

RANDOM TRIANGLE REMOVAL

TOM BOHMAN, ALAN FRIEZE, AND EYAL LUBETZKY

ABSTRACT. Starting from a complete graph on n vertices, repeatedly delete the edges of a uniformly chosen triangle. This stochastic process terminates once it arrives at a triangle-free graph, and the fundamental question is to estimate the final number of edges (equivalently, the time it takes the process to finish, or how many edge-disjoint triangles are packed via the random greedy algorithm). Bollobás and Erdős (1990) conjectured that the expected final number of edges has order $n^{3/2}$. An upper bound of $o(n^2)$ was shown by Spencer (1995) and independently by Rödl and Thoma (1996). Several bounds were given for variants and generalizations (e.g., Alon, Kim and Spencer (1997) and Wormald (1999)), while the best known upper bound for the original question of Bollobás and Erdős was $n^{7/4+o(1)}$ due to Grable (1997). No nontrivial lower bound was available.

Here we prove that with high probability the final number of edges in random triangle removal is equal to $n^{3/2+o(1)}$, thus confirming the $3/2$ exponent conjectured by Bollobás and Erdős and matching the predictions of Gordon, Kuperberg, Patashnik, and Spencer (1996). For the upper bound, for any fixed $\varepsilon > 0$ we construct a family of $\exp(O(1/\varepsilon))$ graphs by gluing $O(1/\varepsilon)$ triangles sequentially in a prescribed manner, and dynamically track the number of all homomorphisms from them, rooted at any two vertices, up to the point where $n^{3/2+\varepsilon}$ edges remain. A system of martingales establishes concentration for these random variables around their analogous means in a random graph with corresponding edge density, and a key role is played by the self-correcting nature of the process. The lower bound builds on the estimates at that very point to show that the process will typically terminate with at least $n^{3/2-o(1)}$ edges left.

1. INTRODUCTION

Consider the following well-known stochastic process for generating a triangle-free graph, and at the same time creating a partial Steiner triple system. Start from a complete graph on n vertices and proceed to repeatedly remove the edges of uniformly chosen triangles. That is, letting $G(0)$ denote the initial graph, $G(i+1)$ is obtained from $G(i)$ by selecting a triangle uniformly at random out of all triangles in $G(i)$ and deleting its 3 edges. The process terminates once no triangles remain, and the fundamental question is to estimate the stopping time

$$\tau_0 = \min\{i : G(i) \text{ is triangle-free}\}.$$

This is equivalent to estimating the number of edges in the final triangle-free graph, since $G(i)$ has precisely $\binom{n}{2} - 3i$ edges by definition. As the triangles removed are mutually edge-disjoint, this process is precisely the random greedy algorithm for triangle packing.

Bollobás and Erdős (1990) conjectured that the expected number of edges in $G(\tau_0)$ has order $n^{3/2}$ (see, e.g., [6, 7]), with the motivation of determining the Ramsey number $R(3, t)$. Behind this conjecture was the intuition that the graph $G(i)$ should be similar to an Erdős-Rényi random graph with the same edge density. Indeed, in the latter random graph with n vertices and $\varepsilon n^{3/2}$ edges there are typically about $\frac{4}{3}\varepsilon^3 n^{3/2}$ triangles, thus, for small ε , deleting all of its triangles one by one would still retain all but a negligible fraction of the edges.

It was shown by Spencer [14] in 1995, and independently by Rödl and Thoma [13] in 1996, that the final number of edges is $o(n^2)$ with high probability (w.h.p.). In 1997, Grable [11] improved this to an upper bound of $n^{11/6+o(1)}$ w.h.p., and further described how similar arguments, using some more delicate calculations, should extend that result to $n^{7/4+o(1)}$. This remained the best upper bound prior to this work. No nontrivial lower bound was available. (See [10] for numerical simulations firmly supporting an answer of $n^{3/2+o(1)}$ to this problem.)

T. B. is supported in part by NSF grant DMS-1001638. A. F. is supported in part by NSF grant DMS-0721878.

Of the various works studying generalizations and variants of the problem, we mention two here. In his paper from 1999, Wormald [15] demonstrated how the differential equation method can yield a nontrivial upper bound on the number of remaining edges during a greedy packing of hyperedges in k -uniform hypergraphs. For the special case $k = 3$, corresponding to triangle packing, this translated to a bound of $n^{2-\frac{1}{57}+o(1)}$. Also in the context of hypergraphs, Alon, Kim and Spencer [1] introduced in 1997 a semi-random variant of the aforementioned process (akin to the Rödl nibble [12] yet with some key differences) which, they showed, finds nearly perfect matchings. Specialized to our setting, that process would result in a collection of edge-disjoint triangles on n vertices that covers all but $O(n^{3/2} \log^{3/2} n)$ of the edges of the complete graph. Alon *et al* [1] then conjectured that the simple random greedy algorithm should match those results, and in particular — generalizing the Bollobás-Erdős conjecture — that applying it to find a maximal collection of k -tuples with pairwise intersections at most $k - 2$ would leave out an expected number of at most $n^{k-1-\frac{1}{k-1}+o(1)}$ uncovered $(k - 1)$ -tuples. They added that “at the moment we cannot prove that this is the case even for $k = 3$ ”, the focus of our work here.

Joel Spencer offered \$200 for a proof that the answer to the problem is $n^{3/2+o(1)}$ w.h.p. ([11, 15]). The main result in this work establishes this precise statement, thus confirming the exponent conjectured by Bollobás and Erdős (1990).

Theorem 1. *Let τ_0 be the number of steps it takes the random triangle removal process to terminate starting from a complete graph on n vertices, and let $E(\tau_0)$ be the edge set of the final triangle-free graph. Then with high probability $\tau_0 = n^2/6 - n^{3/2+o(1)}$, or equivalently, $|E(\tau_0)| = n^{3/2+o(1)}$.*

We prove Theorem 1 by showing that w.h.p. all variables in a collection of $e^{O(1/\varepsilon)}n^2$ random variables, carefully designed to support the analysis, stay close to their respective trajectories throughout the evolution of this stochastic hypergraph process (see §1.1 for details). Establishing concentration for this large collection of variables hinges on their self-correction nature: the further a variable deviates from its trajectory, the stronger its drift is back toward its mean. However, turning this into a rigorous proof is quite challenging given that the various variables interact and errors (deviations from the mean) propagating from other variables may interfere in the attempt of one variable to correct itself. We construct a system of martingales to guarantee that the drift of a variable towards its mean will dominate the errors in our estimates for its peers.

The tools developed here for proving such self-correcting estimates are generic and we believe they will find applications in other settings. In particular, these methods should support an analysis of the random greedy hypergraph matching process as well as that of the corresponding H -free process for a fixed hypergraph H (see §1.2 for more on the triangle-free process).

1.1. Methods. Our starting point for the upper bound is a system of martingales for dynamically tracking an ensemble of random variables consisting of the triangle count and all co-degrees in the graph. The self-correcting behavior of these variables is roughly seen as follows: should the co-degree of a given pair of vertices deviate above/below its average, then more/fewer than average triangles could shrink the co-neighborhood if selected for removal in the next round, and this compensation effect would eventually drive the variable back towards its mean. One can exploit this effect to maintain the concentration of all these variables around their analogous means in a corresponding Erdős-Rényi random graph, as long as the number of edges remains above $n^{7/4+o(1)}$. (In the short note [4] the authors applied this argument to match this upper bound due to Grable.)

It is no coincidence that various methods break precisely at the exponent $7/4$ as it corresponds to the inherent barrier where co-degrees become comparable to the variations in their values that arose earlier in the process. In order to carry out the analysis beyond this barrier, one can for instance enrich the ensemble of tracked variables to address the second level neighborhoods of pairs of vertices. This way one can avoid cumulative worst-case individual errors due to large summations of co-degree variables, en route to co-degree estimates which *improve* as the process evolves. Indeed,




FIGURE 1. A subset of the ensemble of about 2^{30} triangular ladders that are tracked for any pair of roots to sustain the analysis of the upper bound to the point of $n^{3/2+\varepsilon}$ edges for $\varepsilon = \frac{1}{10}$.

these ideas can push this framework to successfully track the process to the point where $n^{5/3+o(1)}$ edges remain. However, beyond that point (corresponding to an edge density of about $n^{-1/3}$) the size of common neighborhoods of triples would become negligible, foiling the analysis. This example shows the benefit of maintaining control over a large family of subgraphs, yet at the same time it demonstrates how for any family of *bounded size* subgraphs the framework will eventually collapse.

In order to prove the $n^{3/2+o(1)}$ upper bound we track an *arbitrarily large* collection of random variables. A natural prerequisite for understanding the evolution of the number of homomorphisms from any small subgraph H to $G(i)$ rooted at a given pair of vertices u, v is to have ample control over homomorphisms from subgraphs obtained by “gluing a triangle” to an edge of H (as the number of those dictates the probability of losing copies of H in the next round on account of losing the mentioned edge). Tracking these would involve additional larger subgraphs etc., and to end this cascade and obtain a system of variables which can self-sustain itself we do the following. To estimate the probability of losing the “last edge” in the largest tracked H , in lieu of counting homomorphisms from a “forward extension” (the result of gluing a new triangle onto the given edge) rooted at u, v , we fix our candidate for the edge in $G(i)$ that would play the role of this last edge and then count homomorphisms rooted at u, v and featuring this specific edge. In other words, we replace the requirement for control over “forward extensions” with control over “backward extensions” from any given edge in $G(i)$. Obviously, rooting the subgraphs at as many as 4 vertices can drastically decrease the expected number of homomorphisms, say to $O(1)$, while our analysis must involve error probabilities that tend to 0 at a super-polynomial rate to account for our polynomial number of variables.

These points lead to a careful definition of the ensemble of graphs we wish to track. We construct what we refer to as *triangular ladders* which, roughly put, are formed by repeatedly taking an existing element in the ensemble and gluing a triangle on one of its last added two edges. We do so until reaching a certain threshold on the number of vertices, namely $O(1/\varepsilon)$ for some fixed $\varepsilon > 0$, amounting to a family of $2^{O(1/\varepsilon)}$ graphs (see Figure 1). Special care needs to be taken so as to avoid certain substructures that would foil our analysis, yet ignoring these for the moment (postponing precise definitions to §4), note that each new step of adding a triangle (a new vertex and two new edges) also increases the number of edges in the ladder under consideration, as long as the edge density is at least $n^{-1/2+\varepsilon}$. When examining a given ladder in this ensemble, its forward extensions will also be in the ensemble, by construction, unless we exceed the size threshold. In the latter case, if the threshold is taken to be suitably large, we can safely appeal to a backward extension argument and wind up with polynomially many homomorphisms to $G(i)$.

The main challenge in this approach is to come up with a canonical definition of the ensemble such that a uniform analysis could be applied to its arbitrarily many elements. Since as usual the errors in one variable propagate to others, we end up with a system of linear constraints on $2^{O(1/\varepsilon)}$ variables. Fortunately, suitable canonical definitions of the main term in each variable as a function of its predecessor in the ensemble make this constraint system one without cyclic dependencies, allowing for a simple solution to it. Altogether, the triangular ladder ensemble can be maintained as long as there are $n^{3/2+\varepsilon}$ edges, and while so it sustains the analysis of the number of triangles.

The lower bound builds on the fact that our analysis of the upper bound yielded asymptotic estimates for all co-degrees up to the point of $n^{3/2+\varepsilon}$ edges, for arbitrarily small $\varepsilon > 0$. Already a weaker bound on the co-degrees, when combined with an analysis of the evolution of the number of triangles, suffices to show that either the final number of edges has order at least $n^{3/2}$ or at some point i there are $|E(i)| \asymp n^{3/2}$ edges and $(\frac{1}{3} - o(1))|E(i)|$ edge-disjoint triangles. These contain a collection of $n^{3/2-O(\varepsilon)}$ edge-disjoint triangles such that for each one there is an edge that will survive in the final graph with probability $n^{-O(\varepsilon)}$. These edge survival events are mutually independent. Altogether this leads to at least $n^{3/2-O(\varepsilon)}$ edges in the final graph w.h.p.

1.2. Comparison with the triangle-free process. A different recipe for obtaining a random triangle-free graph is the so-called “triangle-free process”. In that process, the $\binom{n}{2}$ edges arrive one by one according to a uniformly chosen permutation, and each one is added if and only if no triangles are formed with it in the current graph. It was shown in [8] that the final number of edges in this process is w.h.p. $n^{3/2+o(1)}$, and the correct order of $n^{3/2}\sqrt{\log n}$, along with its ramifications on the Ramsey number $R(3, t)$, was recovered in [3] (see also [5] for generalizations).

Similarly to that process, the triangle removal process studied here can be presented as taking a uniform ordering of the $\binom{n}{3}$ triangles, then sequentially removing the edges of a triangle if and only if all 3 edges presently belong to the graph. The main result in this work shows that this process too culminates in a triangle-free graph with $n^{3/2+o(1)}$ edges.

Despite the high level similarity between the two protocols, the triangle removal process has proven quite challenging already at the level of acquiring the correct exponent of the final number of edges. One easily seen difference is that, in the triangle-free process, at any given point there is a set of forbidden edges (whose addition would form at least one triangle), yet regardless of its structure the next edge to be added is uniformly distributed over the legal ones. As long as the forbidden set of edges is negligible compared to the remaining legal edges, this process therefore mostly adds uniform edges just as in the Erdős-Rényi random graph. In particular, this is the case as long as $o(n^{3/2})$ edges were added (even based on all the edges that arrived rather than those selected). A coupling to the Erdős-Rényi random graph up to that point immediately gives a lower bound of order $n^{3/2}$ on the expected number of edges at the end of the triangle-free process. One could hope for an analogous soft argument to hold for the triangle removal process, where it would translate to an upper bound (as it deletes rather than adds edges).

However, in the triangle removal process studied here, deleting almost all edges save for $n^{3/2+o(1)}$ forms a substantial challenge for the analysis. Already when the edge density is a small constant, still bounded away from 0, the set of forbidden triangles becomes much larger than the set of legal ones. Thus, in an attempt to analyze the process even up to a density of $n^{-\varepsilon}$ for some $\varepsilon > 0$, one is forced to control the geometry of the remaining triangles quite delicately from its very beginning.

2. UPPER BOUND ON THE FINAL NUMBER OF EDGES MODULO CO-DEGREE ESTIMATES

Let (\mathcal{F}_i) be the filtration given by the process. Let $N_u = N_u(i) = \{x \in V_G : xu \in E(i)\}$ denote the neighborhood of a vertex $u \in V_G$, let $N_{u,v} = N_u \cap N_v$ be the co-neighborhood of vertices $u, v \in V_G$ and let $Y_{u,v} = |N_{u,v}|$ denote their co-degree. Our goal is to estimate the number of triangles in $G(i)$ which we denote by $Q(i)$. The motivation behind tracking co-degrees is immediate from the

fact that $Q(i) = \frac{1}{3} \sum_{uv \in E(i)} Y_{u,v}$, yet another important element is their effect on the evolution of Q . By definition,

$$\mathbb{E}[\Delta Q \mid \mathcal{F}_i] = - \sum_{uvw \in Q} (Y_{u,v} + Y_{u,w} + Y_{v,w} - 2) / Q, \quad (2.1)$$

where here and in what follows ΔX is always the one-step change in the variable X , namely

$$\Delta X = X(i+1) - X(i).$$

It is convenient to re-scale the number of steps i and introduce a notion of edge density as follows:

$$t = t(i) = \frac{i}{n^2}, \quad p = p(i, n) = 1 - \frac{6i}{n^2} = 1 - 6t. \quad (2.2)$$

With this notation we have

$$|E(i)| = \binom{n}{2} - 3i = \binom{n}{2} - \frac{1}{2}(1-p)n^2 = \frac{1}{2}(n^2 p - n), \quad (2.3)$$

and we see that, up to the negligible linear term, the number of edges corresponding to the edge density p is in line with that of the Erdős-Rényi random graph $\mathcal{G}(n, p)$. Under the assumption that $G(i)$ indeed resembles $\mathcal{G}(n, p)$ we expect that the number of triangles and co-degrees would satisfy

$$Y_{u,v} \approx np^2 \quad \text{and} \quad Q \approx \frac{1}{6}n^3 p^3.$$

Turning to the evolution of the co-degrees, we see that

$$\mathbb{E}[\Delta Y_{u,v} \mid \mathcal{F}_i] = - \sum_{x \in N_{u,v}} \frac{Y_{u,x} + Y_{v,x} - \mathbb{1}_{\{uv \in E\}}}{Q}, \quad (2.4)$$

and so tight estimates on the co-degrees and Q can already support the analysis of the process. In the short note [4] the authors relied on these variables alone to establish a bound of $O(n^{7/4} \log^{5/4} n)$ on the number of edges that survive to the conclusion of the algorithm. However, once the edge density drops to $p = n^{-1/4}$ we arrive at $Y_{u,v} \approx n^{1/2}$. As the variables $Y_{u,v}$ have order n and standard deviations of order \sqrt{n} while p is constant, we see that at the point $p = n^{-1/4}$ the deviations that developed early in the process cease to be regarded as negligible and a refined analysis is required.

One can achieve better precision for the co-degrees by introducing additional variables aimed at *decreasing* the variation for $Y_{u,v}$ as the process evolves. For instance, through estimates on the number of edges between the common and exclusive neighborhoods for every pair of vertices, along with some additional ingredients exploiting the self-correction phenomenon, one can gain an improvement on the 7/4 exponent. However, eventually these arguments break down, soon after some bounded-size subgraphs become too rare, thus foiling the concentration estimates.

We remedy this via the triangular ladder ensemble, defined in §4 and analyzed in §5. To sustain the analysis up to the point where $n^{3/2+\varepsilon}$ edges remain we must simultaneously control $\exp(O(1/\varepsilon))$ types of random variables, yet our ultimate goal in those sections is simply to estimate all co-degrees $\{Y_{uv} : u, v \in V_G\}$. The following corollary is a special case of Theorem 5.4.

Theorem 2.1. *Define $\zeta = \zeta(p, n)$ to be*

$$\zeta = n^{-1/2} p^{-1} \log n \quad (2.5)$$

and for some (arbitrarily large) absolute constant $\kappa > 0$ let

$$\tau_Q^* = \min \left\{ t : \left| Q / \left(\frac{1}{6} n^3 p^3 \right) - 1 \right| \geq \kappa \zeta^2 \right\}.$$

Then for every $M \geq 3$ w.h.p. the following holds:

$$\left| Y_{u,v} / (np^2) - 1 \right| \leq 3^{3M-1} \zeta$$

for every $u, v \in V_G$ and all t such that $t \leq \tau_Q^$ and $p(t) \geq n^{-1/2+1/M}$.*

The heart of this work is in establishing the above theorem (Sections 3–5), which is complemented by the following result: A multiplicative error of $1 + O(\zeta)$ for all co-degree variables $Y_{u,v}$ can be enhanced to bounds tight up to a factor of $1 + O(\zeta^2)$ on the number of triangles Q .

Theorem 2.2. *Set $\Phi(p, n) = p^{-2/\log n} \log n$ and for some fixed $\alpha > 0$ let*

$$\tau_Y^* = \min \left\{ t : \exists u, v \text{ such that } |Y_{u,v} - np^2| > \alpha n^{1/2} p \Phi \right\}. \quad (2.6)$$

Then w.h.p. the following holds: for all $t \leq \tau_Y^$ and $p(t) \geq n^{-1/2} \log^2 n$ we have*

$$|Q - n^3 p^3 / 6| \leq \alpha^2 n^2 p \Phi^2. \quad (2.7)$$

Proof. Let

$$X(i) = Q - \frac{n^3 p^3}{6}, \quad Z(i) = |X(i)| - \alpha^2 n^2 p \Phi^2.$$

We define below a narrow *critical interval* I_Q that has its upper endpoint at the bound we aim to establish for $|X|$. As long as $|X|$ lies in this interval it is subject to a self-correcting drift, which turns out to be just enough to show that $\mathbb{E}[\Delta Z \mid \mathcal{F}_i] < 0$. That is, as long as $|X| \in I_Q$ the sequence $Z(i)$ is a supermartingale. Standard concentration estimates then imply that $|X|$ is unlikely to ever cross this critical interval, leading to the desired estimate on Q .

Define

$$\Lambda = \log^{-3/2} n$$

and observe that for all $p \geq n^{-1/2} \log^2 n$ we have $n^{-1} p^{-2} \Phi^2 = O(1/\log^2 n) = o(\Lambda)$, whence as long as (2.7) is valid we have $Q = (\frac{1}{6} + o(\Lambda)) n^3 p^3$. This clearly holds at the beginning of the process, and in what follows we may condition on this event as part of showing that the bound (2.7) will be maintained except with an error probability that tends to 0 at a super-polynomial rate.

We begin with estimates on the one-step expected changes of our variables. Recall that due to (2.1) we have

$$\mathbb{E}[\Delta Q \mid \mathcal{F}_i] = - \sum_{uvw \in Q} \frac{Y_{u,v} + Y_{u,w} + Y_{v,w} - 2}{Q} = 2 - \frac{1}{Q} \sum_{uv \in E} Y_{u,v}^2.$$

To bound $\mathbb{E}[\Delta Q \mid \mathcal{F}_i]$ we thus need an estimate on $\sum_{uv \in E} Y_{u,v}^2$, achieved by the next simple lemma.

Lemma 2.3. *Let $a_1, \dots, a_m \in \mathbb{R}$ and suppose that $|a - a_i| \leq \delta$ for all i and some $a \in \mathbb{R}$. Then*

$$\frac{(\sum_i a_i)^2}{m} \leq \sum_i a_i^2 \leq \frac{(\sum_i a_i)^2}{m} + 4m\delta^2.$$

Proof. The lower bound is due to Cauchy-Schwarz. For the upper bound fix $b = \frac{1}{m} \sum_i a_i$ and note that the assumption on the a_i 's implies that $|b - a_i| \leq 2\delta$ for all i . Observe that the convex function $\sum_i a_i^2$ achieves its maximum over the convex set $\{\sum_i a_i = bm\} \cap (\bigcap_i \{|b - a_i| \leq 2\delta\})$ at an extremal point where $a_i = b \pm 2\delta$ for all but at most one i . Hence, $a_i = b - 2\delta$ for $\lfloor m/2 \rfloor$ indices, $a_i = b + 2\delta$ for $\lfloor m/2 \rfloor$ indices and if m is odd there is a single $a_i = b$. In particular

$$\sum_{i=1}^m a_i^2 \leq \lfloor m/2 \rfloor (b - 2\delta)^2 + \lfloor m/2 \rfloor (b + 2\delta)^2 + b^2 \mathbb{1}_{\{m \equiv 1 \pmod{2}\}} \leq mb^2 + 4m\delta^2. \quad \blacksquare$$

Since $\sum_{uv \in E} Y_{u,v} = 3Q$ and $|Y_{u,v} - np^2| < \alpha n^{1/2} p \Phi$ by Eq. (2.6) and the fact that $t \leq \tau_Y^*$, an application of this lemma together with the fact that $|E(i)| = n^2 p / 2 - n / 2$ gives that

$$\frac{1}{Q} \sum_{uv \in E} Y_{u,v}^2 \geq \frac{9Q}{|E|} \geq \frac{18Q}{n^2 p},$$

together with the following upper bound:

$$\begin{aligned} \frac{1}{Q} \sum_{uv \in E} Y_{u,v}^2 &\leq \frac{9Q}{|E|} + 4 \frac{|E|}{Q} \alpha^2 n p^2 \Phi^2 \leq \frac{18Q}{n^2 p (1 - \frac{1}{np})} + \frac{2n^2 p}{(\frac{1}{6} + o(\Lambda)) n^3 p^3} \alpha^2 n p^2 \Phi^2 \\ &\leq \frac{18Q}{n^2 p} + O(p) + (12 + o(\Lambda)) \alpha^2 \Phi^2 = \frac{18Q}{n^2 p} + (12 + O(\Lambda)) \alpha^2 \Phi^2. \end{aligned}$$

(In the last equality we absorbed the $O(p)$ -term into the $O(\Lambda)$ error-term factor of the last expression since $\Lambda \Phi^2 \rightarrow \infty$.) Adding this to our estimate for $\mathbb{E}[\Delta Q \mid \mathcal{F}_i]$ (again absorbing the additive 2 in our error term) while noticing $\Delta(-\frac{1}{6}n^3 p^3) = 3np^2 + O(p/n)$ (recall that $p = p(i, n) = 1 - \frac{6i}{n^2}$) yields

$$\left| \mathbb{E}[\Delta X \mid \mathcal{F}_i] - \frac{18}{n^2 p} \left(-Q + \frac{1}{6} n^3 p^3 \right) \right| \leq (12 + O(\Lambda)) \alpha^2 \Phi^2 + O(p/n) = (12 + O(\Lambda)) \alpha^2 \Phi^2,$$

again incorporating the $O(p/n)$ term into the $O(\Lambda)$ error. Thus, by the definition of X

$$\mathbb{E}[\Delta |X| \mid \mathcal{F}_i] \leq -\frac{18}{n^2 p} |X| + (12 + O(\Lambda)) \alpha^2 \Phi^2.$$

Now assume that in round i_0 the variable $|X|$ rises above $\hat{\alpha}^2 n^2 p \Phi^2$, i.e., it enters the interval

$$I_Q = (\hat{\alpha}^2 n^2 p \Phi^2, \alpha^2 n^2 p \Phi^2) \quad \text{where} \quad \hat{\alpha} = (1 - \log^{-1} n)^{1/2} \alpha.$$

Assume also that up to this point $|X|$ has never reached the upper endpoint of the interval. This will obviously be true the first time, if ever, that $|X|$ enters I_Q . Further let

$$\tau = \min\{i > i_0 : |X(i)| < \hat{\alpha}^2 n^2 p \Phi^2\}.$$

Since $\Delta(-\alpha^2 n^2 p \Phi^2) = (6 + O(\frac{1}{n^2 p})) (1 - \frac{4}{\log n}) \alpha^2 \Phi^2$ and the error-term $O(\frac{1}{n^2 p})$ can be absorbed in the $O(\Lambda)$ -term as $p \geq 1/n$ (with much room to spare), the upper bound on $\mathbb{E}[\Delta |X| \mid \mathcal{F}_i]$ gives

$$\mathbb{E}[\Delta Z \mid \mathcal{F}_i, \tau > i \geq i_0] \leq \left[-3(\hat{\alpha}/\alpha)^2 + 2 + (1 - \frac{4}{\log n}) + O(\Lambda) \right] 6\alpha^2 \Phi^2 \leq -\frac{6 - o(1)}{\log n} \alpha^2 \Phi^2.$$

It follows that $(Z(j \wedge \tau))_{j > i_0}$ is indeed a supermartingale for large n .

