

On the power of random greedy algorithms

He Guo* and Lutz Warnke†

April 15, 2021; revised April 3, 2022

Abstract

In this paper we solve two problems of Esperet, Kang and Thomassé as well as Li concerning (i) induced bipartite subgraphs in triangle-free graphs and (ii) van der Waerden numbers. Each time random greedy algorithms allow us to go beyond the Lovász Local Lemma or alteration method used in previous work, illustrating the power of the algorithmic approach to the probabilistic method.

1 Introduction

The probabilistic method is a widely used tool for proving the existence of hard-to-construct mathematical objects with certain desirable properties: it works by showing that a randomly chosen object has the desired properties with non-zero probability. In classical textbook approaches to the probabilistic method, the underlying random objects are typically generated in a *static* way, e.g., by choosing a graph uniformly at random from a prescribed class of graphs, or by independently including each possible edge.

In this paper we illustrate the power of the algorithmic approach to the probabilistic method, where the random objects are generated step-by-step in a *dynamic* way using a randomized algorithm. To this end we consider two examples from graph theory and additive combinatorics, and show that each time random greedy algorithms allow us to go beyond classical applications of the probabilistic method, i.e., prove existence of mathematical objects with better properties. These algorithmic improvements are key for (i) resolving a problem of Esperet, Kang and Thomassé [8], and (ii) answering a question of Li [24], see Theorems 1.1 and 1.2.

For the two combinatorial examples considered in this paper, previous work used the probabilistic method to show that static random objects can avoid certain forbidden substructures, while maintaining other desired pseudo-random properties. Our technical results show that random greedy algorithms, which by construction avoid these forbidden substructures, create objects with superior pseudo-random properties, see Theorems 1.3 and 1.5. With the benefit of hindsight, earlier work of Rödl [27], Kahn [21], Wormald [34], Spencer [29], Kim [22], Bohman [2], and others [1, 5, 10] can be interpreted similarly. This paper thus reveals the following emerging algorithmic paradigm: one can often take proofs based on classical probabilistic method arguments, and obtain improvements by using an algorithmic approach to the probabilistic method.

1.1 Induced bipartite subgraphs in triangle-free graphs

Our first example is from extremal graph theory, concerning a local refinement of the famous Max Cut problem. Here the history starts in 1988, when Erdős, Faudree, Pach and Spencer [7] introduced the problem of searching for large induced bipartite subgraphs in triangle-free graphs. Around 2018 Esperet, Kang and Thomassé [8] further refined this problem, focusing on induced bipartite subgraphs with large minimum degree. More precisely, for fixed $\eta \in (0, 1)$ they asked to determine the behavior of the parameter $f_\eta(n)$, which is defined as the maximum f such that every n -vertex triangle-free graph with minimum degree at least n^η contains an induced bipartite subgraph with minimum degree at least f . Recent results of Kwan, Letzter, Sudakov and Tran [23] and Cames van Batenburg, de Joannis de Verclos, Kang and Pirot [6] show that

$$f_\eta(n) = \Theta(\max\{\log n, n^{2\eta-1}\}) \quad \text{for fixed } \eta \in (0, 1) \setminus (1/2, 2/3], \quad (1)$$

*Faculty of Mathematics, Technion, Haifa 32000, Israel. E-mail: hguo@campus.technion.ac.il.

†Department of Mathematics, University of California, San Diego, La Jolla CA 92093, USA. E-mail: lwarnke@ucsd.edu. Supported by NSF grant DMS-1703516, NSF CAREER grant DMS-1945481, and a Sloan Research Fellowship.

and also determine $f_\eta(n)$ up to logarithmic factors in the remaining range $\eta \in (1/2, 2/3]$. Illustrating the conceptual punchline of this paper, we use a ‘dynamic’ randomized greedy algorithm to improve existing upper bound constructions [23, 6], which were based on classical probabilistic method tools applied to the binomial random graph $G(n, p)$. This algorithmic improvement allows us to close the logarithmic gap for $\eta \in (1/2, 2/3]$, and determine the order of magnitude of $f_\eta(n)$ for any fixed $\eta \in (0, 1)$. The following result in particular resolves [8, Problem 4.1] of Esperet, Kang and Thomassé up to constant factors.

Theorem 1.1. *For fixed $\eta \in (0, 1)$, we have $f_\eta(n) = \Theta(\max\{\log n, n^{2\eta-1}\})$.*

In comparison to the previous upper bounds [23, 6] based on the probabilistic analysis of $G(n, p)$ via the alteration method or the Lovász Local Lemma, our key improvement stems from the fact that via the so-called semi-random triangle-free process we are able to algorithmically construct pseudo-random triangle-free graphs with higher edge density (see Theorem 1.3), confirming speculations from [8, Section 4] and [6, Section 3].

1.2 Van der Waerden numbers

Our second example is from additive combinatorics, concerning a well-known Ramsey-type parameter for arithmetic progressions. The van der Waerden number $W(r, k)$ is defined as the smallest integer n such that every red and blue coloring of numbers in $[n] := \{1, 2, \dots, n\}$ contains either a monochromatic red r -term arithmetic progression (r -AP) or a monochromatic blue k -AP. The celebrated van der Waerden’s theorem guarantees that $W(r, k)$ is finite for all integers $r, k \geq 2$, making it a natural and interesting problem to determine the asymptotic behavior of $W(r, k)$, see [11, 14]. The ‘off-diagonal’ case, where $r \geq 3$ is fixed and k tends to infinity, was of particular interest to Graham (note that $W(2, k) = \Theta(k)$ holds trivially). Indeed, in the mid 2000s Graham conjectured that $W(3, k) \leq k^{O(1)}$, and mentioned that numerical evidence suggests $W(3, k) = k^{2+o(1)}$, see [12, 13, 16]. Around 2015 Graham even started offering a \$250 reward for his conjecture, see [13, p. 19]. In terms of lower bounds, in 2008 Li and Shu [25] showed that

$$W(r, k) = \Omega((k/\log k)^{r-1}) \quad \text{for fixed } r \geq 3,$$

by applying the Lovász Local Lemma to a random subset of the integers $[n]$. Subsequently, Li raised in 2009 the natural question [24] whether this probabilistic lower bound can be improved via a randomized greedy algorithm that ‘dynamically’ constructs an r -AP free subset of the integers $[n]$. The proof of the following theorem answers Li’s question affirmatively, see also Sections 1.4.2 and 3.

Theorem 1.2. *For fixed $r \geq 3$, we have $W(r, k) = \Omega(k^{r-1}/(\log k)^{r-2})$.*

This result was announced in October 2020, see [17]. While preparing this paper, Green [15] made a breakthrough and showed $W(3, k) \geq k^{(\log k)^{1/3-o(1)}}$ using very different techniques, which in view of $W(r, k) \geq W(3, k)$ disproves the earlier belief that $W(r, k) = k^{O(1)}$ for fixed $r \geq 3$. The best known upper bound $W(3, k) \leq \exp(k^{1-\Omega(1)})$ was obtained by Schoen [28].

1.3 Organization

In Section 1.4 we state our main technical results, which will imply Theorem 1.1 and 1.2 for induced bipartite subgraphs and van der Waerden numbers, respectively. In Sections 2 and 3 we then prove these technical results using an algorithmic approach to the probabilistic method, i.e., by analyzing randomized algorithms that construct pseudo-random triangle-free graphs and r -AP free subsets of the integers, respectively.

1.4 Main technical results

1.4.1 Construction of pseudo-random triangle-free graphs

To prove the upper bound on the parameter $f_\eta(n)$ claimed by Theorem 1.1 for $\eta \in (1/2, 2/3]$, our strategy is to construct a pseudo-random triangle-free graph G_n with $\Theta(n)$ vertices, where pseudo-randomness will intuitively ensure the desired minimum degree properties (in suitable constructions that are based on G_n). Following the conceptual punchline of this paper, we shall construct the desired graph G_n using a semi-random variant of the triangle-free process, which is a randomized greedy algorithm that sequentially adds

more edges to G_n without creating a triangle, see Section 2 for the full details. This algorithmic approach to the probabilistic method is key for obtaining our improved upper bound on $f_\eta(n)$ via the following auxiliary result, since earlier approaches based on the binomial random graph $G(n, p)$ were only able to prove a weaker version of Theorem 1.3, where the triangle-free graph G_n is sparser, i.e., with minimum and maximum degree bounds $\delta(G_n), \Delta(G_n) = \Theta(\sqrt{n})$; see [23, Lemma 5.1] and [6, Theorem 3.1]. Note that we recover these earlier results, by sampling each edge of the graph G_n from Theorem 1.3 independently with probability $1/\sqrt{\log n}$.

