Down-set thresholds

Benjamin Gunby, Xiaoyu He, and Bhargav Narayanan

ABSTRACT. We elucidate the relationship between the threshold and the expectation-threshold of a down-set. Qualitatively, our main result demonstrates that there exist down-sets with polynomial gaps between their thresholds and expectation-thresholds; in particular, the logarithmic gap predictions of Kahn–Kalai and Talagrand (recently proved by Park–Pham and Frankston–Kahn–Narayanan–Park) about up-sets do not apply to down-sets. Quantitatively, we show that any collection $\mathcal{G}$ of graphs on $[n]$ that covers the family of all triangle-free graphs on $[n]$ satisfies the inequality $\sum_{G \in \mathcal{G}} \exp(-\delta e(G^c)/\sqrt{n}) < 1/2$ for some universal $\delta > 0$, and this is essentially best-possible.

1. INTRODUCTION

For a given finite set X, a family $\mathcal{F} \subset 2^X$ is called an *up-set* if it is closed under taking supersets, a *down-set* if it is closed under taking subsets, and *monotone* if it is either an up-set or a down-set.

Our main result demonstrates that there is no analogue for down-set thresholds of the conjectures of Kahn–Kalai [9] (resolved in [11]) and Talagrand [16] (resolved in [7]) regarding up-set thresholds. The following theorem, our main contribution, answers (in the negative) a question that arose in discussions between Kahn and the third author.

Theorem 1.1. *For the down-set $\mathcal{T}_n$ of triangle-free graphs on the vertex set $[n]$, the asymptotic growth rates of the expectation-threshold and the fractional expectation-threshold of $\mathcal{T}_n$ are both between $\sqrt{1/n}$ and $\sqrt{\log n/n}$, while the asymptotic growth rate of the threshold of $\mathcal{T}_n$ is $1/n$.*

The problem of locating thresholds has been a central concern in the study of random discrete structures since the seminal work of Erdős and Rényi [6] on threshold phenomena in random graphs. A great deal of work has since gone into locating thresholds of specific properties of interest; see [4, 8] and the many references therein, for example.

Expectation-thresholds were introduced in [9] as a comparatively easy (and rather general) way of locating thresholds of up-sets, so Theorem 1.1 comes as a bit of a surprise;

the arguments needed to control the expectation-threshold of the aforementioned down-set $\mathcal{T}_n$ are somewhat delicate, in stark contrast to what is needed to locate the threshold of $\mathcal{T}_n$. We shall explain exactly where Theorem 1.1 fits into the general theory of thresholds once we fill in some background, a task to which we now turn.

For a given finite set X and $p \in [0,1]$, we write μ_p for the product measure on the power set 2^X of X given by

$$\mu_p(S) = p^{|S|}(1-p)^{|X \setminus S|}$$

for all $S \subset X$. For a non-trivial (i.e., not 2^X or $\varnothing$) monotone family $\mathcal{F} \subset 2^X$, the *threshold* $p_c(\mathcal{F})$ of $\mathcal{F}$ is the unique p for which $\mu_p(\mathcal{F}) = 1/2$; this is well-defined since $\mu_p(\mathcal{F}) = \sum_{S \in \mathcal{F}} \mu_p(S)$ is strictly increasing in p when $\mathcal{F}$ is a non-trivial up-set, and strictly decreasing in p when $\mathcal{F}$ is a non-trivial down-set.

Following Talagrand [14, 15], we say an up-set $\mathcal{F} \subset 2^X$ is *p-small* if there is a certificate $\mathcal{G} \subset 2^X$ such that $\mathcal{F} \subset \mathcal{G}^\uparrow$, where $\mathcal{G}^\uparrow$ is the increasing family generated by $\mathcal{G}$ (i.e. the family of all sets containing elements of $\mathcal{G}$), and

$$\sum_{S \in \mathcal{G}} p^{|S|} \leq 1/2; \tag{1}$$

in other words, $\mathcal{F}$ is p-small if there is a simple 'first-moment' proof of the fact that $p_c(\mathcal{F}) \geq p$, namely

$$\mu_p(\mathcal{F}) \leq \mu_p(\mathcal{G}^\uparrow) \leq \sum_{S \in \mathcal{G}} p^{|S|} \leq 1/2.$$

Analogously, we say a down-set $\mathcal{F} \subset 2^X$ is *p-small* if there is a certificate $\mathcal{G} \subset 2^X$ such that $\mathcal{F} \subset \mathcal{G}^\downarrow$, where $\mathcal{G}^\downarrow$ is the decreasing family generated by $\mathcal{G}$, and

$$\sum_{S \in \mathcal{G}} (1-p)^{|X \setminus S|} \leq 1/2; \tag{2}$$

again, this means that there is a simple proof of the fact that $p_c(\mathcal{F}) \leq p$, namely

$$\mu_p(\mathcal{F}) \leq \mu_p(\mathcal{G}^\downarrow) \leq \sum_{S \in \mathcal{G}} (1-p)^{|X \setminus S|} \leq 1/2.$$

Now, we define the *expectation-threshold* $q(\mathcal{F})$ of a monotone $\mathcal{F}$ as follows: for an up-set $\mathcal{F}$, this is the *largest* value $p \in [0,1]$ for which $\mathcal{F}$ is p-small, and for a down-set $\mathcal{F}$, this is the *smallest* value $p \in [0,1]$ for which $\mathcal{F}$ is p-small.

Viewing the certificates $\mathcal{G}$ in (1) and (2) as integral maps from 2^X to $\{0,1\}$, the consideration of fractional relaxations of these certificates leads us to the *fractional expectation-threshold* $q_f(\mathcal{F})$ of a monotone $\mathcal{F}$, which is the optimal value $p \in [0,1]$ for which there is a fractional certificate witnessing either the fact that $p_c(\mathcal{F}) \geq p$ for an up-set $\mathcal{F}$, or the fact that $p_c(\mathcal{F}) \leq p$ for a down-set $\mathcal{F}$. Since we shall primarily focus on expectation-thresholds at this stage of the discussion, we defer a careful discussion of these fractional issues to Section 4. Nevertheless, we note that it follows immediately

from these definitions that for any up-set $\mathcal{F}$, we have $q(\mathcal{F}) \leq q_f(\mathcal{F}) \leq p_c(\mathcal{F})$, while for any down-set $\mathcal{F}$, we instead have $p_c(\mathcal{F}) \leq q_f(\mathcal{F}) \leq q(\mathcal{F})$.

