The Annals of Applied Probability
2016, Vol. 26, No. 1, 305–327
DOI: 10.1214/14-AAP1091

THE DENSEST SUBGRAPH PROBLEM IN SPARSE RANDOM GRAPHS[1]

BY VENKAT ANANTHARAM AND JUSTIN SALEZ

University of California, Berkeley and Université Paris Diderot (LPMA)

We determine the asymptotic behavior of the maximum subgraph density of large random graphs with a prescribed degree sequence. The result applies in particular to the Erdős–Rényi model, where it settles a conjecture of Hajek [*IEEE Trans. Inform. Theory* **36** (1990) 1398–1414]. Our proof consists in extending the notion of balanced loads from finite graphs to their local weak limits, using unimodularity. This is a new illustration of the objective method described by Aldous and Steele [In *Probability on Discrete Structures* (2004) 1–72 Springer].

1. Introduction.

Let $G = (V, E)$ be a finite, simple, undirected graph. Write $\vec{E}$ for the set of oriented edges, formed by replacing each edge $\{i, j\} \in E$ with the two oriented edges (i, j) and (j, i). An *allocation* on G is a map $\theta : \vec{E} \to [0, 1]$ satisfying $\theta(i, j) + \theta(j, i) = 1$ for every $\{i, j\} \in E$. The *load* induced by θ at a vertex $o \in V$ is

$$\partial\theta(o) := \sum_{i \sim o} \theta(i, o),$$

where $\sim$ denotes adjacency in G. θ is *balanced* if for every $(i, j) \in \vec{E}$,

$$\partial\theta(i) < \partial\theta(j) \quad \Longrightarrow \quad \theta(i, j) = 0. \tag{1.1}$$

Intuitively, one may think of each edge as carrying a unit amount of load, which has to be distributed over its end-points in such a way that the total load is as balanced as possible across the graph. In that respect, (1.1) is a *local optimality* criterion: modifying the allocation along a single edge cannot further reduce the load imbalance between its end-points. This condition happens to guarantee *global optimality* in a very strong sense. Specifically, the following conditions are equivalent (see [16]):

(i) θ is balanced.

Received January 2014; revised July 2014.

[1]Supported by ARO MURI Grant W911NF-08-1-0233, *Tools for the Analysis and Design of Complex Multi-Scale Networks*, NSF Grant CNS-0910702, and NSF Science & Technology Center Grant CCF-0939370, *Science of Information*.

MSC2010 subject classifications. Primary 60C05, 05C80; secondary 90B15.

Key words and phrases. Maximum subgraph density, load balancing, local weak convergence, objective method, unimodularity, pairing model.

(ii) θ minimizes $\sum_{o \in V} f(\partial\theta(o))$, for some strictly convex $f : [0, \infty) \to \mathbb{R}$.
(iii) θ minimizes $\sum_{o \in V} f(\partial\theta(o))$, for every convex $f : [0, \infty) \to \mathbb{R}$.

In particular, balanced allocations exist on G and they all induce the same loads $\partial\theta : V \to [0, \infty)$. The balanced load $\partial\theta(o)$ induced at a vertex $o \in V$ has a remarkable graph-theoretical interpretation: it measures the *local density* of G at o. Specifically, it was shown in [16] that the vertices receiving the highest load solve the classical densest subgraph problem on G: the value $\max \partial\theta$ coincides with the *maximum subgraph density* of G,

$$\varrho(G) := \max_{\varnothing \subsetneq H \subseteq V} \frac{|E[H]|}{|H|},$$

and the set $H = \operatorname{argmax} \partial\theta$ is precisely the largest set achieving this maximum. Here, $E[H] \subseteq E$ naturally denotes the set of edges with both end-points in H. This surprising connection with a well-known and important graph parameter justifies a deeper study of balanced loads in large graphs. It is convenient to encode the loads induced by a balanced allocation on G into a probability measure on $\mathbb{R}$, called the *empirical load distribution* of G:

$$\mathcal{L}_G = \frac{1}{|V|} \sum_{o \in V} \delta_{\partial\theta(o)}.$$

Motivated by the above connection, Hajek [16] studied the asymptotic behavior of $\mathcal{L}_G$ on the classical *Erdős–Rényi model*, where the graph $G = G_n$ is chosen uniformly at random among all graphs with $m = \lfloor \alpha n \rfloor$ edges on $V = \{1, \dots, n\}$. In the regime where the density parameter $\alpha \geq 0$ is kept fixed while $n \to \infty$, he conjectured that $\mathcal{L}_{G_n}$ should concentrate around a deterministic probability measure $\mathcal{L} \in \mathcal{P}(\mathbb{R})$ and that

$$\varrho(G_n) \xrightarrow[n \to \infty]{\mathbb{P}} \varrho := \sup\{t \in \mathbb{R} : \mathcal{L}([t, +\infty)) > 0\}.$$

