. Its logarithm $\ln t = \ell \ln(2a) - \ln p - b^{\ell - 1}$ diverges to $-\infty$ as ℓ increases, implying $t \to 0$. This means, there exists an ℓ_0 such that $r(\ell) \leq 2^{-\ell}$ for all $\ell \geq \ell_0$.

$$\sum_{\ell=1}^{\infty} r(\ell) \leq \sum_{\ell=1}^{\ell_0} \frac{a^{\ell}}{p \cdot \exp(b^{\ell - 1})} + \sum_{\ell=\ell_0+1}^{\infty} \frac{1}{2^{\ell}} \leq \sum_{\ell=1}^{\ell_0} \frac{a^{\ell}}{p \cdot \exp(b^{\ell - 1})} + 2 = \mathrm{O}(1). \qquad \square$$

Finally, we show that once m is a polynomial factor larger than $1/(1 - p)^n$, the minimization essentially consists of only a single edge, the empty set.

Lemma 25 (Theorem 1 (ii))**.** *There exists a constant $c > 0$, possibly dependent on p, such that for all $m \geq 1/(1 - p)^{n + c \ln n}$, it holds that $1 \leq \mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = 1 + o(1)$.*

Proof. The lower bound is immediate. Suppose we have $m = 1/(1 - p)^{n + f}$ for some non-negative function $f = f(n)$. As soon as the empty set is sampled in one of the m trials, the minimization of $\mathcal{B}_{n,m,p}$ comprises only a single set; otherwise, we fall back to the trivial estimate $|\min(\mathcal{B}_{n,m,p})| \leq m$. Let A denote the event $[\emptyset \in \mathcal{B}_{n,m,p}]$. The law of total expectation together with Corollary 8 implies that

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \mathrm{E}[|\min(\mathcal{B}_{n,m,p})| \mid A] \cdot \mathrm{P}[A] + \mathrm{E}[|\min(\mathcal{B}_{n,m,p})| \mid \neg A] \cdot \mathrm{P}[\neg A]$$

$$\leq \mathrm{P}[A] + m \cdot (1 - (1 - p)^n)^m \leq 1 + \exp(\ln m - m \, (1 - p)^n)$$

$$= 1 + \exp\left(\ln m - (1 - p)^{-f}\right).$$

We have $\ln m = -\ln(1 - p)(n + f)$, where $-\ln(1 - p)$ is a constant strictly larger than 1. Requiring $f \geq c'(\ln n)/(-\ln(1 - p))$ for some arbitrary $c' > 1$ ensures that $\ln m$ is negligible compared to $(1 - p)^{-f}$ and the expression above converges to 1. Setting $c = -c'/\ln(1 - p)$ gives the lemma. $\square$

6.2 The Case $m = 1/(1-p)^n$

There is a mismatch in the upper and lower bounds above in the range of parameters for which they hold. The lower bound has been shown for the full range of $m = 1/(1-p)^{\alpha n} \leq 1/(1-p)^n$, this includes the case where the function $\alpha(n)$ converges to 1. We get from Lemma 23 that $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Omega(1)$ if $\alpha = 1$ (for all n sufficiently large). This is not very surprising as already the the definition of $\min(\mathcal{B}_{n,m,p})$ guarantees the existence of at least one minimal edge. The upper bound in Lemma 24 is more interesting. It only has been proven for $\alpha \leq 1 - \varepsilon$ for any $\varepsilon > 0$. If instead we were to insert, say, $\alpha = 1 - \mathrm{o}(1/n)$, which corresponds to $m = 1/(1-p)^{n-\mathrm{o}(1)}$, we would get $\mathrm{P}[X \leq (1-\alpha)n] = \mathrm{P}[X = 0] = (1-p)^n$ and the result would default to $\mathrm{O}(1)$. We refute this below. Contrarily, we prove a lower bound for $\alpha = 1$ that is stronger than the inherited one. This shows that the binomial characterization breaks down if α converges to 1.