Next consider the one-step variation of Z . Denoting the selected triangle in a given round by uvw , the change in Q following this round is at most $Y_{u,v} + Y_{u,w} + Y_{v,w}$ and in light of our co-degree estimate (2.6) this expression deviates from its expected value of $3np^2$ by at most $3\alpha n^{1/2} p \Phi$. In particular, $|\Delta Z| = O(\sqrt{n} p_0 \log n)$ and letting $p_0 = p(i_0) = 1 - 6i_0/n^2$ this ensures that

$$Z(i_0) \leq (\hat{\alpha}^2 - \alpha^2) n^2 p_0 \Phi^2 + O(\sqrt{n} p_0 \log n) \leq -\frac{1}{2} \alpha^2 n^2 p_0 \log n,$$

where the last inequality holds for large n . With at most $n^2 p_0$ steps remaining until the process terminates, Hoeffding's inequality establishes that for some fixed $c > 0$,

$$\mathbb{P}(\cup_{j \geq i_0} \{Z(j \wedge \tau) \geq 0\}) \leq \exp\left(-c \frac{(n^2 p_0 \log n)^2}{n^2 p_0 (\sqrt{n} p_0 \log n)^2}\right) = e^{-cn/p_0} \leq e^{-cn}.$$

Altogether, w.h.p. $|X(i)| < \alpha^2 n^2 p \Phi^2$ for all i , as required. Thus w.h.p. $|X|$ never leaves I_Q at the upper end, no matter how many times it enters at the lower end. \blacksquare

Remark. For all $p \geq n^{-1/2}$ the function Φ defined in Theorem 2.2 satisfies $1 \leq \Phi/\log n \leq e$, thus the estimates (2.6)–(2.7) precisely feature a relative error of order $n^{-1/2} p^{-1} \log n$ in line with the definition of ζ in Theorem 2.1. Specifically, the conclusion of Theorem 2.1 provides the estimate $|Y_{u,v} - np^2| \leq \alpha n^{1/2} p \Phi$ for $\alpha = 3^{3M-1}$, thus fulfilling the hypothesis of Theorem 2.2 en route to a bound of $|Q/(\frac{1}{6}n^3 p^3) - 1| \leq C\zeta^2$ for $C = (\alpha e)^2$. Altogether, we can apply both theorems in tandem as long as $p \geq n^{-1/2+1/M}$. In particular, at that point (2.7) guarantees that the process is still active with $(\frac{1}{6} + o(1))n^{3/2+3/M}$ triangles, yet there are merely $(\frac{1}{2} - o(1))n^{3/2+1/M}$ edges left. This establishes the upper bound in our main result modulo Theorem 2.1.

3. FROM DEGREES TO EXTENSION VARIABLES

Our goal in this section is to show that strong control on degrees, co-degrees and the number of triangles implies some level of control on arbitrary subgraph counts. Let Y_u denote the degree of the vertex u and as before let $Y_{u,v}$ and Q denote the co-degree of u, v and total number of triangles respectively. Our precise assumption will be that for some absolute constant $C > 0$

$$\begin{cases} |Y_u/(np) - 1| < C\zeta & \text{for all } u \in V_G \\ |Y_{u,v}/(np^2) - 1| < C\zeta & \text{for all } u, v \in V_G \\ |Q/(\frac{1}{6}n^3p^3) - 1| < C\zeta^2 \end{cases} \quad (3.1)$$

where $\zeta = \zeta(n, p) = n^{-1/2}p^{-1} \log n$.

The function ζ begins the process at $\zeta = n^{-1/2+o(1)}$ and then gradually increases with i (or as p decreases) until reaching a value of $n^{-\varepsilon+o(1)}$ for some (arbitrarily small) $\varepsilon > 0$ at the final stage of our analysis. However, in order to emphasize that the value of this ε plays no role in this section we will only make use of the fact that $p(t)$ is such that

$$\log(1/\zeta)/\log \log n \rightarrow \infty. \quad (3.2)$$

Setting

$$\tau_\star = \min\{t > 0 : \text{Eqs. (3.1) or (3.2) are violated}\}, \quad (3.3)$$

we will show that as long as $t \leq \tau_\star$ we can w.h.p. asymptotically estimate the number of copies of any balanced (a precise definition follows) labeled graph with $O(1)$ vertices rooted at a prescribed subset of vertices. These estimates will feature a multiplicative error-term of $1 + O(\zeta^\delta)$, where $\delta > 0$ will only depend on the uniform bound on the sizes of the graphs under consideration.

Definition 3.1. *[extensions and subextensions]* An **extension graph** is a graph $H = (V_H, E_H)$ paired with a strict subset $I_H \subsetneq V_H$ of distinguished vertices which form an independent set in H . The **scaling** of an extension graph is defined to be

$$\mathcal{S}_H = \mathcal{S}_H(n, p) = n^{v_H - \iota_H} p^{e_H},$$

where $e_H = |E_H|$, $v_H = |V_H|$ and $\iota_H = |I_H|$. The **density** of H is defined to be

$$m_H = e_H / (v_H - \iota_H).$$

A **subextension** K of H , denoted $K \subset H$, is a vertex induced subgraph $K = (V_K, E_K)$ that is also an extension graph where $V_K \subset V_H$, $E_K \subset E_H$, and $I_K = I_H$ (thus $V_K \supset I_H$). We say that K is a **proper subextension**, denoted by $K \subsetneq H$, whenever $E_K \subsetneq E_H$. We will denote the trivial subextension $K = (I_H, \emptyset)$ (the edgeless graph on I_H) by $\mathbf{1}$. For any subextension $K \subset H$, define the **quotient** H/K to be the extension graph $H/K = (V_H, E_H \setminus E_K)$ with the distinguished vertex set $I_{H/K} = V_K$. Observe that

$$\mathcal{S}_H = \mathcal{S}_K \mathcal{S}_{H/K}.$$

Let $H = (V_H, E_H)$ be an extension graph associated with the distinguished subset $I_H \subset V_H$ as in Definition 3.1 and set $e_H = |E_H|$, $v_H = |V_H|$ and $\iota_H = |I_H|$. Let $\varphi : I_H \rightarrow V_G$ be an injection. We are interested in tracking the number of copies of H in $G = G(i)$ with $\varphi(I_H)$ playing the role of the distinguished vertices. That is, we are interested in counting the extensions from $\varphi(I_H)$ to copies of H . Formally, an injective map $\psi : V_H \rightarrow V_G$ **extends φ to a copy of H** if ψ is a graph homomorphism that agrees with φ , that is, $\psi|_{I_H} \equiv \varphi$ and $uv \in E_H$ implies that $\psi(u)\psi(v) \in E$. Define

$$\Psi_{H,\varphi} = \#\left\{\psi \in V_H \rightarrow V_G : \psi \text{ extends } \varphi \text{ to a copy of } H\right\}$$

Note that $\Psi_{H,\varphi}$ counts *labeled* copies of H rooted at the given vertex subset $\varphi(I_H)$.

We will track the variable $\Psi_{H,\varphi}$ for extension graphs H that are dense relative to their subextensions. More precisely, we say that H is **balanced** if its density satisfies $m_H \geq m_K$ for any

subextension K of H . Furthermore, H is **strictly balanced** if its density is strictly larger than the densities of all of its subextensions. Our aim is to control the variable $\Psi_{H,\varphi}$ for any balanced extension graph H up to the point where the scaling of H becomes constant, formulated as follows. Let

$$t_H = \min\{t : \mathcal{S}_H(n, p(t)) \leq 1\}, \quad (3.4)$$

where the times t being considered here are, as usual, of the form i/n^2 for integer i . Recalling that $m_H = e_H/(v_H - \iota_H)$ observe that $p(t_H) = n^{-1/m_H} - an^{-2}$ where $0 \leq a < 6$ by (2.2). Furthermore,

$$m_{H_1} \geq m_{H_2} \Leftrightarrow t_{H_1} \leq t_{H_2} \quad (3.5)$$

and in particular any balanced H has $t_H \leq t_K$ for any proper subextension K of H (a strict inequality holds when H is strictly balanced). Thus, the threshold up to which we can track $\Psi_{H,\varphi}$ will indeed be dictated by t_H rather than by t_K for some $K \subsetneq H$.

Finer precision in tracking the variables $\Psi_{H,\varphi}$ will be achievable up to a point slightly earlier than t_H such that \mathcal{S}_H is still reasonably large. To this end it will be useful to generalize the quantity t_H as follows: With (2.5) in mind, for any $r \geq 0$ let

$$t_H^-(r) = \min\{t : \mathcal{S}_H(n, p(t)) \leq 1/\zeta^r\} \quad (3.6)$$

(again selecting among all time-points of the form $t = i/n^2$ for integer values of i). We will establish these more precise bounds on $\Psi_{H,\varphi}$ up to $t_H^-(\delta)$ for some suitably small $\delta > 0$, i.e., up to a point where the scaling \mathcal{S}_H is still super-poly-logarithmic in n . (Recall that $\zeta(t)$ is decreasing in t and yet $\log(1/\zeta)$ tends to ∞ faster than $\log \log n$.)

When examining the times t_H and t_H^- one ought to keep in mind that, for any extension graph H on $O(1)$ vertices, the value of \mathcal{S}_H decreases within a single step of the process by a multiplicative factor of $1 + O(p^{-1}n^{-2})$. In particular, in that case $(1 - O(p^{-1}n^{-2}))\zeta^{-r} \leq \mathcal{S}_H \leq \zeta^{-r}$ at time $t_H^-(r)$.

The main result of this section is the next estimate on the variables $\Psi_{H,\varphi}$, where the proportional error will be in terms of ζ , the varying proportional error we have for the variables $Y_{u,v}$ and Q .

Theorem 3.2. *For every $L > 0$ there is some $0 < \delta < 1$ such that, with probability $1 - O(e^{-\log^2 n})$, for all $t \leq \tau_*$ and every extension graph H on $v_H \leq L$ vertices and injection $\varphi : I_H \rightarrow V_G$ we have:*

- (i) *If H is balanced then $|\Psi_{H,\varphi} - \mathcal{S}_H| \leq \mathcal{S}_H \zeta^\delta$ for all $t \leq t_H^-(\sqrt{\delta})$.*
- (ii) *If H is strictly balanced then $|\Psi_{H,\varphi} - \mathcal{S}_H| \leq (\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^\rho$ with $\rho = 3e^{v_H - \iota_H}$ holds for all $t_H^-(\sqrt{\delta}) \leq t \leq t_H$.*

Before proving Theorem 3.2 we will derive the following corollary which will be applicable to extension graphs that are not necessarily balanced.

Corollary 3.3. *Fix $L > 0$ and let $0 < t \leq \tau_*$. With probability $1 - O(e^{-\log^2 n})$, for every extension graph H on $v_H \leq L$ vertices with $\rho = \rho_H = 3 \exp(v_H - \iota_H)$ and injection $\varphi : I_H \rightarrow V_G$ we have:*

- (1) *If $\mathcal{S}_K(t) \geq 1$ for all subextensions $K \subset H$ then $\Psi_{H,\varphi}(t) \leq (1 + o(1))\mathcal{S}_H(\log n)^\rho$.*
- (2) *If $\mathcal{S}_{H/K}(t) \leq 1$ for all subextensions $K \subset H$ then $\Psi_{H,\varphi}(t) \leq (1 + o(1))(\log n)^\rho$.*

Proof. Assume that the estimates in Theorem 3.2 hold simultaneously for all extension graphs on at most L vertices throughout time t . We may further assume that every vertex $v \in V_H \setminus I_H$ has $\deg(v) \geq 1$. Indeed, for the statement in Part (1) this is w.l.o.g. since adding any isolated vertex to $V_H \setminus I_H$ would contribute a factor of $(1 - o(1))n$ to $\Psi_{H,\varphi}$, balanced by a factor of n to \mathcal{S}_H . For Part (2) this assumption is implied by the hypothesis $\mathcal{S}_{H/K}(t) \leq 1$, as we can take some subextension K with $V_K = V_H \setminus \{v\}$ to obtain that $(p(t))^{\deg(v)} \leq 1/n$.

Define a sequence of strictly increasing sets $I_H = X_0 \subsetneq X_1 \subsetneq X_2 \subsetneq \dots \subsetneq X_k = V_H$ as follows. Associate each candidate for X_i for $i \geq 1$, one that contains X_{i-1} , with the extension graph $K_i = (X_i, E_i)$ where

$$E_i = \{(u, v) \in E_H : u, v \in X_i \text{ and either } u \notin X_{i-1} \text{ or } v \notin X_{i-1}\},$$

i.e., the set of edges of the induced subgraph on X_i that do not have both endpoints in X_{i-1} , and set the distinguished vertex set of K_i to be $I_{K_i} = X_{i-1}$ (guaranteed to be an independent set by our definition of E_i). With this definition in mind, let X_i be the subset of vertices that maximizes m_{K_i} . If more than one such subset exists, let X_i be one of these which has minimal cardinality.

To see that the sequence of X_i 's is strictly increasing, let $i \geq 1$ and consider the potential values for X_i building upon some $X_{i-1} \subsetneq V_H$. Taking $X_i = X_{i-1}$ would yield $m_{K_i} = 0$, strictly smaller than the density that would correspond to $X_i = H$ (otherwise $V_H \setminus X_{i-1}$ would consist of isolated vertices in H , contradicting our minimal degree assumption). Hence, indeed $X_{i-1} \subsetneq X_i$ and overall K_i is strictly balanced by construction. (We note in passing that $K_k = H/K_{k-1}$ and that K_i is only a subextension of H for $i \leq 1$ since $I_H \subsetneq I_{K_i}$ for $i > 1$.)

Motivating these definitions is the fact that for a given sequence X_0, \dots, X_k as above one can recover $\psi \in V_H \rightarrow V_G$ (counted by $\Psi_{H,\varphi}$) iteratively from a sequence of $\psi_i \in V_{K_i} \rightarrow V_G$ (in which $\psi_0 = \varphi$ and $(\psi_i)|_{X_{i-1}} = \psi_{i-1}$ for all i), hence

$$\Psi_{H,\varphi}(t) = \sum_{\psi_1 \in \Psi_{K_1,\varphi}} \sum_{\psi_2 \in \Psi_{K_2,\psi_1}} \cdots \sum_{\psi_{k-1} \in \Psi_{K_{k-1},\psi_{k-2}}} \Psi_{K_k,\psi_{k-1}}(t) \leq \prod_{i=1}^k \max_{\eta} \Psi_{K_i,\eta}(t). \quad (3.7)$$

(The last maximum goes over all injections $\eta : I_{K_i} \rightarrow V_G$ and throughout this section, with a slight abuse of notation, let $\psi_i \in \Psi_{K_i}$ to denote an injective map ψ_i counted by the corresponding Ψ_{K_i} .)

To prove Part (1) of the corollary, observe that $K_1 \subset H$ and so by hypothesis we have $\mathcal{S}_{K_1}(t) \geq 1$, or equivalently, $t \leq t_{K_1}$. For any $2 \leq i \leq k$, the fact that $X_{i-1} \subsetneq X_i$ implies that $m_{K_{i-1}}$ is as large as $m_{K'}$, where K' is the extension graph with vertex set X_i , distinguished vertices $I_{K'} = X_{i-2}$ and all edges of H between vertices of X_i excluding those whose endpoints both lie in X_{i-2} . Since $E_{K'} = E_{i-1} \cup E_i$ we get

$$\frac{|E_{i-1}|}{|X_{i-1}| - |X_{i-2}|} \geq \frac{|E_{i-1}| + |E_i|}{|X_i| - |X_{i-2}|},$$

which readily implies that

$$m_{K_{i-1}} = \frac{|E_{i-1}|}{|X_{i-1}| - |X_{i-2}|} \geq \frac{|E_i|}{|X_i| - |X_{i-1}|} = m_{K_i}.$$

Equivalently (via (3.5)), $t_{K_{i-1}} \leq t_{K_i}$ and by induction we thus have $t \leq t_{K_i}$ for all i . Recalling that K_i is strictly balanced, we can appeal to Theorem 3.2 and derive that for any injection $\eta : I_{K_i} \rightarrow V_G$

$$\Psi_{K_i,\eta}(t) \leq \left(1 + \max\{\zeta^\delta, (\mathcal{S}_{K_i}(t))^{-\sqrt{\delta}}(\log n)^{\rho_i}\}\right) \mathcal{S}_{K_i}(t) \leq (1 + o(1))\mathcal{S}_{K_i}(t)(\log n)^{\rho_i},$$

where $\rho_i = 3 \exp(|X_i| - |X_{i-1}|)$, since $\mathcal{S}_{K_i} \geq 1 - O(1/n^2p)$ (as $t \leq t_{K_i}$) and using (3.2). By definition, $\mathcal{S}_H = \prod_{i=1}^k \mathcal{S}_{K_i}$, and so (3.7) yields that under the hypothesis of Part (1) we have

$$\Psi_{H,\varphi}(t) \leq (1 + o(1))\mathcal{S}_H(t)(\log n)^{\sum_{i=1}^k \rho_i}.$$

The proof of Part (1) is completed by the fact that $\exp(|X_i| - |X_{i-1}|) > 2$ for all i (since the X_i 's are strictly increasing) and in particular

$$\sum_i \rho_i = 3 \sum_i e^{|X_i| - |X_{i-1}|} < 3 \prod_i e^{|X_i| - |X_{i-1}|} = 3e^{v_H - \iota_H} = \rho. \quad (3.8)$$

It remains to prove Part (2). Observe that the hypothesis in this part implies that $\mathcal{S}_H(t) \leq 1$ (by taking $K = \mathbf{1}$), that is we are considering some time point $t \geq t_H$.

To utilize the assumption on the scaling of quotients of H , we go back to the definition of the sequence of X_i 's and argue that, if $H'_i = (V_H, E_H \setminus \cup_{j < i} E_j)$ with $I_{H'_i} = X_{i-1}$ then

$$m_{K_i} \geq m_{H'_i} \quad \text{for all } i \geq 1. \quad (3.9)$$

Indeed, this follows from the fact that selecting $X_i = V_H$ would exactly correspond to having $K_i = H'_i$, and yet X_i is chosen so as to maximize the value of m_{K_i} .

Consider the extension graph $K'_i = (X_{i-1}, \cup_{j < i} E_j)$ (that is, the induced subgraph of H on X_{i-1}) with the distinguished vertex set $I_{K'_i} = I_H$; this is a subextension of H , and we have

$$\mathcal{S}_{H'_i} = n^{v_H - |X_{i-1}|} p^{e_H - \sum_{j < i} |E_j|} = \mathcal{S}_{H/K'_i} \leq 1,$$

where the last inequality follows from our hypothesis on all the quotients of H .

From the inequality $\mathcal{S}_{H'_i} \leq 1$ established above we infer that $t \geq t_{H'}$. At the same time, the combination of (3.5) and (3.9) implies that $t_{H'} \geq t_{K_i}$, and altogether $t \geq t_{K_i}$. Now, since $\Psi_{K_i, \varphi}(t)$ is monotone non-increasing in t and we only aim to bound this quantity from above, it suffices to provide an upper bound on $\Psi_{K_i, \varphi}(t_{K_i})$. To this end, apply Part (ii) of Theorem 3.2 (bearing in mind that $\mathcal{S}_{K_i} = 1 + O(1/n^2 p)$ at time t_{K_i} and that K_i is strictly balanced) to get that for any injection $\eta : I_{K_i} \rightarrow V_G$

$$\Psi_{K_i, \eta}(t_{K_i}) \leq (1 + O(1/n^2 p)) (1 + o(1)) (1 + (\log n)^{\rho_i}), \quad \text{where } \rho_i = 3e^{|X_i| - |X_{i-1}|}.$$

At this point we again appeal to (3.7) to infer that

$$\Psi_{H, \varphi}(t) \leq (1 + o(1)) (\log n)^{\sum_{i=1}^k \rho_i} \leq (1 + o(1)) (\log n)^\rho,$$

where we used the fact that $\sum_i \rho_i \leq \rho$ as seen in (3.8). This concludes the proof. \blacksquare

Proof of Theorem 3.2. Let \mathcal{A}_L denote the set of all balanced extension graphs on at most L vertices and at least one edge (in the edgeless case $\Psi_{H, \varphi} = (1 + O(1/n)) \mathcal{S}_H = (1 + o(\zeta^\delta)) \mathcal{S}_H$ trivially by definition). Since \mathcal{A}_L is finite and $m_H > 0$ for all $H \in \mathcal{A}_L$ we can choose a sufficiently small constant $0 < \delta < 1$ as follows:

$$\delta = \min \left(\left\{ \left| \frac{1}{m_A} - \frac{1}{m_B} \right|^2 : \begin{array}{l} A, B \in \mathcal{A}_L \\ m_A \neq m_B \end{array} \right\} \cup \left\{ \frac{1}{(3e_A)^2} : A \in \mathcal{A}_L \right\} \right). \quad (3.10)$$

Combining the fact that $|m_A^{-1} - m_B^{-1}| \geq \sqrt{\delta}$ whenever $m_A \neq m_B$ together with the facts $p(t_A) = n^{-1/m_A} - O(n^{-2})$ and $p(t_A) \geq n^{-1/2+o(1)}$ (as $t \leq \tau_*$) implies that for any $A, B \in \mathcal{A}_L$ with $t_A < t_B$

$$\mathcal{S}_B(t_A) = \left(n^{m_B^{-1}} p(t_A) \right)^{e_B} = \left(1 + O\left(\frac{1}{n^2 p(t_A)} \right) \right) n^{(m_B^{-1} - m_A^{-1})e_B} \geq (1 - o(1)) n^{\sqrt{\delta} e_B} > \zeta^{-\sqrt{\delta}},$$

where the last inequality used that $e_B \geq 1$ as well as $\zeta^{-1} = o(n)$ (with room to spare). In particular,

$$t_A < t_B \Rightarrow t_A < t_B^-(\sqrt{\delta}) \quad \text{for all } A, B \in \mathcal{A}_L.$$

Contrary to the monotonicity of t_A in m_A as per (3.5), it is generally not true that $t_A^-(r)$ is monotone in m_A . However, the above definition of δ will enable us to still relate these thresholds as we will later argue that if $m_A \geq m_B$ then $t_A^-(\sqrt{\delta}) \leq t_B^-(3 - o(1))\delta$.

Proof of Part (i). Let H and φ be fixed. We will apply the previously used martingale arguments based on critical intervals. We begin with the expected change in $\Psi_{H, \varphi}$. Here the probability of deleting a given copy of H corresponds as usual to the number of triangles resting on each of its edges, and summing these separately can only incur a double count of $O(1)$ since $v_H = O(1)$. Therefore,

$$\begin{aligned} \mathbb{E}[\Delta \Psi_{H, \varphi}] &= - \sum_{\psi \in \Psi_{H, \varphi}} \sum_{uv \in E_H} \frac{Y_{\psi(u)\psi(v)} - O(1)}{Q} = - \frac{\Psi_{H, \varphi} e_H}{Q} (1 + O(\zeta)) n p^2 \\ &= - (6 + O(\zeta)) \frac{\Psi_{H, \varphi} e_H}{n^2 p}. \end{aligned}$$

Now set $X = \Psi_{H,\varphi} - \mathcal{S}_H$. Our critical interval for $|X|$ will be

$$I_\Psi = \left[\left(1 - \frac{\delta}{2e_H}\right) \mathcal{S}_H \zeta^\delta, \mathcal{S}_H \zeta^\delta \right].$$

To evaluate the one-step change in X , we first see that the one-step change in \mathcal{S}_H is deterministically given by

$$\begin{aligned} \Delta \mathcal{S}_H &= n^{v_H - \iota_H} \left((p - 6/n^2)^{e_H} - p^{e_H} \right) = -n^{v_H - \iota_H} \frac{6e_H}{n^2} (p - O(n^{-2}))^{e_H - 1} \\ &= - \left(1 + O(1/n^2 p)\right) \frac{6e_H}{n^2 p} \mathcal{S}_H = - (6 + o(\zeta)) \frac{e_H}{n^2 p} \mathcal{S}_H. \end{aligned} \quad (3.11)$$

When $|X| \in I_\Psi$ we therefore have

$$\begin{aligned} \mathbb{E}[\Delta X] &= - (6 + O(\zeta)) \frac{e_H}{n^2 p} (\mathcal{S}_H + X) + (6 + o(\zeta)) \frac{e_H}{n^2 p} \mathcal{S}_H \\ &= - (6 + o(1)) \frac{e_H}{n^2 p} X + O(\zeta) \frac{e_H}{n^2 p} \mathcal{S}_H = - (6 + o(1)) \frac{e_H}{n^2 p} X, \end{aligned}$$

where the last inequality used the fact that $\mathcal{S}_H \zeta = o(X)$ whenever $|X| \in I_\Psi$ as then we have $|X| \asymp \mathcal{S}_H \zeta^\delta$ for $\delta < 1$. Next we define what would be a supermartingale when $|X| \in I_\Psi$, namely

$$Z = |X| - \mathcal{S}_H \zeta^\delta.$$

As in (3.11), the one-step change in $\mathcal{S}_H \zeta^\delta$ deterministically satisfies

$$\Delta \mathcal{S}_H \zeta^\delta = - \left(1 + O(1/n^2 p)\right) \frac{6(e_H - \delta)}{n^2 p} \mathcal{S}_H \zeta^\delta = - (6 + o(1)) \frac{e_H - \delta}{n^2 p} \mathcal{S}_H \zeta^\delta, \quad (3.12)$$

hence whenever $|X| \in I_\Psi$ we have

$$\begin{aligned} \mathbb{E}[\Delta Z] &= - (6 + o(1)) \frac{e_H}{n^2 p} |X| + (6 + o(1))(e_H - \delta) \frac{\mathcal{S}_H \zeta^\delta}{n^2 p} \\ &\leq - \frac{6 + o(1)}{n^2 p} \mathcal{S}_H \zeta^\delta \left(\left(1 - \frac{\delta}{2e_H}\right) e_H - (e_H - \delta) \right) \leq - \frac{3\delta + o(1)}{n^2 p} \mathcal{S}_H \zeta^\delta < 0, \end{aligned}$$

with the last inequality valid for sufficiently large n . Thus, $Z(t \wedge \tau_\Psi)$ is indeed a supermartingale when τ_Ψ is the stopping time for having $|X|$ exit the interval I_Ψ .

We turn to establish concentration for $Z(t \wedge \tau_\Psi)$. We will invoke the following inequality due to Freedman [9], which was originally stated for martingales yet its proof extends essentially unmodified to supermartingales.

Theorem 3.4 ([9], Thm 1.6). *Let (S_0, S_1, \dots) be a supermartingale w.r.t. a filtration (\mathcal{F}_i) . Suppose that $S_{i+1} - S_i \leq B$ for all i , and write $V_t = \sum_{i=0}^{t-1} \mathbb{E}[(S_{i+1} - S_i)^2 \mid \mathcal{F}_i]$. Then for any $s, v > 0$*

$$\mathbb{P}(\{S_t \geq S_0 + s, V_t \leq v\} \text{ for some } t) \leq \exp \left\{ - \frac{s^2}{2(v + Bs)} \right\}.$$

To qualify an application of this theorem, we must first establish L^∞ and L^2 bounds on ΔZ , which will follow from individually examining the various edges of H whose potential removal at time $t < t_H^-(\sqrt{\delta})$ may decrement the value of $\Psi_{H,\varphi}$.