This paper gets its motivation from [9], where Kahn and Kalai (drawing on a number of important examples in random graph theory) conjectured that $p_c(\mathcal{F}) = O(q(\mathcal{F}) \log |X|)$ for any up-set $\mathcal{F} \subset 2^X$, a conjecture recently proven by Park and Pham [11]. An earlier slight weakening of this result, originally conjectured by Talagrand [16], was proved in [7], where it was shown that $p_c(\mathcal{F}) = O(q_f(\mathcal{F}) \log |X|)$ for any up-set $\mathcal{F} \subset 2^X$. Since decreasing properties also appear quite frequently in the study of random discrete structures, it is natural to ask if there are analogues of the results of [11, 7] for down-sets. Concretely, our primary motivation is the following question that arose from discussions between Kahn and the third author.

Problem 1.2. *Is it true that $q(\mathcal{F})/p_c(\mathcal{F}) \leq (\log |X|)^{O(1)}$ for every down-set $\mathcal{F} \subset 2^X$?*

Of course, down-sets are just complements of up-sets, so one must ask if there is any new content in Problem 1.2, or if it is trivially resolved by the existing machinery for up-sets in [9, 11, 16, 7]. Problem 1.2 is indeed nontrivial, but this merits a short explanation. The results of [11, 7] are all only meaningful for 'large' up-sets $\mathcal{F}$ for which $q(\mathcal{F}) = o(1)$; indeed, if $\mathcal{F}$ is 'small' in the sense of $q(\mathcal{F}) = \Omega(1)$, then $1 \geq p_c(\mathcal{F}) \geq q(\mathcal{F}) = \Omega(1)$, so the Park-Pham theorem for such 'small' $\mathcal{F}$ holds no content. Nevertheless, the following is a natural analogous question for 'small' up-sets in the spirit of the Park-Pham theorem.

Problem 1.3. *Is it true that $q'/p' \leq (\log |X|)^{O(1)}$ for every up-set $\mathcal{F} \subset 2^X$, where $p_c(\mathcal{F}) = 1 - p'$ and $q(\mathcal{F}) = 1 - q'$?*

It is easy to check that Problems 1.2 and 1.3 are equivalent, as every down-set $\mathcal{F}$ can be turned into an up-set $\mathcal{F}' = \{X - A : A \in \mathcal{F}\}$ with $p(\mathcal{F}') = 1 - p(\mathcal{F})$ and $q(\mathcal{F}') = 1 - q(\mathcal{F})$. In this way, all our results on down-sets can also be reformulated in the complementary setting of 'small' up-sets. In turn, Theorem 1.1 demonstrates that the answer to Problem 1.2 (and hence Problem 1.3 as well) is in the negative: there exist down-sets with polynomial gaps between their thresholds and (fractional) expectation-thresholds.

We now turn to a discussion of the sharpness of Theorem 1.1, as well as its proof. For the down-set $\mathcal{T}_n$ of triangle-free graphs on the vertex set $[n]$, it is an easy exercise to show that $p_c(\mathcal{T}_n) = \Theta(1/n)$; see [2], for instance. Thus, the heart of the matter is to establish that $q(\mathcal{T}_n) = \Omega(1/\sqrt{n})$ (along with a similar bound for $q_f(\mathcal{T}_n)$). Therefore, writing $e(H)$ for the number of edges in a graph H, we restrict ourselves to proving the following.

Theorem 1.4. *There exists a universal $\delta > 0$ such that the following holds for all sufficiently large $n \in \mathbb{N}$. If $\mathcal{H}$ is a collection of graphs on $[n]$, and*

$$\sum_{H \in \mathcal{H}} \exp\big((-\delta e(H))/\sqrt{n}\big) < \frac{1}{2}, \tag{3}$$

then there is a triangle-free graph G on $[n]$ sharing at least one edge with each $H \in \mathcal{H}$.

A certificate $\mathcal{G}$ witnessing the expectation-threshold $q(\mathcal{T}_n)$ is just a collection of graphs on $[n]$ that together cover all the triangle-free graphs on $[n]$. Taking $\mathcal{H}$ to be the collection of complements of graphs in $\mathcal{G}$, we see that Theorem 1.4 is equivalent to the fact that $q(\mathcal{T}_n) = \Omega(1/\sqrt{n})$. As we shall see in Section 4, the proof of Theorem 1.4, with very minor alterations, also allows us to prove that $q_f(\mathcal{T}_n) = \Omega(1/\sqrt{n})$, but we focus on $q(\mathcal{T}_n)$ and Theorem 1.4 to keep the presentation simple.

Both Theorems 1.1 and 1.4 are best-possible up to logarithmic factors. As we shall see in Section 5, the machinery of hypergraph containers [3, 12] quickly demonstrates that $q(\mathcal{T}_n) = O(\log n/\sqrt{n})$, and a slightly more careful calculation using the Ramsey number upper bound of $r(3, k) = O(k^2/\log k)$ of Ajtai–Komlós–Szemerédi [1] shows that in fact $q(\mathcal{T}_n) = O(\sqrt{\log n/n})$. This in turn demonstrates that the conclusion of Theorem 1.4 need not hold if the $\sqrt{1/n}$ in the exponent in (3) is replaced with anything growing faster than $\sqrt{\log n/n}$.

We also remark that Theorem 1.4 immediately reproduces the Ramsey number lower bound of $r(3, k) = \Omega(k^2/\log^2 k)$; recall that $r(3, k)$ is the least integer n for which each triangle-free graph on n vertices has independence number at least k. The aforementioned lower bound was proved independently by Erdős [5] and Spencer [13] (but is of course superseded by Kim's [10] tour de force bound of $r(3, k) = \Omega(k^2/\log k)$, which gives the correct asymptotic growth rate). Indeed, for $C > 0$ large enough, taking $\mathcal{H}$ to be the collection of all cliques of size $C\sqrt{n}\log n$ on $[n]$, we see that

$$\sum_{H \in \mathcal{H}} \exp\big(-\delta e(H)n^{-1/2})\big) = \binom{n}{C\sqrt{n}\log n} \exp\big(-\delta C^2\sqrt{n}\log^2 n/2\big)$$

$$\leq \exp\big((C - \delta C^2/2)\sqrt{n}\log^2 n\big)$$

$$< \frac{1}{2},$$

so Theorem 1.4 implies the existence of a triangle-free graph G on the vertex set $[n]$ with no independent set of size $C\sqrt{n}\log n$.