Lemma 26. *For $m = 1/(1-p)^n$, i.e., $\alpha = 1$, it holds that $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Omega(n)$.*

Proof. We extend the idea of the proof of Lemma 25. If none of the m trials produces the empty set, then all distinct sampled singletons are minimal.

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \geq \mathrm{E}[\|\mathcal{B}_{n,m,p}(1,1)\|] \cdot \mathrm{P}[\emptyset \notin \mathcal{B}_{n,m,p}] = n \cdot s_{n,p}(1,m) \cdot (1 - s_{n,p}(0,m)).$$

We show that the product $s_{n,p}(1,m)(1 - s_{n,p}(0,m))$ is bounded away from 0 for all n. From Corollary 8, Lemma 9, and the assumption $m = 1/(1-p)^n$, it follows that

$$s_{n,p}(1,m) = 1 - (1 - p(1-p)^{n-1})^m \geq \frac{mp(1-p)^{n-1}}{1 + mp(1-p)^{n-1}} = p.$$

For the second factor, we have

$$1 - s_{n,p}(0,m) = (1 - (1-p)^n)^m \geq \exp(-m(1-p)^n) \cdot (1 - m(1-p)^{2n}) = \frac{1 - (1-p)^n}{e}. \qquad \square$$

Lemmas 25 and 26 together show that a slight polynomial increase (in n) of the number of trials beyond $1/(1-p)^n$ is enough to push the size of the minimization from at least linear to 1. We leave it as an open problem to give exact bounds for the transitional behavior of $|\min(\mathcal{B}_{n,m,p})|$ around $m = 1/(1-p)^n$, i.e., $\alpha = 1$.

6.3 Phase Transition at $m = 1/(1-p)^{(1-p)n}$

We show that the size of the minimization undergoes a phase transition at $m^* = 1/(1-p)^{(1-p)n}$. This is made explicit in Theorem 2 (restated below) and illustrated in Figure 1. Intuitively, for a small number of edges, the ratio of minimal edges among all edges is constant and thus $|\min(\mathcal{B}_{n,m,p})|$ scales linearly with m. At the transition point, both the size of the multi-hypergraph as well as its minimization are of order $m^* = 1/(1-p)^{(1-p)n} = 2^{(\mathrm{H}(1-p)+p\,\mathrm{ld}\,p)\,n}$. This overlap is indicated in Figure 1b by dashed lines. Beyond that, in the information-theoretic regime, the size of the minimization is instead given by the perplexity of α, it follows $2^{(\mathrm{H}(\alpha)+(1-\alpha)\,\mathrm{ld}\,p)\,n}$ up to polynomial factors. To gauge the transition more accurately, we reuse the correction term $\varphi = \varphi(n, p, 1-\alpha)$ from Section 4. Close to the transition, for $\alpha \approx 1-p$, we have $\varphi \approx 1$. As α moves further away, φ is shrinking and once there is a constant additive gap between α and $1-p$ it is of order $\varphi = \Theta(1/\sqrt{n})$. In summary, our results for the information-theoretic regime show that the growth of $|\min(\mathcal{B}_{n,m,p})|$ continues at first after the transition, but is now only sublinear in m. The minimization peaks at $m = 1/(1-p)^{\frac{n}{1+p}}$ (see Lemma 27) and for larger m, the number of minimal edges is even falling exponentially, although the number of trials further increases. Once m exceeds $1/(1-p)^n$, the minimization collapses under the sheer likelihood of the empty set being sampled.

Theorem 2 (restated). *Let n, m be positive integers and $0 < p < 1$ a non-trivial probability.*

(i) If $m \leq 1/(1-p)^{(1-p)n}$, then $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(m)$.