For each edge $ab \in E_H$ we let $T = T_{ab}$ denote the subextension of H that has the smallest scaling among subextensions that contain ab (and a minimal number of vertices if there is more than one subgraph with the minimum scaling). Notice that $t_H \leq t_T$ since T is a subextension of H . Furthermore, if K is any subextension of T then, being also a subextension of H we again have $t_H \leq t_K$. In particular, $t < t_K$ for all $K \subset T$ and we can invoke Corollary 3.3 (Part (1)) to give

$$\Psi_{T,\varphi} \leq (1 + o(1)) \mathcal{S}_T (\log n)^{\rho_T},$$

where $\rho_T = 3 \exp(v_T - \iota_T)$.

It is worthwhile explaining why invoking Corollary 3.3 at this point is justified. Indeed, one can prove the statements of both parts of the present theorem (and in effect also its corollary) by gradually extending the time range where they all hold. Formally, we consider the stopping time $\tilde{\tau}$ at which one of the assertions of the theorem fails (for any extension graph). As long as $\tilde{\tau} > t$ we may assume the bounds of this theorem (and hence also those of its corollary) as part of our attempt to increase their validity to the next time step, and so on. The bound we will present for $Z(t \wedge \tau_\Psi)$ will remain valid in conjunction with the additional stopping time $\tilde{\tau}$, and similarly for the second part of the theorem (featuring an analogous proof). Finally, analyzing $Z(t \wedge \tau_\Psi \wedge \tilde{\tau})$ and its Part (ii) counterpart is sufficient for proving the theorem since, by definition, at least one of these supermartingales (for the various extension graphs) will need to violate its corresponding large deviation bound in order for $\tilde{\tau}$ to occur.

Consider the effect of removing a given edge $xy \in E_G$ (as part of some triangle xyz) on $\Psi_{H,\varphi}$. Copies of H are lost iff xy corresponds to some edge $ab \in E_H$ in at least one such copy, hence we can bound this amount by adjusting I_H to include a, b , which would be mapped to x, y . Namely, let $H^* = H_{ab}^*$ be the extension graph with $I_{H^*} = I_H \cup \{a, b\}$ whose edges are E_H minus any edges between vertices of I_{H^*} . Further let φ^* be the corresponding injection extending φ to include x, y as the images of a, b (if this is inconsistent with φ then this copy remains intact) to yield

$$|\Delta\Psi_{H,\varphi}| \leq \sum_{ab \in E_H} \Psi_{H_{ab}^*, \varphi^*}.$$

Recalling the above definition of $T = T_{ab}$, let T_{ab}^* similarly include a, b in its distinguished set of vertices. We then get that

$$|\Delta\Psi_{H,\varphi}| \leq \sum_{ab \in E_H} \sum_{\eta \in \Psi_{T_{ab}^*, \varphi^*}} \Psi_{H_{ab}^*/T_{ab}^*, \eta} = \sum_{ab \in E_H} \sum_{\eta \in \Psi_{T_{ab}^*, \varphi^*}} \Psi_{H/T, \eta}, \quad (3.13)$$

where the equality is due to having $V_T = V_{T^*}$ as well as $E(H) \setminus E(T) = E(H^*) \setminus E(T^*)$ as $a, b \in V_T$. Consider this last term $\Psi_{H/T, \eta}$. We claim that

$$\Psi_{H/T, \eta} \leq (1 + o(1)) \mathcal{S}_{H/T} (\log n)^{\rho_{H/T}}. \quad (3.14)$$

Indeed, by the choice of T we have $\mathcal{S}_H \geq \mathcal{S}_T$ (as H is itself could play the role of T), and so $\mathcal{S}_{H/T} = \mathcal{S}_H / \mathcal{S}_T \geq 1$. Moreover, if K is any subextension of H with $ab \in E_K$ and $V_T \subset V_K$ then again $\mathcal{S}_{K/T} \geq 1$. Yet by definition, if $ab \notin E_K$ then $K/T = \bar{K}/T$ where \bar{K} includes the extra edge ab (as both of its endpoints lie in T) and altogether every $K \subset H$ with $V_T \subset V_K$ has $\mathcal{S}_{K/T} \geq 1$. Noticing that every subextension of H/T can be written as K/T for such K now qualifies an application of Part (1) of Corollary 3.3, from which the desired inequality (3.14) follows.

To conclude the L^∞ -bound on ΔZ we need to estimate $\Psi_{T_{ab}^*, \varphi^*}$. We will argue that this quantity is at most $\log^{O(1)} n$. Let U^* be a subextension of T^* , and let U be the modification of U^* into a subextension of T via setting its distinguished vertex set to I_T and adding any edges among these vertices in T (in particular the edge ab). By construction, $\mathcal{S}_{U^*} / \mathcal{S}_U = \mathcal{S}_{T^*} / \mathcal{S}_T$. In addition, since U contains the edge ab then by the choice of T we must have $\mathcal{S}_U \geq \mathcal{S}_T$. Combining the two we get that $\mathcal{S}_{U^*} \geq \mathcal{S}_{T^*}$ and in particular $\mathcal{S}_{T^*/U^*} \leq 1$, i.e., all the quotients of T^* have scaling at most 1. Part (2) of Corollary 3.3 now implies that

$$\Psi_{T_{ab}^*, \varphi^*} \leq (1 + o(1)) (\log n)^{\rho_{T^*}} < (1 + o(1)) (\log n)^{\rho_T}.$$

Combining this with (3.13) and (3.14) we get that

$$|\Delta\Psi_{T, \varphi}| = O(\mathcal{S}_{H/T} (\log n)^{\rho_{H/T} + \rho_T}),$$

while recalling the estimates for $\Delta\mathcal{S}_H$ and $\Delta\mathcal{S}_H \zeta^\delta$ in (3.11), (3.12) we see that they are negligible in comparison (note that $\mathcal{S}_T \leq n^2 p$ since a valid candidate for T is the induced subgraph on $I_H \cup \{a, b\}$,

which contains at least the edge ab) now extends this to the following L^∞ bound on ΔZ .

$$|\Delta Z| = O\left(\mathcal{S}_{H/T}(\log n)^{\rho_{H/T} + \rho_T}\right).$$

For an L^2 bound on the one-step change ΔZ , again consider the effect of removing an edge corresponding to some $ab \in E_H$. For a given $\psi \in \Psi_{T,\varphi}$, the probability that the next drawn triangle will contain this edge is $Y_{\psi(a)\psi(b)}/Q = (6 + o(1))n^{-2}p^{-1}$ due to our estimates on co-degrees and the total number of triangles present. In particular, the probability to remove an edge corresponding to $ab \in E_H$ in any copy specified by some $\psi \in \Psi_{T,\varphi}$ is at most

$$\Psi_{T,\varphi} O(n^{-2}p^{-1}) = O\left(\mathcal{S}_T(\log n)^{\rho_T} n^{-2}p^{-1}\right).$$

As shown above, this can only modify Z by an additive term of $O(\mathcal{S}_{H/T}(\log n)^{\rho_{H/T} + \rho_T})$. Thus, if we enter the critical interval I_Ψ at some time $t_0 = t(i_0)$ then for any $i \geq i_0$,

$$\begin{aligned} \mathbb{E}[(\Delta Z)^2 \mid \mathcal{F}_i] &\leq \sum_{ab \in E_H} O\left(\mathcal{S}_T(\log n)^{\rho_T} n^{-2}p^{-1} \left(\mathcal{S}_{H/T}(\log n)^{\rho_{H/T} + \rho_T}\right)^2\right) \\ &= O\left(\mathcal{S}_H \mathcal{S}_{H/T} n^{-2}p^{-1} (\log n)^{2\rho_{H/T} + 3\rho_T}\right). \end{aligned}$$

The quantities \mathcal{S}_H and $\mathcal{S}_{H/T}$ above are w.r.t. time $t(i)$, yet our scaling is monotone decreasing in i and so we can replace these by the corresponding scaling terms at t_0 . Summing the result over at most $O(n^2 p)$ remaining possible steps gives

$$\sum_{i \geq i_0} \mathbb{E}[(\Delta Z)^2 \mid \mathcal{F}_i] \leq O\left(\mathcal{S}_H(t_0) \mathcal{S}_{H/T}(t_0) (\log n)^{2\rho_{H/T} + 3\rho_T}\right),$$

where $\mathcal{S}_H(t_0)$ is short for $\mathcal{S}_H(n, p(t_0))$. Since $e_H > 0$ one can bound $\sum_{i \geq i_0} \mathcal{S}_H(t) \mathcal{S}_{H/T}(t) p^{-1}(t)$ via integration to obtain that it is $O(n^2 \mathcal{S}_H(t_0) \mathcal{S}_{H/T}(t_0))$, leading to the above inequality.

Compare this last bound for the summation $\sum_i \mathbb{E}[(\Delta Z)^2 \mid \mathcal{F}_i]$ with the product of the L^∞ bound $B = O(\mathcal{S}_{H/T}(\log n)^{\rho_{H/T} + \rho_T})$ and a deviation of $s = \frac{\delta}{3e_H} \mathcal{S}_H \zeta^\delta$. The upper bound on $\sum_i \mathbb{E}[(\Delta Z)^2 \mid \mathcal{F}_i]$ has order at least $Bs\zeta^{-\delta} (\log n)^{\rho_{H/T} + 2\rho_T}$, clearly dominating Bs . Turning to Freedman's inequality and once again applying the fact that $\mathcal{S}_H = \mathcal{S}_T \mathcal{S}_{H/T}$ we get that for some fixed $c > 0$,

$$\mathbb{P}\left(Z(t \wedge \tau_\Psi) \geq Z(t_0) + s \text{ for some } t \geq t_0\right) \leq \exp\left(-c \mathcal{S}_T(t_0) \zeta^{2\delta} (\log n)^{-2\rho_{H/T} - 3\rho_T}\right).$$

We claim that $\mathcal{S}_T(t_0) \geq (1 - o(1))\zeta^{-3\delta}$, and in fact this holds for any $t_0 \leq t \leq t_H^-(\sqrt{\delta})$. If $m_T < m_H$ then $t < t_T^-(\sqrt{\delta})$ since $t \leq t_H < t_T$ and given our definition of δ . Moreover, $\sqrt{\delta} \leq \frac{1}{3}$ (as implied by Eq. (3.10) and the fact that $e_A \geq 1$) thus in this case

$$\mathcal{S}_T \geq (1 - o(1))\zeta^{-\sqrt{\delta}} \geq (1 - o(1))\zeta^{-3\delta}. \quad (3.15)$$

On the other hand, when $m_T = m_H$ we can write

$$\mathcal{S}_T = (n^{1/m_T} p)^{e_T} = (n^{1/m_H} p)^{e_T} = (\mathcal{S}_H)^{e_T/e_H} \geq \mathcal{S}_H^{1/e_H} \geq (1 - o(1))\zeta^{-\sqrt{\delta}/e_H},$$

with the final inequality due to $t \leq t_H^-(\sqrt{\delta})$. By (3.10) we have $\sqrt{\delta} \leq 1/(3e_H)$, hence (3.15) is again valid. Altogether we can conclude that

$$\begin{aligned} \mathbb{P}\left(Z(t \wedge \tau_\Psi) \geq Z(t_0) + s \text{ for some } t \geq t_0\right) &\leq \exp\left(-(c - o(1))\zeta^{-\delta} (\log n)^{-2\rho_{H/T} - 3\rho_T}\right) \\ &\leq \exp(-2 \log^2 n), \end{aligned}$$

where the last inequality is thanks to (3.2). Eq. (3.15) similarly implies that

$$\mathcal{S}_{H/T}(\log n)^{O(1)} = \frac{\mathcal{S}_H}{\mathcal{S}_T}(\log n)^{O(1)} \leq \mathcal{S}_H \zeta^{3\delta - o(1)},$$

that is, the L^∞ -bound on ΔZ , which is valid regardless of whether $|X| \in I_\Psi$, satisfies $B = o(\mathcal{S}_H \zeta^\delta)$. Accounting for the first step whenever $|X|$ enters the interval I_Ψ we get

$$Z(t_0) + s \leq -\frac{\delta}{2e_H} \mathcal{S}_H \zeta^\delta + B + \frac{\delta}{3e_H} \mathcal{S}_H \zeta^\delta < 0$$

for sufficiently large n . In particular, $|X| < \mathcal{S}_H \zeta^\delta$ for all $t \geq t_0$ except with probability $\exp(-2 \log^2 n)$. This probability allows us to absorb a union bound over all $n^{O(1)}$ choices for (H, φ) and t_0 , thus concluding the proof of Part (i) of the theorem.

Proof of Part (ii). Let H and φ be fixed, and let $X = \Psi_{H,\varphi} - \mathcal{S}_H$. We follow the framework of applying a martingale argument within an appropriate critical interval for $|X|$, which would now be

$$I_\Psi = \left[\left(1 - \frac{\sqrt{\delta}}{2}\right) (\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^{\rho_H}, (\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^{\rho_H} \right],$$

where $\rho_H = 3 \exp(v_H - \iota_H)$. From Part (i) of the theorem we get that $|X| \leq \mathcal{S}_H \zeta^\delta$ at time $t = t_H^-(\sqrt{\delta})$. Yet at this time $\mathcal{S}_H \leq \zeta^{-\sqrt{\delta}}$ by definition and so $\mathcal{S}_H^{\sqrt{\delta}} \zeta^\delta \leq 1$. We thus infer that $|X| \leq \mathcal{S}_H \zeta^\delta \leq \mathcal{S}_H^{1-\sqrt{\delta}}$, and in particular $|X| \notin I_\Psi$ at time $t_H^-(\sqrt{\delta})$ provided that n is large enough.

As in the proof of the first part of the theorem,

$$\mathbb{E}[\Delta \Psi_{H,\varphi}] = -(6 + O(\zeta)) \frac{e_H}{n^2 p} \Psi_{H,\varphi}, \quad \Delta \mathcal{S}_H = -(6 + o(\zeta)) \frac{e_H}{n^2 p} \mathcal{S}_H,$$

and so

$$\mathbb{E}[\Delta X] = -(6 + o(1)) \frac{e_H}{n^2 p} X + O(\zeta) \frac{e_H}{n^2 p} \mathcal{S}_H = -(6 + o(1)) \frac{e_H}{n^2 p} X,$$

where the last equality will be justified by showing that $\mathcal{S}_H \zeta = o(|X|)$ when $|X| \in I_\Psi$. Indeed, with $0 < \delta < 1$ a sufficient condition for $\mathcal{S}_H \zeta = o((\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^{\rho_H})$ is to have $\mathcal{S}_H = o(\zeta^{-1} (\log n)^{\rho_H})$, which in turn holds since $t \geq t_H^-(\sqrt{\delta})$ and thanks to the assumption in (3.2) on the decay rate of ζ .

Similarly to (3.11) we have $\Delta(\mathcal{S}_H)^{1-\sqrt{\delta}} = -(6 + o(\zeta))(1 - \sqrt{\delta}) \frac{e_H}{n^2 p} (\mathcal{S}_H)^{1-\sqrt{\delta}}$ deterministically, hence the random variable defined by

$$Z = |X| - (\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^{\rho_H}$$

satisfies the following whenever $|X| \in I_\Psi$ and for any sufficiently large n :

$$\mathbb{E}[\Delta Z] = -(6 + o(1)) \frac{e_H}{n^2 p} |X| + (6 + o(1))(1 - \sqrt{\delta}) \frac{e_H}{n^2 p} (\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^{\rho_H} < 0.$$

Therefore, if $t_0 = t(i_0)$ is the first time where $|X| > (1 - \frac{\sqrt{\delta}}{2})(\mathcal{S}_H)^{1-\sqrt{\delta}} (\log n)^{\rho_H}$ and $\tau_\Psi = \min\{t \geq t_0 : |X(t)| \notin I_\Psi\}$ then $Z(t \wedge \tau_\Psi)$ is a supermartingale.

For L^∞ and L^2 bounds on ΔZ we again wish to analyze the effect of removing some edge $xy \in E_G$ playing the role of some $ab \in E_H$ in a copy of H . Letting $T = T_{ab}$ as before be the subextension of H that has the smallest scaling among subextensions that have $ab \in E_T$, we now claim that in the setting of this part of the theorem we in fact have $T = H$. Indeed, since H is strictly balanced, any $K \subsetneq H$ must have $t_H < t_K$ and thus also $t_H < t_K^-(\sqrt{\delta})$. The hypothesis $t_H^-(\sqrt{\delta}) \leq t \leq t_H$ then implies that $\mathcal{S}_H \leq \zeta^{-\sqrt{\delta}}$ and yet $\mathcal{S}_K > \zeta^{-\sqrt{\delta}}$, precluding the choice $T_{ab} = K$.

Revisiting the argument used to prove the previous part of the theorem, we let H_{ab}^* denote the extension graph that adds a, b to I_H and let φ^* denote a corresponding injection. We then have

$$|\Delta \Psi_{H,\varphi}| \leq \sum_{ab \in E_H} \Psi_{H_{ab}^*, \varphi^*}$$

and yet, the by the minimality of $\mathcal{S}_{H_{ab}^*}$ we must have $\mathcal{S}_{H_{ab}^*/K^*} \leq 1$ for any subextension $K^* \subset H_{ab}^*$ and so Part (2) of Corollary 3.3 implies that

$$\Psi_{H_{ab}^*, \varphi^*} \leq (1 + o(1)) (\log n)^{\rho_{H^*}}.$$

Combining these two equations yields that the potential change in $\Psi_{H,\varphi}$ is at most

$$|\Delta\Psi_{H,\varphi}| = O\left(\max_{ab \in E_H} (\log n)^{\rho_{H^*ab}}\right).$$

Similarly to (3.11), the quantities \mathcal{S}_H and $(\mathcal{S}_H)^{1-\sqrt{\delta}}(\log n)^{\rho_H}$ deterministically change by at most

$$O(\mathcal{S}_H n^{-2} p^{-1} (\log n)^{\rho_H}) = O(\mathcal{S}_H \zeta^2 (\log n)^{2+\rho_H}) = o(1)$$

due to the fact that $\mathcal{S}_H \leq \zeta^{-\sqrt{\delta}}$ for $\delta < 1$ and the decay rate of ζ . Altogether we get

$$|\Delta Z| = O\left(\max\{(\log n)^{\rho_{H^*ab}} : ab \in E_H\}\right) = O\left((\log n)^{(1/e)\rho_H}\right), \quad (3.16)$$

where we used the fact that $\iota_{H^*} > \iota_H$ since at least one of a, b is in $I_{H^*} \setminus I_H$ (as otherwise $ab \notin E_H$ as I_H is an independent set), and so $\rho_H = 3e^{\nu_H - \iota_H} \geq e\rho_{H^*}$. Note that this L^∞ -bound holds in general regardless of whether $|X| \in I_\Psi$. In particular, since $|X| \leq (1 - \frac{\sqrt{\delta}}{2})(\mathcal{S}_H)^{1-\sqrt{\delta}}(\log n)^{\rho_H}$ one step prior to t_0 and can only increase by $O((\log n)^{(1/e)\rho_H})$ at a given step, we have

$$\begin{aligned} |X(t_0)| &= (1 - \frac{\sqrt{\delta}}{2})(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H} + O\left((\log n)^{(1/e)\rho_H}\right) \\ &= (1 - \frac{\sqrt{\delta}}{2} + o(1))(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}, \end{aligned} \quad (3.17)$$

(where the term $(\log n)^{(1/e)\rho_H}$ was absorbed in the higher order $(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}$ -term) and similarly, for all $t_0 \leq t < \tau_\Psi$

$$|X(t)| \leq (1 + o(1))(\mathcal{S}_H(t))^{1-\sqrt{\delta}}(\log n)^{\rho_H}.$$

As such, since the probability to remove an edge that corresponds to some $ab \in E_H$ in a given copy $\psi \in \Psi_{H,\varphi}(t)$ is at most $Y_{\psi(a)\psi(b)}/Q = (6 + o(1))n^{-2}p^{-1}$, the overall probability for this event to occur for some $\psi \in \Psi_{H,\varphi}(t)$ is at most

$$\Psi_{H,\varphi}(t)O(n^{-2}p^{-1}) \leq O\left(\left(\mathcal{S}_H(t) + (\mathcal{S}_H(t))^{1-\sqrt{\delta}}(\log n)^{\rho_H}\right)n^{-2}p^{-1}\right).$$

This implies that for any $t = t(i)$ for $i \geq i_0$ we have that $\mathbb{E}[(\Delta Z)^2 | \mathcal{F}_i]$ is at most

$$\sum_{ab \in E_H} O\left(\left(\mathcal{S}_H(t) + (\mathcal{S}_H(t))^{1-\sqrt{\delta}}(\log n)^{\rho_H}\right)n^{-2}p(t)^{-1}(\log n)^{2e^{-1}\rho_H}\right),$$

and integrating this over the $O(n^2 p(t_0))$ remaining possible steps (again noting that $e_H(1 - \sqrt{\delta}) > 0$ and so $\sum_{i \geq i_0} (\mathcal{S}_H)^{1-\sqrt{\delta}} p^{-1} = O(n^2 (\mathcal{S}_H(t_0))^{1-\sqrt{\delta}})$ as in the proof of Part (i) of this theorem) now gives

$$\begin{aligned} \sum_{i \geq i_0} \mathbb{E}[(\Delta Z)^2 | \mathcal{F}_i] &\leq \sum_{ab \in E_H} O\left(\left(\mathcal{S}_H(t_0) + (\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}\right)(\log n)^{2e^{-1}\rho_H}\right) \\ &= O\left(\left(\mathcal{S}_H(t_0) + (\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}\right)(\log n)^{2e^{-1}\rho_H}\right). \end{aligned} \quad (3.18)$$

Compare this bound with the product of the L^∞ bound given in (3.16) and a deviation value of $s = \frac{\sqrt{\delta}}{4}(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}$. This product has order $(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{(1+e^{-1})\rho_H}$ whereas our bound in (3.18) has order at least $(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{(1+2e^{-1})\rho_H}$, strictly larger. We can thus extract from Freedman's inequality that for some fixed $c > 0$,

$$\mathbb{P}\left(\exists t > t_0 : Z(t \wedge \tau_\Psi) \geq Z(t_0) + s\right) \leq \exp\left(-c \frac{(\mathcal{S}_H(t_0))^{2-2\sqrt{\delta}}(\log n)^{(2-2e^{-1})\rho_H}}{\mathcal{S}_H(t_0) + (\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}}\right)$$

which, according to the larger term between $\{\mathcal{S}_H, (\mathcal{S}_H)^{1-\sqrt{\delta}}(\log n)^{\rho_H}\}$, is at most

$$\exp\left(-c/2 \min\left\{(\mathcal{S}_H)^{1-2\sqrt{\delta}}(\log n)^{(2-2e^{-1})\rho_H}, (\mathcal{S}_H)^{1-\sqrt{\delta}}(\log n)^{(1-2e^{-1})\rho_H}\right\}\right).$$




FIGURE 2. The set of triangular ladders \mathcal{L}_π (labeled by π) for all $|\pi| \leq 4$.

Going back to the definition $\rho_H = 3 \exp(v_H - \iota_H)$, since $v_H - \iota_H \geq 1$ (otherwise H is edgeless) we have $\rho_H \geq 3e$ and thus the exponent of each of the two $\log n$ terms above is at least $3(e-2) > 2$ and we can deduce that for any sufficiently large n we have

$$\mathbb{P}(\exists t > t_0 : Z(t \wedge \tau_\Psi) \geq Z(t_0) + s) \leq \exp(-2 \log^2 n).$$

Finally, recall that $Z(t_0) = -(\frac{\sqrt{\delta}}{2} - o(1))(\mathcal{S}_H(t_0))^{1-\sqrt{\delta}}(\log n)^{\rho_H}$ as guaranteed by (3.17). We can thus conclude that $Z(t) < 0$, that is $|X(t)| < (\mathcal{S}_H(t))^{1-\sqrt{\delta}}(\log n)^{\rho_H}$, for all $t \geq t_0$ and large enough n with probability at least $1 - \exp(-2 \log^2 n)$. A union bound extends this result to all $n^{O(1)}$ choices for t_0 and the pairs (H, φ) with probability at least $1 - \exp(-\log^2 n)$, as required. ■

4. TRIANGULAR LADDERS AND THEIR EXTENSIONS

Definition 4.1. [triangular ladders] Let $\Pi \subset \{e, 0, 1\}^*$ denote the set of all words which do not have 0 in indices 1 and 2, have at most 2 occurrences of e and do not contain the substrings $e0$, $e10$ or $e1e$. For any $\pi \in \Pi$ of length $k \geq 1$, the **triangular ladder** \mathcal{L}_π is the following labeled undirected graph on the vertex set $\{0, 1, 2, \dots, k+1\}$:

- (i) $k = 1$: \mathcal{L}_1 has the two edges $\{02, 12\}$; \mathcal{L}_e has the single edge 12 .
- (ii) $k \geq 2$: let $H = \mathcal{L}_{\pi^-}$ where π^- is the $(|\pi| - 1)$ -prefix of π . Let \mathcal{L}_π be the graph obtained from H as follows: First add the vertex $k+1$ and the edge $(k, k+1)$. In addition, if $\pi(k) = 1$ add the edge $(k-1, k+1)$, whereas if $\pi(k) = 0$ add the edge $(v, k+1)$, where v is the unique neighbor of k in H other than $k-1$.

Observe that vertex 0 is isolated iff $\pi(1) = e$. The set of triangular ladders of length at most 5 is illustrated in Fig. 2. Implicit in the above definition is the fact that by construction, the last vertex of \mathcal{L}_π , i.e., the vertex $k+1$ when $|\pi| = k$, has precisely two neighbors if $\pi(k) \in \{0, 1\}$ (one of which is the vertex k) and exactly one neighbor (the vertex k) when $\pi(k) = e$. In particular, since the substring $e0$ cannot appear in $\pi \in \Pi$, the case $\pi(k) = 0$ in Item (ii) above has $\pi(k-1) \neq e$ and so the vertex k indeed has two neighbors in H . Further note that if $\pi(\ell) = e$ for $\ell > 1$ then the vertex ℓ is a cut-vertex of \mathcal{L}_π . Namely, in this case there is no edge $(a, b) \in E(\mathcal{L}_\pi)$ such that $a < \ell < b$.

Definition 4.2. The **scaling** of a triangular ladder \mathcal{L}_π w.r.t. some given n and p is defined to be $\mathcal{S}_\pi = \mathcal{S}_\pi(n, p) = n^{v_\pi - 2} p^{e_\pi}$, where $e_\pi = |E(\mathcal{L}_\pi)|$ and $v_\pi = |V(\mathcal{L}_\pi)|$.