Theorem 1.3. *There are constants $c, C, C' > 0$ such that for any $0 < \beta < 1/14$ the following holds for any integer $n \geq n_0 = n_0(\beta)$. There exists a triangle-free graph G_n with $v(G_n) \in [n/3, n]$ vertices,*

$$c\sqrt{\beta n \log n} \leq \delta(G_n) \leq \Delta(G_n) \leq C\sqrt{\beta n \log n}, \quad (2)$$

and the property that any induced bipartite subgraph $F \subseteq G_n$ has minimum degree $\delta(F) \leq C' \log n$.

We defer the proof of Theorem 1.3 to Section 2: it is based on a careful refinement¹ of the semi-random triangle-free process analysis of Guo and Warnke [19]. Using Theorem 1.3 we shall in fact establish improved bounds for the more general parameter $g(n, d)$, which denotes the maximum g such that every n -vertex triangle-free graph with minimum degree at least d contains an induced bipartite subgraph with minimum degree at least g . Extending [23, 6], the following result establishes a phase transition of $g(n, d)$ when the minimum degree d is around $\sqrt{n \log n}$, and it also implies Theorem 1.1 since $f_\eta(n) = g(n, n^\eta)$.

Theorem 1.4. *For any fixed $\gamma \in (0, 1)$, we have $g(n, d) = \Theta(\max\{\log d, d^2/n\})$ for all $n^\gamma \leq d \leq n/2$.*

Similar to $f_\eta(n) = g(n, n^\eta)$, the cases $n^\gamma \leq d \leq \sqrt{n}$ and $n^{2/3} \leq d \leq n/2$ of Theorem 1.4 follow from [23]. Furthermore, for $\sqrt{n} \leq d \leq n^{2/3}$ we obtain $g(n, d) = \Omega(\max\{\log d, d^2/n\})$ by combining [23, Theorem 1.3] with the fact that $g(n, d)$ is monotone increasing in d . We now close the gap for $\sqrt{n} \leq d \leq n^{2/3}$ by mimicking the upper bound constructions from [23, 6] using the semi-random triangle-free process based graphs G_n from Theorem 1.3, which have better degree properties than the $G(n, p)$ based graphs used in [23, 6].

Proof of Theorem 1.4 based on Theorem 1.3. Writing $c, C' > 0$ for the constants of Theorem 1.3, let $\beta := 10^{-2}$ and $A := c\sqrt{\beta}/3$. As discussed, it suffices to prove $g(n, d) = O(\max\{\log d, d^2/n\})$ for $\sqrt{n} \leq d \leq n^{2/3}$. To this end we may always assume that n is sufficiently large (whenever necessary), as usual.

We start with the case $\sqrt{n} \leq d \leq A\sqrt{n \log n}$, where we set $\alpha := 2/(c^2\beta)$ and $n' := \lceil \alpha d^2 / \log n \rceil$. Note that $n^{2/3} \ll n' \leq \lceil \alpha A^2 n \rceil \leq n/2$ and $n' \ll d^2$. By taking the disjoint union of $\lfloor n/n' \rfloor$ copies of $G_{n'}$, we obtain a triangle-free graph H_n with $v(H_n) = \lfloor n/n' \rfloor \cdot v(G_{n'}) \in [n/6, n]$ vertices and minimum degree

$$\delta(H_n) = \delta(G_{n'}) \geq c\sqrt{\beta n' \log n'} \geq \sqrt{c^2\beta\alpha \cdot d^2 \cdot 2/3} > d.$$

Furthermore, every induced bipartite subgraph $F \subseteq H_n$ is a disjoint union of induced bipartite subgraphs from copies of $G_{n'}$ and thus has minimum degree at most $\delta(F) \leq C' \log n' \leq 2C' \log d$. By ‘blowing up’ each vertex of H_n into an independent set of suitable sizes between one and six (i.e., after replacing each vertex of H_n by an independent set, we add a complete bipartite graph between every pair of independent sets that correspond to an edge in H_n), we thus obtain an n -vertex triangle-free graph $G_{n,d}$ with $\delta(G_{n,d}) \geq \delta(H_n) \geq d$, where furthermore every induced bipartite subgraph $F \subseteq G_{n,d}$ has minimum degree at most $\delta(F) \leq 6 \cdot 2C' \log d$ (by analogous disjoint reasoning as before), establishing that $g(n, d) = O(\log d)$.

Finally, in the remaining case $A\sqrt{n \log n} \leq d \leq n^{2/3}$ we set $\alpha := c^2\beta/18$ and $n' := \lfloor \alpha(n/d)^2 \log n \rfloor$. Note that $n^{2/3} \ll n' \leq \alpha n/A^2 \leq n/2$. By ‘blowing up’ each vertex of $G_{n'}$ into an independent set of size $\lfloor n/n' \rfloor$, we obtain a triangle-free graph H_n with $v(H_n) = \lfloor n/n' \rfloor \cdot v(G_{n'}) \in [n/6, n]$ vertices and minimum degree

$$\delta(H_n) = \lfloor n/n' \rfloor \cdot \delta(G_{n'}) \geq \frac{n}{2n'} \cdot c\sqrt{\beta n' \log n'} \geq \sqrt{\frac{c^2\beta n^2 \log(n^{2/3})}{4n'}} \geq \sqrt{\frac{c^2\beta \cdot d^2 \cdot 2/3}{4\alpha}} > d.$$

Furthermore, every induced bipartite subgraph $F \subseteq H_n$ has minimum degree at most $\delta(F) \leq \lfloor n/n' \rfloor \cdot C' \log n' \leq 2\alpha^{-1}C' \cdot d^2/n$. By blowing up each vertex of H_n into an independent set of suitable sizes between one and six, we then obtain an n -vertex triangle-free graph $G_{n,d}$ that establishes $g(n, d) = O(d^2/n)$. \square

¹In the proof of Theorem 1.3 we mainly use the semi-random variant instead of the standard triangle-free process [2, 9, 4] for technical convenience, since in Section 2 this leads to conceptually simpler proofs of certain pseudo-random properties (i.e., where the necessary refinements of existing proofs require less technical work for the semi-random variant).

1.4.2 Construction of pseudo-random r -AP free sets of integers

To prove the lower bound on the van der Waerden number $W(r, k)$ claimed by Theorem 1.2, our strategy is to construct a large subset $I \subseteq [n]$ of the integers that is r -AP free and pseudo-random, where pseudo-randomness will intuitively ensure that $[n] \setminus I$ is k -AP free for fairly large $k = k(n)$. For technical reasons, it will be convenient to work with the field $\mathbb{Z}/N\mathbb{Z}$ for a prime number N , where a set of numbers $\{a_1, \dots, a_r\} \subseteq \mathbb{Z}/N\mathbb{Z}$ is formally called an r -term arithmetic progression (r -AP) in $\mathbb{Z}/N\mathbb{Z}$ if $|\{a_1, \dots, a_r\}| = r$ and $a_i \equiv_N a_1 + (i-1)d$ for some $d \not\equiv_N 0$. Following the conceptual punchline of this paper, we shall construct the desired pseudo-random r -AP free subset $I \subseteq \mathbb{Z}/N\mathbb{Z}$ using the so-called random greedy r -AP free process, which is a randomized greedy algorithm that step-by-step adds more random numbers to I without creating an r -AP, see Section 3 for the full details. This algorithmic approach to the probabilistic method is key for obtaining our improved lower bound on $W(r, k)$ via the following result, since earlier approaches based on random subsets the integers were only able to prove Theorem 1.5 with the weaker parameter choice $k = \Theta(N^{1/(r-1)} \log N)$, see [25]. Note that we recover these earlier results, by observing that the set $I \subseteq \mathbb{Z}/N\mathbb{Z}$ from Theorem 1.5 also satisfies $|I \cap K| \geq 1$ for all k -APs K in $\mathbb{Z}/N\mathbb{Z}$ of size $k \geq k_N$.

Theorem 1.5. *For any fixed $r \geq 3$, there are constants $C, N_0 > 0$ such that the following holds for any prime number $N \geq N_0$. There exists a set $I \subseteq \mathbb{Z}/N\mathbb{Z}$ which (i) is r -AP free in $\mathbb{Z}/N\mathbb{Z}$ and (ii) satisfies $|I \cap K| \geq 1$ for all k -APs K in $\mathbb{Z}/N\mathbb{Z}$ of size $k = k_N := \lceil C(N/\log N)^{1/(r-1)} \log N \rceil$.*

Proof of Theorem 1.2 based on Theorem 1.5. Assuming that k is sufficiently large (as we may), we pick the largest prime number $N \geq \max\{2, N_0\}$ satisfying $k \geq k_N = \lceil C(N/\log N)^{1/(r-1)} \log N \rceil$. Using Bertrand's postulate it follows that $n := N - 1 = \Theta(k^{r-1}/(\log k)^{r-2})$. For $I \subseteq \mathbb{Z}/N\mathbb{Z}$ as given by Theorem 1.5, we color $I \cap [n]$ red and $[n] \setminus I$ blue. Properties (i)–(ii) of Theorem 1.5 and $k \geq k_N$ ensure that this coloring contains no red r -APs or blue k -APs in $[n]$, since any AP in $[n]$ corresponds to an AP in $\mathbb{Z}/N\mathbb{Z}$ (and any blue k -AP contains all numbers of at least one blue k_N -AP). It follows that $W(r, k) > n = \Theta(k^{r-1}/(\log k)^{r-2})$. \square

We defer the proof of Theorem 1.5 to Section 3: it is based on the differential equation method and results of Bohman and Bennett [1] for the so-called random greedy independent set algorithm. Noteworthily, in our analysis we need to ensure that all of the polynomially many k -APs are ‘hit’ by the set I produced by the r -AP process. This is in great contrast to the analysis of the H -free process arising in graph Ramsey theory, where one typically needs to ensure that an exponential number of substructures are hit [2, 3, 9, 4, 30, 31, 26].