(ii) *For any $\varepsilon > 0$ and all $1/(1-p)^{(1-p)n} \leq m \leq 1/(1-p)^{(1-\varepsilon)n}$, i.e., all $1 - p \leq \alpha \leq 1 - \varepsilon$,*

$$\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(1) \cdot \min\left(1, \frac{1}{(\alpha - (1-p))\sqrt{(1-\alpha)n}}\right) \cdot 2^{(\mathrm{H}(\alpha) + (1-\alpha)\log_2 p)\,n}$$

$$= \Theta(1) \cdot \min\left(1, \frac{1}{(\alpha - (1-p))\sqrt{(1-\alpha)n}}\right) \cdot \left(\frac{p^{1-\alpha}}{(1-\alpha)^{1-\alpha}\,\alpha^\alpha}\right)^n.$$

Proof. The first statement covers the linear regime of $m \leq 1/(1-p)^{(1-p)n}$. By the binomial characterization in Theorem 1, we have $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(m) \cdot \mathrm{P}[X \leq i^*]$, with $X \sim \mathrm{Bin}(n,p)$. It is thus enough to verify that $\mathrm{P}[X \leq i^*] = \mathrm{P}[X \leq (1-\alpha)n]$ does not converge to 0. This follows easily from $\alpha \leq 1 - p$ and pn being the median of the binomial distribution.

In the remainder, we treat the information-theoretic regime of all m such that $1/(1-p)^{(1-p)} \leq m \leq 1/(1-p)^{(1-\varepsilon)n}$ for some fixed $\varepsilon > 0$. In particular, $(1-\alpha)n$ is now smaller than $\mathrm{E}[X]$. Recall the definition of function φ from Section 4, inserting $1 - \alpha$ for x gives

$$\varphi = \varphi(n, p, 1 - \alpha) = \min\left(1, \frac{1}{(\alpha - (1-p))\sqrt{(1-\alpha)n}}\right).$$

Our improved Chernoff–Hoeffding bound (Theorem 3) implies $\mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta(m) \cdot \varphi \cdot 2^{-\mathrm{D}(1-\alpha \,\|\, p)\,n}$. Expressing the (the power of) the divergence in terms of the perplexity, and using $\mathrm{H}(1-\alpha) = \mathrm{H}(\alpha)$ as well as $m = 1/(1-p)^{\alpha n}$, shows that this is equal to

$$\Theta(1) \cdot \varphi \cdot m\, 2^{-\mathrm{D}(1-\alpha \,\|\, p)\,n} = \Theta(1) \cdot \varphi \cdot \frac{1}{(1-p)^{\alpha n}}\, 2^{\mathrm{H}(1-\alpha)n}\, p^{(1-\alpha)n}\,(1-p)^{\alpha n}$$

$$= \Theta(1) \cdot \varphi \cdot 2^{(\mathrm{H}(\alpha) + (1-\alpha)\,\mathrm{ld}\,p)\,n}. \qquad \square$$

The behavior of the minimization for increasing m suggests that there is a sweet spot in the information-theoretic regime where the expected number of minimal edges is maximum. We apply Theorem 2 to calculate this maximum.

Lemma 27. *The maximum expected number of minimal edges is $\max_{m \geq 1}\ \mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] = \Theta((1+p)^n/\sqrt{n})$, it is attained for $m = 1/(1-p)^{\frac{n}{1+p}}$.*

Proof. We first verify that that the maximum indeed sits in the information-theoretic regime. Observe that $1/(1-p)^{1-p} < 1 + p$ holds for all non-trivial probabilities p. This can be seen from $1/(1-p)^{1-p}$ being strictly concave on the open unit interval and $1 + p$ being its tangent line at $p = 0$. Theorem 2 (i) shows that the sample sizes in the linear regime are too small to lead to the claimed bound of $\Theta^*((1+p)^n)$. Moreover, $2^{\mathrm{H}(\alpha) + (1-\alpha)\,\mathrm{ld}\,p}$ as a function of α is continuous and concave, it converges to $2^{\mathrm{H}(1-p) + p\,\mathrm{ld}\,p} = 1/(1-p)^{(1-p)}$ as $\alpha \searrow 1 - p$ (for any p). Hence, there exists some $\delta > 0$ small enough such that $2^{\mathrm{H}(1-p+\delta) + (p-\delta)\,\mathrm{ld}\,p} < 1 + p$. Similarly, the function converges to 1 from above as $\alpha \nearrow 1$. Let $\varepsilon > 0$ be such that $2^{\mathrm{H}(1-\varepsilon) + \varepsilon\,\mathrm{ld}\,p} < 1 + p$. In other words, any bound exponential in $1 + p$ must come from an $\alpha \in [1 - p + \delta, 1 - \varepsilon]$, we are in the setting of Theorem 2 (ii).