FIGURE 3. Fans in triangular ladders: The vertices highlighted in red form 4-fans at vertices 2, 4 and 5 (note the exceptional structure of the fan at $a = 4$ due to having $\pi(a - 2) = \mathbf{e}$).

The scaling \mathcal{S}_π will correspond asymptotically to the expected number of labeled copies of \mathcal{L}_π rooted at two given vertices (matching the labels 0 and 1) along our process. Recalling the notion of extension graphs (see Definition 3.1), \mathcal{S}_π is simply the scaling of the extension graph obtained from \mathcal{L}_π via the distinguishing vertex set $I = \{0, 1\}$. Observe that if $|\pi| = k$ and π contains $0 \leq s \leq 2$ copies of \mathbf{e} then $\mathcal{S}_\pi = n^k p^{2k-s}$.

Our focus will be on tracking a special subset of the family of all triangular ladders, whose formulation will require the following definition.

Definition 4.3 (*f-fan*). For $f \geq 3$, an *f-fan* at some vertex $a > 0$ in a triangular ladder \mathcal{L}_π is a sequence of vertices that has one of the following two forms:

- (1) $\pi(a - 2) \neq \mathbf{e}$ and vertices $a + 1, a + 2, \dots, a + f + 1$ are all adjacent to the vertex a .
- (2) $\pi(a - 2) = \mathbf{e}$ and vertices $a - 1, a + 1, \dots, a + f$ are all adjacent to the vertex a .

Equivalently, \mathcal{L}_π has an *f-fan* at $a > 0$ with $\pi(a - 2) \neq \mathbf{e}$ if $\pi(a + 1) = 1$ and $\pi(a + i) = 0$ for $2 \leq i \leq f$. Similarly, \mathcal{L}_π has an *f-fan* at $a > 0$ with $\pi(a - 2) = \mathbf{e}$ if $\pi(a + 1) = 1$ and $\pi(a + i) = 0$ for $2 \leq i \leq f - 1$.

Note that there cannot be an *f-fan* at 0, nor can there be one at a vertex a whenever $\pi(a) = \mathbf{e}$. Figure 3 illustrates the structure of *f-fans* as defined above.

Definition 4.4. [*M-bounded triangular ladders*] For $M \geq 3$, define \mathcal{B}_M to be the set of all $\pi \in \Pi$ that do not contain an *M-fan* and in addition satisfy one of the following conditions:

- (a) π contains no copy of \mathbf{e} and has length at most $3M - 1$, or
- (b) π contains 1 copy of \mathbf{e} , has length at most $2M$ and does not have $\pi(2M) = \mathbf{e}$, or
- (c) π contains 2 copies of \mathbf{e} , has length at most $M + 1$ and does not have $\pi(M + 1) = \mathbf{e}$.

If $\pi \in \mathcal{B}_M$ we say that \mathcal{L}_π is an *M-bounded triangular ladder*.

It is easily seen that $\mathcal{B}_M \subset \mathcal{B}_{M'}$ for any $M \leq M'$. The following additional observation demonstrates some of the reasons for the precise choice of constants in the definition of \mathcal{B}_M .

Observation 4.5. The family \mathcal{B}_M is closed under the prefix operation. Indeed, let σ be a proper prefix of some $\pi \in \mathcal{B}_M$. Precluding an *M-fan* from π clearly also precludes it from σ . While σ may contain fewer copies of \mathbf{e} , these would meet relaxed length criteria in Definition 4.4, hence $\sigma \in \mathcal{B}_M$.

Having defined the family of graphs we wish to track throughout the process, we proceed to a classification of the edges in \mathcal{L}_π for $\pi \in \mathcal{B}_M$. This will in turn be used to study the effect of removing such edges as part of triangles eliminated by our process.

Definition 4.6 (boundary edges). Let \mathcal{L}_π be an *M-bounded triangular ladder* for some $M \geq 3$ and let $yz \in E(\mathcal{L}_\pi)$ be such that $0 < y < z$ and $\pi(y) \neq \mathbf{e}$. Let

$$\pi' = \begin{cases} \pi_{z-1} \circ 0 & \text{if } y < z - 1 \\ \pi_{z-1} \circ 1 & \text{if } y = z - 1, \end{cases}$$




FIGURE 4. Examples of M -bounded triangular ladders for $M = 5$. Outer boundary edges w.r.t. \mathcal{B}_5 are highlighted in red, side boundary edges are green and initial edges are brown.

where π_ℓ denotes the ℓ -prefix of π and \circ denotes string concatenation. We say yz is a **boundary edge** of \mathcal{L}_π w.r.t. \mathcal{B}_M iff $\pi' \notin \mathcal{B}_M$. In other words, an edge is a boundary edge if the ladder obtained by removing all vertices larger than z and then attaching a new vertex adjacent to y, z no longer belongs to the family \mathcal{B}_M . We classify a boundary edge yz into one of two categories:

- **outer boundary edges:** $z = |\pi| + 1$ (last vertex in \mathcal{L}_π) and in addition $|\pi|$ is maximal w.r.t. its number of copies of \mathbf{e} (i.e., $|\pi|$ is $3M - 1$, $2M$ or $M + 1$ if it has 0, 1 or 2 such copies resp.).
- **side boundary edges:** all other boundary edges. Necessarily $y < z - 1$, there is at most 1 copy of \mathbf{e} in π , the ladder \mathcal{L}_π has an $(M - 1)$ -fan at y whilst $\mathcal{L}_{\pi'}$ has an M -fan at y .

To see that indeed side boundary edges are accompanied by π of the form specified in the above definition ($y < z - 1$, at most 1 copy of \mathbf{e} and an M -fan in π') argue as follows. First, we claim that if π has 2 copies of \mathbf{e} then \mathcal{L}_π cannot have side boundary edges. To show this, recall that in this case $|\pi| \leq M + 1$ and there are 2 copies of \mathbf{e} in π_M . After excluding the case $|\pi| = M + 1$, $z = M + 2$ (corresponding to an outer boundary edge) we are left with either $z = M + 1$ or $z < M + 1$. In the former case, $|\pi'| = M + 1$ and π' retains both \mathbf{e} copies of π in its M -prefix, thus $\pi' \in \mathcal{B}_M$ by Item (c) of Definition 4.4. In the latter, $|\pi'| \leq M$, and either it contains 2 copies of \mathbf{e} — here $\pi' \in \mathcal{B}_M$ as per the value of $|\pi'|$ — or it contains 1 copy of \mathbf{e} and again $\pi' \in \mathcal{B}_M$ since $\mathcal{L}_{\pi'}$ does not contain an M -fan (any f -fan in π' would have π' be at least $f + 2$ symbols long). We have thus established that π can have at most 1 copy of \mathbf{e} . Next, recall that by definition π has no M -fan. The exclusion of the cases $|\pi| = 2M$, $z = |\pi| + 1$ with 1 copy of \mathbf{e} and $|\pi| = 3M - 1$, $z = |\pi| + 1$ with no copies of \mathbf{e} (corresponding to outer boundary) implies that $|\pi'| \leq |\pi|$, and yet $\pi' \notin \mathcal{B}_M$ (since yz is a boundary edge), hence it must contain an M -fan. Comparing π to π' now implies that $\pi' = \pi_{z-1} \circ 0$ so as to accommodate a new M -fan (missing from π_{z-1}), therefore indeed $y < z - 1$. Finally, it further follows from this analysis that

$$z = \begin{cases} y + M & \text{if } \pi(y - 2) \neq \mathbf{e}, \\ y + M - 1 & \text{if } \pi(y - 2) = \mathbf{e}. \end{cases} \quad (4.1)$$

We complement the definition of boundary edges by categorizing the each of the remaining edges into one of two types:

Definition 4.7 (non-boundary edge classes). *Let \mathcal{L}_π be an M -bounded triangular ladder and let yz be a non-boundary edge of \mathcal{L}_π . We say that yz is an **initial edge** if $y < z$ and $\pi(y) = \mathbf{e}$, and otherwise we say that yz is an **interior edge**.*

Figure 4 depicts the family of 3-bounded triangular ladders, highlighting, outer boundary edges, side boundary edges and initial edges in these ladders w.r.t. \mathcal{B}_3 .

Remark. While f -fans and M -bounded triangular ladders were both defined for values of $f, M \geq 3$, one could extend these notions to $f = M = 2$. However, going back to Definition 4.3, a 2-fan at vertex 3 occurs whenever π begins with $\mathbf{e}111$, as opposed to f -fans for $f > 2$ which are accompanied by non-overlapping patterns of the form $10\dots 0$. In light of this anomaly at $M = 2$, and given that our application would have M be arbitrarily large, our attention will be restricted to $f, M \geq 3$.

Definition 4.8 (backward extension). Let \mathcal{L}_π be a triangular ladder that contains some outer boundary edge yz . The **backward extension from yz to \mathcal{L}_π** , denoted by $B_{\pi,yz}$, is defined as the graph obtained by deleting from \mathcal{L}_π all edges within the distinguished vertex set $I = \{0, 1, y, z\}$.

Note that we have the following scalings for these extension graphs (these will be derived later).

$$\mathcal{S}_{B_{\pi,yz}} = \begin{cases} n^{3M-3}p^{6M-3} & \text{if } |\pi| = 3M - 1, \\ n^{2M-2}p^{4M-2} & \text{if } |\pi| = 2M, \\ n^{M-1}p^{2M-1} & \text{if } |\pi| = M + 1. \end{cases} \quad (4.2)$$

As in Definition 3.1, the density of a backward extension graph H is equal to $m_H = e_H/(v_H - 4)$, as the distinguished vertex set of H is $I = \{0, 1, y, z\}$ for some outer boundary edge yz . To qualify an application of Theorem 3.2 on backward extension graphs we need to verify that they are balanced (that is, the density of a backward extension graph H is as large as the density of any subextension K of H).

Lemma 4.9. Let \mathcal{L}_π be an M -bounded triangular ladder for some $M \geq 3$, and let yz be an outer boundary edge of \mathcal{L}_π . Then the backward extension graph $B_{\pi,yz}$ is balanced.

Proof. With Definition 4.6 in mind, without loss of generality we have $z = |\pi| + 1$ where $|\pi|$ is either $3M - 1$, $2M$ or $M + 1$ depending on whether it has 0, 1 or 2 occurrences of \mathbf{e} , respectively. We will first prove the statement of the lemma for the case where $\pi(1) = \mathbf{e}$.

Suppose for the sake of contradiction that $B = B_{\pi,yz} = (V_B, E_B)$ is not balanced, and let $C = (V_C, E_C)$ be a maximal proper subextension of B such that $m_C > m_B$. Let a, b be vertices such that $[a, b)$ is the first interval of vertices missing from V_C , that is, $j \in V_C$ for $1 \leq j < a$ and for $j = b$ whereas $j \notin V_C$ for $a \leq j < b$. Observe that $2 \leq a < b \leq z = |V_B|$ since $\{1, y, z\} \subset V_C \subsetneq V_B$. By the maximality assumption, the extension graph we get by adding $a, \dots, b - 1$ (and all incident edges) to C has density at most m_B . Let e_0 be the number of edges within $a, \dots, b - 1$ or connecting these vertices to V_C in B (or equivalently in \mathcal{L}_π , since $\{0, 1, y, z\} \subset V_C$ and so the two edge-sets are equal). Since $m_C > m_B$ we now get

$$\frac{e_0}{b - a} < m_B. \quad (4.3)$$

As each vertex that joins a triangular ladder has either one or two neighbors among the previous vertices, one of which is always the preceding vertex, we also have

$$e_0 \geq 2(b - a) + 1 - \#\{a \leq j \leq b - 1 : \pi(j - 1) = \mathbf{e}\}. \quad (4.4)$$

In this observation we included the edge $(b - 1, b)$ (the $+1$ term in the right hand side) yet disregarded another potential edge joining b to some $\gamma < b - 1$ (which is present iff $\pi(b - 1) \neq \mathbf{e}$). We call γb the ‘extra’ edge, and a careful consideration of when γb exists will play a role below. Consider now 2 cases depending on the structure of π .

(1) π contains one copy of \mathbf{e} :

Here we have $|V_B| = |V(\mathcal{L}_\pi)| = 2M + 2$ and $e_\pi = 2(|V_B| - 2) - 1 = 4M - 1$. We claim that $1y, 1z \notin E(\mathcal{L}_\pi)$ and thus the extension graph B satisfies $E_B = E(\mathcal{L}_\pi) \setminus \{yz\}$. Indeed, as the pattern $\mathbf{e}10$ is forbidden from any $\pi \in \Pi$, the neighborhood of the vertex 1 is always a subset of $\{2, 3\}$. Since the form of π dictates that $z = 2M + 1$ and that \mathcal{L}_π has no M -fan, in particular $y \geq M + 1 \geq 4$ and so neither y nor z are neighbors of 1, as claimed. Altogether, $|E_B| = e_\pi - 1 = 4M - 2$ and

$$m_B = \frac{|E_B|}{|V_B| - 4} = 2 + \frac{1}{M - 1}. \quad (4.5)$$

Further note that the extra edge γb with $\gamma < b - 1$ is present since $b - 1 \geq a \geq 2$ and so $\pi(b - 1) \neq \mathbf{e}$. Several cases are now possible.

Case (1.1): $a > 2$. Then the symbol \mathbf{e} does not appear in $\pi(a-1), \dots, \pi(b-2)$. It follows from (4.3),(4.4),(4.5) that $b-a > M-1$, and as $3 \leq a < b \leq 2M+1$ we get

$$M \leq b-a \leq 2M-2. \quad (4.6)$$

Examine the extra edge γb . By Definition 4.4, the ladder \mathcal{L}_π does not contain an M -fan, implying that $\gamma \geq b-M$ (this inequality is tight if there is an $(M-1)$ -fan at γ). Combining this with the fact $b-a \geq M$ yields $a \leq \gamma < b-1$, hence we can adjust (4.4) by the extra edge γb to have $e_0 \geq 2(b-a)+2$. By (4.6) we now get

$$\frac{e_0}{b-a} \geq 2 + \frac{2}{b-a} \geq 2 + \frac{1}{M-1} = m_B, \quad (4.7)$$

contradicting (4.3).

Case (1.2): $a = 2$ and $b > 3$. Note that here $b < 2M+1$ as otherwise $V_C = \{0, 1, 2M+1\}$ contradicting the fact that $\{0, 1, y, z\} \subset V_C$. Since $\pi(1) = \mathbf{e}$, we have a contribution of -1 from the last term in (4.4). On the other hand, since $b > 3$, the extra edge γb must feature $2 \leq \gamma \leq b-2$ (the vertex 1 has only $\{2, 3\}$ as neighbors in particular excluding $b \geq 4$). Taking this into account we can modify (4.4) into $e_0 \geq 2(b-a)+1$. At this point we can immediately preclude the case $b+1 \in V_C$. Indeed, if $b+1 \in V_C$ then this vertex has 2 neighbors among previous vertices (as $\pi(b) \neq \mathbf{e}$), one of which must be among $2, \dots, b-1$, thus $e_0 \geq 2(b-a)+2$. At the same time, having $b+1 \leq |V_B| = 2M+1$ implies that $b-a \leq 2M-2$ and so inequality (4.7) is again valid, contradicting (4.3). Assume therefore that $b+1 \notin V_C$ and let b' be the smallest index greater than b such that $b' \in V_C$. Let e_1 be the number of edges among $2, \dots, b-1, b+1, \dots, b'-1$ plus those joining this set with V_C . In place of (4.3) we now write

$$\frac{e_1}{b'-3} < m_B = 2 + \frac{1}{M-1},$$

and accounting for the edge 12, two edges going backwards from each of the vertices $3 \leq j \leq b'-1$ and one edge going back from b' we arrive at $e_1 \geq 2b'-4$. This implies that $e_1/(b'-3) \geq 2 + 2/(b'-3) \geq m_B$, a contradiction.

Case (1.3): $a = 2$ and $b = 3$. Here, the first interval missing from V_C is simply the singleton $\{2\}$. Recall that $|V_B| = 2M+2 \geq 8$ and that π begins with the prefix $\mathbf{e}11$ (thus the vertex 2 is adjacent to $\{1, 3, 4\}$ to the very least in \mathcal{L}_π). Clearly, if $4 \in V_C$ then $e_0/(b-a) \geq 3 > m_B$ in contradiction with (4.3). Assume therefore that $4 \notin V_C$ and let b' be the smallest index greater than 3 such that $b' \in V_C$. Again let e_1 be the number of edges among the vertex $2, 4, \dots, b'-1$ plus the number of edges joining these vertices with V_C . Instead of (4.3) we have

$$\frac{e_1}{b'-3} < m_B = 2 + \frac{1}{M-1}. \quad (4.8)$$

Counting the edges 12, 23, two edges going backward from each $4 \leq j \leq b'-1$ as well as the edge $(b'-1, b')$ we get $e_1 \geq 2b'-5$. Therefore, it follows from (4.8) that $b' > M+2$. This case will be concluded by examining the extra edge $\gamma'b'$ such that $\gamma' < b'-1$ (present since $\pi(b'-1) \neq \mathbf{e}$). Crucially, \mathcal{L}_π does not contain an M -fan at the vertex 3 and so $\gamma' \neq 3$ (if f is the maximal f -fan at 3 then by definition the neighbors of vertex 3 are precisely $\{2, 4, \dots, 3+f\}$, a set precluding b'). Since $b' \geq M+3 \geq 6$ we further have $\gamma' \neq 1$ and overall the extra edge $\gamma'b'$ further increments e_1 yielding $e_1 \geq 2b'-4$. Since $b' \leq 2M+1$ this gives $e_1/(b'-3) \geq 2 + 2/(b'-3) \geq 2 + 1/(M-1)$, contradicting (4.8).

(2) π contains two copies of \mathbf{e} :

Here $|V_B| = |V(\mathcal{L}_\pi)| = M+3$ and $e_\pi = 2(|V_B|-2) - 2 = 2M$. Again we claim that $1y, 1z \notin E(\mathcal{L}_\pi)$ and hence $E_B = E(\mathcal{L}_\pi) \setminus \{yz\}$. To see this, observe that when π begins with \mathbf{ee} the vertex 1 has only 2 as its neighbor, which has a subset of $\{1, 3, 4\}$ as its neighbors due to the forbidden pattern $\mathbf{e}10$. Therefore, in this case having $z = M+2 \geq 5$ implies that $y > 2$ and $1y \notin E(\mathcal{L}_\pi)$. Alternatively, if π has 2 copies of \mathbf{e} and yet $\pi(2) \neq \mathbf{e}$, then the index $j > 1$

such that $\pi(j) = \mathbf{e}$ in fact satisfies $j > 3$ (and hence is not a neighbor of 1) due to the forbidden patterns $\mathbf{e}0$ and $\mathbf{e}1\mathbf{e}$, and moreover $j < M + 1$ due to Definition 4.4. The fact that the vertex j is a cut-vertex of \mathcal{L}_π now leads to the conclusion that $|E_B| = e_\pi - 1 = 2M - 1$, thus

$$m_B = \frac{|E_B|}{|V_B| - 4} = 2 + \frac{1}{M - 1}.$$

Case (2.1): There are no copies of \mathbf{e} in $\{\pi(a - 1), \dots, \pi(b - 2)\}$. Observe that $a \geq 3$ since $\pi(1) = \mathbf{e}$. By (4.4) we have

$$e_0/(b - a) \geq 2 + 1/(b - a) \geq 2 + 1/(M - 1) \quad (4.9)$$

(where we used the fact that $b - a \leq M - 1$ as $3 \leq a < b \leq M + 2$), contradicting (4.3).

Case (2.2): $a \geq 3$ and the symbol \mathbf{e} appears (once) in $\{\pi(a - 1), \dots, \pi(b - 2)\}$. Here there are no later occurrences of \mathbf{e} and so the extra edge γb (with $\gamma < b - 1$) is present. Moreover, $\gamma \geq \ell$ where $a - 1 \leq \ell \leq b - 2$ is the aforementioned index of \mathbf{e} in π , since, crucially, vertex ℓ is a cut-vertex of \mathcal{L}_π . It now follows that when $\ell \geq a$ the extra edge γb contributes to the count of e_0 and gives $e_0 \geq 2(b - a) + 1$, which as shown above (see (4.9)) contradicts (4.3). Conversely, if $\pi(a - 1) = \mathbf{e}$ then the forward neighbors of $a - 1$ are confined to the subset $\{a, a + 1\}$. Consequently, as long as $b > a + 1$ we again get that $\gamma \geq a$ and the extra edge γb again leads to the sought inequality $e_0 \geq 2(b - a) + 1$ that yields a contradiction to (4.3). It remains to deal with the case $\pi(a - 1) = \mathbf{e}$ and $b = a + 1$, where we claim that $b < z$. To see this, observe that if b were the last vertex of C (vertex $z = M + 2$) then its neighbors would be $\{a - 1, a\}$. However, since $y \neq a$ (by the fact that $a \notin V_C$ whereas y must belong to C as it is a distinguished vertex) we would then get $y = a - 1$ and $\pi(y) = \mathbf{e}$, contradicting the fact that yz is a boundary edge (see Definition 4.6). It now follows that $|\pi| \geq b = a + 1$ and since $\pi(a - 1) = \mathbf{e}$ we must have $\pi(a) = \pi(a + 1) = 1$ by the definition of Π (see Definition 4.1). In particular, the edge $(a, a + 2)$ belongs to $E(\mathcal{L}_\pi)$ (and also to B , recalling that a is not a distinguished vertex). To conclude this case argue as follows. Either $a + 2 \in V_C$, in which case $e_0 = 3$ due to the neighbors $\{a - 1, a + 1, a + 2\}$ of the vertex a and we have $e_0/(b - a) = 3 > 2 + 1/(M - 1)$ as needed. Otherwise, letting b' be the smallest index larger than $b = a + 1$ such that $b' \in V_C$ (which must exist since $z = M + 3 \in V_C$), we can consider the set $\{a, a + 2, \dots, b' - 1\}$ and denote by e_1 the number of edges incident to it in C . We must have $e_1/(b' - a - 1) < m_B$ and yet $e_1 \geq 2(b' - a - 1) + 1$ since it includes two edges going backward from each of $a + 2, \dots, b' - 1$, the two edges $(a - 1, a), (a, a + 1)$ and finally the edge $(b' - 1, b')$. This yields $e_1/(b' - a - 1) \geq 2 + 1/(b' - a - 1)$, and as $3 \leq a < b' \leq M + 2$ we get that $e_1/(b' - a - 1) \geq 2 + 1/(M - 1) = m_B$, contradiction.

Case (2.3): $a = 2$ and $b > 3$. If the second occurrence of \mathbf{e} is not in $\{\pi(2), \dots, \pi(b - 1)\}$ then the extra edge γb contributes to e_0 (we are assured that $1b \notin E(\mathcal{L}_\pi)$ since $b > 3$) and again we obtain $e_0 \geq 2(b - a) + 1$ to contradict (4.3). (Note that $b - a \leq M - 1$ in this case — hence the contradiction — due to the assumption that π features a second copy of \mathbf{e} in one of the indices $\{b, \dots, |\pi| - 1\}$, in fact implying that $b - a \leq M - 2$). Assume therefore that the second occurrence of \mathbf{e} does appear in the sequence. Observe that the assertion $b < z = M + 2$ still holds since otherwise $V_C = \{1, z\}$ leaving out the distinguished vertex y . Now, if $b + 1 \in V_C$ then this vertex has a neighbor in $2, \dots, b - 1$, and the corresponding edge cancels the contribution of the extra copy of \mathbf{e} to the right hand side of (4.4). This again gives $e_0 \geq 2(b - a) + 1$ and leads to the same contradiction (recall that $b \leq M + 1$ and so $b - a \leq M - 1$). It remains to deal with the case $b + 1 \notin V_C$. Again let b' be the minimal index such that $b < b' \in V_C$ and let e_1 count the number of edges incident to $2, \dots, b - 1, b + 1, \dots, b' - 1$, we must have $e_1/(b' - 3) < m_B$. However, e_1 includes two edges going backward from all but two of these $b' - 3$ vertices (those corresponding to copies of \mathbf{e} that contribute just one backward edge), the extra edge γb and the edges $(b, b - 1)$ and $(b' - 1, b')$. Overall, $e_1 \geq 2(b' - 3) + 1 \geq 2 + 1/(M - 1) = m_B$ as before.

Case (2.4): $a = 2$ and $b = 3$. Recall that the last vertex in V_C is $z = M + 2 \geq 5$, and examine the vertex 4. Regardless of whether π begins with $\mathbf{e}11$ or $\mathbf{e}\mathbf{e}11$, the vertex 2 has $\{1, 3, 4\}$ as its neighbors in $E(\mathcal{L}_\pi)$ and so if $4 \in V_C$ we immediately get that $e_0/(b-a) = 3 > m_B$, contradiction. Assume therefore that $4 \notin V_C$, let $b' > 4$ be the minimal index beyond 4 with $b' \in V_C$ (bearing in mind that $z \in V_C$ by definition) and let e_1 be the number of edges incident to $2, 4, \dots, b' - 1$ in V_C .

If the symbol \mathbf{e} does not appear in $\{\pi(4), \dots, \pi(b' - 2)\}$ (it is forbidden from occurring at $\pi(3)$ as per the definition of Π) then $e_1 \geq 2b' - 5$, since the vertices $4, \dots, b' - 1$ together send $2(b' - 4)$ edges backwards, accompanied by the edges $12, 23$ and $(b' - 1, b')$. Thus, in this case $e_1/(b' - 3) \geq 2 + 1/(b' - 3) \geq m_B$, as desired. Finally, if the symbol \mathbf{e} does appear in $\{\pi(4), \dots, \pi(b' - 2)\}$ then the extra edge $\gamma'b'$ (satisfying $\gamma' < b' - 1$) is present. Moreover, if we let $4 \leq j \leq b' - 2$ be the index such that $\pi(j) = \mathbf{e}$ then $\gamma' \geq j$ since the vertex j is a cut-vertex in \mathcal{L}_π . As such, the edge $\gamma'b'$ contributes to e_1 , thus canceling the negative contribution of the extra \mathbf{e} symbol. That is, again $e_1 \geq 2b' - 5$, eliminating this case.

All cases above led to a contradiction, completing the proof of the lemma for ladders beginning with \mathbf{e} . It remains to prove the balance property for ladders beginning with 1, which will follow from similar arguments yet here the vertex 0 will play a delicate role.