Using a nonrigorous analogy with the case of finite trees, Hajek even proposed a description of $\mathcal{L}$ and ϱ in terms of the solutions to a distributional fixed-point equation which will be given later. In this paper, we establish this conjecture together with its analogue for various other sparse random graphs, using the unifying framework of *local weak convergence*.

2. Local weak convergence.

This section gives a brief account of the theory of local weak convergence. For more details, we refer to the seminal paper [6] and to the surveys [2, 3].

Rooted graphs. A *rooted graph* (G, o) is a graph $G = (V, E)$ together with a distinguished vertex $o \in V$, called the *root*. We let $\mathcal{G}_\star$ denote the set of all locally finite connected rooted graphs considered up to *rooted isomorphism*, that is, $(G, o) \equiv (G', o')$ if there exists a bijection $\gamma : V \to V'$ that preserves roots ($\gamma(o) = o'$) and adjacency ($\{i, j\} \in E \iff \{\gamma(i), \gamma(j)\} \in E'$). We write $[G, o]_h$ for the (finite) rooted subgraph induced by the vertices lying at graph-distance at most $h \in \mathbb{N}$ from o. The distance

$$\mathrm{DIST}\big((G, o), (G', o')\big) := \frac{1}{1+r} \qquad \text{where } r = \sup\big\{h \in \mathbb{N} : [G, o]_h \equiv \big[G', o'\big]_h\big\},$$

turns $\mathcal{G}_\star$ into a complete separable metric space; see [2].

Local weak limit. Let $\mathcal{P}(\mathcal{G}_\star)$ denote the set of Borel probability measures on $\mathcal{G}_\star$, equipped with the topology of weak convergence [7]. Given a finite graph $G = (V, E)$, let $\mathcal{U}(G)$ denote the law induced on $\mathcal{G}_\star$ by rooting G at a uniformly chosen vertex $o \in V$ and restricting G to the connected component of o. If $\{G_n\}_{n \geq 1}$ is a sequence of finite graphs such that $\{\mathcal{U}(G_n)\}_{n \geq 1}$ admits a limit $\mu \in \mathcal{P}(\mathcal{G}_\star)$, we call μ the *local weak limit* of $\{G_n\}_{n \geq 1}$ and write

$$G_n \xrightarrow[n \to \infty]{\text{LWC}} \mu.$$

Edge-rooted graphs. Let $\mathcal{G}_{\star\star}$ denote the set of locally finite connected graphs with a distinguished oriented edge, taken up to the natural isomorphism relation and equipped with the natural distance. With any function $f : \mathcal{G}_{\star\star} \to \mathbb{R}$ is naturally associated a function $\partial f : \mathcal{G}_\star \to \mathbb{R}$, defined by

$$\partial f(G, o) = \sum_{i \sim o} f(G, i, o).$$

Dually, with any $\mu \in \mathcal{P}(\mathcal{G}_\star)$ is naturally associated a nonnegative measure $\vec{\mu}$ on $\mathcal{G}_{\star\star}$, defined by the following relation: for any Borel $f : \mathcal{G}_{\star\star} \to [0, \infty)$,

$$\int_{\mathcal{G}_{\star\star}} f \, d\vec{\mu} = \int_{\mathcal{G}_\star} (\partial f) \, d\mu.$$

Note that $\vec{\mu}(\mathcal{G}_{\star\star}) = \deg(\mu)$, where $\deg(\mu)$ is the average degree of the root:

$$\deg(\mu) := \int_{\mathcal{G}_\star} \deg(G, o) \, d\mu(G, o).$$

Unimodularity. Given $f : \mathcal{G}_{\star\star} \to \mathbb{R}$, we define its *reversal* $f^* : \mathcal{G}_{\star\star} \to \mathbb{R}$ by

$$f^*(G, i, o) = f(G, o, i).$$

It was shown in [2] that any $\mu \in \mathcal{P}(\mathcal{G}_\star)$ arising as the local weak limit of some sequence of finite graphs satisfies the symmetry