The proof of Theorem 1.4 proceeds by generalising of Erdős' lower bound for $r(3, k)$. Erdős built triangle-free graphs with large independence number by starting with an Erdős–Rényi random graph $\Gamma = G(n, p)$ with $p \asymp n^{-1/2}$ and picking a maximal triangle-free subgraph $G \subset \Gamma$; he showed that with high probability, such a G has an edge in every set of vertices of size at least $\sqrt{n}\log n$. We shall show, more generally,

that for any choice of $\mathcal{H}$ satisfying (3), it is likely for a similarly constructed triangle-free G to share an edge with every $H \in \mathcal{H}$.

One of the technical tools that we need to execute the above strategy is a large-deviations estimate for the number of triangles completed by *any* fixed set of edges in the random graph. This large-deviations estimate, and the mechanism we use to 'preserve independence' in the proof of this estimate (which involves passing from random graphs to random directed graphs and back) might both be of independent interest; the details appear in Sections 2 and 3.

This paper is organised as follows. The proof of Theorem 1.4, modulo the large-deviations estimate mentioned above, is given in Section 2, and the beef follows in Section 3 where the requisite large-deviations estimate is established. The fractional analogue of our main result is stated and sketched in Section 4. We present the constructions demonstrating the sharpness of our results in Section 5. Finally, we close in Section 6 with a discussion of some open problems.

2. Proof of the main result

Writing $G(n, p)$ for the Erdős–Rényi random graph, the proof of Theorem 1.4 follows from a large-deviations estimate for the number of triangles in $G(n, p)$ 'closed' by the addition of *any* given set of m edges. To state this lemma, we need a little bit of notation: given a graph Γ, we write $e(\Gamma)$ for the number of edges of Γ, $\Delta(\Gamma)$ for the maximum degree of Γ, and Γ^2 for the square of Γ, which is the graph on the same vertex set as Γ whose edges are all pairs of vertices with a common neighbour in Γ. At the heart of our argument is the following fact.

Lemma 2.1. *There exist universal* $\gamma, \varepsilon > 0$ *such that the following holds for all sufficiently large* $m, n \in \mathbb{N}$. *Let* H *be a graph on* $[n]$ *with* m *edges. If* $\Gamma \sim G(n, p)$ *with* $p \le \varepsilon n^{-1/2}$, *then*

$$\mathbb{P}\Big(e(H \cap \Gamma^2) \ge 3m/4 \bigwedge \Delta(\Gamma) \le 2pn \Big) \le 15 \exp\big(-\gamma m n^{-1/2}\big).$$

We defer the proof of Lemma 2.1 to the next section, but briefly mention why the lemma demands that we work within the event $\{\Delta(\Gamma) < 2pn\}$; by considering the case where H is a star with $m = \Theta(n)$ edges, we see that $\mathbb{P}(e(H \cap \Gamma^2) \ge 3m/4) = p^{\Theta(n)}$, thus demonstrating that the bound in Lemma 2.1 does not hold uniformly in H for $\mathbb{P}(e(H \cap \Gamma^2) \ge 3m/4)$. We now describe how to deduce Theorem 1.4 from Lemma 2.1.

Proof of Theorem 1.4. Take $\gamma, \varepsilon > 0$ to be the constants in Lemma 2.1, and put $p = \varepsilon n^{-1/2}$ and $\Gamma = G(n, p)$. For $H \in \mathcal{H}$, we say that an edge $uv \in E(H)$ is *H-good* if $uv \notin E((\Gamma \setminus H)^2)$, i.e., uv is an edge of H that does not form any triangle uvw where $uw, vw \in E(\Gamma \setminus H)$.

5

We claim that if Γ contains an H-good edge, then every maximal triangle-free subgraph of Γ intersects H. Indeed, suppose for contradiction that Γ contains an H-good edge e, and that G is a maximal triangle-free subgraph of Γ with $G \cap H = \varnothing$. We see that $G \cup \{e\}$ is also triangle-free, since any triangle thereof must contain e, but this would violate the H-goodness of e. This would imply that $G \cup \{e\}$ is also a triangle-free subgraph of Γ, violating the maximality of G.

It remains to show that with positive probability, Γ contains an H-good edge for every $H \in \mathcal{H}$. This proves the theorem, as any maximal triangle-free subgraph G of Γ then demonstrates the desired result. To this end, we define three types of bad events, and show that they are all unlikely to happen. Let Z be the event that $\Delta(\Gamma) > 2pn$. For $H \in \mathcal{H}$, let Y_H be the event that $e(H \cap \Gamma^2) \geq 3e(H)/4$, and let X_H be the event that Γ contains no H-good edges.

Writing $\overline{E}$ for the complement of an event E, we have

$$\mathbb{P}\left(\bigvee_{H \in \mathcal{H}} X_H\right) \leq \mathbb{P}(Z) + \sum_{H \in \mathcal{H}} \mathbb{P}\left(Y_H \wedge \overline{Z}\right) + \sum_{H \in \mathcal{H}} \mathbb{P}\left(X_H \wedge \overline{Y_H}\right), \tag{4}$$

since either Z occurs, some Y_H occurs without Z, or else some X_H occurs without the corresponding Y_H. We now treat the three terms on the right hand side of (4) separately.

For the first, it is easily seen that $\mathbb{P}(Z) \to 0$ as $n \to \infty$. For the second term, we apply Lemma 2.1 to obtain

$$\mathbb{P}\left(Y_H \wedge \overline{Z}\right) \leq 15 \exp\left(-\gamma e(H) n^{-1/2}\right). \tag{5}$$

For the third and final term in (4), note that $\mathbb{P}\left(X_H \wedge \overline{Y_H}\right) \leq \mathbb{P}\left(X_H \mid \overline{Y_H}\right)$. We can understand this conditional probability of X_H by exposing the randomness of $\Gamma = G(n, p)$ in two stages. First, we expose the edges of Γ outside H and count the number of H-good edges in H. We then expose the edges of Γ in H, noting that when doing this, each edge of H has an independent p chance of lying inside Γ. Thus,

$$\begin{aligned}
\mathbb{P}\left(X_H \mid \overline{Y_H}\right) &\leq \mathbb{P}\left(X_H \mid e(H \setminus (\Gamma \setminus H)^2) \geq e(H)/4\right) \\
&\leq (1-p)^{e(H)/4} \\
&\leq \exp\left(-\frac{\varepsilon}{4} e(H) n^{-1/2}\right).
\end{aligned} \tag{6}$$