2 Semi-random triangle-free process

In this section we prove Theorem 1.3 by showing that a semi-random variant of the so-called triangle-free process typically finds a triangle-free graph $G_n \subseteq K_n$ with the desired properties. Intuitively, this process starts with an empty graph, and then iteratively adds a large number of carefully chosen edges (instead of just adding a single edge as in the original triangle-free process) such that the resulting graph stays triangle-free.

2.1 More details and heuristics

The formal details of the semi-random triangle-free process given in [19, Section 2] are rather involved, so here we shall only introduce those aspects that are important for the upcoming arguments of this paper (keeping the notation from [19] to minimize notational differences). The semi-random process starts with

$$E_0 = T_0 := \emptyset \quad \text{and} \quad O_0 := E(K_n), \tag{3}$$

and the rough plan is to step-by-step build up a ‘random’ set of edges $E_i \subseteq E(K_n)$, a triangle-free edge subset $T_i \subseteq E_i$, and a set of ‘open’ edges $O_i \subseteq E(K_n) \setminus E_i$, each of which can still be added to E_i without creating triangles. More precisely, in step $i + 1 \geq 1$ of the semi-random triangle-free process we sample a random edge subset $\Gamma_{i+1} \subseteq O_i$, where each edge $e \in O_i$ is included independently with probability

$$p := \sigma/\sqrt{n} \quad \text{for} \quad \sigma := (\log n)^{-2}, \tag{4}$$

and update the random set of edges by setting

$$E_{i+1} := E_i \cup \Gamma_{i+1}. \tag{5}$$

To determine the new triangle-free edge subset $T_{i+1} \subseteq T_i \cup \Gamma_{i+1}$, the idea is to delete a suitable set $D_{i+1} \subseteq \Gamma_{i+1}$ of edges from Γ_{i+1} with $|\Gamma_{i+1} \setminus D_{i+1}| \approx |\Gamma_{i+1}|$, such that

$$T_{i+1} := T_i \cup (\Gamma_{i+1} \setminus D_{i+1}) \quad (6)$$

remains triangle-free, see [19, (13)–(14) in Section 2.1] for the precise definition of D_{i+1} (this construction intuitively works since only few new triangles are created in $E_i \cup \Gamma_{i+1}$ due to the fact that Γ_{i+1} is fairly small). To determine the new open edge set $O_{i+1} \subseteq O_i \setminus \Gamma_{i+1}$, we certainly have to remove the set C'_{i+1} of ‘newly closed’ edges, which simply contains all edges $e \in O_i$ that form a triangle with some two edges of $E_{i+1} = E_i \cup \Gamma_{i+1}$. For technical reason we also remove an extra random edge subset $S_{i+1} \subseteq O_i$ and set

$$O_{i+1} := O_i \setminus (\Gamma_{i+1} \cup C'_{i+1} \cup S_{i+1}), \quad (7)$$

see [19, (15)–(20) in Section 2.1] for the precise definition of $C'_{i+1} \cup S_{i+1}$ (the removal of extra edges is a technical twist that intuitively makes it easier to prove certain concentration statements).

Stopping this iterative construction after $I \approx n^\beta$ steps, the pseudo-random intuition from [19, Section 2] suggests that, with respect to various edge statistics, the resulting n -vertex triangle-free graph

$$H := ([n], T_I) \quad \text{with} \quad I := \lceil n^\beta \rceil \quad (8)$$

heuristically resembles a binomial random graph $G(n, \rho)$ with edge probability

$$\rho := \sqrt{\beta(\log n)/n}, \quad (9)$$

but with the notable exception that it by construction contains no triangles (such a random graph would typically contain many triangles). This heuristic makes it plausible that $G_n = H$ satisfies the degree properties claimed by Theorem 1.3, since routine arguments show that the random graph $G(n, \rho)$ typically has these degree properties. To keep the modifications of [19] minimal, we shall in fact find an induced subgraph $G_n \subseteq H$ with the desired degree properties (this extra step is convenient but not necessary, see Remark 2.1).

2.2 Setup and proof of Theorem 1.3

We now turn to the technical details of our proof of Theorem 1.3, which extends [19, Sections 2–3]. Here our setup is guided by the pseudo-random heuristic discussed in [19, Section 2.2], which loosely suggests that

$$\mathbb{P}(e \in E_i) \approx \pi_i/\sqrt{n} \quad \text{and} \quad \mathbb{P}(e \in O_i) \approx q_i, \quad (10)$$

where the parameters π_i and q_i defined in [19, Section 2.3] satisfy the technical properties

$$\pi_i := \sigma + \sum_{0 \leq j < i} \sigma q_j, \quad 0 < q_i \leq 1 = q_0 \quad \text{and} \quad \pi_I/\sqrt{n} = (1 + o(1))\rho, \quad (11)$$

see [19, Section 2.3 and Lemma 17] for the full details, which formally justify that we may indeed use the parameters π_i and q_i and their properties (11) in this paper (in contrast to the heuristic approximations (10), which we of course may not use in our proofs). In particular, to get a handle on the number of edges between large sets of vertices, consistent with (10)–(11) we introduce the pseudo-random events

$$\mathcal{T}_I^* := \left\{ |T_I(A, B)| \geq (1 - \delta)|A||B|\rho \text{ for all disjoint } A, B \subseteq [n] \text{ with } |A| = |B| = s \right\}, \quad (12)$$

$$\mathcal{T}_I^+ := \left\{ |T_I(A, B)| \leq (1 + \delta)2s|B|\rho \text{ for all disjoint } A, B \subseteq [n] \text{ with } 1 \leq |A| = |B| \leq 2s \right\}, \quad (13)$$

where we write $S(A, B) := \{ab \in S : a \in A, b \in B\}$ for the set of edges from S that go between A and B , and use the carefully chosen (see [19, Section 2.3 and Theorem 9]) size parameter

$$s := \lceil D(\log n)/\rho \rceil \quad \text{with} \quad D := 108/\delta^2 \quad \text{and} \quad \delta := 1/10. \quad (14)$$

To eventually get a handle on the maximum degree, we also introduce the auxiliary event

$$\mathcal{N}_{\leq I} := \left\{ |N_{\Gamma_i}(v)| \leq 2\sigma q_{i-1}\sqrt{n} \text{ for all } v \in [n] \text{ and } 0 < i \leq I \right\}, \quad (15)$$

writing $N_S(v) := \{w \in [n] : vw \in S\}$ for the set of neighbors of v in a given edge set S .

Results of Guo and Warnke, see [19, Theorem 9], imply that

$$\mathbb{P}(\mathcal{T}_I^* \cap \mathcal{N}_{\leq I}) \geq 1 - n^{-\omega(1)}. \quad (16)$$

As we shall show next, Theorem 1.3 then follows from the claim

$$\mathbb{P}(\mathcal{T}_I^+) \geq 1 - o(1), \quad (17)$$

whose stochastic domination based proof we defer to Section 2.3.