- (a) π contains no copies of \mathbf{e} : Here $v_\pi = 3M + 1$ and $e_\pi = 6M - 2$, and therefore $|V_B \setminus I_B| = 3M - 3$ (here we used the fact that $y \neq 0, 1$ — and thus $|I_B| = 4$ — since $z = 3M$ and forbidding an M -fan implies that $y \geq z - M = 2M$) and $|E_B| = 6M - 3$. Therefore

$$m_B = \frac{6M - 3}{3M - 3} = 2 + \frac{1}{M - 1}.$$

As before, if we consider the first interval $\{a, \dots, b-1\}$ missing from V_C we have $e_0/(b-a) < m_B$. The lack of copies of \mathbf{e} implies that $e_0 \geq 2(b-a) + 1$ (with the additive 1 accounting for the edge $(b-1, b)$ as usual) and so $b-a > M-1$. This in turn implies that $b-a \geq M$ and so the extra edge γb lands in the interval $[a, b)$ (since $\gamma \geq b-M$ due to the forbidden M -fans). As such, we cannot possibly have $b = z$ since both of its neighbors are missing from V_C (whereas y must belong to V_C as a distinguished vertex). Therefore, $b \leq z - 1 = 3M - 1$ and, since $a \geq 2$ (again recall that 1 is distinguished) we have $b-a \leq 3M - 3$. Furthermore, this entitles us to consider the vertex $b+1$. One more edge incident to this interval is needed to secure the contradiction: if $b+1 \in V_C$ it immediately provides such an edge, since its two neighbors are $\gamma' < b$ and $\gamma' \geq \gamma \geq a$. In this case we get

$$\frac{e_0}{b-a} \geq 2 + \frac{3}{b-a} \geq 2 + \frac{3}{3M-3} = m_B,$$

as required. On the other hand, if $b+1 \notin V_C$ we can again look at the next missing interval $\{b+1, \dots, b'-1\}$ where $b' \in V_C$ satisfies $b' \leq z = 3M$ (hence $b' - a \leq 3M - 2$). The edges $(b-1, b), \gamma b, (b'-1, b')$ now give the required contradiction since

$$\frac{e_1}{b'-a-1} \geq 2 + \frac{3}{b'-a-1} \geq 2 + \frac{3}{3M-3} = m_B.$$

- (b) π contains 1 copy of \mathbf{e} : We have $v_\pi = 2M + 2$ and $e_\pi = 4M - 1$, thus $|V_B \setminus I_B| = 2M - 2$ (again $y \neq 0, 1$ since $z = 2M + 1$ and $y \geq z - M = M + 1$) and $|E_B| = 4M - 2$, and as before we obtain

$$m_B = \frac{4M - 2}{2M - 2} = 2 + \frac{1}{M - 1}.$$

If $\mathbf{e} \notin \{\pi(a-1), \dots, \pi(b-2)\}$ then exactly as before we have $e_0 \geq 2(b-a) + 1$, implying that $b-a \geq M$ and so, if present, the extra edge γb has an endpoint in $[a, b)$. At the same time, $b \leq z - 1 = 2M$ and so $b-a \leq 2M - 2$ and we conclude that $e_0/(b-a) \geq 2 + 2/(2M - 2) = m_B$, contradiction. However, γb may not exist — this is the case iff $\pi(b-1) = \mathbf{e}$. Again then

necessarily $b < z$ and we can turn to the vertex $b + 1$. If $b + 1 \in V_C$, the enforced pattern $\mathbf{e}11$ implies that $(b - 1, b + 1)$ belongs to E_B and, counting this in place of γb we get $e_0/(b - a) \geq 2 + 2/(2M - 2)$ as before. If $b + 1 \notin V_C$ then joining $[a, b)$ to the next missing interval $[b + 1, b')$ gives $e_1/(b' - a - 1) \geq 2 + 2/(b' - a - 1)$ on account of the edges $(b - 1, b)$ and $(b' - 1, b')$, and the desired inequality $e_1/(b' - a - 1) \geq m_B$ follows from the fact that $b' - a \leq z - 2 = 2M - 1$.

If $\mathbf{e} \in \{\pi(a - 1), \dots, \pi(b - 2)\}$ we can only say that $e_0 \geq 2(b - a)$ since the contribution of the edge $(b - 1, b)$ is negated by the shortage of one edge due to the \mathbf{e} symbol. Nevertheless, if the extra edge γb is such that $\gamma \in [a, b)$ then $e_0 \geq 2(b - a) + 1$ and the argument used above implies that $b - a \geq M$ and $b \leq z - 1 = 2M$. We now consider $b + 1$: it has a neighbor γ' with $\gamma \leq \gamma' < b$, in particular satisfying $\gamma' \in [a, b)$. Thus, if $b + 1 \in V_C$ we are done by having $e_0/(b - a) \geq 2 + 2/(b - a) \geq m_B$. If $b + 1 \notin V_C$ then together with the next missing interval we have $e_1/(b' - a - 1) \geq 2 + 2/(b' - a - 1)$ (accounting for the edges γb , $(b - 1, b)$ and $(b' - 1, b')$ and subtracting one edge for the \mathbf{e} symbol), and the fact that $b' - a \leq z - 2 = 2M - 1$ now gives the sought inequality $e_1/(b' - a - 1) \geq 2 + 1/(M - 1) = m_B$.

Finally, if there is an \mathbf{e} symbol in $\{\pi(a - 1), \dots, \pi(b - 2)\}$ and yet the extra edge γb leaps over the interval $[a, b)$ then necessarily $\pi(a - 1) = \mathbf{e}$ (since having $\pi(j) = \mathbf{e}$ makes j a cut-vertex) and $\gamma = a - 1$. This can only occur when $b = a + 1$, and since neither of its neighbors can play the role of y (having $a \notin V_C$ and $\pi(\gamma) = \mathbf{e}$) we can infer that $b < z$. The symbol \mathbf{e} must be followed by 11 (in lieu of any additional \mathbf{e} symbols in π) and so the edge $(a, b + 1)$ belongs to \mathcal{L}_π (and to B , since $a \neq y$). Altogether, $e_0 = 3$ on account of the neighbors $\{a - 1, a + 1, a + 2\}$ of a and so $e_0/(b - a) = 3 > m_B$.

- (c) π contains 2 copies of \mathbf{e} : Here $v_\pi = M + 3$ and $e_\pi = 2M$, leading to $|V_B \setminus I_B| = M - 1$ (now $z = M + 2$ and $y \geq z - M = 2$) and $|E_B| = 2M - 1$. We get that

$$m_B = \frac{2M - 1}{M - 1} = 2 + \frac{1}{M - 1}.$$

Let $0 \leq s \leq 2$ denote the number of \mathbf{e} symbols in $\{\pi(a - 1), \dots, \pi(b - 2)\}$.

If $s = 0$ then $e_0 \geq 2(b - a) + 1$ and so $b - a > M - 1$. In this case either $\gamma \in [a, b)$ or the edge γb does not exist, either way precluding b from being equal to z , thus $b - a \leq (z - 1) - 2 = M - 1$, contradiction.

If $s = 1$ then a-priori we only have $e_0 \geq 2(b - a)$. If the extra edge γb falls in $[a, b - 2]$ then $e_0/(b - a) \geq 2 + 1/(b - a)$ and again $b - a > M - 1$, while at the same time $b < z$ and so $b - a \leq M - 1$. If the extra edge leaps over $[a, b)$ then necessarily $\pi(a - 1) = \mathbf{e}$ and $b = a + 1$. As before, here we must have 3 edges incident to a and so $e_0/(b - a) = 3 > m_B$. If there is no extra edge γb then $\pi(b - 1) = \mathbf{e}$ and $b \leq z - 1 = M + 1$, thus we can again appeal to $b + 1$ and either get an extra edge $(b - 1, b + 1)$ from the pattern $\mathbf{e}11$ whenever $b + 1 \in V_C$ (leading to $e_0/(b - a) \geq 2 + 1/(b - a) \geq 2 + 1/(M - 1) = m_B$), or join $[a, b)$ with the next missing interval $[b + 1, b')$ and for an extra edge $(b' - 1, b')$ (leading to $e_1/(b' - a - 1) \geq 2 + 1/(b' - a - 1) \geq m_B$).

It remains to treat $s = 2$, in which case clearly $\gamma \in [a, b)$ and so $b < z$. If $b + 1 \in V_C$ then it has a neighbor γ' with $\gamma \leq \gamma' < b$ and we arrive at $e_0 \geq 2(b - a) + 1$ (accounting for $(b - 1, b)$, γb , $(\gamma', b + 1)$ while subtracting 2 for the \mathbf{e} symbols). We have $b - a \leq (z - 1) - 2 = M - 1$ and so $e_0/(b - a) \geq 2 + 1/(M - 1) = m_B$, as needed. The final case has $b + 1 \notin V_C$, and combining $\{a, \dots, b - 1\}$ with the next missing interval $\{b + 1, \dots, b' - 1\}$ now gives $e_1 \geq 2(b' - a - 1) + 1$, where we counted the edges $(b - 1, b)$, γb , $(b' - 1, b')$ and subtracted two edges for \mathbf{e} symbols. We have $b' - a \leq z - 2 = M$ and $e_1/(b' - a - 1) \geq 2 + 1/(M - 1) = m_B$, concluding the proof. ■

Let $M \geq 3$ be some arbitrarily large constant. We aim to establish control over the number of labeled copies of every M -bounded triangular ladder along the triangle removal process as long as

$$p \geq p_M := n^{-\frac{1}{2} + \frac{1}{M}}. \quad (4.10)$$

Recalling (4.2) observe that $\mathcal{S}_H(p) \geq n$ for any backward extension H as long as we have $p \geq p_M$, with room to spare (more precisely, $\mathcal{S}_H(p_M) \geq n^{3/2-1/M}$ when the corresponding ladder \mathcal{L}_π has $|\pi| = M + 1$ and is even larger when $|\pi| = 2M$ or $|\pi| = 3M - 1$). Since ζ as defined in (2.5) satisfies $\zeta^{-1} = o(\sqrt{n})$, we have $\mathcal{S}_H > \zeta^{-\sqrt{\delta}}$ for any $t(p)$ with $p \geq p_M$ and any $0 < \delta < 1$. In particular, in this regime $t \leq t_H^-(\sqrt{\delta})$ for any backward extension H , thus we may invoke Part (i) of Theorem 3.2 to control the variables $\Psi_{H,\varphi}$ for all the extension graphs in \mathcal{B}_M (each of which has at most $L = 3M + 1$ vertices). Further note that at $t(p_M)$ we have $\zeta(t) = n^{-1/M+o(1)}$, thus the decay rate of ζ satisfies Eq. (3.2) and we can conclude:

Corollary 4.10. *For every $M \geq 3$ there is some $0 < \delta(M) < 1$ such that w.h.p. the following holds: For every backward extension graph $B = B_{\pi,yz}$ with $\pi \in \mathcal{L}_\pi$ and any corresponding injection φ we have $|\Psi_{B,\varphi} - \mathcal{S}_B| \leq \mathcal{S}_B \zeta^\delta$ as long as $t \leq \tau_\star$ and $p(t) \geq p_M$.*

Backward extensions are only defined for outer boundary edges. For all other edges (side boundary, initial or interior) we will need to keep track of the following.

Definition 4.11 (forward extension). *Let \mathcal{L}_π be an M -bounded triangular ladder and let $yz \in E(\mathcal{L}_\pi)$ with $y < z$ be an edge that is not an outer boundary edge. The **forward extension from yz to \mathcal{L}_π** , denoted by $F_{\pi,yz}$, is the graph obtained by deleting from \mathcal{L}_π all edges between vertices in the following set I of distinguished vertices:*

- *If yz is an initial edge then $I = \{0, 1, \dots, y\} \cup \{z\}$.*
- *If yz is a side boundary edge then $I = \{0, 1, \dots, y\} \cup \{z\}$ in case $\pi(y-2) \neq \mathbf{e}$ and otherwise $I = \{0, 1, \dots, y-2\} \cup \{y, z\}$.*
- *If yz is an interior edge then $I = \{0, 1, \dots, z\}$.*

Similarly to the case of backward extension graphs, estimates to within a multiplicative factor of $1 + O(\zeta^\delta)$ over forward extension graphs can be readily derived from Theorem 3.2.

Corollary 4.12. *Let $M \geq 3$ and let $0 < \delta < 1$ be the constant from Corollary 4.10. Then w.h.p. for every forward extension graph $F = F_{\pi,yz}$ with $\pi \in \mathcal{B}_M$ and any corresponding injection φ we have $\Psi_{F,\varphi} = (1 + O(\zeta^\delta))\mathcal{S}_F$ as long as $t \leq \tau_\star$ and $p(t) \geq p_M$. Furthermore, when yz is not a side boundary edge the stronger bound $\Psi_{F,\varphi} = (1 + O(\zeta))\mathcal{S}_F$ holds throughout this regime.*

Proof. Suppose first that yz is an *interior* edge of \mathcal{L}_π . We wish to count all possible mappings $\psi \in \Psi_{F,\varphi}$ by iteratively exposing the images of $z+1, \dots, |V(\mathcal{L}_\pi)|$ in V_G . Following Definition 4.1 of triangular ladders we see that every vertex j is attached to either 1 or 2 vertices among its predecessors $\{0, 1, \dots, j-1\}$ (depending on whether or not we then encounter the symbol \mathbf{e} in π). Therefore, given $\psi(1), \dots, \psi(j-1)$ for some $j > z$, the number of possibilities for $\psi(j)$ is dictated by either Y_u for some $u \in \{\psi(1), \dots, \psi(j-1)\}$ or by $Y_{u,v}$ for some u, v in this set. Using (3.1) to control these variables up to a relative error of $1 + O(\zeta)$ and iterating this expansion we get

$$\Psi_{F,\varphi} = (1 + O(|V(\mathcal{L}_\pi)|\zeta))\mathcal{S}_F = (1 + O(\zeta))\mathcal{S}_F. \quad (4.11)$$

This argument remains valid of course when yz is an *initial* edge with $y = z - 1$, since again we begin with the distinguished set $I = \{0, 1, \dots, z\}$ whose images are predetermined by φ . On the other hand, if yz is an initial edge with $y < z - 1$ then necessarily $y = z - 2$ (to see this observe that $\pi(y) = \mathbf{e}$ by Definition 4.7 and thus, since π does not contain the substring $\mathbf{e}10$, the potential forward neighbors of y are confined to the set $\{y+1, y+2\}$) and similarly the neighbors of $z-1$ in \mathcal{L}_π are precisely y, z . Thus, the number of possibilities for $\psi(z-1)$ is given by $Y_{\psi(y)\psi(z)}$ and from this point on we can iteratively expose the images of $z+1, \dots, |V(\mathcal{L}_\pi)|$ as above to establish (4.11).

It remains to deal with the case where yz is a *side boundary* edge, which as we next describe demonstrates the necessity for a careful analysis of f -fans. We focus on the case $\pi(y-2) \neq \mathbf{e}$, noting that the case $\pi(y-2) = \mathbf{e}$ will follow by a nearly identical argument. Recall that the distinguished vertex set is now $I = \{0, 1, \dots, y\} \cup \{z\}$. By Definition 4.6, here we have $y < z - 1$, there is at

most 1 occurrence of \mathbf{e} in π and the ladder $\mathcal{L}_{\pi_{z-1}}$ contains an $(M-1)$ -fan at y . Therefore, if we attempt to construct a copy $\psi \in \Psi_{F,\varphi}$ iteratively as before, upon reaching the vertex $z-1$ we will find that its neighbors in \mathcal{L}_π are precisely $\{y, z-2, z\}$ and so, once the images of these vertices are preset, $\psi(z-1)$ is given by some variable N_{uvw} which we may very well have no control over (our approximation within a relative error of $1 + O(\zeta)$ for these variables ceases to be valid once $p < n^{-1/3+o(1)}$ as the triple co-neighborhoods become merely poly-logarithmic in size).

To remedy this issue we will rely on Theorem 3.2. Let K be the subextension of F that has $V_K = \{0, 1, \dots, z\}$. Since $I_K = I_F = \{0, 1, \dots, y\} \cup \{z\}$, each vertex $j \in \{y+2, \dots, z-1\}$ has $\{y, j-1, j+1\}$ as its neighbors in K . We can temporarily think of these edges as being directed by letting each such vertex j send one edge to $j+1$ and another to y , while the vertex $y+1$ does the same and sends an additional edge to some $\ell < y$. With this view it becomes clear that $e_K = 2(v_K - \iota_K) + 1$ and furthermore that K is a (strictly) balanced extension graph. To conclude the proof, note that $v_K - \iota_K = z - y - 1$ whereas $z = y + M$ since z is the final vertex of an $(M-1)$ -fan at y (see Eq. (4.1) bearing in mind that we are in the case $\pi(y-2) \neq \mathbf{e}$). Plugging this into the expression for $v_K - \iota_K$ we obtain that $\mathcal{S}_K = n^{M-1}p^{2M-1}$. Hence, as long as $p \geq p_M$ we have $\mathcal{S}_K \geq n^{(3M-2)/(2M)}$. This translates into a lower bound of at least $n^{7/6}$ by the assumption $M \geq 3$ and, since $\zeta^{-1} = o(\sqrt{n})$, this implies in turn that $t < t_K^-(\sqrt{\delta})$ for any $0 < \delta < 1$. We can therefore appeal to Part (i) of Theorem 3.2 and get that

$$|\Psi_{K,\varphi} - \mathcal{S}_K| \leq \mathcal{S}_K \zeta^\delta$$

where $\delta = \delta(M)$ is precisely the same constant as in Corollary 4.10. From this point we can complete K into a copy of F by iteratively enumerating over the options for $z+1, \dots, |V(\mathcal{L}_\pi)|$ as before, each time at a cost of a multiplicative factor of $1 + O(\zeta)$. Altogether we deduce that

$$|\Psi_{F,\varphi} - \mathcal{S}_F| \leq (\zeta^\delta + O(\zeta))\mathcal{S}_F = O(\zeta^\delta)\mathcal{S}_F,$$

as required. Essentially the same argument extends this estimate also to the case of a side boundary edge yz with $\pi(y-2) = \mathbf{e}$. Now $I_F = \{0, 1, \dots, y-2\} \cup \{y, z\}$ yet $z = y+M-1$ (see (4.1)) and a closer look reveals that the vertices $y-1, y+1, \dots, z$ forming the $(M-1)$ -fan play an identical role to that assumed by vertices $y+1, \dots, z$ in the previous case (e.g., each can be viewed as sending one edge to its successor in this list and another to y , the first vertex $y-1$ sends one additional to $\ell = y-2$ and so on). Therefore, the subextension K is again strictly balanced and has $e_K = 2(v_K - \iota_K) + 1$. Having $v_K - \iota_K = z - y$ and $z = y + M - 1$ we again get $\mathcal{S}_K = n^{M-1}p^{2M-1} \geq n^{7/6}$ for all $p \geq p_M$, and the rest of the argument holds unmodified. \blacksquare

5. TRACKING TRIANGULAR LADDERS

5.1. Expected change. For a labeled undirected graph H and $u, v \in V_G$ write $\Psi_{H,uv} = \Psi_{H,uv}(t)$ to denote the random variable $\Psi_{H,\varphi}$ where φ maps $0 \mapsto u$ and $1 \mapsto v$. If $H = \mathcal{L}_\pi$ we will use the abbreviated form $\Psi_{\pi,uv}$, and as usual we use v_π and e_π to denote $|V(\mathcal{L}_\pi)|$ and $|E(\mathcal{L}_\pi)|$, respectively. Our aim in this subsection is to estimate the expected change in variables $\Psi_{\pi,uv}$ within one step of the process for any M -bounded triangular ladder.

Theorem 5.1. *Let $M \geq 3$ and let $0 < \delta < 1$ be the constant from Corollaries 4.10 and 4.12. For any $\pi \in \mathcal{B}_M$ let π^- denote its $(|\pi| - 1)$ -prefix and define*

$$X_{\pi,uv} = \Psi_{\pi,uv} - \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv}$$

where $\Psi_{\emptyset,uv} := 1$. Let $\zeta = n^{-1/2}p^{-1} \log n$ and for some absolute constant $\kappa > 0$ let

$$\tau_\zeta^* = \tau_\zeta^*(\kappa) = \min \left\{ t : \left| Q / \left(\frac{1}{6} n^3 p^3 \right) - 1 \right| \geq \kappa \zeta^2 \right\}.$$

Let i be some time such that $t(i) \leq \tau_Q^*$ and $p(i) \geq p_M$. Suppose that $\Psi_{\pi,uv} = (1 + o(1))\mathcal{S}_\pi$ holds at time i for all $\pi \in \mathcal{B}_M$ and all $u, v \in V_G$. Then

$$\mathbb{E}[\Delta X_{\pi,uv} \mid \mathcal{F}_i] = -(6 - o(1))\frac{e_\pi}{n^2 p} X_{\pi,uv} + \Xi_0(\pi) + \Xi_1(\pi) + \frac{O(\zeta^{1+\delta})}{n^2 p} \mathcal{S}_\pi \quad (5.1)$$

for every $\pi \in \mathcal{B}_M$ and every $u, v \in V_G$, where the correction term $\Xi_1(\pi)$ is given by

$$\Xi_1(\pi) = -\frac{6 + o(1)}{n^3 p^3} X_{\pi \circ 1, uv} \mathbb{1}\{(|\pi|, |\pi| + 1) \text{ is an initial/interior edge}\},$$

whereas if there exists $y < |\pi|$ such that $(y, |\pi| + 1) \in E(\mathcal{L}_\pi)$ then

$$\begin{aligned} \Xi_0(\pi) = & -\frac{6 + o(1)}{n^2 p} X_{\pi, uv} \mathbb{1}\{(y, |\pi| + 1) \text{ is an initial edge}\} \\ & -\frac{6 + o(1)}{n^3 p^3} X_{\pi \circ 0, uv} \mathbb{1}\{(y, |\pi| + 1) \text{ is an interior edge}\} \\ & -\frac{6 + o(1)}{n^3 p^3} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi_{y-1} \circ \mathbf{e}}} X_{\pi_{y-1} \circ \mathbf{e}1, uv} \mathbb{1}\{(y, |\pi| + 1) \text{ is a side boundary edge, } \pi(y-2) \neq \mathbf{e}\} \\ & -\frac{6 + o(1)}{n^3 p^3} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi_{y-2} \circ \mathbf{e}}} X_{\pi_{y-2} \circ \mathbf{e}1, uv} \mathbb{1}\{(y, |\pi| + 1) \text{ is a side boundary edge, } \pi(y-2) = \mathbf{e}\} \end{aligned}$$

and otherwise $\Xi_0(\pi) = 0$.

Our proof will crucially rely on the properties of triangular ladders in order to express the number of copies of \mathcal{L}_π that rest on a given edge in terms of the total number of other triangular ladders. To demonstrate this we use the following simple lemma of [2] whose proof is provided for completeness.

Lemma 5.2. *Suppose $(x_i)_{i \in I}$ and $(y_i)_{i \in I}$ are real numbers so that $|x_i - x| \leq \delta_1$ and $|y_i - y| < \delta_2$ for some $x, y \in \mathbb{R}$ and every $i \in I$. Then*

$$\left| \sum_{i \in I} x_i y_i - \frac{1}{|I|} \left(\sum_{i \in I} x_i \right) \left(\sum_{i \in I} y_i \right) \right| \leq 2|I|\delta_1\delta_2.$$

Proof. By the triangle inequality, $|\sum_{i \in I} (x_i - x)(y_i - y)| \leq |I|\delta_1\delta_2$, and at the same time it is easy to see that $\sum_{i \in I} (x_i - x)(y_i - y)$ is precisely equal to

$$\sum_{i \in I} x_i y_i - \frac{1}{|I|} \left(\sum_i x_i \right) \left(\sum_i y_i \right) + |I| \left(\frac{\sum_i x_i}{|I|} - x \right) \left(\frac{\sum_i y_i}{|I|} - y \right).$$

The hypothesis on the x_i 's and y_i 's implies that the last term above is at most $|I|\delta_1\delta_2$ in absolute value, from which the desired result follows. \blacksquare

Prior to commencing with the proof of Theorem 5.1 we wish to illustrate the role that the backward/forward extensions play in the analysis of the variables $\Psi_{\pi,uv}$.

Example 5.3 (expected one-step change in $\Psi_{\mathbf{ee}}$). In general, in order to estimate the one-step change in $\Psi_{\pi,uv}$ we examine the edges of the given ladder and for each one consider the probability of hitting this edge (as part of some labeled copy of the given ladder rooted at $u, v \in V_G$) by the current triangle as well as the number of copies of the ladder that it participates in. For $\mathcal{L}_{\mathbf{ee}}$ we write separate summations for the two edges 12 and 23 to get

$$\begin{aligned} \mathbb{E}[\Delta \Psi_{\mathbf{ee}, uv} \mid \mathcal{F}_i] &= - \sum_{a \in N_u} \frac{Y_{u,a}(Y_a - 1)}{Q} - \sum_{\substack{a \in N_u \\ b \in N_a \setminus \{u\}}} \frac{Y_{a,b} - \mathbb{1}_{\{ua \in E\}}}{Q} \\ &= - \sum_{a \in N_u} \frac{Y_{u,a}(Y_a - 1)}{Q} - \frac{\Psi_{\mathbf{ee}1, uv}}{Q} + O\left(\frac{1}{n^3 p^3}\right) \mathcal{S}_{\mathbf{ee}}, \end{aligned}$$

where the indicators $\mathbb{1}_{\{ua \in E\}}$ treat the double counting of a triangle that contains both these edges, and the last equality rewrote their sum as $O(\mathcal{S}_{ee}/(n^3 p^3))$ via the hypothesis on the degrees and Q . We now apply Lemma 5.2 with $I = N_u$, the sequence $x_a = Y_{u,a}$ with $x = np^2$ and $\delta_1 = np^2 O(\zeta)$ and the sequence $y_a = Y_a$ with $y = np$ and $\delta_2 = np O(\zeta)$ (the variables $\Psi_{\pi,uv}$ count labeled graphs, thus for instance $\sum_{a \in N_u} Y_{u,a} = \Psi_{e1,uv}$). Since $|I| \delta_1 \delta_2 / Q = O(\zeta^2 p)$ while $\mathcal{S}_{ee} = n^2 p^2$ this gives

$$\mathbb{E}[\Delta \Psi_{ee,uv} \mid \mathcal{F}_i] = -\frac{\Psi_{e1,uv} \Psi_{ee,uv}}{Y_u Q} - \frac{\Psi_{ee1,uv}}{Q} + \frac{O(\zeta^2)}{n^2 p} \mathcal{S}_{ee}.$$

(The $O(\frac{1}{n^3 p^3})$ -term was absorbed by the $O(\frac{\zeta^2}{n^2 p})$ -term as the latter has order $\frac{\log^2 n}{n^3 p^3}$ by Def. (2.5).)