Now, we fix $\delta = \min(\varepsilon/4, \gamma)/5$. We know from (3) that

$$\sum_{H \in \mathcal{H}} \exp\left(-\delta e(H) n^{-1/2}\right) < \frac{1}{2};$$

in particular, each individual term in the sum is less than $1/2$. Thus,

$$\sum_{H \in \mathcal{H}} \exp\left(-5\delta e(H)n^{-1/2}\right) < \frac{1}{16} \sum_{H \in \mathcal{H}} \exp\left(-\delta e(H)n^{-1/2}\right) < \frac{1}{32}.$$

Putting this together with (4), (5), and (6), we find that

$$\mathbb{P}\left(\bigvee_{H \in \mathcal{H}} X_H\right) \leq \sum_{H \in \mathcal{H}} \left(15 \exp\left(-\gamma e(H)n^{-1/2}\right) + \exp\left(-\frac{\varepsilon}{4}e(H)n^{-1/2}\right)\right) + o(1)$$

$$\leq 16 \sum_{H \in \mathcal{H}} \exp\left(-5\delta e(H)n^{-1/2}\right) + o(1) < 1$$

for all $n \in \mathbb{N}$ sufficiently large, as desired.

We have shown that with positive probability, Γ contains an H-good edge for every $H \in \mathcal{H}$. Thus, with positive probability, any maximal triangle-free subgraph G of Γ will share an edge with every $H \in \mathcal{H}$, as desired. $\qed$

3. Large-deviations for closed triangles

Our proof of Lemma 2.1 needs us, amongst other things, to show that for a fixed graph H on $[n]$ and a random set $U \subset [n]$ containing each vertex with probability p, the number of edges in the induced subgraph $H[U]$ is concentrated. In particular, we need the following technical bound on the exponential moment of $e(H[U])$ in the case where H is bipartite.

Lemma 3.1. *Let H be a bipartite graph on $[n]$ with m edges, and let U be a random subset of $[n]$ such that each vertex is chosen independently with the same probability $p \in [0, 1]$. Let X be the random variable counting the number of edges in $H[U]$, and let Z be the indicator random variable of the event $|U| \leq 5pn$. Then*

$$\mathbb{E}\left[\exp\left(\frac{ZX}{5pn}\right)\right] \leq \exp\left(\frac{pm}{n}\right).$$

Proof. Let $V_L \cup V_R$ be a bipartition of the vertex set of H, and let $U_L = U \cap V_L$ and $U_R = U \cap V_R$. We expose the random set U by first exposing the random set U_L and then exposing U_R.

For $i \in V_R$, let d_i be the random variable counting the number of neighbours of i in U_L, and let X_i be the indicator random variable of the event $i \in U_R$, so that $X = \sum_{i \in V_R} d_i X_i$. Notice that d_i depends only on U_L and not on U_R, so conditional on any exposure of U_L, the random variables $\{d_i X_i : i \in V_R\}$ are independent. Therefore, taking Z_L to be the indicator of the event $|U_L| \leq 5pn$, and noting that Z_L is $\{0, 1\}$-valued and that $Z \leq Z_L$, we get

$$\mathbb{E}\left[\exp\left(\frac{ZX}{5pn}\right)\right] \leq \mathbb{E}\left[\exp\left(\frac{Z_L X}{5pn}\right)\right] = \mathbb{E}\left[1 - Z_L + Z_L \exp\left(\frac{X}{5pn}\right)\right]$$

7

$$= \mathbb{E}_{U_L}\left[1 - Z_L + \mathbb{E}_{U_R}\left[Z_L \exp\left(\frac{X}{5pn}\right)\right]\right]$$

$$= \mathbb{E}_{U_L}\left[1 - Z_L + Z_L\mathbb{E}_{U_R}\left[\exp\left(\frac{\sum_{i\in V_R} d_i X_i}{5pn}\right)\right]\right]$$

$$= \mathbb{E}_{U_L}\left[1 - Z_L + Z_L \prod_{i\in V_R} \mathbb{E}_{X_i}\left[\exp\left(\frac{d_i X_i}{5pn}\right)\right]\right]$$

$$= \mathbb{E}_{U_L}\left[1 - Z_L + Z_L \prod_{i\in V_R}\left(1 - p + p\exp\left(\frac{d_i}{5pn}\right)\right)\right], \qquad (7)$$

the last equality holding since X_i is a Bernoulli random variable with $\mathbb{P}(X_i = 1) = p$.

We now state a simple estimate that we will use to control (7): for all $0 \le x \le 1$, we have

$$e^x - 1 \le (e - 1)x; \qquad (8)$$

indeed, equality holds at $x = 0$ and $x = 1$ and the right hand side is linear, so this bound follows from convexity of $e^x - 1$.

Returning to (7), if $Z_L = 1$, then $d_i \le |U_L| \le 5pn$; thus (8) with $x = d_i/5pn$ implies that either $Z_L = 0$ or

$$\exp\left(\frac{d_i}{5pn}\right) - 1 \le \frac{(e - 1)d_i}{5pn} \le \frac{d_i}{2pn}.$$

Therefore,

$$Z_L \prod_{i\in V_R}\left[1 - p + p\exp\left(\frac{d_i}{5pn}\right)\right] \le Z_L \prod_{i\in V_R}\left(1 + \frac{d_i}{2n}\right) \le Z_L\exp\left(\frac{\sum_{i\in V_R} d_i}{2n}\right),$$

which when plugged into (7) gives

$$\mathbb{E}\left[\exp\left(\frac{ZX}{5pn}\right)\right] \le \mathbb{E}_{U_L}\left[1 - Z_L + Z_L\exp\left(\frac{\sum_{i\in V_R} d_i}{2n}\right)\right] \le \mathbb{E}_{U_L}\left[\exp\left(\frac{\sum_{i\in V_R} d_i}{2n}\right)\right]. \quad (9)$$

Now, for $j \in V_L$, let $a_j = \deg_H(j)$ and let Y_j be the indicator random variable of the event $j \in U_L$, and note that

$$\sum_{i\in V_R} d_i = \sum_{j\in V_L} a_j Y_j$$

since both sides count the number of edges between U_L and V_R. As the random variables $\{Y_j : j \in V_L\}$ are independent, we get

$$\mathbb{E}_{U_L}\left[\exp\left(\frac{\sum_{i\in V_R} d_i}{2n}\right)\right] = \mathbb{E}_{U_L}\left[\exp\left(\frac{\sum_{j\in V_L} a_j Y_j}{2n}\right)\right]$$

$$= \prod_{j\in V_L} \mathbb{E}_{Y_j}\left[\exp\left(\frac{a_j Y_j}{2n}\right)\right]$$

$$= \prod_{j \in V_L} \left(1 - p + p \exp\left(\frac{a_j}{2n}\right)\right), \tag{10}$$

the last equality holding since Y_j is a Bernoulli random variable with $\mathbb{P}(Y_i = 1) = p$.