Proof of Theorem 1.3 assuming inequality (17). Combining (16)–(17) we infer that $\mathbb{P}(\mathcal{T}_I^* \cap \mathcal{N}_{\leq I} \cap \mathcal{T}_I^+) > 0$ for all sufficiently large n , so by the probabilistic method we may henceforth fix a graph $H = ([n], T_I)$ for which the event $\mathcal{T}_I^* \cap \mathcal{N}_{\leq I} \cap \mathcal{T}_I^+$ holds. We then construct the induced triangle-free subgraph $G_n \subseteq H$ by iteratively deleting vertices of degree at most $\delta/4 \cdot n\rho$, and now verify that it has the claimed properties, starting with the degree bound (2). Noting $e_H(A, B) = |T_I(A, B)|$, the event \mathcal{T}_I^* implies, via a double-counting argument for $e(H)$ analogous to the proof of [19, Theorem 5], that the number of edges of H is at least

$$e(H) = \frac{\sum_{A \subseteq [n]: |A|=s} \sum_{B \subseteq [n] \setminus A: |B|=s} |T_I(A, B)|}{2 \binom{n-2}{s-1} \binom{n-s-1}{s-1}} \geq \frac{\binom{n}{s} \binom{n-s}{s} \cdot (1-\delta)s^2\rho}{2 \binom{n-2}{s-1} \binom{n-s-1}{s-1}} = (1-\delta) \binom{n}{2} \rho. \quad (18)$$

Furthermore, by the recursive definition (6) of the edge set $T_I \subseteq \bigcup_{0 \leq i < I} \Gamma_{i+1}$, using the properties (11) of π_i we infer, for all sufficiently large $n \geq n_0(\beta)$, that the event $\mathcal{N}_{\leq I}$ implies the maximum degree bound

$$\Delta(H) = \max_{v \in [n]} |N_{T_I}(v)| \leq \max_{v \in [n]} \sum_{0 \leq i < I} |N_{\Gamma_{i+1}}(v)| \leq \sum_{0 \leq i < I} 2\sigma q_i \sqrt{n} \leq 2\pi_I \sqrt{n} \leq (2+\delta)\rho n.$$

By construction of $G_n \subseteq H$, using $\delta = 1/10$ we thus infer, for all sufficiently large $n \geq n_0(\beta)$, that

$$v(G_n) \geq \frac{2e(G_n)}{\Delta(G_n)} \geq \frac{2[e(H) - n \cdot \delta/4 \cdot n\rho]}{\Delta(H)} \geq \frac{2(1-2\delta) \binom{n}{2} \rho}{(2+\delta)n\rho} > \frac{n}{3},$$

and so the claimed degree bound (2) follows with $c := \delta/4$ and $C := 2 + \delta$.

Next, suppose that $F \subseteq G_n$ is an induced bipartite subgraph with two parts A and B , where we may assume that $|A| \geq |B| \geq 1$. Since $F \subseteq G_n$ and $G_n \subseteq H$ are both induced subgraphs, we have

$$e_F(A, B) = e_{G_n}(A, B) = e_H(A, B) = |T_I(A, B)|.$$

Furthermore, since A and B are both independent sets in F , we have $|B| \leq |A| \leq \alpha(F) \leq \alpha(H) \leq 2 \cdot s$, where the last inequality holds because the event \mathcal{T}_I^* implies that in H there is at least one edge between any two disjoint s -vertex sets. Using a double counting argument similar to (18), the event \mathcal{T}_I^+ then implies that

$$|T_I(A, B)| = \frac{\sum_{A' \subseteq A: |A'|=|B|} |T_I(A', B)|}{\binom{|A|-1}{|B|-1}} \leq \frac{\binom{|A|}{|B|} \cdot (1+\delta)2s|B|\rho}{\binom{|A|-1}{|B|-1}} = (1+\delta)2s|A|\rho.$$

The definitions (14) of $s \approx D(\log n)/\rho$ and $\delta = 1/10$ give $(1+\delta)2s\rho \leq 3D \log n$ for sufficiently large n . By averaging it follows that $\delta(F) \leq e_F(A, B)/|A| \leq 3D \log n$, completing the proof with $C' := 3D$. \square

Remark 2.1. For any $\delta > 0$ one can in fact show that the minimum and maximum degree of the n -vertex graph $H = ([n], T_I)$ satisfy $(1-\delta)n\rho \leq \delta(H) \leq \Delta(H) \leq (1+\delta)n\rho$ with high probability (by adapting [19, Sections 3.1–3.5]), which would allow us to directly use $G_n = H$ in the above proof of Theorem 1.3. However, the coarser bounds used above suffice for our purposes, and require less technical modifications of [19].

2.3 Pseudo-randomness: deferred proof of inequality (17)

This subsection is devoted to the deferred proof of inequality (17), i.e., $\mathbb{P}(\mathcal{T}_I^+) \geq 1 - o(1)$. To this end we shall adapt the strategy from [19, Sections 3.4–3.5] to our setting, i.e., use estimates on the number of open edges $|O_i(A, B)|$ to eventually get a handle on the total number of added edges $|T_I(A, B)|$.

Turning to the details, let \mathcal{S} denote the set of all pairs of vertex disjoint $A, B \subseteq [n]$ with $1 \leq |A| = |B| \leq 2s$. To keep the changes to [19] minimal, for each pair $(A, B) \in \mathcal{S}$ we enlarge A to A^+ by adding the lexicographic first $2s - |A|$ vertices from $[n] \setminus (A \cup B)$. Note that the vertex set A^+ is determined by A . Consistent with the heuristic approximations (10), we then introduce the ‘open edges’ related pseudo-random events

$$\tilde{\mathcal{Q}}_i^+ := \left\{ |O_i(A^+, B)| \leq q_i |A^+| |B| \text{ for all } (A, B) \in \mathcal{S} \right\} \quad \text{and} \quad \tilde{\mathcal{Q}}_{\leq I}^+ := \bigcap_{0 \leq i \leq I} \tilde{\mathcal{Q}}_i^+. \quad (19)$$

Note that $1 \leq |B| \leq |A^+| = 2s$ for all pairs $(A, B) \in \mathcal{S}$. Furthermore, there are at most n^{2j} pairs $(A, B) \in \mathcal{S}$ with $|B| = j$. With these two key properties in mind, the proof of [19, Lemma 24] carries over to the pairs (A^+, B) virtually unchanged (that proof merely exploits that $|A|$ is large, and only uses $|A| = |B|$ to control the final union bound estimate over all pairs (A, B) of vertex subsets), giving

$$\max_{0 \leq i < I} \mathbb{P}(\neg \tilde{\mathcal{Q}}_{i+1}^+ \mid \tilde{\mathcal{Q}}_{\leq i}^+ \cap \mathfrak{X}_{\leq i}) \leq \sum_{(A, B) \in \mathcal{S}} n^{-\omega(|B|)} \leq \sum_{1 \leq j \leq 2s} n^{2j - \omega(j)} \leq n^{-\omega(1)},$$

where $\mathfrak{X}_{\leq i}$ is a ‘good’ event determined by $(O_j, E_j, T_j, \Gamma_j, S_j)_{0 \leq j \leq i}$ that is formally defined in [19, Section 2.4]; here we shall only use that the event $\mathfrak{X}_{\leq i+1}$ implies $\mathfrak{X}_{\leq i}$, and that $\mathbb{P}(\neg \mathfrak{X}_{\leq I}) \leq n^{-\omega(1)}$ by [19, Theorem 9]. In view of $q_0 = 1$, see (11), it is straightforward to check that the event $\tilde{\mathcal{Q}}_0^+ = \tilde{\mathcal{Q}}_{\leq 0}^+$ always holds. Since the event $\mathfrak{X}_{\leq I}$ implies $\mathfrak{X}_{\leq i}$ for all $0 \leq i < I$, using $\tilde{\mathcal{Q}}_{\leq i+1}^+ = \tilde{\mathcal{Q}}_{i+1}^+ \cap \tilde{\mathcal{Q}}_{\leq i}^+$ and $I \approx n^\beta$ it follows that

$$\begin{aligned} \mathbb{P}(\neg \tilde{\mathcal{Q}}_{\leq I}^+) &\leq \mathbb{P}(\neg \mathfrak{X}_{\leq I}) + \mathbb{P}(\neg \tilde{\mathcal{Q}}_{i+1}^+ \cap \tilde{\mathcal{Q}}_{\leq i}^+ \cap \mathfrak{X}_{\leq i} \text{ for some } 0 \leq i < I) \\ &\leq \mathbb{P}(\neg \mathfrak{X}_{\leq I}) + \sum_{0 \leq i < I} \mathbb{P}(\neg \tilde{\mathcal{Q}}_{i+1}^+ \mid \tilde{\mathcal{Q}}_{\leq i}^+ \cap \mathfrak{X}_{\leq i}) \leq (I+1) \cdot n^{-\omega(1)} \leq n^{-\omega(1)}. \end{aligned} \quad (20)$$

Turning to the total number of added edges $|T_I(A, B)|$ for $(A, B) \in \mathcal{S}$, using $A \subseteq A^+$ and $T_I \subseteq E_I$ together with the recursive definition (5) of the edge set $E_I = \bigcup_{0 \leq i < I} \Gamma_{i+1}$, it follows that

$$|T_I(A, B)| \leq |E_I(A^+, B)| = \sum_{0 \leq i < I} |O_i(A^+, B) \cap \Gamma_{i+1}|. \quad (21)$$

Recall that the event $\tilde{\mathcal{Q}}_i^+$ implies $|O_i(A^+, B)| \leq q_i |A^+| |B|$, and that $\Gamma_{i+1} \subseteq O_i$ is the random subset where each edge $e \in O_i$ is included independently with probability p . Combining these properties, by mimicking the stochastic domination arguments from the proof of [19, Claim 30] it then follows that

$$\mathbb{P}(|E_I(A^+, B)| \geq t \text{ and } \tilde{\mathcal{Q}}_{\leq I}^+) \leq \mathbb{P}(Z^+ \geq t) \quad \text{with} \quad Z^+ \stackrel{d}{=} \text{Bin}\left(\sum_{0 \leq i < I} \lfloor q_i |A^+| |B| \rfloor, p\right).$$