$$\int_{\mathcal{G}_{\star\star}} f \, d\vec{\mu} = \int_{\mathcal{G}_{\star\star}} f^* \, d\vec{\mu}, \tag{2.1}$$

for any Borel $f : \mathcal{G}_{\star\star} \to [0, \infty)$. A measure $\mu \in \mathcal{P}(\mathcal{G}_\star)$ satisfying (2.1) is called *unimodular*, and the set of such measures is denoted by $\mathcal{U}$. The property (2.1) may be viewed as an infinite analogue of the trivial identity

$$\sum_{o \in V} \sum_{i \sim o} f(i, o) = \sum_{o \in V} \sum_{i \sim o} f(o, i),$$

valid for any finite graph $G = (V, E)$ and any $f : \vec{E} \to \mathbb{R}$.

Marks on oriented edges. It will sometimes be convenient to work with *networks*, that is, graphs equipped with a map from $\vec{E}$ to some fixed complete separable metric space Ξ. The above definitions extend naturally; see [2].

Unimodular Galton–Watson trees. Let $\pi = \{\pi_k\}_{k \geq 0}$ be a probability distribution on $\mathbb{N}$ with finite, nonzero mean. A *unimodular Galton–Watson tree* with degree distribution π is a random rooted tree obtained by a Galton–Watson branching process where the root has offspring distribution π and all descendants have the size-biased offspring distribution $\widehat{\pi} = \{\widehat{\pi}_k\}_{k \geq 0}$, where

$$\widehat{\pi}_k = \frac{(k+1)\pi_{k+1}}{\sum_i i \pi_i}. \tag{2.2}$$

The resulting law is unimodular and is denoted by $\mathrm{UGWT}(\pi)$. Such trees play a distinguished role in the local weak convergence theory, as they are the limits of many natural sequences of random graphs, including the Erdős–Rényi model and more generally, random graphs with prescribed degrees.

The pairing model. Given a sequence $\mathbf{d} = \{d(i)\}_{1 \leq i \leq n}$ of nonnegative integers whose sum is even, the pairing model [8, 18] generates a random graph $\mathbb{G}[\mathbf{d}]$ on $V = \{1, \dots, n\}$ as follows: $d(i)$ *half-edges* are attached to each $i \in V$, and the $2m = d(1) + \cdots + d(n)$ half-edges are paired uniformly at random to form m edges. Loops and multiple edges are then removed (a few variants exist—see [28]—but they are equivalent for our purpose). Now, consider a degree sequence $\mathbf{d}_n = \{d_n(i)\}_{1 \leq i \leq n}$ for each $n \geq 1$ and assume that

$$\forall k \in \mathbb{N}, \qquad \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\{d_n(i) = k\}} \xrightarrow[n \to \infty]{} \pi_k, \tag{2.3}$$

for some probability distribution $\pi = \{\pi_k\}_{k \geq 0}$ on $\mathbb{N}$ with finite, nonzero mean. Under the additional assumption that

$$\sup_{n \geq 1} \left\{ \frac{1}{n} \sum_{i=1}^{n} d_n^2(i) \right\} < \infty,$$

the local weak limit of $\{\mathbb{G}[\mathbf{d}_n]\}_{n \geq 1}$ is $\mu := \mathrm{UGWT}(\pi)$ almost surely; see [9].

3. Results. Our first main result is that the notion of balanced allocations can be extended from finite graphs to their local weak limits, in such a way that the induced loads behave continuously with respect to local weak convergence. Let us define a *Borel allocation* as a measurable function $\Theta : \mathcal{G}_{\star\star} \to [0, 1]$ such that $\Theta + \Theta^* = 1$, and call it *balanced* on a given $\mu \in \mathcal{U}$ if for $\vec{\mu}$-almost-every $(G, i, o) \in \mathcal{G}_{\star\star}$,

$$\partial\Theta(G, i) < \partial\Theta(G, o) \quad \Longrightarrow \quad \Theta(G, i, o) = 0.$$

This definition is the natural analogue of (1.1) when finite graphs are replaced by unimodular measures. We then have the following result.