There are suitable constants $C, C' > 0$ such that

$$C \cdot \varphi \cdot 2^{(\mathrm{H}(\alpha) + (1-\alpha)\,\mathrm{ld}\,p)\,n} \leq \mathrm{E}[|\min(\mathcal{B}_{n,m,p})|] \leq C' \cdot \varphi \cdot 2^{(\mathrm{H}(\alpha) + (1-\alpha)\,\mathrm{ld}\,p)\,n}.$$

As α is bounded away from both $1 - p$ and 1, we have $\varphi = \Theta(1/\sqrt{n})$ and the hidden constants are independent of α (but possibly depending on δ and ε). We are thus left to determine the extremum of the exponential part. Let

$$g(\alpha, p) = \mathrm{H}(\alpha) + (1-\alpha)\,\mathrm{ld}\,p = -\alpha\,\mathrm{ld}\,\alpha - (1-\alpha)\,\mathrm{ld}(1-\alpha) + (1-\alpha)\,\mathrm{ld}\,p$$

be the exponent to base 2^n of the expression above. Its partial derivative

$$\frac{\partial}{\partial \alpha} g(\alpha, p) = \mathrm{ld}\left(\frac{1-\alpha}{\alpha}\right) - \mathrm{ld}\,p$$

has a single zero in the interval $[1 - p + \delta, 1 - \varepsilon]$ at $\alpha^* = 1/(1 + p)$, resulting in an exponent of

$$g(\alpha^*, p) = -\frac{1}{1+p} \operatorname{ld}\left(\frac{1}{1+p}\right) - \frac{p}{1+p} \operatorname{ld}\left(\frac{p}{1+p}\right) + \frac{p}{1+p} \operatorname{ld} p = \operatorname{ld}(1+p). \qquad \square$$

7 Conclusion

We calculated the number of minimal edges of maximum-entropy multi-hypergraphs with expected edge size pn and discovered a phase transition with respect to the total number of edges. As long as m is at most $m^* = 1/(1 - p)^{(1-p)n}$, the expected size of the minimization is linear in m. Once the sample size increases beyond that point, the minimization is instead governed by the entropy of α where $m = 1/(1 - p)^{\alpha n}$. The minimization continues to grow sublinearly until m reaches $1/(1 - p)^{\frac{n}{1+p}}$, from there on, it decays rapidly. Raising m above $1/(1 - p)^n$ finally ensures that only the empty set is minimal.

The Chernoff–Hoeffding theorem played an integral role in verifying these results. We tightened the tail bounds on the binomial distribution and provided explicit upper and lower bounds on the constants involved. We are convinced that this improvement can help other researchers in probability well beyond the scope of this paper.

A possible extension of our work would be to generalize Theorem 1 to the case where α converges to 1 as n increases. It is also interesting to investigate hypergraph models for which the sampling probability p is no longer constant, but tends to 0. This would bring the study of random hypergraphs closer to that of sparse random graphs. Finally, a promising direction in the light of our original motivation of relational databases is to allow different sample probabilities per vertex as well as dependencies between the elements. This requires the maximum-entropy model to incorporate additional constraints.

Acknowledgments

References

[1] Z. Abedjan, L. Golab, F. Naumann, and T. Papenbrock. *Data Profiling.* Synthesis Lectures on Data Management. Morgan & Claypool Publishers, San Rafael, CA, USA, 2018.

[2] K. Anand and G. Bianconi. Entropy Measures for Networks: Toward an Information Theory of Complex Topologies. *Physical Review E*, 80:045102, 2009.