To finish, we see that since $a_j \leq |V_R| \leq n$, the bound (8) with $x = a_j/2n$ gives

$$\exp\left(\frac{a_j}{2n}\right) - 1 \leq \frac{(e-1)a_j}{2n} \leq \frac{a_j}{n}.$$

Substituting this bound into (10), then substituting the result into (9), and using the fact that $\sum_{j \in V_L} a_j = m$, we get

$$\mathbb{E}\left[\exp\left(\frac{ZX}{5pn}\right)\right] \leq \prod_{j \in V_L} \left(1 - p + p \exp\left(\frac{a_j}{2n}\right)\right)$$
$$\leq \prod_{j \in V_L} \left(1 + \frac{pa_j}{n}\right)$$
$$\leq \exp\left(\frac{p \sum_{j \in V_L} a_j}{n}\right) = \exp\left(\frac{pm}{n}\right),$$

completing the proof. $\qquad\qquad\square$

To prove Lemma 2.1, it will be more convenient to work with random directed graphs. This model gives us some additional independence that is crucial in proving the directed analogue of Lemma 2.1 below. We shall subsequently show that Lemma 2.1 follows from its directed analogue.

We need a little more notation. We write $\vec{G}(n, p)$ for the random directed graph (or digraph, for short) on $[n]$, where each directed edge (u, v) appears independently with probability p, and pairs of anti-parallel edges (u, v) and (v, u) are allowed to be simultaneously present. For a digraph D, we write $\Delta(D)$ for its maximum out-degree, and $\widehat{D}$ for the undirected graph whose edges are pairs of distinct vertices in D with a common in-neighbour. The directed analogue of Lemma 2.1 is as follows.

Lemma 3.2. *There exist universal $\gamma', \varepsilon' > 0$ such that the following holds for all sufficiently large $m, n \in \mathbb{N}$. Let H be a bipartite graph on $[n]$ with m edges. If $D \sim \vec{G}(n, p)$ with $p \leq \varepsilon' n^{-1/2}$, then*

$$\mathbb{P}\left(e(H \cap \widehat{D}) \geq m/16 \bigwedge \Delta(D) \leq 5pn - 1\right) \leq \exp\left(-\gamma' m n^{-1/2}\right).$$

Proof. Let $D_\circ$ be the random directed graph generated in the same way as D except that self-loops are also chosen with probability p. Then in the natural coupling between D and $D_\circ$ given by removing the self-loops of $D_\circ$, we have both

$$e(H \cap \widehat{D}_\circ) \geq e(H \cap \widehat{D})$$

and
$$\Delta(D_\circ) \leq \Delta(D) + 1,$$
whence it follows that
$$\mathbb{P}\Big(e(H \cap \widehat{D}) \geq m/16 \bigwedge \Delta(D) \leq 5pn - 1\Big) \leq \mathbb{P}\Big(e(H \cap \widehat{D}_\circ) \geq m/16 \bigwedge \Delta(D_\circ) \leq 5pn\Big).$$

For each $v \in [n]$, let U_v be the out-neighbourhood of v in $D_\circ$, and let X_v be the random variable counting the number of edges of H contained in U_v. Since any edge $e \in E(\widehat{D}_\circ)$ is contained in U_v for some $v \in [n]$, we have
$$e(H \cap \widehat{D}_\circ) \leq \sum_{v \in [n]} X_v.$$

Next, let Z be the indicator random variable of the event that $\Delta(D_\circ) \leq 5pn$. For each $v \in [n]$, let Z_v be the indicator random variable of the event that $|U_v| \leq 5pn$. Clearly, $Z \leq Z_v$ for all $v \in [n]$, so it follows that

$$\mathbb{P}\Big(e(H \cap \widehat{D}_\circ) \geq m/16 \bigwedge \Delta(D_\circ) \leq 5pn\Big) \leq \mathbb{P}\left(Z \sum_{v \in [n]} X_v \geq m/16\right)$$
$$\leq \mathbb{P}\left(\sum_{v \in [n]} Z_v X_v \geq m/16\right).$$

Note that the random variables $\{Z_v X_v : v \in [n]\}$ are independent, as $Z_v X_v$ only depends on the out-neighbourhood U_v of v, and all of these are independent; this decoupling is why we find it crucial to work with random *directed* graphs.

Now, U_v is a random subset of $[n]$ with each vertex chosen independently with probability p, so by Lemma 3.1, we have
$$\mathbb{E}\left[\exp\left(\frac{Z_v X_v}{5pn}\right)\right] \leq \exp\left(\frac{pm}{n}\right)$$
for every $v \in [n]$. Therefore, using an exponential moment bound, we get

$$\mathbb{P}\left(\sum_{v \in [n]} Z_v X_v \geq m/16\right) = \mathbb{P}\left(\exp\left(\frac{\sum_{v \in [n]} Z_v X_v}{5pn}\right) \geq \exp\left(\frac{m}{80pn}\right)\right)$$
$$\leq \frac{\mathbb{E}\left[\exp\left(\frac{\sum_{v \in [n]} Z_v X_v}{5pn}\right)\right]}{\exp\left(\frac{m}{80pn}\right)}$$
$$= \frac{\prod_{v \in [n]} \mathbb{E}\left[\exp\left(\frac{Z_v X_v}{5pn}\right)\right]}{\exp\left(\frac{m}{80pn}\right)}$$

$$\le \exp\left(\frac{m}{pn}\left(p^2 n - \frac{1}{80}\right)\right).$$

For $p \le (20\sqrt{n})^{-1}$, say, we have $p^2 n - 1/80 \le -1/100$, so

$$\mathbb{P}\left(\sum_{v \in [n]} Z_v X_v \ge m/16\right) \le \exp\left(-\frac{m}{100pn}\right) \le \exp\left(-\frac{m}{5\sqrt{n}}\right).$$

Thus, Lemma 3.2 is seen to hold with $\gamma' = 1/5$ and $\varepsilon' = 1/20$. $\qquad\square$

Finally, the proof of Lemma 2.1 follows from Lemma 3.2 and Markov's inequality by constructing a suitable coupling.