THEOREM 1. *Let* $\mu \in \mathcal{U}$ *be such that* $\deg(\mu) < \infty$. *Then*,

Existence and optimality. *There is a Borel allocation* Θ_0 *that is balanced on* μ, *and for any Borel allocation* Θ *the following are equivalent*:

(i) Θ *is balanced on* μ.
(ii) Θ *minimizes* $\int f \circ \partial\Theta \, d\mu$ *for some strictly convex* $f : [0, \infty) \to \mathbb{R}$.
(iii) Θ *minimizes* $\int f \circ \partial\Theta \, d\mu$ *for every convex* $f : [0, \infty) \to \mathbb{R}$.
(iv) $\partial\Theta = \partial\Theta_0$, μ*-almost-everywhere*.

Continuity. *For any sequence* $\{G_n\}_{n \geq 1}$ *of finite graphs*,

$$\Bigl(G_n \xrightarrow[n\to\infty]{\text{LWC}} \mu\Bigr) \quad \Longrightarrow \quad \Bigl(\mathcal{L}_{G_n} \xrightarrow[n\to\infty]{\mathcal{P}(\mathbb{R})} \mathcal{L}_\mu\Bigr),$$

where $\mathcal{L}_\mu$ *is the law of the random variable* $\partial\Theta_0 \in L^1(\mathcal{G}_\star, \mu)$.

Variational characterization. *The stop-loss transform of* $\mathcal{L}_\mu$ *is given by*

$$\Psi_{\mathcal{L}_\mu}(t) = \max_{\substack{f : \mathcal{G}_\star \to [0,1] \\ \text{Borel}}} \left\{ \frac{1}{2} \int_{\mathcal{G}_{\star\star}} \widehat{f} \, d\vec{\mu} - t \int_{\mathcal{G}_\star} f \, d\mu \right\}, \qquad t \geq 0,$$

where $\widehat{f}(G, i, o) := f(G, o) \wedge f(G, i)$.

Recall that the *stop-loss transform* of a nonnegative integrable random variable X (in fact, of its law $\mathcal{L}$) is the function $\Psi_X = \Psi_{\mathcal{L}}$ defined by

$$\Psi_X(t) = \mathbb{E}\bigl[(X - t)^+\bigr] = \int_{\mathbb{R}} (x - t)^+ \, d\mathcal{L}(x), \qquad t \geq 0.$$

This function plays a central role in the theory of *convex ordering*, due to the classical equivalence between the following conditions (see, e.g., [26]):

(i) $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$ for every convex function $f : [0, \infty) \to \mathbb{R}$.
(ii) $\Psi_X \leq \Psi_Y$ and $\Psi_X(0) = \Psi_Y(0)$.

In particular, if $\Psi_X = \Psi_Y$ then X and Y have the same Laplace transforms. This shows that $\Psi_{\mathcal{L}}$ characterizes $\mathcal{L}$. Consequently, the above variational problem completely determines the limiting empirical load distribution.

Our second main result is an explicit resolution of this variational problem in the important case where $\mu = \mathrm{UGWT}(\pi)$, for an arbitrary degree distribution $\pi = \{\pi_k\}_{k\geq 0}$ on $\mathbb{N}$ with finite, nonzero mean. Throughout the paper, we let $[x]_0^1$ denote the closest point to $x \in \mathbb{R}$ in the interval $[0, 1]$, that is,

$$[x]_0^1 := \begin{cases} 0, & \text{if } x \leq 0, \\ x, & \text{if } x \in [0,1], \\ 1, & \text{if } x \geq 1. \end{cases}$$

Given $t \in \mathbb{R}$ and $Q \in \mathcal{P}([0,1])$, we let $F_{\pi,t}(Q) \in \mathcal{P}([0,1])$ denote the law of

$$[1 - t + \xi_1 + \cdots + \xi_{\widehat{D}}]_0^1,$$

where $\widehat{D}$ follows the size-biased distribution $\widehat{\pi}$ defined at (2.2), and where $\{\xi_k\}_{k\geq 1}$ are i.i.d. with law Q, independent of $\widehat{D}$. As conjectured by Hajek [16], the value of $\Psi_{\mathcal{L}_\mu}(t)$ turns out to be controlled by the solutions to the distributional fixed point equation $Q = F_{\pi,t}(Q)$. The latter can be solved numerically; see [16] for detailed tables in the case where π is Poisson.

THEOREM 2. *When $\mu = \mathrm{UGWT}(\pi)$, we have for every $t \in \mathbb{R}$:*

$$\Psi_{\mathcal{L}_\mu}(t) = \max_{Q = F_{\pi,t}(Q)} \left\{ \frac{\mathbb{E}[D]}{2} \mathbb{P}(\xi_1 + \xi_2 > 1) - t\mathbb{P}(\xi_1 + \cdots + \xi_D > t) \right\},$$

where $D \sim \pi$ and where $\{\xi_k\}_{k\geq 1}$ are i.i.d. with law Q, independent of D. The maximum is over all choices of $Q \in \mathcal{P}([0,1])$ subject to $Q = F_{\pi,t}(Q)$.