[3] N. A. Arafat, D. Basu, L. Decreusefond, and S. Bressan. Construction and Random Generation of Hypergraphs with Prescribed Degree and Dimension Sequences. *CoRR*, abs/2004.05429, 2020. ArXiv preprint, to appear at DEXA 2020.

[4] R. B. Ash. *Information Theory.* Dover Books on Mathematics. Dover Publications, Mineola, NY, USA, 1990. Reprint of the Interscience Publishers 1965 edition.

[5] G. Badkobeh, P. K. Lehre, and D. Sudholt. Black-box Complexity of Parallel Search with Distributed Populations. In *Proceedings of the 2015 Conference on Foundations of Genetic Algorithms (FOGA)*, pp. 3–15, 2015.

[6] M. Ball, A. Rosen, M. Sabin, and P. N. Vasudevan. Average-case fine-grained hardness. In H. Hatami, P. McKenzie, and V. King, editors, *Proceedings of the 49th Annual Symposium on Theory of Computing, (STOC)*, pp. 483–496, 2017.

[7] M. Behrisch, A. Coja-Oghlan, and M. Kang. The Order of the Giant Component of Random Hypergraphs. *Random Structures and Algorithms*, 36:149–184, 2010.

[8] M. Behrisch, A. Coja-Oghlan, and M. Kang. Local Limit Theorems for the Giant Component of Random Hypergraphs. *Combinatorics, Probability and Computing*, 23:331—366, 2014.

[9] C. Berge. *Hypergraphs - Combinatorics of Finite Sets*, Vol. 45 of *North-Holland Mathematical Library*. North-Holland, Amsterdam, Netherlands, 1989.

[10] L. Bertossi. *Database Repairing and Consistent Query Answering*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, San Rafael, CA, USA, 2011.

[11] G. Bianconi. The Entropy of Randomized Network Ensembles. *Europhysics Letters*, 81: 28005, 2007.

[12] T. Bläsius, T. Friedrich, and M. Schirneck. The Parameterized Complexity of Dependency Detection in Relational Databases. In *Proceedings of the 11th International Symposium on Parameterized and Exact Computation (IPEC)*, pp. 6:1–6:13, 2016.

[13] T. Bläsius, T. Friedrich, and M. Schirneck. The Minimization of Random Hypergraphs. In *Proceedings of the 28th Annual European Symposium on Algorithms (ESA)*, pp. 21:1–21:15, 2020.

[14] T. Bleifuß, S. Kruse, and F. Naumann. Efficient Denial Constraint Discovery with Hydra. *Proceedings of the VLDB Endowment*, 11:311–323, 2017.

[15] T. Bläsius, T. Friedrich, J. Lischeid, K. Meeks, and M. Schirneck. Efficiently Enumerating Hitting Sets of Hypergraphs Arising in Data Profiling. In *Proceedings of the 21st Meeting on Algorithm Engineering and Experiments (ALENEX)*, pp. 130–143, 2019.

[16] B. Bollobás. *Random Graphs*. Studies in Advanced Mathematics. Cambridge University Press, Cambridge, UK, 2 edition, 2001.

[17] B. Bollobás. A Probabilistic Proof of an Asymptotic Formula for the Number of Labelled Regular Graphs. *European Journal of Combinatorics*, 1:311–316, 1980.

[18] M. Borassi, P. Crescenzi, and M. Habib. Into the Square: On the Complexity of Some Quadratic-time Solvable Problems. *Electronic Notes in Theoretical Computer Science*, 322: 51–67, 2016.

[19] A. A. Bruen and M. A. Forcinito. *Cryptography, Information Theory, and Error-Correction: A Handbook for the 21st Century*. Wiley-Interscience, New York, NY, USA, 2004.

[20] P. S. Chodrow. Configuration Models of Random Hypergraphs. *Journal of Complex Networks*, 8, 2020. Article number cnaa018.