Using $p = \sigma/\sqrt{n}$, the properties (11) of π_i and $|A^+| = 2s$, similar to [19, Section 3.5] we have

$$\mu^+ := \mathbb{E}Z^+ \sim \sum_{0 \leq i < I} \sigma q_i / \sqrt{n} \cdot |A^+| |B| = (\pi_I - \sigma) / \sqrt{n} \cdot |A^+| |B| \sim 2s |B| \rho.$$

Using the definitions (14) of the parameter $s \approx D(\log n)/\rho$ and the constant $D = 108/\delta^2$, for sufficiently large n it follows that $(1 + \delta)2s|B|\rho \geq (1 + \delta/2)\mu^+$ and $\delta^2\mu^+/12 > 12|B|\log n$, say. Similar to [19, (97)–(98)], standard Chernoff bounds such as [20, Theorem 2.1] thus routinely give

$$\mathbb{P}(|E_I(A^+, B)| \geq (1 + \delta)2s|B|\rho \text{ and } \tilde{\mathcal{Q}}_{\leq I}^+) \leq \mathbb{P}(Z^+ \geq (1 + \delta/2)\mu^+) \leq \exp(-\delta^2\mu^+/12) \leq n^{-12|B|}.$$

Recalling that the vertex set A^+ is determined by A , and that there are at most most n^{2j} pairs $(A, B) \in \mathcal{S}$ with $|B| = j$, in view of inequality (21) it then follows via a standard union bound argument that

$$\mathbb{P}(\neg \mathcal{T}_I^+ \cap \tilde{\mathcal{Q}}_{\leq I}^+) \leq \sum_{(A, B) \in \mathcal{S}} n^{-12|B|} \leq \sum_{1 \leq j \leq 2s} n^{2j - 12j} = o(n^{-9}),$$

which together with (20) implies $\mathbb{P}(\mathcal{T}_I^+) \geq 1 - o(1)$. This completes the proof of inequality (17) and thus Theorem 1.3, as discussed. \square

3 Random greedy r -AP free process

In this section we prove Theorem 1.5 by showing that the random greedy r -AP free process typically finds an r -AP free subset $I \subseteq \mathbb{Z}/N\mathbb{Z}$ with the desired properties. Intuitively, this process starts with an empty set $I = \emptyset$, and then iteratively adds new random numbers from $\mathbb{Z}/N\mathbb{Z}$ such that the resulting set I stays r -AP free. More formally, fixing $r \geq 3$, the random greedy r -AP free process starts with

$$I(0) := \emptyset \quad \text{and} \quad S(0) := \mathbb{Z}/N\mathbb{Z}. \quad (22)$$

Here $I(i)$ denotes the growing r -AP free set found after i steps, and $S(i)$ denotes the set of ‘available’ numbers in $\mathbb{Z}/N\mathbb{Z} \setminus I(i)$, i.e., that can be added to $I(i)$ without creating an r -AP. In step $i + 1 \geq 1$ of the random greedy r -AP free process, we then choose $x_{i+1} \in S(i)$ uniformly at random and update the r -AP free set and available set via

$$I(i+1) := I(i) \cup \{x_{i+1}\}, \quad (23)$$

$$S(i+1) := S(i) \setminus (\{x_{i+1}\} \cup Y_{x_{i+1}}(i)), \quad (24)$$

where we write $Y_{x_{i+1}}(i)$ for the set of numbers that become ‘unavailable’ when x_{i+1} is added, i.e.,

$$Y_x(i) := \{y \in S(i) \setminus \{x\} : \text{there is } A \in \mathcal{A}_{N,r} \text{ such that } x, y \in A \text{ and } A \setminus \{x, y\} \subseteq I(i)\}, \quad (25)$$

in which $\mathcal{A}_{N,\ell}$ is a shorthand for the collection of all ℓ -APs in $\mathbb{Z}/N\mathbb{Z}$.

3.1 Proof strategy

In this subsection we discuss our proof strategy for Theorem 1.5. To this end, let us first record the basic observation that each number $x \in \mathbb{Z}/N\mathbb{Z}$ is contained in exactly

$$D := r|\mathcal{A}_{N,r}|/N = \Theta(N) \quad (26)$$

many r -APs $A \in \mathcal{A}_{N,r}$. Our strategy is then to analyze the random greedy r -AP free process for

$$m := \xi \cdot ND^{-\frac{1}{r-1}} (\log N)^{\frac{1}{r-1}} \quad (27)$$

steps, and show that the r -AP free set $I := I(m) \subseteq \mathbb{Z}/N\mathbb{Z}$ typically satisfies $I \cap K \neq \emptyset$ for all k -APs $K \in \mathcal{A}_{N,k}$ of size

$$k := 9\xi^{-1} \cdot (D/\log N)^{1/(r-1)} \log N = \Theta((N/\log N)^{1/(r-1)} \log N), \quad (28)$$

deferring the choices of the sufficiently small constants $0 < \xi, \delta < 1/(2r)$. As usual, we are henceforth treating both m and k as integers (since rounding has an asymptotically negligible effect on our arguments).

The outlined proof strategy is consistent with the pseudo-random heuristic that $I = I(m)$ resembles a random m -element subset of $\mathbb{Z}/N\mathbb{Z}$. Indeed, noting $km = 9N \log N$, this heuristic suggests that

$$\mathbb{P}(I \cap K = \emptyset) \approx \frac{\binom{N-k}{m}}{\binom{N}{m}} = \prod_{0 \leq j < k} \left(1 - \frac{m}{N-j}\right) \leq \exp\left(-\frac{km}{N}\right) \ll N^{-2},$$

which is small enough to employ a union bound argument over the at most N^2 many k -APs $K \in \mathcal{A}_{N,k}$. In (37) below and Section 3.3 we will essentially make this heuristic reasoning rigorous, albeit in a slightly roundabout way (via several pseudo-random events and the differential equation method).

3.2 Setup and proof of Theorem 1.5

We now turn to the technical details of our proof of Theorem 1.5, which require some setup. In order to relate the discrete steps of the process to continuous trajectories, we introduce the convenient scaling

$$t_i := i/M \quad \text{with} \quad M := ND^{-\frac{1}{r-1}}. \quad (29)$$

To get a handle on all k -APs $K \in \mathcal{A}_{N,k}$, we denote the number of available numbers in K by

$$S_K(i) := S(i) \cap K. \quad (30)$$

Henceforth using the shorthand $X = (a \pm b)x$ for $X \in [(a-b)x, (a+b)x]$ to avoid clutter (as usual), we then define $\mathcal{K}_{\leq j}$ as the pseudo-random event that for all $0 \leq i \leq j$ we have

$$|S_K(i)| = (1 \pm e(t_i))kq(t_i) \quad \text{for all } K \in \mathcal{A}_{N,k}, \quad (31)$$

and similarly define $\mathcal{S}_{\leq j}$ as the pseudo-random event that for all $0 \leq i \leq j$ we have

$$|S(i)| = (1 \pm D^{-\delta})Nq(t_i) \quad \text{and} \quad \max_{x \in S(i)} |Y_x(i) - s_2(t_i)| \leq D^{\frac{1}{r-1}-\delta}, \quad (32)$$

using the deterministic functions

$$q(t) := e^{-t^{r-1}}, \quad s_2(t) := (r-1)D^{\frac{1}{r-1}}t^{r-2}q(t) \quad \text{and} \quad e(t) := e^{5(t+t^{r-1})} \cdot D^{-\delta}. \quad (33)$$

Note that, by choosing $\xi = \xi(r, \delta) > 0$ small enough compared to $r, \delta > 0$, we may assume that for all steps $0 \leq i \leq m$ we have $0 \leq t_i \leq t_m = m/M = \xi(\log N)^{\frac{1}{r-1}}$ as well as

$$0 < D^{-\delta} \leq e(t) = o(1) \quad \text{and} \quad 0 \leq t \leq D^{o(1)} \quad \text{for} \quad 0 \leq t \leq t_m. \quad (34)$$

Results of Bohman and Bennett (which require N to be prime), see [1, Section 4], imply that for sufficiently² small $\xi, \delta > 0$ we have

$$\mathbb{P}(\neg \mathcal{S}_{\leq m}) \leq \exp(-N^{\Omega(1)}). \quad (35)$$

Using estimates (32) and (34), we see that the event $\mathcal{S}_{\leq m}$ implies $\min_{0 \leq i \leq m} |S(i)| = |S(m)| \geq Nq(t_m)/2 > 0$ for all sufficiently large n , which in turn ensures that the random greedy r -AP free process does not terminate before step m (since the process is always able to select a new number in each of the first m steps). As we shall show next, Theorem 1.5 then follows easily from the claim

$$\mathbb{P}(\neg \mathcal{G}_{\leq m}) = o(1) \quad \text{for} \quad \mathcal{G}_{\leq i} := \mathcal{S}_{\leq i} \cap \mathcal{K}_{\leq i}, \quad (36)$$

whose differential equation method based proof we defer to Section 3.3.