We stress that many of our applications of Lemma 5.2 have this common form. We are computing the expected change in some variable $\Psi_{\pi,uv}$ due to the loss of a particular edge in the triangular ladder. We compute this as a sum over such edges $yz \in E(\mathcal{L}_\pi)$. For each one we break up the ladder \mathcal{L}_π into a smaller subgraph K that contains u, v and the edge yz , then proceed to (i) count the number of labeled copies of K in G rooted at u, v (thereby yielding an estimate on the probability of removing such an edge in the current round), denoted by some $\Psi_{K,uv}$, and (ii) count the number of extensions from K to \mathcal{L}_π in G (giving the total number of copies of \mathcal{L}_π eliminated by removing the edge under consideration), denoted by some $\Psi_{\mathcal{L}_\pi/K}$. Observe that, by definition, $\sum_{\psi \in \Psi_{K,uv}} \Psi_{\mathcal{L}_\pi/K, \psi} = \Psi_{\pi,uv}$. Moreover, given the injection $\psi : K \rightarrow V_G$ the probability to remove $\psi(y)\psi(z)$ is equal to $Y_{\psi(y)\psi(z)}/Q$, thus we arrive at the expression

$$\sum_{\psi \in \Psi_{K,uv}} \frac{Y_{\psi(y)\psi(z)}}{Q} \Psi_{\mathcal{L}_\pi/K, \psi}.$$

Applying Lemma 5.2 on this summation we now get a main term which is a product of $\Psi_{\pi, \psi}$ (the result of $\sum_{\psi} \Psi_{\mathcal{L}_\pi/K, \psi}$) and $\Psi_{K',uv}$, where K' is obtained from K by attaching to it a new vertex incident to the edge yz (the result of $\sum_{\psi} Y_{\psi(y)\psi(z)}$). We now see the reason behind our definition of the family of triangular ladders, since if K is a triangular ladder then K' will commonly also belong to our family of tracked variables. As an error term in this analysis we will sustain the approximation error in $Y_{\psi(y)\psi(z)}/Q$, i.e., an $O(\zeta/(n^2 p))$ term, multiplied by $\Psi_{K,uv}$ and then by the approximation error for $\Psi_{\mathcal{L}_\pi/K}$. Usually we will choose K so that \mathcal{L}_π/K will be the forward extension $F_{\pi, yz}$ introduced in the previous section, whose corresponding variable $\Psi_{F, \varphi}$ was estimated there to within an error of $O(\zeta^\delta) \mathcal{S}_F$ uniformly over the injection φ . Overall, we get the desired $O(\zeta^{1+\delta}/(n^2 p)) \Psi_{\pi,uv}$ error-term.

We are now ready to prove the main result of this subsection.

Proof of Theorem 5.1. By hypothesis we control the variables $Y_u = \Psi_{e,uv}$, $Y_{u,v} = \Psi_{1,uv}$ and Q to within a multiplicative factor of $1 + O(\zeta)$. Via Corollaries 4.10 and 4.12 this guarantees estimates to within a factor of $1 + O(\zeta^\delta)$ in backward and forward extension variables $B_{\pi, yz}$ and $F_{\pi, yz}$. (Forward extensions of initial/interior edges feature a better approximation of $1 + O(\zeta)$.)

As mentioned above, we will analyze the different possible edges $yz \in E(\mathcal{L}_\pi)$ whose images in $E(G)$ (under any mapping $\psi \in \Psi_{\pi,uv}$) may be removed in the upcoming step, and the possible impact this event would have on $X_{\pi,uv}$.

Given \mathcal{F}_i , for each $\psi \in \Psi_{\pi,uv}$ we define the events

$$\mathcal{E}_\psi = \bigcup \{ \mathcal{E}_\psi^{yz} : yz \in E(\mathcal{L}_\pi) \} \quad \text{where} \quad \mathcal{E}_\psi^{yz} = \{ \psi(y)\psi(z) \text{ is removed at time } i \}, \quad (5.2)$$

and for each edge $yz \in E(\mathcal{L}_\pi)$ further set

$$\Delta_{yz} \Psi_{\pi,uv} = - \sum_{\psi \in \Psi_{\pi,uv}} \mathbb{1}_{\mathcal{E}_\psi^{yz}}.$$

According to these definitions, $\Delta\Psi_{\pi,uv} = -\sum_{\psi \in \Psi_{\pi,uv}} \mathbb{1}_{\mathcal{E}_\psi}$. At the same time, we crucially have

$$\mathbb{P}(\mathcal{E}_\psi \mid \mathcal{F}_i) = \sum_{yz \in E(\mathcal{L}_\pi)} \mathbb{P}(\mathcal{E}_\psi^{yz} \mid \mathcal{F}_i) + O(1/Q) \quad (5.3)$$

since $\mathbb{P}(\mathcal{E}_\psi^{yz}, \mathcal{E}_\psi^{y'z'} \mid \mathcal{F}_i) \leq 1/Q$ for any distinct edges $yz, y'z' \in E(\mathcal{L}_\pi)$ (the event that both are removed identifies a single triangle in G), and there are $O(e_\pi^2) = O(1)$ such pairs. It follows that

$$\mathbb{E}[\Delta\Psi_{\pi,uv} \mid \mathcal{F}_i] = \sum_{yz \in E(\mathcal{L}_\pi)} \mathbb{E}[\Delta_{yz}\Psi_{\pi,uv}] + O(\Psi_{\pi,uv}/Q) = \sum_{yz \in E(\mathcal{L}_\pi)} \mathbb{E}[\Delta_{yz}\Psi_{\pi,uv}] + \frac{o(\zeta^2)}{n^2p}\Psi_{\pi,uv}, \quad (5.4)$$

using the fact that $Q \asymp n^3p^3$ whereas $\zeta^{-2}n^2p \asymp n^3p^3 \log^2 n$. Our ultimate goal in this theorem is to estimate the one-step change in $X_{\pi,uv} = \Psi_{\pi,uv} - (\mathcal{S}_\pi/\mathcal{S}_{\pi^-})\Psi_{\pi^-,uv}$. As usual we write the change from time i to time $i+1$ as

$$\Delta X_{\pi,uv} = \Delta\Psi_{\pi,uv} - \frac{\mathcal{S}_\pi(i+1)}{\mathcal{S}_{\pi^-}(i+1)}\Delta\Psi_{\pi^-,uv} - \Psi_{\pi^-,uv}\Delta\left(\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\right).$$

Since we have $\mathcal{S}_\pi/\mathcal{S}_{\pi^-} = np^\theta$ where $\theta = 1$ if the last symbol in π is \mathbf{e} and otherwise $\theta = 2$, the change in the last term is deterministically given by

$$\Psi_{\pi^-,uv}\Delta\left(\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\right) = -\theta np^{\theta-1} \left(\frac{6}{n^2}\right) \left(1 - O\left(\frac{1}{n^2p}\right)\right) \Psi_{\pi^-,uv} = -\frac{6\theta - o(\zeta^3)}{n^2p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv},$$

relying in the last equality on the fact that $\zeta^{-3} = n^{\frac{3}{2}-o(1)}p^3 = o(n^2p)$. Adding the similar deterministic estimate $(\mathcal{S}_\pi/\mathcal{S}_{\pi^-})(i) = (1 - o(\zeta^3))(\mathcal{S}_\pi/\mathcal{S}_{\pi^-})(i-1)$ we can rewrite the equation for $\Delta X_{\pi,uv}$ as

$$\Delta X_{\pi,uv} = \Delta\Psi_{\pi,uv} - (1 + o(\zeta^3))\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\Delta\Psi_{\pi^-,uv} + \frac{6\theta - o(\zeta^3)}{n^2p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv}. \quad (5.5)$$

Finally, bearing in mind that θ counts the number of edges that are incident to the vertex $|\pi|+1$, we can break up the expected change in $\Delta\Psi_{\pi,uv}$ and $\Delta\Psi_{\pi^-,uv}$ as per equation with (5.4) and get

$$\begin{aligned} \mathbb{E}[\Delta X_{\pi,uv} \mid \mathcal{F}_i] &= \sum_{\substack{yz \in E(\mathcal{L}_\pi) \\ y < z \leq |\pi|}} \left(\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] - (1 - o(\zeta^3))\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\mathbb{E}[\Delta_{yz}\Psi_{\pi^-,uv} \mid \mathcal{F}_i] \right) \\ &\quad + \sum_{\substack{yz \in E(\mathcal{L}_\pi) \\ y < z = |\pi|+1}} \left(\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6 - o(\zeta^3)}{n^2p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv} \right) + \frac{o(\zeta^2)}{n^2p}\Psi_{\pi,uv}. \end{aligned}$$

(Notice that the first sum goes over all edges of \mathcal{L}_{π^-} , leaving θ edges of $E(\mathcal{L}_\pi) \setminus E(\mathcal{L}_{\pi^-})$ for the second sum. This allows us to divide the last term of (5.5) to each of these latter summands.)

By assumption, $(\mathcal{S}_\pi/\mathcal{S}_{\pi^-})\Psi_{\pi^-,uv} = (1 + o(1))\Psi_{\pi,uv}$, and similarly,

$$\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\mathbb{E}[\Delta_{yz}\Psi_{\pi^-,uv} \mid \mathcal{F}_i] = -\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \sum_{\varphi \in \Psi_{\pi^-,uv}} \frac{Y_{\varphi(y)\varphi(z)}}{Q} \asymp \frac{1}{n^2p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv} \asymp \frac{1}{n^2p}\Psi_{\pi,uv},$$

so the $o(\zeta^3)$ terms in the last identity for $\mathbb{E}[\Delta X_{\pi,uv} \mid \mathcal{F}_i]$ can be absorbed in the $o(\zeta^2)$ term, yielding:

$$\begin{aligned} \mathbb{E}[\Delta X_{\pi,uv} \mid \mathcal{F}_i] &= \sum_{\substack{yz \in E(\mathcal{L}_\pi) \\ y < z \leq |\pi|}} \left(\mathbb{E}[\Delta_{yz} \Psi_{\pi,uv} \mid \mathcal{F}_i] - \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \mathbb{E}[\Delta_{yz} \Psi_{\pi^-,uv} \mid \mathcal{F}_i] \right) \\ &+ \sum_{\substack{yz \in E(\mathcal{L}_\pi) \\ y < z = |\pi|+1}} \left(\mathbb{E}[\Delta_{yz} \Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6}{n^2 p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv} \right) + \frac{o(\zeta^2)}{n^2 p} \Psi_{\pi,uv}. \end{aligned} \quad (5.6)$$

This identity becomes particularly useful in light of the fact that if $yz \in E(\mathcal{L}_{\pi^-})$ then yz has the same edge classification in both \mathcal{L}_π and \mathcal{L}_{π^-} . (Indeed, the criterion for classifying yz as a boundary edge is a function of π_{z-1} and the values of y, z , whereas initial edges are identified by $\pi(y)$.) We now proceed to estimate the appropriate quantities in (5.6) for each $yz \in E(\mathcal{L}_\pi)$ depending on the type of the edge under consideration.

- *Initial edges:* If yz is an initial edge with $y < z$ then $\pi(y) = \mathbf{e}$ and $z \in \{y+1, y+2\}$. Consider the case $z = y+1$ and examine π_y , the length- y prefix of π . We will count the number of (labeled, rooted at u, v) copies of \mathcal{L}_{π_y} to estimate the number of edges that can play the role of yz in G . Given some mapping $\varphi \in \Psi_{\pi_y,uv}$, in the event that we remove the edge $\varphi(y)\varphi(z)$ the number of affected copies $\psi \in \Psi_{\pi,uv}$ (i.e., those agreeing with φ on y, z) is simply the number of forward extensions $F = F_{\pi,yz}$ from φ to the full copy of \mathcal{L}_π . Thus,

$$\mathbb{E}[\Delta_{yz} \Psi_{\pi,uv} \mid \mathcal{F}_i] = - \sum_{\varphi \in \Psi_{\pi_y,uv}} \frac{Y_{\varphi(y)\varphi(z)} \Psi_{F,\varphi}}{Q}.$$

For a fixed $\varphi \in \Psi_{\pi_y,uv}$, the probability of removing the image of yz is $Y_{\varphi(y)\varphi(z)}/Q$, whereas the sum over $\varphi \in \Psi_{\pi_y,uv}$ of $Y_{\varphi(y)\varphi(z)}$ is $\Psi_{\pi_y \circ 1,uv}$. We claim that $\pi_y \circ 1 \in \mathcal{B}_M$. For $y = 1$ indeed $\mathbf{e}1 \in \mathcal{B}_M$. For $y > 1$, clearly π has either $\mathfrak{C}_\pi \in \{1, 2\}$ copies of \mathbf{e} . Going back to the definition of \mathcal{B}_M we see that, as $\pi \in \mathcal{B}_M$, necessarily $|\pi| \leq \ell$ where $\ell = 2M$ if $\mathfrak{C}_\pi = 1$ and $\ell = M+1$ if $\mathfrak{C}_\pi = 2$. Moreover, $y \neq \ell$ so as to avoid $\pi(\ell) = \mathbf{e}$, thus $y < \ell$ and it follows that $\pi_y \circ 1$ has length at most ℓ , at most \mathfrak{C}_π copies of \mathbf{e} and a last symbol of 1, thus $\pi_y \circ 1 \in \mathcal{B}_M$ as claimed.

As stated before, $\sum_\varphi \Psi_{F,\varphi}$ gives $\Psi_{\pi,uv}$ itself. Altogether, an application of Lemma 5.2 gives

$$\mathbb{E}[\Delta_{yz} \Psi_{\pi,uv} \mid \mathcal{F}_i] = - \frac{\Psi_{\pi_y \circ 1,uv}}{\Psi_{\pi_y,uv} Q} \Psi_{\pi,uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi,uv}, \quad (5.7)$$

where in the notation of that lemma we substituted $|I| = \Psi_{\pi_y,uv}$, $\delta_1 = O(\zeta)/(n^2 p)$ and $\delta_2 = O(\zeta)\mathcal{S}_F$ uniformly over the injection φ thanks to Corollary 4.12 (note that yz is not a side boundary edge), thus $|I|\delta_1\delta_2 = O(\zeta^2/(n^2 p)) \sum_{\varphi \in \Psi_{\pi_y,uv}} \Psi_{F,\varphi} = O(\zeta^2/(n^2 p)) \Psi_{\pi,uv}$.

Now consider the case of an initial edge yz where $z = y+2$. If we wanted to use argument analogous to the one above we would examine the ladder $\mathcal{L}_{\pi_{z-1}}$. However, in this case the sum over $\varphi \in \Psi_{\pi_{z-1},uv}$ of $Y_{\varphi(y)\varphi(z)}$ would give $\Psi_{\pi_{z-1} \circ 0,uv}$, and $\pi_{z-1} \circ 0 \notin \mathcal{B}_M$ since $\pi(y) = \mathbf{e}$ and the substrings $\mathbf{e}10$ and $\mathbf{e}0$ are forbidden. A closer look at $\mathcal{L}_{\pi_{z-1}}$ reveals that the vertices $y+1, z = y+2$ play identical roles in this graph — they are connected to one another and each shares an edge with y . By switching their labels we can examine the graph \mathcal{L}_{π_y} where the role of z will be played by the isomorphic $y+1$. This gives

$$\mathbb{E}[\Delta_{yz} \Psi_{\pi,uv} \mid \mathcal{F}_i] = - \sum_{\varphi \in \Psi_{\pi_y,uv}} \frac{Y_{\varphi(y)\varphi(y+1)} \Psi_{F_{\pi,yz},\varphi'}}{Q},$$

where φ' above stands for the composition of φ and the map $z \mapsto y+1$. One should note that the above equality used the fact that $F_{\pi,yz}$ has the distinguished vertex set $\{1, \dots, y\} \cup \{z\}$ thus vertex $y+1$ that was not mapped in φ' is indeed added to form a full copy of \mathcal{L}_π . Henceforth

the analysis of this summation can be carried out exactly as in the above case of $z = y + 1$, and we conclude that Eq. (5.7) holds also for the case of $z = y + 2$.

Having established (5.7) for all initial edges, we now wish to estimate the contribution to $\mathbb{E}[\Delta X \mid \mathcal{F}_i]$ for such edges. This will depend on y, z as per the following three cases:

(1) $y < z \leq |\pi|$: From (5.7), applied both to $\Psi_{\pi,uv}$ and to $\Psi_{\pi^-,uv}$, we have

$$\begin{aligned} \mathbb{E}[\Delta_{yz} \Psi_{\pi,uv} \mid \mathcal{F}_i] - \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \mathbb{E}[\Delta_{yz} \Psi_{\pi^-,uv} \mid \mathcal{F}_i] &= -\frac{\Psi_{\pi_y \circ 1, uv}}{\Psi_{\pi_y, uv} Q} X_{\pi, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv} \\ &= -\frac{6 - o(1)}{n^2 p} X_{\pi, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv}, \end{aligned} \quad (5.8)$$

where the last equality applied our hypothesis that $\Psi_{\sigma, uv} = (1 + o(1))\mathcal{S}_\sigma$ for any $\sigma \in \mathcal{B}_M$ (namely for π_y and $\pi_y \circ 1$).

(2) $y = |\pi|$ and $z = |\pi| + 1$: Here $\pi_y = \pi$, thus the main term in the right-hand-side of (5.7) simplifies into $-\Psi_{\pi \circ 1, uv}/Q$. Writing $\Psi_{\pi \circ 1, uv} = X_{\pi \circ 1, uv} + (\mathcal{S}_{\pi \circ 1}/\mathcal{S}_\pi)\Psi_{\pi, uv}$ we get

$$\begin{aligned} \mathbb{E}[\Delta_{yz} \Psi_{\pi, uv} \mid \mathcal{F}_i] + \frac{6}{n^2 p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-, uv} &= -\frac{X_{\pi \circ 1, uv} + np^2 \Psi_{\pi, uv}}{Q} + \frac{6}{n^2 p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv} \\ &= -\frac{6}{n^2 p} X_{\pi, uv} - \frac{6 - o(1)}{n^3 p^3} X_{\pi \circ 1, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv}. \end{aligned} \quad (5.9)$$

(The final error term above absorbed a $1 + O(\zeta^2)$ factor due to the approximation of Q .)

(3) $y = |\pi| - 1$ and $z = |\pi| + 1$: Here $\pi_y = \pi^-$ and $\pi_y \circ 1 = \pi$ (we must have $\pi(y + 1) = 1$ to allow such an edge yz). The identity $\Psi_{\pi, uv} = X_{\pi, uv} + (\mathcal{S}_\pi/\mathcal{S}_{\pi^-})\Psi_{\pi^-, uv}$ now yields

$$\begin{aligned} \mathbb{E}[\Delta_{yz} \Psi_{\pi, uv} \mid \mathcal{F}_i] + \frac{6}{n^2 p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-, uv} &= -\frac{X_{\pi, uv} + np^2 \Psi_{\pi^-, uv}}{\Psi_{\pi^-, uv} Q} \Psi_{\pi, uv} + \frac{6}{n^2 p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv} \\ &= -\frac{12 + o(1)}{n^2 p} X_{\pi, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv}, \end{aligned} \quad (5.10)$$

where we applied the facts $Q = (\frac{1}{6} + O(\zeta^2))n^3 p^3$ and $\Psi_{\pi, uv}/\Psi_{\pi^-, uv} = (1 + o(1))np^2$.

- *Interior edges*: First consider an interior edge yz where $y < z \leq |\pi|$. Here we let $\mathcal{L}_{\pi_{z-1}}$ be the intermediate graph containing yz . As before, for each $\varphi \in \Psi_{\pi_{z-1}, uv}$ the probability of choosing a triangle that contains $\varphi(y)\varphi(z)$ is $Y_{\varphi(y)\varphi(z)}/Q$, and the number of extensions in $\Psi_{\pi, uv}$ that are lost with this choice is the number of forward extensions from the fixed φ to $F = F_{\pi, yz}$. Therefore, as in the argument for initial edges we get

$$\mathbb{E}[\Delta_{yz} \Psi_{\pi, uv} \mid \mathcal{F}_i] = - \sum_{\varphi \in \Psi_{\pi_{z-1}, uv}} \frac{Y_{\psi(y)\psi(z)} \Psi_{F, \varphi}}{Q} = -\frac{\Psi_{\pi_{z-1} \circ b_{yz}, uv}}{\Psi_{\pi_{z-1}, uv} Q} \Psi_{\pi, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv},$$

where $b_{yz} = 0$ if $y < z - 1$ whereas $b_{yz} = 1$ if $y = z - 1$. Observe that $\pi_{z-1} \circ b_{yz} \in \mathcal{B}_M$ by the characterization of yz as an interior edge, hence we have control over $\Psi_{\pi_{z-1} \circ b_{yz}, uv}$ and so

$$\begin{aligned} \mathbb{E}[\Delta_{yz} \Psi_{\pi, uv} \mid \mathcal{F}_i] - \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \mathbb{E}[\Delta_{yz} \Psi_{\pi^-, uv} \mid \mathcal{F}_i] &= -\frac{\Psi_{\pi_{z-1} \circ b_{yz}, uv}}{\Psi_{\pi_{z-1}, uv} Q} X_{\pi, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv} \\ &= -\frac{6 - o(1)}{n^2 p} X_{\pi, uv} + \frac{O(\zeta^2)}{n^2 p} \Psi_{\pi, uv}. \end{aligned} \quad (5.11)$$

Now consider an interior edge yz with $y < z = |\pi| + 1$. Again let $b_{yz} = 0$ if $y < z - 1$ and $b_{yz} = 1$ if $y = z - 1$. Since here $\pi_{z-1} = \pi$ there is no need to enumerate over extensions from an

intermediate graph to \mathcal{L}_π and we simply have, due to our $1 + O(\zeta^2)$ control over Q ,

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] = -\frac{\Psi_{\pi ob_{yz},uv}}{Q} = -\frac{X_{\pi ob_{yz},uv} + np^2\Psi_{\pi,uv}}{Q} = -(6 + O(\zeta^2))\frac{X_{\pi ob_{yz},uv}}{n^3p^3} - \frac{6 + O(\zeta^2)}{n^2p}\Psi_{\pi,uv}.$$

In particular,

$$\begin{aligned} \mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6}{n^2p}\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\Psi_{\pi^-,uv} &= \mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6}{n^2p}(\Psi_{\pi,uv} - X_{\pi,uv}) \\ &= -\frac{6}{n^2p}X_{\pi,uv} - \frac{6 - o(1)}{n^3p^3}X_{\pi ob_{yz},uv} + \frac{O(\zeta^2)}{n^2p}\Psi_{\pi,uv}. \end{aligned} \quad (5.12)$$

Observe that this last equation provides a large self-correcting term, and yet we see that errors in $X_{\pi o0,uv}$ and $X_{\pi o1,uv}$ will influence the error in $X_{\pi,uv}$.

- *Side boundary edges with a standard fan:* Let yz be a side boundary edge such that $y < z$ and $\pi(y-2) \neq \mathbf{e}$ (the case $\pi(y-2) = \mathbf{e}$, where the fan captured by the corresponding side boundary edge has an exceptional structure, will be treated separately).

Recall that in this case there are $\mathfrak{C}_\pi \in \{0, 1\}$ copies of \mathbf{e} in π and $|\pi| \leq 3M - 1 - (M - 1)\mathfrak{C}_\pi$. Moreover, $y > 0$ and $z = y + M$ following Eq. (4.1), and $z \leq 3M - 1 - (M - 1)\mathfrak{C}_\pi$ as otherwise yz would be classified as an outer boundary edge rather than a side boundary edge. From these facts we infer that $\pi_{y-1} \circ \mathbf{e}1 \in \mathcal{B}_M$ since this string contains $\mathfrak{C}_\pi + 1$ copies of \mathbf{e} , its last symbol is 1 (as opposed to \mathbf{e}) and it has length $y + 1 = z - M + 1 \leq 3M - 1 - (M - 1)(\mathfrak{C}_\pi + 1)$.

Similar to our argument above for initial edges (the case of an initial edge yz with $z = y + 2$), we will let $y + 1$ in the graph $\mathcal{L}_{\pi_{y-1} \circ \mathbf{e}}$ play the role of z . To be precise, for each $\varphi \in \Psi_{\pi_{y-1} \circ \mathbf{e}, uv}$ let φ' be the composition of φ and the map $z \mapsto y + 1$. By the definition of the forward extension $F = F_{\pi, yz}$ as having $I_F = \{0, 1, \dots, y\} \cup \{z\}$, the variable $\Psi_{F, \varphi'}$ exactly captures the number of ways to extend φ' to a full copy of \mathcal{L}_π , thus

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] = - \sum_{\varphi \in \Psi_{\pi_{y-1} \circ \mathbf{e}, uv}} \frac{Y_{\varphi(y)\varphi(y+1)}\Psi_{F, \varphi'}}{Q}.$$

Since $\sum_{\varphi} Y_{\varphi(y)\varphi(y+1)} = \Psi_{\pi_{y-1} \circ \mathbf{e}1}$ while $\Psi_{F, \varphi'} = (1 + O(\zeta^\delta))\mathcal{S}_F$ for all φ' by Corollary 4.12, we can conclude from Lemma 5.2 that

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] = -\frac{\Psi_{\pi_{y-1} \circ \mathbf{e}1, uv}}{\Psi_{\pi_{y-1} \circ \mathbf{e}, uv}Q}\Psi_{\pi,uv} + \frac{O(\zeta^{1+\delta})}{n^2p}\Psi_{\pi,uv}. \quad (5.13)$$

We now turn to the effect this has on the one-step change in ΔX . When $y < z \leq |\pi|$ we have

$$\begin{aligned} \mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] - \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\mathbb{E}[\Delta_{yz}\Psi_{\pi^-,uv} \mid \mathcal{F}_i] &= -\frac{\Psi_{\pi_{y-1} \circ \mathbf{e}1, uv}}{\Psi_{\pi_{y-1} \circ \mathbf{e}, uv}Q}X_{\pi,uv} + \frac{O(\zeta^{1+\delta})}{n^2p}\Psi_{\pi,uv} \\ &= -\frac{6 - o(1)}{n^2p}X_{\pi,uv} + \frac{O(\zeta^{1+\delta})}{n^2p}\Psi_{\pi,uv}. \end{aligned} \quad (5.14)$$

Alternatively, if $y < z = |\pi| + 1$, since the main term of the right-hand-side of (5.13) equals

$$\frac{X_{\pi_{y-1} \circ \mathbf{e}1, uv} + np^2\Psi_{\pi_{y-1} \circ \mathbf{e}, uv}}{\Psi_{\pi_{y-1} \circ \mathbf{e}, uv}Q}\Psi_{\pi,uv} = \frac{6 + o(1)}{n^3p^3}\frac{\Psi_{\pi,uv}}{\Psi_{\pi_{y-1} \circ \mathbf{e}, uv}}X_{\pi_{y-1} \circ \mathbf{e}1, uv} + \frac{6 + O(\zeta^2)}{n^2p}\Psi_{\pi,uv}$$

(the error terms above are due solely to the approximation of Q) it follows that

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6}{n^2p}\frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}}\Psi_{\pi^-,uv} = -\frac{6}{n^2p}X_{\pi,uv} - \frac{6 + o(1)}{n^3p^3}\frac{\Psi_{\pi,uv}}{\Psi_{\pi_{y-1} \circ \mathbf{e}, uv}}X_{\pi_{y-1} \circ \mathbf{e}1, uv} + \frac{O(\zeta^{1+\delta})}{n^2p}\Psi_{\pi,uv}. \quad (5.15)$$

- *Side boundary edges with an exceptional fan:* Let yz where $y < z$ is a side boundary edge such that $\pi(y-2) = \mathbf{e}$. To tackle the expected change in this scenario we will modify the argument from the previous case to incorporate estimates on homomorphisms from $\mathcal{L}_{\pi_{y-2}\circ\mathbf{e}}$ and $\mathcal{L}_{\pi_{y-2}\circ\mathbf{e}1}$.