[21] C. Cooper, A. M. Frieze, and W. Pegden. On the Rank of a Random Binary Matrix. *Electronic Journal of Combinatorics*, 26, 2019. Article number P4.12.

[22] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications and Signal Processing. Wiley-Interscience, New York, NY, USA, 2nd edition, 2006.

[23] H. Cramér. Sur un nouveau théorème-limite de la théorie des probabilités. In *Actualités scientifiques et industrielles*, Vol. 763, pp. 5–23, 1938. Colloque consacré à la théorie des probabilités. (On a New Limit Theorem in Probability.) In French.

[24] M. Dalirrooyfard, A. Lincoln, and V. Vassilevska Williams. New Techniques for Proving Fine-Grained Average-Case Hardness. *CoRR*, abs/2008.06591, 2020. ArXiv preprint, to appear at FOCS 2020.

[25] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*, Vol. 38 of *Stochastic Modelling and Applied Probability*. Springer, Berlin and Heidelberg, Germany, 2nd edition, 2010.

[26] J. Demetrovics, G. O. H. Katona, D. Miklos, O. Seleznjev, and B. Thalheim. Asymptotic Properties of Keys and Functional Dependencies in Random Databases. *Theoretical Computer Science*, 190:151–166, 1998.

[27] J.-D. Deuschel and D. W. Stroock. *Large Deviations*, Vol. 137 of *Pure and Applied Mathematics*. Academic Press, Boston, MA, USA, 1989.

[28] B. Doerr. Probabilistic Tools for the Analysis of Randomized Optimization Heuristics. In B. Doerr and F. Neumann, editors, *Theory of Evolutionary Computation*, chapter 1. Springer International, Basel, Switzerland, 2020.

[29] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 2009.

[30] P. Erdős and A. Rényi. On Random Graphs I. *Publicationes Mathematicae Debrecen*, 6: 290–297, 1959.

[31] J. A. Fill. Convergence Rates Related to the Strong Law of Large Numbers. *Annals of Probability*, 11:123–142, 1983.

[32] V. Froese, R. van Bevern, R. Niedermeier, and M. Sorge. Exploiting Hidden Structure in Selecting Dimensions That Distinguish Vectors. *Journal of Computer and System Sciences*, 82:521–535, 2016.

[33] J. Gao, R. Impagliazzo, A. Kolokolova, and R. R. Williams. Completeness for First-Order Properties on Sparse Structures With Algorithmic Applications. *Transactions on Algorithms*, 15:23:1–23:35, 2018.

[34] D. Garlaschelli and M. I. Loffredo. Maximum Likelihood: Extracting Unbiased Information from Complex Networks. *Physical Review E*, 78:015101, 2008.

[35] E. N. Gilbert. Random Graphs. *Annals of Mathematical Statistics*, 30:1141–1144, 1959.

[36] Y. Grandvalet and Y. Bengio. Entropy Regularization. In O. Chapelle, B. Schölkopf, and A. Zien, editors, *Semi-Supervised Learning*, chapter 9. MIT Press, Cambridge, MA, USA, 2006.

[37] P. Harremoës. Binomial and Poisson Distributions as Maximum Entropy Distributions. *Transactions on Information Theory*, 47:2039–2041, 2001.

[38] W. Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[39] R. v. d. Hofstad. *Random Graphs and Complex Networks*, Vol. 1 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, UK, 2016.

[40] E. T. Jaynes. Information Theory and Statistical Mechanics. *Phyical Review Series II*, 106: 620–630, 1957.

[41] E. T. Jaynes. Information Theory and Statistical Mechanics II. *Phyical Review Series II*, 108:171–190, 1957.

[42] D. M. Kane and R. R. Williams. The Orthogonal Vectors Conjecture for Branching Programs and Formulas. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, pp. 48:1–48:15, 2018.

[43] M. Karoński and T. Łuczak. The Phase Transition in a Random Hypergraph. *Journal of Computational and Applied Mathematics*, 142:125–135, 2002.