By analogy with the case of finite graphs, we define the maximum subgraph density of a measure $\mu \in \mathcal{U}$ with $\deg(\mu) < \infty$ as the essential supremum of the random variable $\partial\Theta_0$ constructed in Theorem 1. In other words,

$$\varrho(\mu) := \sup\{t \in \mathbb{R} : \Psi_{\mathcal{L}_\mu}(t) > 0\}.$$

In light of Theorem 1, it is natural to seek a continuity principle of the form

$$\left(G_n \xrightarrow[n\to\infty]{\mathrm{LWC}} \mu\right) \quad \Longrightarrow \quad \left(\varrho(G_n) \xrightarrow[n\to\infty]{} \varrho(\mu)\right). \tag{3.1}$$

However, a moment of thought shows that the graph parameter $\varrho(G)$ is too sensitive to be captured by local weak convergence. Indeed, if $|V(G_n)| \to \infty$, then adding a large but fixed clique to G_n will arbitrarily boost the value of $\varrho(G_n)$ without affecting the local weak limit of $\{G_n\}_{n\geq 1}$. Similarly, for random graphs with a prescribed asymptotic degree distribution $\pi \in \mathcal{P}(\mathbb{N})$, we expect (3.1) to fail when π has heavy tail, due to the presence of extremely dense subgraphs with negligible size. Understanding the maximum subgraph density in that regime remains an interesting open question. Our third main result establishes (3.1) in the light-tail regime.

THEOREM 3. *Consider a sequence* $\{\mathbf{d}_n\}_{n\geq 1}$ *of degree sequences that approach some distribution* $\pi = \{\pi_k\}_{k\geq 0}$ *in the sense of* (2.3). *Assume that* $\pi_0 + \pi_1 < 1$ *and that* π *has light tail*, *that is*, *for some* $\theta > 0$,

$$\sup_{n\geq 1}\left\{\frac{1}{n}\sum_{i=1}^{n} e^{\theta d_n(i)}\right\} < \infty. \tag{3.2}$$

Then $\varrho(\mathbb{G}[\mathbf{d}_n]) \xrightarrow[n\to\infty]{\mathbb{P}} \varrho(\mu)$, *where* $\mu = \mathrm{UGWT}(\pi)$.

In particular, the result applies to the Erdős–Rényi random graph $\mathbb{G}_n$ with n vertices and $m = \lfloor \alpha n \rfloor$ edges. Indeed, the conditional law of $\mathbb{G}_n$ given its (random) degree sequence $\mathbf{d_n}$ is precisely that of $\mathbb{G}[\mathbf{d_n}]$, and $\{\mathbf{d_n}\}_{n\geq 1}$ satisfies a.s. the conditions (2.3) and (3.2) with $\pi = \mathrm{Poisson}(2\alpha)$. Therefore, Theorems 1, 2 and 3 settle the conjectures of [17] and validate the numerical tables for ϱ given therein. We note that the quantity ϱ depends monotonically and continuously on the connectivity parameter α. More precisely, it is not hard to show that for any positive $\alpha < \beta$,

$$1 \leq \frac{\varrho(\beta)}{\varrho(\alpha)} \leq \frac{\beta}{\alpha}.$$

Also, since a graph G is *k-orientable* ($k \in \mathbb{N}$) if and only if $\varrho(G) < k$, Theorem 3 extends recent results on the k-orientability of the Erdős–Rényi random graph [11, 12]. See [13, 15, 20, 21] for various generalizations.

4. Proof ingredients and related work.

The objective method. This work is a new illustration of the general principles exposed in the *objective method* by Aldous and Steele [3]. The latter provides a powerful framework for the unified study of sparse random graphs and has already led to several remarkable results. Two prototypical examples are the celebrated $\zeta(2)$ limit in the random assignment problem due to Aldous [4], and the asymptotic enumeration of spanning trees in large graphs by Lyons [23]. Since then, the method has been successfully applied to various other combinatorial enumeration/optimization problems on graphs, including (but not limited to) [10, 14, 19–21, 24, 25, 27].