Proof of Theorem 1.5 assuming inequality (36). For any k -AP $K \in \mathcal{A}_{N,k}$ in $\mathbb{Z}/N\mathbb{Z}$, whenever the event $\mathcal{G}_{\leq i}$ holds, by combining the concentration bounds (31)–(32) with the error estimate (34) we infer that

$$\mathbb{P}(x_{i+1} \notin S(i) \cap K \mid \mathcal{F}_i) = 1 - \frac{|S_K(i)|}{|S(i)|} \leq 1 - \frac{\frac{1}{2}kq(t_i)}{2Nq(t_i)} = 1 - \frac{k}{4N},$$

where \mathcal{F}_i denotes the natural filtration associated with the process after i steps (which intuitively keeps track of the ‘history’ of the process, i.e., all the information available up to and including step i). Since the event $\mathcal{G}_{\leq m}$ implies the event $\mathcal{G}_{\leq i}$ for all $0 \leq i \leq m$, using $km = 9N \log N$ it routinely follows that

$$\mathbb{P}(I(m) \cap K = \emptyset \quad \text{and} \quad \mathcal{G}_{\leq m}) \leq \prod_{0 \leq i \leq m-1} \left(1 - \frac{k}{4N}\right) \leq \exp\left(-\frac{km}{4N}\right) \ll N^{-2}. \quad (37)$$

Taking a union bound over the at most N^2 many k -APs $K \in \mathcal{A}_{N,k}$ in $\mathbb{Z}/N\mathbb{Z}$ then completes the proof of Theorem 1.5 with $I := I(m)$, since $\mathbb{P}(\neg \mathcal{G}_{\leq m}) = o(1)$ by the assumed inequality (36). \square

²For example, the explicit choices $\delta = 1/(40r^2)$, $\xi = \delta/500$ satisfy all constraints of this paper and [1].

3.3 Dynamic concentration: deferred proof of inequality (36)

This subsection is devoted to the deferred proof of inequality (36), i.e., $\mathbb{P}(-\mathcal{G}_{\leq m}) = o(1)$, which in view of the probability estimate (35) and the definition of the event $\mathcal{K}_{\leq i}$ requires us to establish the dynamic concentration estimate (31) for $|S_K(i)|$. To this end, following the differential equation method approach to dynamic concentration [34, 2, 32], for all k -APs $K \in \mathcal{A}_{N,k}$, steps $0 \leq i \leq m$ and sign-parameters $\sigma \in \{+, -\}$ we introduce the auxiliary random variables

$$X_K^\sigma(i) := \sigma[|S_K(i)| - kq(t_i)] - kq(t_i)e(t_i). \quad (38)$$

The point is that the desired estimate (31) follows when the inequalities $X_K^+(i) \leq 0$ and $X_K^-(i) \leq 0$ both hold. In the following we shall use supermartingale arguments to establish these inequalities, by analyzing the (expected and worst-case) one-step changes of $X_K^\sigma(i)$ and $|S_K(i)|$.

3.3.1 Expected one-step changes

We start by estimating the expected value of the one-step changes $\Delta S_K(i) := |S_K(i+1)| - |S_K(i)|$ of the number of available numbers in any k -AP $K \in \mathcal{A}_{N,k}$, assuming that $0 \leq i < m$ and $\mathcal{G}_{\leq i}$ hold. Note that $|S_K(i)|$ is monotone decreasing. Furthermore, a number $x \in S_K(i)$ is removed from the set of available numbers if the process chooses a number x_{i+1} from $Y_x(i) \cup \{x\}$. Since $x_{i+1} \in S(i)$ is chosen uniformly at random, using the estimates (31)–(32) implied by $\mathcal{G}_{\leq i}$ it follows that

$$\mathbb{E}(\Delta S_K(i) \mid \mathcal{F}_i) = - \sum_{x \in S_K(i)} \frac{|Y_x(i)| \pm 1}{|S(i)|} = \frac{-[1 \pm e(t_i)]kq(t_i) \cdot [s_2(t_i) \pm 2D^{\frac{1}{r-1}-\delta}]}{[1 \pm D^{-\delta}]Nq(t_i)}. \quad (39)$$

Recalling that $0 < D^{-\delta} \leq e(t_i) = o(1)$ by (34) and that $q(t) = e^{-t^{r-1}}$ by (33), using $s_2(t_i)/D^{\frac{1}{r-1}} = (r-1)t_i^{r-2}q(t_i) = -q'(t_i)$ and $D^{\frac{1}{r-1}}/N = 1/M$ it follows that

$$\begin{aligned} \mathbb{E}(\Delta S_K(i) \mid \mathcal{F}_i) &= - \left(1 \pm 4e(t_i)\right) \left(s_2(t_i) \pm 2D^{\frac{1}{r-1}-\delta}\right) \frac{k}{N} \\ &= \frac{kq'(t_i)}{M} \pm \left(4(r-1)t_i^{r-2} \cdot q(t_i)e(t_i) + 4D^{-\delta}\right) \frac{k}{M}. \end{aligned} \quad (40)$$

In preparation for the upcoming supermartingale arguments, we now show that the expected values of the one-step changes $\Delta X_K^\sigma(i) := |X_K^\sigma(i+1)| - |X_K^\sigma(i)|$ of the auxiliary variables are negative, again assuming that $0 \leq i < m$ and $\mathcal{G}_{\leq i}$ hold. Set $f(t) := q(t)e(t)$. Recalling the shorthand $t_i = i/M$ and the definition (38) of $X_K^\sigma(i)$, by applying Taylor's theorem with remainder to the functions q and f , it follows that

$$\begin{aligned} \mathbb{E}(\Delta X_K^\sigma(i) \mid \mathcal{F}_i) &= \sigma \left[\mathbb{E}(\Delta S_K(i) \mid \mathcal{F}_i) - k[q(t_{i+1}) - q(t_i)] \right] - k[f(t_{i+1}) - f(t_i)] \\ &= \sigma \left[\mathbb{E}(\Delta S_K(i) \mid \mathcal{F}_i) - \frac{kq'(t_i)}{M} \right] - \frac{kf'(t_i)}{M} + O\left(\max_{0 \leq t \leq t_m} \frac{k(|q''(t)| + |f''(t)|)}{M^2} \right). \end{aligned} \quad (41)$$

Using (40) we see that in (41) the main $kq'(t_i)/M$ term cancels up to second order terms. In the following we shall show that the main error term $-kf'(t_i)/M$ is large enough to make the expected change (41) negative. Indeed, noting $f(t) = q(t)e(t) \geq D^{-\delta}$, we have

$$f'(t_i) = (5 + 4(r-1)t_i^{r-2})q(t_i)e(t_i) \geq 5D^{-\delta} + 4(r-1)t_i^{r-2} \cdot q(t_i)e(t_i).$$

Furthermore, $D = \Theta(N)$ and $\delta \leq 1/(2r)$ imply $M = N/D^{\frac{1}{r-1}} \gg D^{2\delta}$. Recalling that $q(t) \leq 1$, $f(t) \leq e(t) \ll 1$ and $0 \leq t \leq D^{o(1)}$, see (34), it then routinely follows that

$$\frac{|q''(t)| + |f''(t)|}{M} \leq \frac{O(\sum_{0 \leq j \leq 2r} t^j) \cdot [q(t) + f(t)]}{M} \leq \frac{D^{o(1)}}{D^{2\delta}} \ll D^{-\delta}.$$

Inserting these estimates and (40) into the expected one-step changes (41) of $X_K^\sigma(i)$, it follows that

$$\mathbb{E}(\Delta X_K^\sigma(i) \mid \mathcal{F}_i) \leq -(1 - o(1))kD^{-\delta}/M < 0. \quad (42)$$

3.3.2 Bounds on the one-step changes

We next bound the expected values of the one-step changes $|\Delta S_K(i)| = ||S_K(i+1)| - |S_K(i)||$, assuming that $0 \leq i < m$ and $\mathcal{G}_{\leq i}$ hold. Since $|S_K(i)|$ is step-wise decreasing, by combining $s_2(t_i) \leq rD^{\frac{1}{r-1}}t_i^{r-2}$ with the first estimate of the expected one-step changes (40), using $e(t_i) = o(1)$ and $0 \leq t_i \leq D^{o(1)}$ it follows that

$$\mathbb{E}(|\Delta S_K(i)| \mid \mathcal{F}_i) = -\mathbb{E}(\Delta S_K(i) \mid \mathcal{F}_i) \leq O\left(D^{\frac{1}{r-1}}t_i^{r-2} + D^{\frac{1}{r-1}-\delta}\right) \cdot \frac{k}{N} \ll kD^{\frac{1}{r-1}+\delta/2}/N. \quad (43)$$