Proof of Lemma 2.1. Let H be an arbitrary graph on $[n]$ with m edges. Choose a bipartite subgraph H' of H with $m' \ge m/2$ edges. Let $\gamma', \varepsilon' > 0$ be the constants in Lemma 3.2, suppose $p \le \varepsilon' n^{-1/2}$, and pick p' such that $2p' - (p')^2 = p$. Then $p' \le \varepsilon' n^{-1/2}$ satisfies the conditions of Lemma 3.2, and for $m, n \in \mathbb{N}$ sufficiently large, Lemma 3.2 applies to H' and p', telling us that if $D \sim \vec{G}(n, p')$, then

$$\mathbb{P}\left(e(H' \cap \widehat{D}) \ge m'/16 \bigwedge \Delta(D) \le 5p'n - 1\right) \le \exp\left(-\gamma' m' n^{-1/2}\right).$$

Let $\Gamma \sim G(n, p)$ be an undirected Erdős–Rényi random graph. By our choices of p and p', we can couple $\Gamma \sim G(n, p)$ and $D \sim \vec{G}(n, p')$ such that two vertices in Γ are adjacent if and only if there exists at least one directed edge between them in D. In particular, $\Delta(D) \le \Delta(\Gamma)$, so

$$\Delta(\Gamma) \le 2pn \implies \Delta(D) \le 5p'n - 1.$$

Since H' contains at least $m' \ge m/2$ edges of H, we also have

$$e(H \cap \Gamma^2) \ge 3m/4 \implies e(H' \cap \Gamma^2) \ge m'/2.$$

Now, if an edge lies in Γ^2, then by the coupling above, it has at least a $1/4$ chance of lying in $\widehat{D}$. Thus, for any fixed graph Γ_0 on $[n]$ with $e(H' \cap \Gamma_0^2) \ge m'/2$, we have

$$\mathbb{E}[e(H' \cap \widehat{D}) \mid \Gamma = \Gamma_0] \ge m'/8.$$

Since $e(H' \cap \widehat{D})$ is a random variable supported on $[0, m']$, Markov's inequality applied to $X = m' - e(H' \cap \widehat{D})$ yields

$$\mathbb{P}\left(e(H' \cap \widehat{D}) < m'/16 \mid \Gamma = \Gamma_0\right) = \mathbb{P}(X \ge 15m'/16 \mid \Gamma = \Gamma_0) \le \frac{7m'/8}{15m'/16} = \frac{14}{15},$$

implying that for any Γ_0 with $e(H \cap \Gamma_0^2) \ge 3m/4$, we have

$$\mathbb{P}\left(e(H' \cap \widehat{D}) \ge m'/16 \mid \Gamma = \Gamma_0\right) \ge \frac{1}{15}.$$

Summing this estimate over all possible Γ_0, it follows that

$$\mathbb{P}\Big(e(H \cap \Gamma^2) \geq 3m/4 \bigwedge \Delta(\Gamma) \leq 2pn\Big)$$
$$\leq 15\mathbb{P}\Big(e(H' \cap \widehat{D}) \geq m'/16 \bigwedge \Delta(D) \leq 5p'n - 1\Big)$$
$$\leq 15\exp\big(-\gamma' m' n^{-1/2}\big)$$
$$= 15\exp\big(-\gamma m n^{-1/2}\big),$$

where $\gamma = \gamma'/2$; this completes the proof. $\qquad\square$

4. FRACTIONAL RELAXATIONS

We now turn to the fractional analogue of Theorem 1.4; to state this carefully, we need some definitions.

Following Talagrand [14, 15] once again, we say an up-set $\mathcal{F} \subset 2^X$ is *weakly p-small* if there is a certificate $a : 2^X \to \mathbb{R}_{\geq 0}$ such that for each $S \in \mathcal{F}$, we have

$$\sum_{T \subset S} a(T) \geq 1,$$

and

$$\sum_{T \subset X} a(T)p^{|T|} \leq 1/2;$$

thus, if $\mathcal{F}$ is weakly p-small, then we have a simple proof, as before, of the fact that $p_c(\mathcal{F}) \geq p$. Analogously, we say a down-set $\mathcal{F} \subset 2^X$ is *weakly p-small* if there is a certificate $a : 2^X \to \mathbb{R}_{\geq 0}$ such that for each $S \in \mathcal{F}$, we have

$$\sum_{T \supset S} a(T) \geq 1,$$

and

$$\sum_{T \subset X} a(T)(1-p)^{|X \setminus T|} \leq 1/2;$$

thus, if $\mathcal{F}$ is weakly p-small, then we again have a simple proof of the fact that $p_c(\mathcal{F}) \leq p$.

Now, we define the *fractional expectation-threshold* $q_f(\mathcal{F})$ of a monotone $\mathcal{F}$ as follows: for an up-set $\mathcal{F}$, this is the *largest* value $p \in [0,1]$ for which $\mathcal{F}$ is weakly p-small, and for a down-set $\mathcal{F}$, this is the *smallest* value $p \in [0,1]$ for which $\mathcal{F}$ is weakly p-small.

The statement that $q_f(\mathcal{T}_n) = \Omega(1/\sqrt{n})$ is equivalent to the following modification of Theorem 1.4.

Theorem 4.1. *There exists a universal $\delta > 0$ such that the following holds for all sufficiently large $n \in \mathbb{N}$. If $\mathcal{H}$ is the collection of all graphs on $[n]$ and $a : \mathcal{H} \to \mathbb{R}_{\geq 0}$*

12

satisfies

$$\sum_{H \in \mathcal{H}} a(H) \exp\big((-\delta e(H))/\sqrt{n}\big) < \frac{1}{2},$$

then there is a triangle-free graph G on $[n]$ such that

$$\sum_{H \in \mathcal{G}(G)} a(H) < 1,$$

where $\mathcal{G}(G) = \{H \in \mathcal{H} : G \cap H = \varnothing\}$.