Recall that now there is precisely 1 occurrence of \mathbf{e} in π (at index $y-2$) and so $|\pi| \leq 2M$. We further have $y \geq 3$ and $z = y + M - 1$ following (4.1), and $z \leq 2M$ so as not to make yz an outer boundary edge. Thus $y = z - (M-1) \leq M+1$ and we see that $\pi_{y-2} \circ \mathbf{e}1 \in \mathcal{B}_M$ as it has 2 copies of \mathbf{e} and ends with a 1 at index $y \leq M+1$ (of course $\pi_{y-2} \circ \mathbf{e} \in \mathcal{B}_M$ as well).

For each $\varphi \in \Psi_{\pi_{y-2}\circ\mathbf{e}}$ let φ' be the composition of φ and the map that sends $y \mapsto y-1$ and $z \mapsto y$. By definition, the forward extension $F_{\pi,yz}$ has $I_F = \{0, 1, \dots, y-2\} \cup \{y, z\}$ and it is easy to verify that as in the previous case (where $\pi(y-2) \neq \mathbf{e}$) we have

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] = - \sum_{\varphi \in \Psi_{\pi_{y-2}\circ\mathbf{e},uv}} \frac{Y_{\varphi(y-1)\varphi(y)}\Psi_{F,\varphi'}}{Q}.$$

The rest of the argument proceeds exactly as in the case $\pi(y-2) \neq \mathbf{e}$ with the single exception that π_{y-2} replaces all occurrences of π_{y-1} . Overall we obtain that (5.14) is valid whenever $y < z \leq |\pi|$, whereas in the situation of $y < z = |\pi| + 1$ we have

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6}{n^2p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv} = -\frac{6}{n^2p} X_{\pi,uv} - \frac{6+o(1)}{n^3p^3} \frac{\Psi_{\pi,uv}}{\Psi_{\pi_{y-2}\circ\mathbf{e},uv}} X_{\pi_{y-2}\circ\mathbf{e}1,uv} + \frac{O(\zeta^{1+\delta})}{n^2p} \Psi_{\pi,uv}. \quad (5.16)$$

- *Outer boundary edges:* Let yz with $z = |\pi| + 1$ be an outer boundary edge. By definition, extending \mathcal{L}_π by adding a new vertex incident to y, z leads us out of \mathcal{B}_M , foiling the argument we used for the other edge classes. In lieu of counting mappings of subgraphs that include yz and then seeking estimates on their forward extensions to \mathcal{L}_π , we will first directly identify the images of y, z and then estimate the number of backward extensions to \mathcal{L}_π . Formally, let $B = B_{\pi,yz}$ be the backward extension from yz to \mathcal{L}_π and for a given ordered pair (a, b) with $ab \in E(G)$ and $\{a, b\} \cap \{u, v\} = \emptyset$ let φ^{ab} be the injection that maps $0 \mapsto u, 1 \mapsto v, y \mapsto a$ and $z \mapsto b$. Then

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] = - \sum_{\substack{(a,b) : \{a,b\} \cap \{u,v\} = \emptyset \\ ab \in E(G)}} \frac{Y_{ab}\Psi_{B,\varphi^{ab}}}{Q}.$$

Summing Y_{ab} over the $2(|E(G)| - Y_u - Y_v + \mathbb{1}_{\{uv \in E\}})$ ordered pairs (a, b) as above amounts to $6Q - 2\Psi_{\mathbf{e}1,uv} - 2\Psi_{\mathbf{e}1,vu} + \mathbb{1}_{\{uv \in E\}}Y_{uv}$, with the terms $\Psi_{\mathbf{e}1,uv}$ and $\Psi_{\mathbf{e}1,vu}$ accounting for triangles incident to u (playing the role of either a or b , hence the factor of 2) and those incident to v . The corresponding summation over $\Psi_{B,\varphi}$ as usual gives $\Psi_{\pi,uv}$, and we infer from Corollary 4.10 that

$$\begin{aligned} \mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] &= -\frac{(3Q - \Psi_{\mathbf{e}1,uv} - \Psi_{\mathbf{e}1,vu} + O(Y_{uv}))\Psi_{\pi,uv}}{(|E(G)| - Y_u - Y_v + O(1))Q} + \frac{O(\zeta^{1+\delta})}{n^2p} \Psi_{\pi,uv} \\ &= -\frac{6}{n^2p} \Psi_{\pi,uv} + \frac{O(\zeta^{1+\delta})}{n^2p} \Psi_{\pi,uv}, \end{aligned}$$

where we used the fact that $\Psi_{\mathbf{e}1,uv}/Q$ as well as $Y_u/|E(G)|$ and $Y_v/|E(G)|$ are both $O(1/n)$, while $Y_{uv}/Q = O(1/n^2p) = o(1/n)$, to omit all the terms $\Psi_{\mathbf{e}1,uv}, \Psi_{\mathbf{e}1,vu}, Y_u, Y_v, Y_{uv}$ at an additive cost of $O(\frac{1}{n}) \frac{\Psi_{\pi,uv}}{n^2p} = o(\zeta^2) \frac{\Psi_{\pi,uv}}{n^2p}$, and similarly, we replaced the quantity $|E(G)| = \frac{1}{2}(n^2p - n)$ simply by $\frac{1}{2}n^2p$ at an additive cost of $O(\frac{1}{np}) \frac{\Psi_{\pi,uv}}{n^2p} = o(\zeta^2) \frac{\Psi_{\pi,uv}}{n^2p}$.

It now follows that for any outer boundary edge yz ,

$$\mathbb{E}[\Delta_{yz}\Psi_{\pi,uv} \mid \mathcal{F}_i] + \frac{6}{n^2p} \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv} = -\frac{6}{n^2p} X_{\pi,uv} + \frac{O(\zeta^{1+\delta})}{n^2p} \Psi_{\pi,uv}. \quad (5.17)$$

To conclude the proof, observe that every edge type contributes the term $-(6+o(1))\frac{1}{n^2p}X_{\pi,uv}$ to the summation in (5.6), and correction terms accumulate when $z = |\pi| + 1$. Outer boundary edges do not contribute any additional terms. When $y = z - 1$, Eqs. (5.9),(5.12) reveal that both initial edges and interior edges contribute an extra term of $-(6+o(1))\frac{1}{n^3p^3}X_{\pi\circ 1,uv}$, whereas side boundary edges cannot have $y = z - 1$ for any $M \geq 3$. When $y < z - 1$, initial edges contribute the extra term $-(6+o(1))\frac{1}{np^2}X_{\pi,uv}$ as per (5.10), interior edges contribute the extra term $-(6+o(1))\frac{1}{n^3p^3}X_{\pi\circ 1,uv}$ as per (5.12) and side boundary edges contribute the term featured in (5.15). These correction terms match the definition of Ξ_0, Ξ_1 in the statement of the theorem, thus completing the proof. \blacksquare

5.2. Concentration. In this subsection we prove the following theorem:

Theorem 5.4. *Let $M \geq 3$ and let $0 < \delta < 1$ be the constant from Corollaries 4.10 and 4.12. For any π in \mathcal{B}_M let π^- denote its $(|\pi| - 1)$ -prefix and define*

$$X_{\pi,uv} = \Psi_{\pi,uv} - \frac{\mathcal{S}_\pi}{\mathcal{S}_{\pi^-}} \Psi_{\pi^-,uv},$$

where $\Psi_{\emptyset,uv} := 1$. Let $\zeta = n^{-1/2}p^{-1} \log n$ and for some absolute constant $\kappa > 0$ let

$$\tau_Q^* = \tau_Q^*(\kappa) = \min \left\{ t : \left| Q / \left(\frac{1}{6}n^3p^3 \right) - 1 \right| \geq \kappa \zeta^2 \right\}.$$

Then w.h.p., for every $\pi \in \mathcal{B}_M$, every $u, v \in V_G$ and all t such that $t \leq \tau_Q^*$ and $p(t) \geq p_M$

$$|X_{\pi,uv}| \leq \omega_\pi \zeta \mathcal{S}_\pi,$$

where $\omega_\pi = 3^{3M-|\pi|-(M-1)\mathfrak{c}_\pi}$ with \mathfrak{c}_π denoting the number of copies of \mathfrak{e} in π .

Proof. Let $\bar{\tau}$ be the stopping time where for some $\pi \in \mathcal{B}_M$ and $u, v \in V_G$ we have $|X_{\pi,uv}| > \omega_\pi \zeta \mathcal{S}_\pi$. To show that w.h.p. we do not encounter $\bar{\tau}$ as long as $t \leq \tau_*$ and $p(t) \geq p_M$, consider a pair π, uv where at some step i_0 the variable $|X_{\pi,uv}|$ has just entered its critical interval

$$I_\pi = \left[\left(1 - \frac{1}{4e_\pi} \right) \omega_\pi \zeta \mathcal{S}_\pi, \omega_\pi \zeta \mathcal{S}_\pi \right].$$

We define

$$Z_{\pi,uv} = |X_{\pi,uv}| - \omega_\pi \zeta \mathcal{S}_\pi,$$

and will now argue that $Z_{\pi,uv}$ is a supermartingale as long as $|X_{\pi,uv}| \in I_\pi$ and in addition we have not yet encountered $\bar{\tau} \wedge \tau_*$ and are above the final probability threshold p_M . Indeed, since $\Delta(\zeta \mathcal{S}_\pi) = -\frac{6+o(1)}{n^2p}(e_\pi - 1)\zeta \mathcal{S}_\pi$, in this setting Theorem 5.1 implies that

$$\begin{aligned} \mathbb{E}[\Delta Z_{\pi,uv} \mid \mathcal{F}_i] &\leq \frac{6+o(1)}{n^2p} \left((e_\pi - 1)\omega_\pi \zeta \mathcal{S}_\pi - e_\pi |X_{\pi,uv}| \right) + |\Xi_0| + |\Xi_1| + \frac{O(\zeta^{1+\delta})}{n^2p} \mathcal{S}_\pi \\ &\leq -\frac{6+o(1)}{n^2p} \frac{3}{4} \omega_\pi \zeta \mathcal{S}_\pi + |\Xi_0| + |\Xi_1| + \frac{O(\zeta^{1+\delta})}{n^2p} \mathcal{S}_\pi. \end{aligned} \quad (5.18)$$

Examining the definitions of Ξ_0, Ξ_1 from that theorem, we see that the contribution of Ξ_0 when the mentioned edge $(y, |\pi| + 1)$ is an initial edge translates in our setting to $-\frac{6+o(1)}{n^2p}|X_{\pi,uv}|$, strictly negative for large enough n . As we look for an upper bound on $\Xi_0 + \Xi_1$, it is thus legitimate to ignore this case. However, in all other cases with these correction terms, we have no guarantee that $X_{\sigma,uv}$ for $\sigma \neq \pi$ shares the same sign as $X_{\pi,uv}$. Worst case estimates for these variables (recalling the bounds on their absolute values, available throughout $t < \bar{\tau}$) give

$$|\Xi_1| \leq \frac{6+o(1)}{n^3p^3} \omega_{\pi\circ 1} \zeta \mathcal{S}_{\pi\circ 1} = \frac{6+o(1)}{n^2p} \omega_{\pi\circ 1} \zeta \mathcal{S}_\pi,$$

and similarly,

$$\frac{|\mathbb{E}_0|}{\frac{6}{n^2p}\zeta\mathcal{S}_\pi} \leq \begin{cases} (1+o(1))\omega_{\pi_{y-1}\circ\mathbf{e}1} & \text{if } \mathcal{L}_\pi \text{ has a side boundary edge } yz \text{ with } z = |\pi| + 1, \pi(y-2) \neq \mathbf{e} \\ (1+o(1))\omega_{\pi_{y-2}\circ\mathbf{e}1} & \text{if } \mathcal{L}_\pi \text{ has a side boundary edge } yz \text{ with } z = |\pi| + 1, \pi(y-2) = \mathbf{e} \\ (1+o(1))\omega_{\pi\circ 0} & \text{if } \mathcal{L}_\pi \text{ has an interior edge } yz \text{ with } y < z - 1 \text{ and } z = |\pi| + 1 \end{cases}$$

Returning to (5.18), when \mathcal{L}_π does not contain a side boundary edge yz with $z = |\pi| + 1$,

$$\begin{aligned} \mathbb{E}[\Delta Z_{\pi,uv} \mid \mathcal{F}_i] &\leq -\frac{6+o(1)}{n^2p} \left(\frac{3}{4}\omega_\pi - \omega_{\pi\circ 0} - \omega_{\pi\circ 1} \right) \zeta\mathcal{S}_\pi + \frac{O(\zeta^{1+\delta})}{n^2p} \mathcal{S}_\pi \\ &= -\frac{1/2+o(1)}{n^2p} \omega_\pi \zeta\mathcal{S}_\pi + \frac{O(\zeta^{1+\delta})}{n^2p} \mathcal{S}_\pi, \end{aligned}$$

where we set $\omega_\sigma = 0$ if $\sigma \notin \mathcal{B}_M$ and the equality stemmed from the fact that for any $b \in \{0, 1\}$, since π and $\pi \circ b$ share the same number of \mathbf{e} symbols, our definition of the constants $(\omega_\sigma)_{\sigma \in \mathcal{B}_M}$ implies that $\omega_{\pi\circ b} = \frac{1}{3}\omega_\pi$.

Similarly, when yz with $z = |\pi| + 1$ is a side boundary edge of \mathcal{L}_π we can let $r = y - 1$ if $\pi(y-2) \neq \mathbf{e}$ and $r = y - 2$ otherwise, and obtain that

$$\mathbb{E}[\Delta Z_{\pi,uv} \mid \mathcal{F}_i] \leq -\frac{6+o(1)}{n^2p} \left(\frac{3}{4}\omega_\pi - \omega_{\pi\circ 1} - \omega_{\pi_{r\circ\mathbf{e}1}} \right) \zeta\mathcal{S}_\pi + \frac{O(\zeta^{1+\delta})}{n^2p} \mathcal{S}_\pi.$$

By definition of side boundary edges, π has at most a single copy of \mathbf{e} , thus $\omega_\pi = 3^{3M-|\pi|-(M-1)\mathbf{e}_\pi}$ and $\omega_{\pi_{r\circ\mathbf{e}1}} = 3^{3M-(r+2)-(M-1)(\mathbf{e}_\pi+1)} = 3^{2M-1-r-(M-1)\mathbf{e}_\pi}$. Furthermore, revisiting (4.1) we have $y = z - M = |\pi| + 1 - M$ if $\pi(y-2) \neq \mathbf{e}$ (and then $r = |\pi| - M$) while $y = z - (M-1) = |\pi| - M + 2$ if $\pi(y-2) = \mathbf{e}$ (and then again $r = |\pi| - M$). In both cases we therefore get that $\omega_{\pi_{r\circ\mathbf{e}1}} = \frac{1}{3}\omega_\pi = \omega_{\pi\circ 1}$ and so again we obtain that

$$\mathbb{E}[\Delta Z_{\pi,uv} \mid \mathcal{F}_i] \leq -\frac{1/2+o(1)}{n^2p} \omega_\pi \zeta\mathcal{S}_\pi + \frac{O(\zeta^{1+\delta})}{n^2p} \mathcal{S}_\pi.$$

In particular, $\mathbb{E}[\Delta Z_{\pi,uv} \mid \mathcal{F}_i] < 0$ for large enough n , as claimed.

We are left with the task of bounding the one-step changes ΔZ given \mathcal{F}_i in L^∞ and L^2 norm. Revisiting the events \mathcal{E}_ψ and \mathcal{E}_ψ^{yz} defined in (5.2), observe that when aiming for an upper bound on $|\Delta\Psi_{\pi,uv}|$ (as opposed to tight asymptotics, where (5.3) incurred an approximation error) one always has

$$|\Delta\Psi_{\pi,uv}| = \sum_{\psi \in \Psi_{\pi,uv}} \mathbb{1}_{\mathcal{E}_\psi} \leq \sum_{\psi \in \Psi_{\pi,uv}} \sum_{yz \in E(\mathcal{L}_\pi)} \mathbb{1}_{\mathcal{E}_\psi^{yz}} = \sum_{yz \in E(\mathcal{L}_\pi)} |\Delta_{yz}\Psi_{\pi,uv}|.$$

In particular, for the sake of an upper bound on $|\Delta Z|$ and $\mathbb{E}[(\Delta Z)^2 \mid \mathcal{F}_i]$ it will suffice to analyze the effect of each edge $yz \in E(\mathcal{L}_\pi)$ separately.

Let $yz \in E(\mathcal{L}_\pi)$ with $y < z$ and let $H = (V(\mathcal{L}_\pi), E(\mathcal{L}_\pi))$ be the extension graph featuring $I_H = \{0, 1\}$. As argued in the proof of Theorem 3.2, let $T = T(yz) \subset H$ be the subextension of H with minimal scaling out of all subextensions that contain the edge yz (and, if multiple such subextensions exist, with minimal cardinality out of those).

Consider some subextension of the quotient H/T . This graph can be written as K/T for some subextension $K \subset H$ that contains all edges of T , thus $\mathcal{S}_K \geq \mathcal{S}_T$ by the minimality of T , and equivalently $\mathcal{S}_{K/T} \geq 1$. Part (1) of Corollary 3.3 now implies that for some $c_1 > 0$ and any injection $\varphi : T \rightarrow V_G$ we have

$$\Psi_{H/T, \varphi} \leq \mathcal{S}_{H/T} (\log n)^{c_1}. \quad (5.19)$$

Notice that at the same time any subextension $K \subset H$ has $\mathcal{S}_K \geq 1$. This follows from the fact that, by Definition 4.1, every vertex $1 < i \leq |\pi| + 1$ shares at most 2 edges with the vertices $\{j : j < i\}$, thus $e_K \leq 2(v_K - \iota_K)$ and $\mathcal{S}_K \geq n^{(v_K - \iota_K)/M} \geq 1$ due to the hypothesis $p \geq p_M$. Since every $K \subset T$

has $\mathcal{S}_K \geq 1$ (being also a subextension of H) we may again appeal to Part (1) of Corollary 3.3 to obtain that for some $c_2 > 0$ and any $u, v \in V_G$ we have

$$\Psi_{T,uv} \leq \mathcal{S}_T(\log n)^{c_2}. \quad (5.20)$$

A final ingredient we need is to control the number of extensions to a copy of T given the image of the edge yz . Formally, let T_0 be the induced graph on the vertex set $\{0, 1, y, z\}$ and consider the quotient $T^* = T/T_0$ (that is, the extension graph obtained from T by increasing the distinguished vertex set to $\{0, 1, y, z\}$ and removing any edges among these vertices). The choice of T implies that every $K^* \subsetneq T^*$ has $\mathcal{S}_{K^*} > \mathcal{S}_{T^*}$ (as otherwise the result of moving the vertices y, z from I_K to $V_K \setminus I_K$ and adding all edges of T among $\{0, 1, y, z\}$ would yield a graph K with $\mathcal{S}_K \leq \mathcal{S}_T$ and a strictly smaller vertex set). Equivalently, $\mathcal{S}_{T^*/K^*} \leq 1$ for all $K^* \subset T^*$ and so, by Part (2) of Corollary 3.3, we deduce that for some $c_3 > 0$ and any injection $\eta : \{0, 1, y, z\} \rightarrow V_G$,

$$\Psi_{T/T_0,\eta} \leq (\log n)^{c_3}. \quad (5.21)$$

We now translate these facts to L^∞ and L^2 bounds on $\Delta_{yz}X_{\pi,uv}$ given \mathcal{F}_i . The analysis will differ depending on whether or not $y \in \{0, 1\}$.

- (i) In case $y > 1$: Since T contains the edge yz we have $\mathcal{S}_T \geq n^2p^4$ already due to the vertices y, z (as argued above, each vertex $i > 1$ that belongs to V_T sends at most two edges backwards).

For the L^∞ bound we look at the maximum number of extensions $\psi \in \Psi_{\pi,uv}$ all of which map yz to a particular edge in $E(G)$. Equivalently, we are examining the maximum possible value of $\Psi_{H/T_0,\eta}$ over all injections $\eta : V(T_0) \rightarrow V_G$ that map $0 \rightarrow u$ and $1 \rightarrow v$ with the additional property of preserving all edges of the induced graph of H on $\{0, 1, y, z\}$. Each such extension ψ factors into extensions $T_0 \hookrightarrow T$ and $T \hookrightarrow H$, thus an application of (5.21) and (5.19) gives

$$|\Delta_{yz}\Psi_{\pi,uv}| \leq \max_{\eta} \Psi_{H/T_0,\eta} \leq \max_{\eta} \sum_{\varphi \in \Psi_{T/T_0,\eta}} \Psi_{H/T,\varphi} \leq \mathcal{S}_{H/T}(\log n)^{c_1+c_3}. \quad (5.22)$$

Since $\mathcal{S}_{H/T} = \mathcal{S}_H/\mathcal{S}_T$ and $\mathcal{S}_T \geq n^2p^4$ we further have

$$|\Delta_{yz}\Psi_{\pi,uv}| \leq \max_{\eta} \Psi_{H/T_0,\eta} \leq \frac{(\log n)^{c_1+c_3}}{n^2p^4} \mathcal{S}_H. \quad (5.23)$$

For the L^2 bound, observe that the set of edges in $E(G)$ that can play the role of yz for some copy of H rooted at uv is clearly contained in $\{\psi(y)\psi(z) : \psi \in \Psi_{T,uv}\}$. Using (5.20), the probability that such an edge is chosen in the next round is therefore at most

$$\sum_{\psi \in \Psi_{T,uv}} \frac{Y_{\psi(y)\psi(z)}}{Q} \leq (6 + o(1)) \frac{\mathcal{S}_T(\log n)^{c_2}}{n^2p}.$$

By (5.22), combined with the usual fact that $\mathcal{S}_{H/T} = \mathcal{S}_H/\mathcal{S}_T$, it follows that

$$\mathbb{E}[(\Delta_{yz}\Psi_{\pi,uv})^2 \mid \mathcal{F}_i] \leq (6 + o(1)) \frac{(\mathcal{S}_H)^2(\log n)^{c_2+2(c_1+c_3)}}{\mathcal{S}_T n^2p} \leq \frac{(\mathcal{S}_H)^2(\log n)^C}{n^4p^5},$$

where the last inequality featured some suitably large constant $C > 0$ and relied on the fact that $\mathcal{S}_T \geq n^2p^4$. Integrating this over a maximum of n^2p remaining steps gives

$$\sum_{i > i_0} \mathbb{E}[(\Delta_{yz}\Psi_{\pi,uv})^2 \mid \mathcal{F}_i] \leq \sum_{i > i_0} \frac{(\mathcal{S}_H)^2(\log n)^C}{n^4p(i)^5} \leq \frac{(\mathcal{S}_H)^2}{n^2p(i_0)^4} (\log n)^C. \quad (5.24)$$

- (ii) In case $y \in \{0, 1\}$ and $\pi(1) = 1$:

- Suppose $z = 2$. Here $V(T) = \{0, 1, 2\}$ and $E(T) = \{02, 12\}$ (any additional vertex would increase \mathcal{S}_T). The number of extensions $\psi \in \Psi_{\pi, uv}$ which map yz to a prescribed edge in $E(G)$ is exactly the number of forward extensions from yz to π . Since $F = F_{\pi, yz}$ has $V(T)$ as its distinguished vertices, $\mathcal{S}_F = \mathcal{S}_H/(np^2)$ and Corollary 4.12 gives

$$\Delta_{yz}\Psi_{\pi, uv} = O(\mathcal{S}_H/(np^2)). \quad (5.25)$$

We claim that the probability of choosing a triangle incident to yz is $O(p/n)$. Indeed, if $y = 0$ then we are looking at triangles incident to u and to some neighbor of v (to assume the role of the image of z), and there are $N_{u,v} = O(np^2)$ candidates for the image of z each sharing $O(np^2)$ triangles with u . The same holds symmetrically if $y = 1$ and we deduce that

$$\sum_{i>i_0} \mathbb{E}[(\Delta_{yz}\Psi_{\pi, uv})^2 \mid \mathcal{F}_i] \leq \sum_{i>i_0} O\left(\frac{p}{n} \frac{\mathcal{S}_H^2}{n^2 p^4}\right) = O\left(\frac{\mathcal{S}_H^2}{np(i_0)^2}\right). \quad (5.26)$$

- Suppose $z \geq 3$. Here necessarily $y = 1$ and either $\pi_z = 110\dots 0$ or $\pi_z = 11$ (since $\pi(2) \neq 0$). Note that if $V(T) = \{0, 1, 2, \dots, z\}$ then $\mathcal{S}_T \geq n^2 p^4$ already due to two backward edges sent from 2 and two such edges from z . Thus, the argument for $y > 1$ (Case (i)) applies in this case as well and implies the bounds in (5.23) and (5.24) for the L^∞ and L^2 norms, resp. Otherwise, we must have $V(T) = \{0, 1, z\}$ and $E(T) = \{1z\}$. (Indeed, including any $w > z$ in $V(T)$ would clearly increase \mathcal{S}_T . If $V(T) \cap A \subsetneq A$ for $A = \{2, \dots, z-1\}$ then of course $z > 3$ and so $\pi_z = 110\dots 0$. Thus, the vertex with minimal index in $V(T) \cap A$ sends 2 edges backwards and its removal from T would decrease \mathcal{S}_T , contradicting the minimality of T .)