[44] G. O. H. Katona. Testing Functional Connection Between Two Random Variables. In A. N. Shiryaev, S. R. S. Varadhan, and E. L. Presman, editors, *Prokhorov and Contemporary Probability Theory*, pp. 335–348. Springer, Berlin and Heidelberg, Germany, 2013. Festschrift.

[45] G. O. H. Katona. Random Databases with Correlated Data. In A. Düsterhöft, M. Klettke, and K.-D. Schewe, editors, *Conceptual Modelling and Its Theoretical Foundations: Essays Dedicated to Bernhard Thalheim on the Occasion of His 60th Birthday*, pp. 29–35. Springer, Berlin and Heidelberg, Germany, 2012. Festschrift.

[46] M. J. Kearns and U. V. Vazirani. *An Introduction to Computational Learning Theory.* MIT Press, Cambridge, MA, USA, 1994.

[47] H. K. Kesavan. Jaynes' Maximum Entropy Principle. In C. A. Floudas and P. M. Pardalos, editors, *Encyclopedia of Optimization*, pp. 1779–1782. Springer, Boston, MA, USA, 2009.

[48] B. Klar. Bounds on Tail Probabilities of Discrete Distributions. *Probability in the Engineering and Informational Sciences*, 14:161–171, 2000.

[49] E. H. Lieb and J. Yngvason. The Physics and Mathematics of the Second Law of Thermodynamics. *Physics Reports*, 310:1–96, 1999.

[50] E. Livshits, A. Heidari, I. F. Ilyas, and B. Kimelfeld. Approximate Denial Constraints. *Proceedings of the VLDB Endowment*, 13:1682–1695, 2020.

[51] M. Mitzenmacher and E. Upfal. *Probability and Computing.* Cambridge University Press, New York, NY, USA, 2nd edition, 2017.

[52] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, Cambridge, UK, 1995.

[53] M. E. J. Newman. Scientific Collaboration Networks. I. Network Construction and Fundamental Results. *Physical Review E*, 64:016131, 2001.

[54] M. E. J. Newman. *Networks: An Introduction.* Oxford University Press, New York, NY, USA, 2010.

[55] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition.* Cambridge University Press, Cambridge, UK, 2010.

[56] R. O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, Cambridge, UK, 2014.

[57] P. S. Oliveto and C. Witt. Improved Time Complexity Analysis of the Simple Genetic Algorithm. *Theoretical Computer Science*, 605:21 – 41, 2015.

[58] T. Papenbrock, J. Ehrlich, J. Marten, T. Neubert, J.-P. Rudolph, M. Schönberg, J. Zwiener, and F. Naumann. Functional Dependency Discovery: An Experimental Evaluation of Seven Algorithms. *Proceedings of the VLDB Endowment*, 8:1082–1093, 2015.

[59] J. Park and M. E. J. Newman. The Statistical Mechanics of Networks. *Physical Review E*, 70:066117, 2004.

[60] Y. V. Prokhorov. Асимптотическое поведение биномиального распределения. Успехи математических наук, 8:135–142, 1953. (Asymptotic Behavior of the Binomial Distribution.) In Russian.

[61] F. Saracco, R. Di Clemente, A. Gabrielli, and T. Squartini. Randomizing Bipartite Networks: The Case of the World Trade Web. *Scientific Reports*, 5:10595, 2015.

[62] J. Schmidt-Pruzan and E. Shamir. Component Structure in the Evolution of Random Hypergraphs. *Combinatorica*, 5:81–94, 1985.

[63] C. E. Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27:379–423, 1948.

[64] E. V. Slud. Distribution Inequalities for the Binomial Law. *The Annals of Probability*, 5: 404–412, 1977.

[65] R. R. Williams. A New Algorithm for Optimal 2-Constraint Satisfaction and Its Implications. *Theoretical Computer Science*, 348:357–365, 2005.

[66] K. A. Zweig. *Network Analysis Literacy: A Practical Approach to the Analysis of Networks*. Lecture Notes in Social Networks. Springer, Vienna, Austria, 2014.