Proof. This can be proved in a manner similar to Theorem 1.4. In particular, let $\Gamma = G(n,p)$, let G be any maximal triangle-free subgraph of Γ, and define X_H, Y_H, and Z as in the proof of Theorem 1.4. Recall that if $\overline{X_H}$ holds, then $G \cap H \neq \varnothing$. Thus

$$\mathbb{E}\left[\sum_{H \in \mathcal{G}(G)} a(H) \,\bigg|\, \overline{Z}\right] = \sum_{H \in \mathcal{H}} a(H) \mathbb{P}\big(G \cap H = \varnothing \,|\, \overline{Z}\big)$$

$$\leq \sum_{H \in \mathcal{H}} a(H) \mathbb{P}\big(X_H \,|\, \overline{Z}\big)$$

$$= \sum_{H \in \mathcal{H}} a(H) \frac{\mathbb{P}\big(X_H \wedge \overline{Z}\big)}{\mathbb{P}\big(\overline{Z}\big)}.$$

By the bounds on $\mathbb{P}(X_H \wedge \overline{Y_H})$ and $\mathbb{P}(Y_H \wedge \overline{Z})$ in the proof of Theorem 1.4, we have

$$\mathbb{P}\big(X_H \wedge \overline{Z}\big) \leq 16 \exp\big(-ce(H)n^{-1/2}\big)$$

for some absolute constant $c > 0$. Of course, this probability is bounded above by 1, so since $\mathbb{P}(\overline{Z}) = 1 - o(1)$, we get

$$\mathbb{E}\left[\sum_{H \in \mathcal{G}(G)} a(H) \,\bigg|\, \overline{Z}\right] \leq (1 + o(1)) \sum_{H \in \mathcal{H}} a(H) \min\big(16 \exp\big(-ce(H)n^{-1/2}\big), 1\big).$$

Putting $\delta = c/8$, we then have

$$\min\big(16 \exp\big(-ce(H)n^{-1/2}\big), 1\big) \leq \sqrt{2} \exp\big(-\delta e(H)n^{-1/2}\big),$$

as the right hand side is the geometric mean of one copy of $16 \exp\big(-ce(H)n^{-1/2}\big)$ and seven copies of 1. Thus, if

$$\sum_{H \in \mathcal{H}} a(H) \exp\big((-\delta e(H))/\sqrt{n}\big) < \frac{1}{2},$$

then

$$\mathbb{E}\left[\sum_{H \in \mathcal{G}(G)} a(H) \,\bigg|\, \overline{Z}\right] \leq \big(\sqrt{2} + o(1)\big) \sum_{H \in \mathcal{H}} a(H) \exp\big((-\delta e(H))/\sqrt{n}\big)$$

$$\leq \frac{1}{\sqrt{2}} + o(1) < 1,$$

proving that there is some triangle-free graph G on $[n]$ for which

$$\sum_{H \in \mathcal{G}(G)} a(H) < 1,$$

as required. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

5. Uniform covers

Our goal in this section is to collect together some constructions that yield good upper bounds on the expectation-threshold of $\mathcal{T}_n$ (and consequently, its fractional expectation-threshold as well, since $q_f(\mathcal{T}_n) \leq q(\mathcal{T}_n)$).

For $0 \leq m \leq \binom{n}{2} - n^2/4$, let $f(m, n)$ be the smallest integer k such that there exists a family $\mathcal{H}$ of k graphs on $[n]$ with $e(H) = \binom{n}{2} - m$ for all $H \in \mathcal{H}$ such that every triangle-free graph is contained in some $H \in \mathcal{H}$.

Notice that $f(m, n)$ exists for all m in the given range, as all triangle-free graphs have at most $n^2/4$ edges. The relationship between the expectation-threshold $q(\mathcal{T}_n)$ and the function f is as follows.

Proposition 5.1. *For all $0 \leq m \leq \binom{n}{2} - n^2/4$, we have $q(\mathcal{T}_n) \leq \log(2f(m, n))/m$.*

Proof. Take a family $\mathcal{H}$ of $f(m, n)$ graphs, each with $\binom{n}{2} - m$ edges, that covers $\mathcal{T}_n$. The probability that $G(n, p)$ is contained in one of these graphs is $(1 - p)^m$. Setting $p = q(\mathcal{T}_n)$, we must have

$$f(m, n)(1 - p)^m \geq 1/2$$

by the minimality of the expectation-threshold; noting that $(1 - p)^m \leq e^{-pm}$, the conclusion follows. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

We now collect together estimates for $f(m, n)$ that hold in various regimes.

Proposition 5.2. *We have the following bounds on $f(m, n)$ for all sufficiently large $n \in \mathbb{N}$.*

(1) *For all $0 \leq m \leq \binom{n}{2} - n^2/4$, we have*
$$f(m, n) = \exp\big(\Omega\big(\max\big(m/\sqrt{n}, \sqrt{m}\big)\big)\big).$$

(2) *There exists a universal $c > 0$ such that if $m < cn \log n$, then*
$$f(m, n) \leq \exp\big(2\sqrt{m} \log n\big).$$

(3) *For all $0 < c < 1/4$, there exists $C = C(c) > 0$ such that*
$$f\big(cn^2, n\big) \leq \exp\big(Cn^{3/2} \log n\big).$$

Before turning to the proof of these estimates, notice that Proposition 5.1 and Item 2 in the above proposition combine to show that

$$q_f(\mathcal{T}_n) \le q(\mathcal{T}_n) = O\left(\frac{\sqrt{n \log n} \log n}{n \log n}\right) = O\left(\sqrt{\frac{\log n}{n}}\right),$$

as was earlier claimed.

Proof of Proposition 5.2. We start with Item 1. The first inequality here, i.e., $f(m,n) = \exp(\Omega(m/\sqrt{n}))$, comes from Proposition 5.1, and our proof of the fact that $q(\mathcal{T}_n) = \Omega(1/\sqrt{n})$. For the other bound, consider a random complete bipartite graph G in which each vertex is independently sent to either partition class of G with probability $1/2$. For any graph H with $e(H) = \binom{n}{2} - m$, let T be a spanning forest of H^c (i.e. the union of spanning trees of each connected component), and note that $e(T) \ge v(T)/2 = \Omega(\sqrt{m})$. It is not difficult to see that the edges of T are contained in G with independent probabilities of $1/2$, so

$$\mathbb{P}(G \subset H) \le \mathbb{P}(G \cap T = \varnothing) = 2^{-e(T)} = \exp\left(-\Omega\left(\sqrt{m}\right)\right).$$

Since such a G is always triangle-free, if $\mathcal{H}$ is a family of graphs, each with $\binom{n}{2} - m$ edges, that covers $\mathcal{T}_n$, a union bound yields $|\mathcal{H}| = \exp(\Omega(\sqrt{m}))$, and the second inequality follows.