Lack of correlation decay. In the problem considered here, there is a major obstacle to a direct application of the objective method: the balanced load at a vertex is *not* determined by the local environment around that vertex. For example, every vertex of a d-regular graph with girth h has the same h-neighborhood as the root of a d-regular tree with height h. However, the induced load is $\frac{d}{2}$ in the first case and $1 - \frac{1}{(d-1)^{h-1}d}$ in the second. This long-range dependence gives rise to nonuniqueness issues when trying to extend the notion of balanced loads to infinite graphs. We refer to [17] for a detailed study of this phenomenon—therein called *load percolation*—as well as several fascinating questions.

Relaxation. To overcome the lack of correlation decay, we introduce a suitable relaxation of the balancing condition (1.1), which we call ε-balancing. Remarkably enough, any positive value of the perturbative parameter ε suffices to annihilate the long-range dependences described above. This allows us to define a unique ε-balanced Borel allocation $\Theta_\varepsilon : \mathcal{G}_{\star\star} \to [0, 1]$ and to establish the continuity of the induced load $\partial\Theta_\varepsilon : \mathcal{G}_\star \to [0, \infty)$ with respect to local convergence (Section 5). We then use unimodularity to prove that, as the perturbative parameter ε tends to 0, Θ_ε converges in a certain sense to a balanced Borel allocation Θ_0 (Section 6). This quickly leads to a proof of Theorem 1 (Section 7). In spirit, the role of the perturbative parameter $\varepsilon > 0$ is comparable to that of the temperature in [10], although no Gibbs–Boltzmann measure is involved in the present work.

Recursion on trees. As many other graph-theoretical problems, load balancing has a simple recursive structure when considered on trees. Indeed, once the value of the allocation along a given edge $\{i, j\}$ has been fixed, the problem naturally decomposes into two independent sub-problems, corresponding to the two disjoint subtrees formed by removing $\{i, j\}$. Note, however, that in the resulting sub-problems the loads of i and j must be shifted by a suitable amount to take into account the contribution of the removed edge. The precise effect of this shift on the loads induced at i and j defines what we call the *response functions* of the two subtrees (Section 8). Those response functions satisfy a recursion (Section 9). Recursions on trees automatically give rise to distributional fixed-point equations when specialized to Galton–Watson trees. Such equations are a common ingredient in the objective method; see [5]. This leads to the proof of Theorem 2 (Section 10).

Dense subgraphs in the pairing model. Finally, the proof of Theorem 3 (Section 11) relies on a property of random graphs with a prescribed degree sequence that might be of independent interest: under the exponential moment assumption (3.2), we show that dense subgraphs must be extensively large with high probability. Our argument is based on the first-moment method. See Proposition 11.1 for the precise statement, and [22], Lemma 6, for a result in the same direction.

5. ε-balancing. Throughout this section, $G = (V, E)$ is a locally finite graph and $\varepsilon > 0$ is a fixed parameter. An allocation θ on G is called *ε-balanced* if for every $(i, j) \in \vec{E}$,

$$\theta(i, j) = \left[\frac{1}{2} + \frac{\partial\theta(i) - \partial\theta(j)}{2\varepsilon}\right]_0^1. \tag{5.1}$$

This condition can be viewed as a relaxation of (1.1). Its interest lies in the fact that it fixes the nonuniqueness issue on infinite graphs.

PROPOSITION 5.1 (Existence, uniqueness and monotony). *If G has bounded degrees, then there is a unique ε-balanced allocation θ on G. If moreover $E' \subseteq E$, then the ε-balanced allocation θ' on $G' = (V, E')$ satisfies $\partial\theta' \leq \partial\theta$.*

PROOF. The set K of all allocations on G is clearly a compact convex subset of the locally convex space $\mathbb{R}^{\vec{E}}$ equipped with the topology of coordinate-wise convergence. Moreover, the mapping $K \ni \theta \mapsto \theta' \in K$ defined by

$$\theta'(i,j) = \left[\frac{1}{2} + \frac{\partial\theta(i) - \partial\theta(j)}{2\varepsilon}\right]_0^1$$

is continuous. It must therefore admit a fixed point, by the Schauder–Tychonoff fixed-point theorem (see, e.g., [1], Theorem 8.2). This proves existence. Now, consider $E' \subseteq E$ and let θ, θ' be ε-balanced allocations on G, G', respectively. Fix $o \in V$ and set

$$I := \{i \in V : \{i,o\} \in E', \theta'(i,o) > \theta(i,o)\}.$$

Clearly,

$$\partial\theta'(o) - \partial\theta(o) \le \sum_{i \in I} (\theta'(i,o) - \theta(i,o)).$$