Turning to the worst-case one-step changes of $|S_K(i)|$, we introduce the auxiliary event

$$\mathcal{N}_{\leq j} := \left\{ \max_{x \in S(i)} |Y_x(i) \cap K| \leq D^{\frac{1}{r-1}-3\delta} \text{ for all } K \in \mathcal{A}_{N,k} \text{ and } 0 \leq i \leq j \right\}. \quad (44)$$

Recalling the reasoning leading to (39), the crux is that when $\mathcal{N}_{\leq i}$ holds, then we have

$$|\Delta S_K(i)| \leq 1 + \max_{x \in S(i)} |Y_x(i) \cap K| \leq 2D^{\frac{1}{r-1}-3\delta}. \quad (45)$$

We now claim that the auxiliary event $\mathcal{N}_{\leq m}$ typically holds, i.e., more precisely that

$$\mathbb{P}(\neg \mathcal{N}_{\leq m} \text{ and } \mathcal{S}_{\leq m}) \leq \exp(-N^{\Omega(1)}). \quad (46)$$

Turning to the proof details, with an eye on $|Y_x(i) \cap K|$ let

$$\mathcal{I} = \mathcal{I}(K, x) := \{W : |W| = r-2, W \cup \{x, y\} \in \mathcal{A}_{N,r} \text{ for some } y \in K\}. \quad (47)$$

Note that $|\mathcal{I}| \leq kr^2$, as there are at most r^2 many r -APs containing two distinct numbers $\{x, y\}$. Let

$$N_{K,x} := \sum_{W \in \mathcal{I}} Y_W \quad \text{with} \quad Y_W := \mathbb{1}_{\{W \subseteq I(m) \text{ and } \mathcal{S}_{\leq m}\}}. \quad (48)$$

Since $\{x\} \cup W$ contains $r-1 \geq 2$ elements, by similar reasoning as for $|\mathcal{I}|$ it follows that

$$\max_{0 \leq i \leq m} |Y_x(i) \cap K| \cdot \mathbb{1}_{\{\mathcal{S}_{\leq m}\}} \leq N_{K,x} \cdot r^2. \quad (49)$$

We shall bound $N_{K,x}$ via the following Chernoff-type upper tail estimate for combinatorial random variables with ‘controlled dependencies’, which is a convenient corollary of [33, Theorem 7 and Remarks 9–10].

Lemma 3.1. *Let $(Y_\alpha)_{\alpha \in \mathcal{I}}$ be a finite family of variables with $Y_\alpha \in [0, 1]$ and $\sum_{\alpha \in \mathcal{I}} \lambda_\alpha \leq \mu$, where $(\lambda_\alpha)_{\alpha \in \mathcal{I}}$ satisfies $\mathbb{E}(\prod_{i \in [s]} Y_{\alpha_i}) \leq \prod_{i \in [s]} \lambda_{\alpha_i}$ for all $(\alpha_1, \dots, \alpha_s) \in \mathcal{I}^s$ with $\alpha_i \cap \alpha_j = \emptyset$ for $i \neq j$. Set $Y := \sum_{\alpha \in \mathcal{I}} Y_\alpha$. If $\max_{\alpha \in \mathcal{I}} |\{\beta \in \mathcal{I} : \alpha \cap \beta \neq \emptyset\}| \leq C$, then $\mathbb{P}(Y \geq z) \leq (e\mu/z)^{z/C}$ for all $z > \mu$.*

With an eye on $\mathbb{E}N_{K,x}$, we first record the basic observation that when $\mathcal{S}_{\leq m}$ holds, then in every step $i \leq m$ there are at least $|S(i)| \geq |S(m)| \gg ND^{-\delta/4}$ available numbers, say. For any set U of numbers from $\mathbb{Z}/N\mathbb{Z}$ a straightforward adaptation of the proof of [3, Lemma 4.1] (which proceeds by taking a union bound over all possible steps where the numbers of U could appear) then ensures that

$$\mathbb{P}(U \subseteq I(m) \text{ and } \mathcal{S}_{\leq m}) \leq m^{|U|} \cdot \left(\frac{1}{ND^{-\delta/4}}\right)^{|U|} \leq \pi^{|U|} \quad \text{with} \quad \pi := D^{-\frac{1}{r-1}+\delta/2}, \quad (50)$$

where we used that $m/N = D^{-\frac{1}{r-1}}(\log N)^{O(1)} \ll D^{-\frac{1}{r-1}+\delta/4}$, say. In particular, for any sequence of sets $(W_1, \dots, W_s) \in \mathcal{I}^s$ satisfying $W_i \cap W_j = \emptyset$ for $i \neq j$, using the definition (48) of Y_W and (50) it follows that

$$\mathbb{E}\left(\prod_{i \in [s]} Y_{W_i}\right) = \mathbb{P}\left(\bigcup_{i \in [s]} W_i \subseteq I(m) \text{ and } \mathcal{S}_{\leq m}\right) \leq \pi^{s(r-2)} = \prod_{i \in [s]} \lambda_{W_i} \quad \text{with} \quad \lambda_W := \pi^{r-2}.$$

Furthermore, combining $|\mathcal{I}| \leq kr^2$ with (50) and the definition (28) of k , it also follows that

$$\sum_{W \in \mathcal{I}} \lambda_W \leq kr^2 \cdot \pi^{r-2} = 9\xi^{-1}r^2(\log n)^{1-\frac{1}{r-1}}D^{\frac{3-r}{r-1}}D^{(r-2)\delta/2} \ll D^{\frac{1}{r-1}}D^{-4\delta} =: \mu.$$

To estimate the associated C -parameter of Lemma 3.1, note that any set $W \in \mathcal{I}$ satisfies

$$|\{W' \in \mathcal{I} : W \cap W' \neq \emptyset\}| \leq \sum_{w \in W} \sum_{A \in \mathcal{A}_{N,r} : \{x,w\} \subseteq A} 2^{|A|} \leq r \cdot r^2 \cdot 2^r =: C.$$

Using inequality (49), by invoking Lemma 3.1 with $z := \mu D^\delta / r^2 \geq D^{\Omega(1)}$ it follows that

$$\mathbb{P}\left(\max_{0 \leq i \leq m} |Y_x(i) \cap K| \geq D^{\frac{1}{r-1}} D^{-3\delta} \text{ and } \mathcal{S}_{\leq m}\right) \leq \mathbb{P}(N_{K,x} \geq z) \leq (e\mu/z)^{z/C} \leq \exp(-N^{\Omega(1)}). \quad (51)$$

Taking a union bound over all of the at most $N \cdot N^2 = N^{O(1)}$ possible pairs (x, K) then establishes the claimed inequality (46).

3.3.3 Supermartingale arguments

We are now ready to prove $\mathbb{P}(-\mathcal{G}_{\leq m}) \leq \exp(-N^{\Omega(1)})$, by showing that $X_K^\sigma(i) \geq 0$ is extremely unlikely. Here our main probabilistic tool is the following supermartingale inequality [18, Lemma 19], which allows us to exploit that $X_K^\sigma(i)$ is defined (38) as the sum of a random variable and a deterministic function.

Lemma 3.2. *Let $(S_i)_{i \geq 0}$ be a supermartingale adapted to the filtration $(\mathcal{F}_i)_{i \geq 0}$. Assume that $S_i = X_i + D_i$, where X_i is \mathcal{F}_i -measurable and D_i is $\mathcal{F}_{\max\{i-1, 0\}}$ -measurable. Writing $\Delta X_i := X_{i+1} - X_i$, assume that $\max_{i \geq 0} |\Delta X_i| \leq C$ and $\sum_{i \geq 0} \mathbb{E}(|\Delta X_i| \mid \mathcal{F}_i) \leq V$. Then, for all $z > 0$,*

$$\mathbb{P}(S_i \geq S_0 + z \text{ for some } i \geq 0) \leq \exp\left(-\frac{z^2}{2C(V+z)}\right). \quad (52)$$

Turning to the details, we define the stopping time T as the minimum of m and the first step $i \geq 0$ where the ‘good’ event $\mathcal{G}_{\leq i} \cap \mathcal{N}_{\leq i}$ fails. For brevity, set $i \wedge T := \min\{i, T\}$. Recalling the definition (36) of the event $\mathcal{G}_{\leq m} = \mathcal{S}_{\leq m} \cap \mathcal{K}_{\leq m}$, by the discussion below (38) it follows that

$$\begin{aligned} \mathbb{P}(-\mathcal{G}_{\leq m}) &= \mathbb{P}(-\mathcal{S}_{\leq m} \cup \neg \mathcal{N}_{\leq m}) + \mathbb{P}(\neg \mathcal{K}_{\leq m} \cap \mathcal{S}_{\leq m} \cap \mathcal{N}_{\leq m}) \\ &\leq \mathbb{P}(\neg \mathcal{S}_{\leq m}) + \mathbb{P}(\neg \mathcal{N}_{\leq m} \text{ and } \mathcal{S}_{\leq m}) + \sum_{\sigma \in \{+, -\}} \sum_{K \in \mathcal{A}_{N,k}} \mathbb{P}\left(X_K^\sigma(i \wedge T) \geq 0 \text{ for some } i \geq 0\right). \end{aligned} \quad (53)$$