For the L^∞ bound, we can control the number of forward extensions $F_{\pi, yz}$ from yz to π via Corollary 4.12. If yz happens to be a side boundary edge (in our setting this will happen iff z concludes an $(M-1)$ -fan at vertex 1) then $I_F = V(T) = \{0, 1, z\}$ thus $\mathcal{S}_F = \mathcal{S}_H/(np)$ and

$$\Delta_{yz}\Psi_{\pi, uv} = O(\mathcal{S}_H/(np)).$$

For all other edge-types $I_F = \{0, 1, \dots, z\}$ and we must examine the number of extensions from $V(T) = \{0, 1, z\}$ to the induced graph of H on I_F . Specifically, Let R be the extension graph of \mathcal{L}_{π_z} with the distinguished vertex set $\{0, 1, z\}$. Recalling that $\pi_z = 11\dots 0$, there are $z-2$ vertices and $2(z-1)-1 = 2z-3$ edges in R . Observe that R is strictly balanced, since we can view any vertex $1 < j < z$ is accountable for two exclusive directed edges going from it to 1 and $j+1$ (the maximal density is attained when all such j 's are present). Finally, $\mathcal{S}_R \geq 1$ as otherwise adding the vertices to $\{2, \dots, z-1\}$ to $V(T)$ would give an extension graph containing yz with a scaling of $\mathcal{S}_T \mathcal{S}_R < \mathcal{S}_T$, violating the definition of T . That is,

$$n^{z-2} p^{2z-3} \geq 1.$$

As before, we count extensions $\psi \in \Psi_{\pi, uv}$ which map yz to a prescribed edge $e = vw$ via the product of $\Psi_{R, uvw}$ and $\max_\eta \Psi_{F_{\pi, yz}, \eta}$. The latter is asymptotically $\mathcal{S}_F = \mathcal{S}_H/(np\mathcal{S}_R)$ by Corollary 4.12. For the former we apply Theorem 3.2, yielding an upper bound of $(1+o(1))\mathcal{S}_R$ as long as $t \leq t_R^-(\sqrt{\delta})$ and the upper bound $\mathcal{S}_R + \mathcal{S}_R^{1-\sqrt{\delta}}(\log n)^\rho$ for some fixed $\rho = \rho_R > 0$ for $t_R^-(\sqrt{\delta})t \leq t_R$. We have $\mathcal{S}_R \geq 1$, hence indeed $t \leq t_R$ and the combined upper bound of $\mathcal{S}_R + \max\{\mathcal{S}_R, \mathcal{S}_R^{1-\sqrt{\delta}}(\log n)^\rho\}$ is a valid one for $\Psi_{R, uvw}$. Moreover, since \mathcal{S}_R is decreasing we see that this bound is $2\mathcal{S}_R$ up to the first point where $\mathcal{S}_R \leq (\log n)^{\rho/\sqrt{\delta}}$ (namely, $p = n^{-\frac{z-2}{2z-3}+o(1)} \leq n^{-2/3+o(1)}$ for any $z \geq 3$). Beyond that point we can simply use the bound $2(\log n)^c$ for $c = (1 + \frac{1-\sqrt{\delta}}{\delta})\rho$. (Altogether,

$$\Delta_{yz}\Psi_{\pi, uv} \leq (2+o(1)) \max\{\mathcal{S}_R, \log^c n\} \frac{\mathcal{S}_H}{np\mathcal{S}_R} = O\left(\frac{\mathcal{S}_H}{np^{5/4}}\right), \quad (5.27)$$

where the last inequality used the aforementioned fact that the term $\log^c n$ can dominate \mathcal{S}_R only starting at $p \leq n^{-2/3+o(1)}$, whilst it is then easily absorbed by the extra factor of $p^{-1/4}$. The L^∞ bound in (5.27) is of larger order than the $O(\mathcal{S}_H/(np))$ bound given for the case where yz is a side boundary edge, thus we can safely use it for any edge-type.

For the L^2 bound, note that the probability of selecting a triangle incident to v (supporting the removal of yz) is $\sum_{w \in N_v} N_{vw}/Q = O(1/n)$. Summing the one-step expected variances in the usual way yields

$$\sum_{i>i_0} \mathbb{E}[(\Delta_{yz} \Psi_{\pi,uv})^2 \mid \mathcal{F}_i] \leq \sum_{i>i_0} O\left(\frac{\mathcal{S}_H^2}{n^3 p(i)^{5/2}}\right) = O\left(\frac{\mathcal{S}_H^2}{np(i_0)^{3/2}}\right). \quad (5.28)$$

(iii) In case $\pi(1) = \mathbf{e}$ and $y = 1$: By definition yz is an initial edge, and since the corresponding forward extension $F_{\pi,yz}$ has $I_F = \{0, 1, z\}$ we can invoke Corollary 4.12 and infer that

$$\Delta_{yz} \Psi_{\pi,uv} \leq \max_{\eta} \Psi_{F_{\pi,yz},\eta} = O(\mathcal{S}_F) = O(\mathcal{S}_H/(np)) \quad (5.29)$$

and (again by the fact that the probability of choosing a triangle incident to v is $O(1/n)$)

$$\sum_{i>i_0} \mathbb{E}[(\Delta_{yz} \Psi_{\pi,uv})^2 \mid \mathcal{F}_i] \leq \sum_{i>i_0} O\left(\frac{\mathcal{S}_H^2}{n^3 p(i)^2}\right) \leq O\left(\frac{\mathcal{S}_H^2}{np(i_0)}\right). \quad (5.30)$$

We now combine the bounds above for analogous estimates on ΔZ . Comparing (5.23), (5.25), (5.27) and (5.29) we see that for all π and all edges yz in \mathcal{L}_π we have

$$\Delta_{yz} \Psi_{\pi,uv} = O\left(\frac{\mathcal{S}_H}{np^2}\right) \quad (5.31)$$

whereas a comparison of (5.24), (5.26), (5.28) and (5.30) shows that similarly we always have

$$\sum_{i>i_0} \mathbb{E}[(\Delta_{yz} \Psi_{\pi,uv})^2 \mid \mathcal{F}_i] \leq O\left(\frac{\mathcal{S}_H^2}{np^2}\right). \quad (5.32)$$

By summing these expressions over all $yz \in E(H)$ we arrive at the corresponding upper bounds on $\Delta \Psi_{\pi,uv}$ (of the same order). Furthermore, if $z \leq |\pi|$ we may repeat this analysis for \mathcal{L}_{π^-} , and since the obtained bounds will be in terms of \mathcal{S}_{π^-} while the variable $X_{\pi,uv}$ features $(\mathcal{S}_\pi/\mathcal{S}_{\pi^-})\Psi_{\pi^-,uv}$ we deduce that these bounds extend the aforementioned bounds to $X_{\pi,uv}$ itself. It is furthermore easy to verify that the change in Z due to the term $\zeta \mathcal{S}_\pi$, changing deterministic with p by an additive term of relative order $O(1/(n^2 p))$, is negligible in comparison with these bounds. (Indeed, it is a factor of $n/(p\zeta)$ smaller than our Lipschitz bound.) Altogether, the L^∞ and cumulative L^2 bounds hold for the supermartingale Z and we can now exploit these to obtain a large deviation estimate.

Observe that the L^∞ bound in (5.31) is $o(\zeta \mathcal{S}_H)$ since $n^{1/2} p$ tends to infinity (it is at least as large as $n^{1/M}$). Thus, the first step of Z whenever $|X_{\pi,uv}|$ enters the critical interval I_π is negligible in terms of the order of this window and we have $Z(i_0) = (-1/(4e_\pi) + o(1))\omega_\pi \zeta \mathcal{S}_\pi$.

Set $s = (5e_\pi)^{-1} \omega_\pi \zeta \mathcal{S}_\pi$ as the target deviation for our supermartingale Z . The product of the uniform L^∞ bound and s has order $\zeta(\mathcal{S}_\pi)^2/(np^2)$. This is of lower order compared to the cumulative L^2 -bound in (5.32) since $\zeta \rightarrow 0$ given the fact that $np^2 \geq n^{1/M}$ throughout the range $p \geq p_M$.

We may now invoke Freedman's inequality and get that for some fixed $c > 0$,

$$\mathbb{P}\left(\exists t > t_0 : Z(t \wedge \tau_* \wedge \bar{\tau}) \geq Z(t_0) + s\right) \leq \exp\left(-\frac{c(\zeta \mathcal{S}_\pi)^2}{\mathcal{S}_\pi^2/(np^2)}\right) = e^{-cnp^2 \zeta^2} = e^{-c(\log n)^2}.$$

Having arrived at an error probability that decays to 0 at a super-polynomially rate, a union bound over the ensemble of $(\Psi_{\pi,uv})$ variables establishes the required estimates for all $\pi \in \mathcal{B}_M$. \blacksquare

6. LOWER BOUND ON THE FINAL NUMBER OF EDGES

Our lower bound of $n^{3/2-o(1)}$ on the final number of edges will rely solely on the fact that for any fixed $\varepsilon > 0$ we have asymptotically tight estimates for all the co-degrees in the process at time t such that $p(t) = n^{-1/2+\varepsilon}$.

Theorem 6.1. *Suppose that for some fixed $0 < \varepsilon < \frac{1}{6}$, all co-degrees satisfy $Y_{u,v} = (1 + o(1))np^2$ throughout $p \geq p_0 = n^{-1/2+\varepsilon}$. Then w.h.p. the final number of edges is at least $n^{3/2-6\varepsilon-o(1)}$.*

Remark 6.2. The above result was formulated in accordance with our upper bound analysis, which achieves the required estimate over the co-degrees as long as the number of edges is of order $n^{3/2+\varepsilon}$ for an arbitrarily small fixed $\varepsilon > 0$. The proof argument of Theorem 6.1 in fact yields a stronger result, namely that an asymptotic upper bound of np^2 over all co-degrees throughout the point of $n^{3/2}h$ edges, for some $h(n) < n^{1/6}$ that grows to ∞ with n , implies that the final number of edges is w.h.p. at least $cn^{3/2}h^{-6}$ for some absolute $c > 0$.

An important ingredient in the proof is the following straightforward lemma which provides an upper bound on Q valid well beyond the point where our co-degree estimates break down. An estimate of this nature, albeit weaker, was provided in [4]. Towards the end of the process, what begins as a second order term in this bound will become the main term, hence the importance of estimating it sharply.

Lemma 6.3. *Let τ_c denote the minimal time where $Y_{u,v} > 2(np^2 + n^{1/3})$ for some $u, v \in V_G$. With high probability, for all $t \leq \tau_c$ such that $p(t) \geq n^{-3/5}$, the number of triangles satisfies*

$$Q - \frac{1}{6}(n^3p^3 + n^2p) \leq n^{7/3}p^2.$$

Proof. By a slight abuse of notation we let Q denote both the collection of all triangles in G and its cardinality. By definition,

$$\mathbb{E}[\Delta Q \mid \mathcal{F}_i] = - \sum_{uvw \in Q} \frac{Y_{u,v} + Y_{u,w} + Y_{v,w} - 2}{Q} = 2 - \frac{1}{Q} \sum_{uv \in E} Y_{u,v}^2.$$

Since $\sum_{uv \in E} Y_{u,v} = 3Q$, applying Cauchy-Schwarz shows that $\sum_{uv \in E} Y_{u,v}^2 \geq 9Q^2/|E|$ and thus

$$\mathbb{E}[\Delta Q \mid \mathcal{F}_i] \leq 2 - \frac{9Q}{|E|} = 2 - \frac{9Q}{\frac{1}{2}(n^2p - n)} \leq 2 - \frac{18}{n^2p}Q.$$

Define

$$\Upsilon = \frac{1}{6}(n^3p^3 + n^2p) + n^{7/3}p^2$$

and observe that the one-step change in this variable is deterministically given by

$$\Delta(\Upsilon) = - \left(1 + O\left(\frac{1}{n^2p}\right)\right) \frac{6}{n^2} \left(\frac{1}{2}n^3p^2 + \frac{1}{6}n^2 + 2n^{7/3}p\right) = - \left(3np^2 + 1 + 12n^{1/3}p\right) + O(p/n).$$

Therefore, if we set

$$Z = Q - \Upsilon \quad , \quad I_Q = \left[\Upsilon - \frac{1}{6}n^{7/3}p^2, \Upsilon\right]$$

it follows that whenever $Q \in I_Q$ we have

$$\mathbb{E}[\Delta Z \mid \mathcal{F}_i] \leq 2 - \left(3np^2 + 3 + 15n^{1/3}p\right) + \left(3np^2 + 1 + 12n^{1/3}p\right) + O(p/n) = -3n^{1/3}p + O(p/n) < 0,$$

where the last inequality holds for any sufficiently large n .

Let i_0 be the first time where $Q \in I_Q$ and write $t_0 = t(i_0)$ and $p_0 = p(i_0)$. For any $t_0 < \tau_c$

$$|\Delta Q| \leq 3 \max_{u,v} Y_{u,v} = O(np_0^2 + n^{1/3}),$$

and observing that this bound is of order strictly lower than $n^{7/3}p_0^2$ for all $p_0 \geq 1/n$ we obtain that

$$Z(i_0) = -\left(\frac{1}{6} - o(1)\right)n^{7/3}p_0^2.$$

From this point the process has at most $n^2p_0/2$ steps until its conclusion, therefore we may substitute $s = \frac{1}{8}n^{7/3}p_0^2$ and deduce from Hoeffding's inequality that for some fixed $c > 0$

$$\mathbb{P}(\exists i > i_0 : Z(i) \geq Z(i_0) + s) \leq \exp\left(-c \frac{s^2}{n^2p_0(np_0^2 + n^{1/3})^2}\right) = \exp\left(-c \frac{n^{8/3}p_0^3}{(np_0^2 + n^{1/3})^2}\right).$$

If $p_0 \geq n^{-1/3}$ then the last exponent has order $n^{2/3}/p_0 \geq n^{2/3}$, and otherwise it has order $n^2p_0^3 \geq n^{1/5}$ for any $p_0 \geq n^{-3/5}$. A union bound over i_0 now completes the proof. \blacksquare

Proof of Theorem 6.1. Fix $0 < \varepsilon < \frac{1}{6}$ and let i_0 be such that

$$p_0 = p(i_0) = n^{-1/2+\varepsilon}.$$

Condition on \mathcal{F}_{i_0} and assume that up to this point we had $Y_{u,v} = (1 + o(1))n^2p$ for all $u, v \in V_G$.

For two distinct triangles $uvw, xyz \in Q$ we write $uvw \sim xyz$ to denote that they share an edge. Further denote the number of distinct triangles incident to some $uvw \in Q$ by

$$B_{uvw} = \#\{xyz \in Q : xyz \sim uvw\} = Y_{u,v} + Y_{u,w} + Y_{v,w} - 3.$$

Observe that our hypothesis on the co-degrees implies that $B_{uvw} = n^{2\varepsilon+o(1)}$ for all Q triangles at time i_0 . In addition, we claim that if $uvw \in Q(j-1)$ then at the end of round j either $uvw \notin Q(j)$ or $B_{uvw}(j-1) - B_{uvw}(j) \leq 4$, and in the latter case the co-degree of at most 2 edges in uvw has changed. Indeed, if the triangle selected in round j is either uvw itself or some triangle $xyz \sim uvw$ then at once $uvw \notin Q(j)$, whereas removing a vertex disjoint triangle xyz does not affect B_{uvw} . Consider therefore the case of removing a triangle xyz sharing exactly one vertex in common with uvw , that is, w.l.o.g. $x = u$ and $\{y, z\} \cap \{v, w\} = \emptyset$. In this case, each of the edges uy, uz may participate in at most one triangle involving the edge uv and at most another involving the edge uw , thus at most 4 triangles can be deducted from the count of B_{uvw} in this case whereas $Y_{v,w}$ remains unchanged. (This last scenario is tight whenever both y and z belong to $N_{u,v} \cap N_{u,w}$.)

Next we will consider a random subset \mathcal{X} of all triangles in $Q(i_0)$. Initialize this set by $\mathcal{X}(i_0) = \emptyset$ and let it evolve until the culmination of the triangle removal process, denoting its final value by \mathcal{X}_* . To formulate its evolution rule we further maintain an auxiliary set \mathcal{Y} , also initialized by $\mathcal{Y}(i_0) = \emptyset$. At round $j > i_0$ do as follows:

- Process the triangles $Q(j)$ in an arbitrary order.
- If the current triangle $uvw \in Q(j)$ has $uvw \notin \mathcal{Y}$ and in addition this is the first round j after which some edge uv in this triangle has no other triangles incident to it ($Y_{u,v} = 1$) then:
 - Add uvw to \mathcal{X} .
 - Add all triangles in $\mathcal{N}_2(uvw)$ to \mathcal{Y} , where $\mathcal{N}_1(uvw) = \{xyz : xyz \sim uvw\}$ denotes all triangles incident to uvw and $\mathcal{N}_2(uvw) = \mathcal{N}_1(uvw) \cup \{x'y'z' : x'y'z' \sim xyz \in \mathcal{N}_1(uvw)\}$.

(Note that at no point do we ever delete triangles from \mathcal{X} or from \mathcal{Y} .) As explained above, removing some triangle xyz can affect triangles incident to at most two edges of uvw . Therefore, if j is the first round at the end of which uvw has some edge with co-degree 1, and uvw is intact at the end of this round, then at least one of its other edges has co-degree greater than 1 and so $B_{uvw} > 0$ upon the insertion of uvw to \mathcal{X} .

With the above explanation in mind, consider some time $i_1 > i_0$. At time i_0 , every triangle $uvw \in Q(i_0)$ had all of its co-degrees equal $n^{2\varepsilon+o(1)} > 1$. If $uvw \in Q(i_1)$ has $B_{uvw} = 0$ then either $uvw \in \mathcal{X}$ or necessarily there exists some $xyz \in \mathcal{X}$ such that at some point $i_0 < j \leq i_1$ we had $uvw \in \mathcal{N}_2(xyz)$. By the co-degrees hypothesis we know that already at time i_0 we had

$|\mathcal{N}_2(xyz)| \leq n^{4\varepsilon+o(1)}$ for all $xyz \in Q(i_0)$, and since co-degrees only decrease throughout our process, we can conclude that

$$|\mathcal{X}(i_1)| \geq |\{uvw \in Q(i_1) : B_{uvw} = 0\}| n^{-4\varepsilon-o(1)}. \quad (6.1)$$

Let i_1 be such that $p_1 = p(i_1) = \frac{1}{\sqrt{n \log n}}$. At this point $|E| = \frac{1}{2}(n^2 p_1 - n) = (\frac{1}{2} + o(1))n^{3/2}(\log n)^{-1}$ while $n^3 p_1^3 = O(n^{3/2}(\log n)^{-3})$. Obviously, if the process concludes before time i_1 we have at least $|E(i_1)| = n^{3/2-o(1)}$ edges in the final graph and there is nothing left to prove. Assume therefore that the process is still active at time i_1 . By Lemma 6.3 (noting that our assumption on the co-degrees satisfies the hypothesis $t(i_1) < \tau_c$ in that lemma for any sufficiently large n since $\varepsilon < \frac{1}{6}$) we get

$$Q(i_1) \leq \frac{1}{6}n^{3/2}(\log n)^{-3} + \frac{1}{6}n^{3/2}(\log n)^{-1} + n^{4/3} = (\frac{1}{3} + o(1))|E(i_1)|. \quad (6.2)$$

Clearly, if $|\{uv \in E(i_1) : Y_{u,v} = 0\}| \geq \delta|E(i_1)|$ for some arbitrarily small fixed $\delta > 0$, that is, a positive fraction of the edges have no triangles incident to them, then at least $\delta|E(i_1)| = n^{3/2-o(1)}$ edges will survive the removal process and we can conclude the proof. Assume therefore that there are at most $o(|E(i_1)|)$ such edges. In particular,

$$Q(i_1) = \frac{1}{3} \sum_{uv \in E(i_1)} Y_{u,v} \geq \frac{1}{3} |\{uv \in E(i_1) : Y_{u,v} \neq 0\}| \geq (\frac{1}{3} - o(1))|E(i_1)|.$$

Furthermore, together with (6.2), this implies that almost all triangles are edge disjoint. A simple way to see this is to order the $Q(i_1)$ triangles arbitrarily and then sequentially add their edges to the initial edge set $\{uv \in E(i_1) : Y_{u,v} = 0\}$, eventually arriving at the edge set of $G(i_1)$. If for some fixed $\delta > 0$ at least $\delta Q(i_1)$ triangles share any edges with the triangles preceding them in this procedure then the total number of edges would be at most $(3 - \delta)Q(i_1) + o(|E(i_1)|)$, which by (6.2) is at most $(1 - \frac{\delta}{3} + o(1))|E(i_1)|$, contradiction.

Examining the above statement we see that one can exclude $o(|E(i_1)|)$ triangles and be left with a set \mathcal{A} of $(\frac{1}{3} - o(1))|E(i_1)|$ edge disjoint triangles, i.e., $B_{uvw} = 0$ for every $uvw \in \mathcal{A}$. Re-adding the excluded triangles may now increase the value of B_{uvw} for at most $o(|E(i_1)|)$ distinct triangles $uvw \in \mathcal{A}$ (an excluded triangle impacts at most 3 triangles in \mathcal{A}). Altogether,

$$\#\{uvw \in Q(i_1) : B_{uvw} = 0\} = (\frac{1}{3} - o(1))|E(i_1)| = n^{3/2-o(1)},$$

which when added to (6.1) now implies that

$$|\mathcal{X}_*| \geq |\mathcal{X}(i_1)| \geq n^{3/2-4\varepsilon-o(1)}. \quad (6.3)$$

We now focus our attention on which edges among the triangles of \mathcal{X}_* belongs to the final outcome of the process. To this end, it will be useful to consider an equivalent formulation of the process starting from time i_0 . Instead of drawing a uniform triangle in every step, consider a uniform permutation σ over the triangles $Q(i_0)$, and the deterministic process that examines these triangles sequentially and each time removes the 3 edges of the triangle under consideration iff it is still intact (i.e., iff all of its edges still belong to the current graph). Clearly this process is merely a time-rescaled version of the usual triangle removal process, hence it suffices to analyze this one for establishing some properties of the final output.

For each $uvw \in \mathcal{X}$ define an event \mathcal{E}_{uvw} which will depend only on the internal ordering of the triangles $\mathcal{N}_2(uvw)$ in σ . The formulation of the event will depend on the constellation of triangles incident to uvw upon its addition to \mathcal{X}_* . Let j denote the round at the end of which uvw was added to \mathcal{X}_* . By definition, at the end of this round there was some edge of uvw that had no other triangles incident to it, that is, w.l.o.g. we had $Y_{u,v} = 1$. Choose some arbitrary $xyz \in Q(j)$ which at the end of round j satisfied $xyz \in \mathcal{N}_1(uvw)$, while recalling that by construction $B_{uvw}(j) > 0$ thus such a triangle necessarily exists. Define \mathcal{E}_{uvw} to be the event that σ arranges xyz prior to all triangles $\mathcal{N}_1(xyz)$ (including uvw).

First, recall that for any xyz as above $\mathcal{N}_1(xyz) \leq n^{2\varepsilon+o(1)}$ by the co-degree hypothesis, thus $\mathbb{P}(\mathcal{E}_{uvw}) \geq n^{-2\varepsilon-o(1)}$. Second, observe that on the event \mathcal{E}_{uvw} the triangle xyz under consideration will be examined for possible removal prior to examining any other $x'y'z' \sim xyz$ from $Q(j)$, thus the process will delete the edges $\{xy, xz, yz\}$ at that time. Only later will $uvw \in \mathcal{N}_1(xyz)$ be processed, and at that time no action will be taken due to its missing edge (the edge in common with xyz). In particular, the edge uv with $Y_{u,v}(j) = 1$, whose only incident triangle at time j was uvw , will belong to the final output.

Since the triangles in \mathcal{X}_* are edge disjoint, the final graph contains at least $\sum \mathbb{1}_{\mathcal{E}_{uvw}}$ edges. Moreover, the occurrence of \mathcal{E}_{uvw} is purely a function of the internal ordering according to σ between $xyz \in \mathcal{N}_1(uvw)$ and $\mathcal{N}_1(xyz) \subseteq \mathcal{N}_2(uvw)$. By construction, we exclude all triangles in $\mathcal{N}_2(uvw)$ from possible future inclusion in \mathcal{X}_* (by adding them to \mathcal{Y}), thus $\sum \mathbb{1}_{\mathcal{E}_{uvw}}$ stochastically dominates a binomial random variable with parameters $\text{Bin}(|\mathcal{X}_*|, n^{-2\varepsilon-o(1)})$. By (6.3) and standard estimates for the binomial distribution this is at least $n^{3/2-6\varepsilon-o(1)}$ w.h.p., as required. ■

ACKNOWLEDGMENTS

We thank Noga Alon, Jeff Kahn, Peter Keevash and Joel Spencer for fruitful discussions, as well as an anonymous referee for a thorough review.

REFERENCES

- [1] N. Alon, J.-H. Kim, and J. Spencer, *Nearly perfect matchings in regular simple hypergraphs*, Israel J. Math. **100** (1997), 171–187.
- [2] P. Bennett and T. Bohman, *A natural barrier in random greedy hypergraph matching*. Manuscript.
- [3] T. Bohman, *The triangle-free process*, Adv. Math. **221** (2009), no. 5, 1653–1677.
- [4] T. Bohman, A. Frieze, and E. Lubetzky, *A note on the random greedy triangle-packing algorithm*, J. Comb. **1** (2010), no. 3-4, 477–488.
- [5] T. Bohman and P. Keevash, *The early evolution of the H -free process*, Invent. Math. **181** (2010), no. 2, 291–336.
- [6] B. Bollobás, *The Life and Work of Paul Erdős*, Wolf Prize in mathematics. Vol. 1 (S. S. Chern and F. Hirzebruch, eds.), World Scientific Publishing Co. Inc., River Edge, NJ, 2000, pp. 292–315.
- [7] B. Bollobás, *To prove and conjecture: Paul Erdős and his mathematics*, Amer. Math. Monthly **105** (1998), no. 3, 209–237.
- [8] P. Erdős, S. Suen, and P. Winkler, *On the size of a random maximal graph*, Random Structures Algorithms **6** (1995), no. 2-3, 309–318.
- [9] D. A. Freedman, *On tail probabilities for martingales*, Ann. Probability **3** (1975), 100–118.
- [10] D. M. Gordon, G. Kuperberg, O. Patashnik, and J. H. Spencer, *Asymptotically optimal covering designs*, J. Combin. Theory Ser. A **75** (1996), no. 2, 270–280.
- [11] D. A. Grable, *On random greedy triangle packing*, Electron. J. Combin. **4** (1997), Research Paper 11, 19 pp.
- [12] V. Rödl, *On a packing and covering problem*, European J. Combin. **6** (1985), no. 1, 69–78.
- [13] V. Rödl and L. Thoma, *Asymptotic packing and the random greedy algorithm*, Random Structures Algorithms **8** (1996), no. 3, 161–177.
- [14] J. Spencer, *Asymptotic packing via a branching process*, Random Structures Algorithms **7** (1995), no. 2, 167–172.
- [15] N. C. Wormald, *The differential equation method for random graph processes and greedy algorithms*, Lectures on approximation and randomized algorithms (M. Karonski and H.J. Prömel, eds.), Advanced Topics in Mathematics, Polish Scientific Publishers PWN, Warsaw, 1999, pp. 73–155.

TOM BOHMAN

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA 15213, USA.

E-mail address: `tbohman@math.cmu.edu`

ALAN FRIEZE

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA 15213, USA.

E-mail address: `alan@random.math.cmu.edu`

E. LUBETZKY

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012, USA.

E-mail address: `eyal@courant.nyu.edu`