On the other hand, since the map $x \mapsto [\frac{1}{2} + \frac{x}{2\varepsilon}]_0^1$ is nondecreasing and $\frac{1}{2\varepsilon}$-Lipschitz, our assumption on θ, θ' implies that for every $i \in I$,

$$\theta'(i,o) - \theta(i,o) \le \frac{1}{2\varepsilon}(\partial\theta'(i) - \partial\theta(i) - \partial\theta'(o) + \partial\theta(o)).$$

Injecting this into the above inequality and rearranging, we obtain

$$\begin{aligned} \partial\theta'(o) - \partial\theta(o) &\le \frac{1}{|I| + 2\varepsilon} \sum_{i \in I} (\partial\theta'(i) - \partial\theta(i)) \\ &\le \frac{\Delta}{\Delta + 2\varepsilon} \max_{i \in I} (\partial\theta'(i) - \partial\theta(i)), \end{aligned} \tag{5.2}$$

where Δ denotes the maximum degree in G. Now, observe that $\partial\theta, \partial\theta'$ are $[0, \Delta]$-valued, so that $M := \sup_V (\partial\theta' - \partial\theta)$ is finite. Property (5.2) forces $M \le 0$, which proves the monotony $E' \subseteq E \Longrightarrow \partial\theta' \le \partial\theta$. In particular, $E' = E$ implies $\partial\theta' = \partial\theta$, which in turns forces $\theta' = \theta$, thanks to (5.2). □

We now remove the bounded-degree assumption as follows. Fix $\Delta \in \mathbb{N}$, and consider the truncated graph $G^\Delta = (V, E^\Delta)$ obtained from G by isolating all nodes having degree more than Δ, that is,

$$E^\Delta = \{\{i,j\} \in E : \deg(G,i) \vee \deg(G,j) \le \Delta\}.$$

By construction, G^Δ has degree at most Δ, and we let $\Theta_\varepsilon^\Delta(G,i,j)$ denote the mass sent along $(i,j) \in \vec{E}$ in the unique ε-balanced allocation on G^Δ, with the understanding that $\Theta_\varepsilon^\Delta(G,i,j) = 0$ if $\{i,j\} \notin E^\Delta$. By uniqueness, this quantity depends only on the isomorphism class of the edge-rooted graph (G,i,j), so that we have a well-defined map $\Theta_\varepsilon^\Delta : \mathcal{G}_{\star\star} \to [0,1]$. By an immediate induction on $r \in \mathbb{N}$, the local contraction (5.2) yields

$$[G,o]_r \equiv [G',o']_r \quad \Longrightarrow \quad \left|\partial\Theta_\varepsilon^\Delta(G,o) - \partial\Theta_\varepsilon^\Delta(G',o')\right| \le \Delta\left(1+\frac{2\varepsilon}{\Delta}\right)^{-r}.$$

Since the map $x \mapsto [\frac{1}{2} + \frac{x}{2\varepsilon}]_0^1$ is $\frac{1}{2\varepsilon}$-Lipschitz, it follows that

$$[G,i,j]_r \equiv [G',i',j']_r \quad \Longrightarrow \quad \left|\Theta_\varepsilon^\Delta(G,i,j) - \Theta_\varepsilon^\Delta(G',i',j')\right| \le \frac{\Delta}{2\varepsilon}\left(1+\frac{2\varepsilon}{\Delta}\right)^{-r}.$$

Thus, the map $\Theta_\varepsilon^\Delta$ is equicontinuous. Now, the sequence of sets $\{E_\Delta\}_{\Delta\ge 1}$ increases to E, so the monotony in Proposition 5.1 guarantees that $\{\partial\Theta_\varepsilon^\Delta\}_{\Delta\ge 1}$ converges pointwise on $\mathcal{G}_\star$. Moreover, any given $\{i,j\} \in E$ belongs to E^Δ for large enough Δ, and the definition of ε-balancing yields

$$\Theta_\varepsilon^\Delta(G,i,j) = \left[\frac{1}{2} + \frac{\partial\Theta_\varepsilon^\Delta(G,i) - \partial\Theta_\varepsilon^\Delta(G,j)}{2\varepsilon}\right]_0^1.$$