For any $K \in \mathcal{A}_{N,k}$, we initially have $S_K(0) = |K| = k$. By definition (38) of $X_K^\sigma(i)$ we thus have

$$X_K^\sigma(0 \wedge T) = X_K^\sigma(0) = \sigma[|S_K(0)| - k] - k\epsilon(0) = -kD^{-\delta}. \quad (54)$$

Note that the estimates in Sections 3.3.1–3.3.2 apply for $0 \leq i \leq T-1$ (since then $0 \leq i \leq m-1$ and $\mathcal{G}_{\leq i} \cap \mathcal{N}_{\leq i}$ hold), which in particular implies that the expected one-step changes of $X_K^\sigma(i)$ satisfy (42), and that the one-step changes $|\Delta S_K(i)|$ satisfy the worst case bound (45) and the expectation bound (43). Recalling (54), the stopped sequence $S_i := X_K^\sigma(i \wedge T)$ thus is a supermartingale with $S_0 = -kD^{-\delta}$, to which Lemma 3.2 can be applied with $X_i = \sigma|S_K(i \wedge T)|$, $C = O(D^{\frac{1}{r-1}-3\delta})$ and $V = m \cdot kD^{\frac{1}{r-1}+\delta/2}/N = O(kD^{2\delta/3})$. Invoking inequality (52) with $z = kD^{-\delta}$, using the definition (28) of k and $D = \Theta(N)$ it follows that

$$\mathbb{P}\left(X_K^\sigma(i \wedge T) \geq 0 \text{ for some } i \geq 0\right) \leq \exp(-\Omega(kD^{\delta/3}/D^{\frac{1}{r-1}})) \leq \exp(-N^{\Omega(1)}). \quad (55)$$

Inserting (55) and $|\mathcal{A}_{N,k}| \leq N^2$ into (53), then $\mathbb{P}(-\mathcal{G}_{\leq m}) \leq \exp(-N^{\Omega(1)})$ follows from (35) and (46), which completes the proof of inequality (36) and thus Theorem 1.5, as discussed. \square

3.4 Generalization of Theorem 1.5

We close by recording that a minor variant of our proof yields the following generalization of Theorem 1.5.

Theorem 3.3. *For any fixed $r \geq 3$ and $c > 0$, there are constants $C, N_0 > 0$ such that the following holds for any prime number $N \geq N_0$, setting $k = k_N := \lceil C(N/\log N)^{1/(r-1)} \log N \rceil$. For any family $\mathcal{B}_{N,k}$ of k -element subsets of $\mathbb{Z}/N\mathbb{Z}$ satisfying $|\mathcal{B}_{N,k}| \leq N^c$ there exists a set $I \subseteq \mathbb{Z}/N\mathbb{Z}$ which (i) is r -AP free in $\mathbb{Z}/N\mathbb{Z}$ and (ii) satisfies $|I \cap K| \geq 1$ for all $K \in \mathcal{B}_{N,k}$.*

Proof-Sketch. Our differential equation method based proof of inequality (36) immediately carries over to $\mathcal{B}_{N,k}$, since besides $|\mathcal{A}_{N,k}| \leq N^{O(1)}$ we did not use any AP-specific properties (note that in the union bound arguments below (51) and (55) there is plenty of elbow room). Furthermore, after changing the constant 9 in the definition (28) of k to $4c + 1$, say, we see that the probability estimate (37) holds with N^{-2} replaced by N^{-c} , which ensures that the union bound argument establishing Theorem 1.5 carries over to $\mathcal{B}_{N,k}$. \square

Acknowledgements. We thank the referees for helpful suggestions.

References

- [1] P. Bennett and T. Bohman. A note on the random greedy independent set algorithm. *Rand. Struct. Algor.* **49** (2016), 479–502.
- [2] T. Bohman. The triangle-free process. *Adv. Math.* **221** (2009), 1653–1677.
- [3] T. Bohman and P. Keevash. The early evolution of the H -free process. *Invent. Math.* **181** (2010), 291–336.
- [4] T. Bohman and P. Keevash. Dynamic concentration of the triangle-free process. *Rand. Struct. Algor.* **58** (2021), 221–293.
- [5] T. Bohman and L. Warnke. Large girth approximate steiner triple systems. *J. Lond. Math. Soc.* **100** (2019), 895–913.
- [6] W. Cames van Batenburg, R. de Joannis de Verclos, R. Kang, and F. Pirot. Bipartite induced density in triangle-free graphs. *Electron. J. Combin.* **27** (2020), Paper 2.34, 19 pp.
- [7] P. Erdős, R. Faudree, J. Pach, and J. Spencer. How to make a graph bipartite. *J. Combin. Theory Ser. B* **45** (1988), 86–98.
- [8] L. Esperet, R. Kang, and S. Thomassé. Separation choosability and dense bipartite induced subgraphs. *Combin. Probab. Comput.* **28** (2019), 720–732.
- [9] G. Fiz Pontiveros, S. Griffiths, and R. Morris. The triangle-free process and the Ramsey number $R(3, k)$. *Mem. Amer. Math. Soc.* **263** (2020), no. 1274.
- [10] S. Glock, D. Kühn, A. Lo, and D. Osthus. On a conjecture of Erdős on locally sparse Steiner triple systems. *Combinatorica* **40** (2020), 363–403.
- [11] W.T. Gowers. A new proof of Szemerédi’s theorem. *Geom. Funct. Anal.* **11** (2001), 465–588.
- [12] R. Graham. On the growth of a van der Waerden-like function. *Integers* **6** (2006), A29, 5 pp.
- [13] R. Graham and S. Butler. *Rudiments of Ramsey theory*. 2nd ed., Amer. Math. Soc., Providence (2015).
- [14] B. Green. Arithmetic progressions in sumsets. *Geom. Funct. Anal.* **12** (2002), 584–597.
- [15] B. Green. New lower bounds for van der Waerden numbers. Preprint (2021). [arXiv:2102.01543](https://arxiv.org/abs/2102.01543)
- [16] B. Green and J. Wolf. A note on Elkin’s improvement of Behrend’s construction. In *Additive number theory*, Springer (2010), pp. 141–144.
- [17] H. Guo. On the power of random greedy algorithms. *Talk at AMS Fall Southeastern Sectional Meeting (virtual)*, October 10, 2020. Abstract available at: http://www.ams.org/amsmtgs/2281_abstracts/1161-05-280.pdf
- [18] H. Guo, K. Patton, and L. Warnke. Prague dimension of random graphs. Preprint (2020). [arXiv:2011.09459](https://arxiv.org/abs/2011.09459)
- [19] H. Guo and L. Warnke. Packing nearly optimal Ramsey $R(3, t)$ graphs. *Combinatorica* **40** (2020), 63–103.
- [20] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. Wiley-Interscience, New York (2000).
- [21] J. Kahn. Asymptotically good list-colorings. *J. Combin. Theory Ser. A* **73** (1996), 1–59.
- [22] J.H. Kim. The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$. *Rand. Struct. Algor.* **7** (1995), 173–207.
- [23] M. Kwan, S. Letzter, B. Sudakov, and T. Tran. Dense induced bipartite subgraphs in triangle-free graphs. *Combinatorica* **40** (2020), 283–305.
- [24] Y. Li. A random process on graphs. *Talk at Random graphs and complex networks Conference (Shanghai Jiaotong University)*, November 14, 2009. Slides available at: http://math.sjtu.edu.cn/institution/C_S/09-11-15/4-LYS.pdf
- [25] Y. Li and J. Shu. A lower bound for off-diagonal van der Waerden numbers. *Adv. in Appl. Math.* **44** (2010), 243–247.
- [26] M. Piccollelli. The final size of the C_ℓ -free process. *SIAM J. Discrete Math.* **28** (2014), 1276–1305.
- [27] V. Rödl. On a packing and covering problem. *European J. Combin.* **6** (1985), 69–78.
- [28] T. Schoen. A subexponential upper bound for van der Waerden numbers $W(3, k)$. *Electron. J. Combin.* **28** (2021), Paper 2.34, 10pp.
- [29] J. Spencer. Asymptotic packing via a branching process. *Rand. Struct. Algor.* **7** (1995), 167–172.
- [30] L. Warnke. The C_ℓ -free process. *Rand. Struct. Algor.* **44** (2014), 490–526.
- [31] L. Warnke. When does the K_4 -free process stop? *Rand. Struct. Algor.* **44** (2014), 355–397.
- [32] L. Warnke. On Wormald’s differential equation method. *Combin. Probab. Comput.*, to appear. [arXiv:1905.08928](https://arxiv.org/abs/1905.08928)

- [33] L. Warnke. On the missing log in upper tail estimates. *J. Combin. Theory Ser. B* **140** (2020), 98–146.
- [34] N. Wormald. Differential equations for random processes and random graphs. *Ann. Appl. Probab.* **5** (1995), 1217–1235.