Next, we turn to Item 2. Ajtai, Komlós and Szemerédi [1] showed that $r(3,n) = O(n^2/\log n)$, i.e., there is a universal $c > 0$ such that $r(3, 2\sqrt{m}) \le n$ for all $m < cn \log n$. This implies that we can cover $\mathcal{T}_n$ by taking all graphs of the form $K_n - B$ where B is a clique of size $2\sqrt{m}$. There are $\binom{n}{2\sqrt{m}}$ such graphs, and they each have $\binom{2\sqrt{m}}{2} \ge m$ non-edges, so this shows that

$$f(m,n) \le \binom{n}{2\sqrt{m}} \le \exp\left(2\sqrt{m} \log n\right).$$

Finally, for Item 3, a standard application of the machinery of hypergraph containers (as in [12], for example) shows that for all $\varepsilon > 0$, there exists $C > 0$ such that for all $n \in \mathbb{N}$, there is a family of $\exp(Cn^{3/2} \log n)$ graphs on $[n]$, each with at most $(\frac{1}{4} + \varepsilon)n^2$ edges, that collectively contain all triangle-free graphs on $[n]$; the claimed bound follows by adding edges to each of these graphs until each of them has exactly $\binom{n}{2} - cn^2$ edges. $\square$

6. FURTHER QUESTIONS

Our results leave open the problems of determining the asymptotic growth rates of $q(\mathcal{T}_n)$ and $q_f(\mathcal{T}_n)$. We have shown that both these thresholds are $\Omega(1/\sqrt{n})$ and $O(\sqrt{\log n/n})$, but getting more precise estimates remains an interesting open problem.

Problem 6.1. *What are the asymptotic growth rates of $q(\mathcal{T}_n)$ and $q_f(\mathcal{T}_n)$?*

Towards Problem 6.1, notice that Proposition 5.2 effectively bounds $f(m, n)$ when $m = O(n \log n)$ and $m = \Theta(n^2)$, i.e., up to a logarithmic factor in the exponent, $f(m, n)$ behaves like $\exp(\sqrt{m})$ when $m = O(n \log n)$, and like $\exp(n^{3/2})$ when $m = \Theta(n^2)$. This motivates the following question.

Problem 6.2. *What is the behaviour of $f(m, n)$ for $n \log n \ll m \ll n^2$? Is the bound $f(m, n) = \exp(\Omega(m/\sqrt{n}))$ sharp in this regime (possibly up to a logarithmic factor in the exponent)?*

Another natural question, whose answer might help resolve Problem 6.1, is as follows.

Problem 6.3. *For $m = O(n \log n)$ and $m = \Theta(n^2)$, can we eliminate the logarithmic gaps from our bounds on $f(m, n)$?*

Finally, we remark that our machinery is specialised to the case of triangle-free graphs, and does not prove analogous results for larger cliques; this prompts the following question.

Problem 6.4. *For $r \geq 4$ and $\mathcal{F}_{r,n}$ the down-set of all K_r-free graphs on $[n]$, what are the asymptotic growth rates of $q(\mathcal{F}_{r,n})$ and $q_f(\mathcal{F}_{r,n})$?*

Acknowledgements

References

1. M. Ajtai, J. Komlós, and E. Szemerédi, *A note on Ramsey numbers*, J. Combin. Theory Ser. A **29** (1980), 354–360. 4, 15

2. N. Alon and J. H. Spencer, *The probabilistic method*, fourth ed., Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2016. 3

3. J. Balogh, R. Morris, and W. Samotij, *Independent sets in hypergraphs*, J. Amer. Math. Soc. **28** (2015), 669–709. 4

4. B. Bollobás, *Random graphs*, Cambridge Studies in Advanced Mathematics, vol. 73, Cambridge University Press, Cambridge, 2001. 1

5. P. Erdős, *Graph theory and probability II*, Canad. J. Math. **13** (1961), 346–352. 4

6. P. Erdős and A. Rényi, *On the evolution of random graphs*, Magyar Tud. Akad. Mat. Kutató Int. Közl. **5** (1960), 17–61. 1

7. K. Frankston, J. Kahn, B. Narayanan, and J. Park, *Thresholds versus fractional expectation-thresholds*, Ann. of Math. **194** (2021), 475–495. 1, 3

8. S. Janson, T. Łuczak, and A. Rucinski, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000. 1

9. J. Kahn and G. Kalai, *Thresholds and expectation thresholds*, Combin. Probab. Comput. **16** (2007), 495–502. 1, 3

10. J. H. Kim, *The Ramsey number $R(3,t)$ has order of magnitude $t^2/\log t$*, Random Structures Algorithms **7** (1995), 173–207. 4

11. J. Park and H. T. Pham, *A proof of the Kahn-Kalai conjecture*, Preprint, arXiv:2203.17207. 1, 3

12. D. Saxton and A. Thomason, *Hypergraph containers*, Invent. Math. **201** (2015), 925–992. 4, 15

13. J. Spencer, *Ramsey's theorem – a new lower bound*, J. Combin. Theory Ser. A **18** (1975), 108–115. 4

14. M. Talagrand, *Are all sets of positive measure essentially convex?*, Geometric aspects of functional analysis (Israel, 1992–1994), Oper. Theory Adv. Appl., vol. 77, Birkhäuser, Basel, 1995, pp. 295–310. 2, 12

15. _____, *The generic chaining*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2005. 2, 12

16. _____, *Are many small sets explicitly small?*, Proceedings of the 2010 ACM International Symposium on Theory of Computing, 2010, pp. 13–35. 1, 3

Department of Mathematics, Rutgers University, Piscataway, NJ 08854, USA
Email address: bg570@connect.rutgers.edu

Department of Mathematics, Princeton University, Princeton, NJ 08540, USA
Email address: xiaoyuh@princeton.edu

Department of Mathematics, Rutgers University, Piscataway, NJ 08854, USA
Email address: narayanan@math.rutgers.edu