Consequently, the pointwise limit $\Theta_\varepsilon := \lim_{\Delta\to\infty}\Theta_\varepsilon^\Delta$ exists in $[0,1]^{\mathcal{G}_{\star\star}}$. It clearly satisfies $\Theta_\varepsilon + \Theta_\varepsilon^* = 1$ and it is Borel as the pointwise limit of continuous maps. Thus, it is a Borel allocation. Letting $\Delta \to \infty$ above yields

$$\Theta_\varepsilon(G,i,j) = \left[\frac{1}{2} + \frac{\partial\Theta_\varepsilon(G,i) - \partial\Theta_\varepsilon(G,j)}{2\varepsilon}\right]_0^1. \tag{5.3}$$

6. The $\varepsilon \to 0$ limit. We now send the perturbative parameter ε to 0, and show that Θ_ε converges in a certain sense to a balanced Borel allocation Θ_0. Fix $\mu \in \mathcal{U}$ with $\deg(\mu) < \infty$. We write $\|f\|_p$ for the norm in both $L^p(\mu)$ and $L^p(\vec{\mu})$: which is meant should be clear from the context. Note that by unimodularity, we have for any Borel allocation Θ,

$$\|\Theta\|_1 = \int_{\mathcal{G}_{\star\star}} \Theta\, d\vec{\mu} = \int_{\mathcal{G}_{\star\star}} \frac{\Theta+\Theta^*}{2}\, d\vec{\mu} = \frac{\deg(\mu)}{2}. \tag{6.1}$$

PROPOSITION 6.1 (Existence of a balanced Borel allocation). *The limit $\Theta_0 := \lim_{\varepsilon\to 0}\Theta_\varepsilon$ exists in $L^2(\vec{\mu})$ and is a balanced Borel allocation on μ.*

PROOF. We will establish the following property: for $0 < \varepsilon \le \varepsilon'$,

$$\|\Theta_{\varepsilon'} - \Theta_\varepsilon\|_2^2 \le \|\Theta_\varepsilon\|_2^2 - \|\Theta_{\varepsilon'}\|_2^2. \tag{6.2}$$

In particular, $\|\Theta_\varepsilon\|_2^2 \ge \|\Theta_{\varepsilon'}\|_2^2$ so $\lim_{\varepsilon\to 0}\uparrow \|\Theta_\varepsilon\|_2^2$ exists. Consequently, the right-hand side tends to 0 as $\varepsilon, \varepsilon' \to 0$, hence so does the left-hand side. This provides a Cauchy criterion in $L^2(\vec{\mu})$ for $\{\Theta_\varepsilon\}_{\varepsilon>0}$, ensuring the existence of $\Theta_0 = \lim_{\varepsilon\to 0}\Theta_\varepsilon$. The rest of the claim follows, since Borel allocations are closed in $L^2(\vec{\mu})$ and letting $\varepsilon \to 0$ in (5.3) shows that Θ_0 is balanced on μ. In order to prove (6.2), let us first assume that

$$\mu\big(\{(G,o) : \deg(G,o) \le \Delta\}\big) = 1, \tag{6.3}$$

for some $\Delta \in \mathbb{N}$. This ensures that $f \in L^2(\vec{\mu})$, where

$$f(G,i,o) := \partial\Theta_\varepsilon(G,o) + \varepsilon\Theta_\varepsilon(G,i,o).$$

A straightforward manipulation of (5.3) shows that

$$f(G,i,o) > f(G,o,i) \quad \Longrightarrow \quad \Theta_\varepsilon(G,i,o) = 0.$$

This implies $\Theta_\varepsilon f + \Theta_\varepsilon^* f^* = f \wedge f^*$. On the other hand, $f \wedge f^* \le \Theta_{\varepsilon'} f + \Theta_{\varepsilon'}^* f^*$ since $\Theta_{\varepsilon'} + \Theta_{\varepsilon'}^* = 1$. Thus, $\Theta_\varepsilon f + \Theta_\varepsilon^* f^* \le \Theta_{\varepsilon'} f + \Theta_{\varepsilon'}^* f^*$. Integrating against $\vec{\mu}$ and invoking unimodularity, we get $\langle \Theta_\varepsilon - \Theta_{\varepsilon'}, f\rangle_{L^2(\vec{\mu})} \le 0$ or more explicitly,

$$\langle \partial\Theta_\varepsilon - \partial\Theta_{\varepsilon'}, \partial\Theta_\varepsilon\rangle_{L^2(\mu)} + \varepsilon\langle \Theta_\varepsilon - \Theta_{\varepsilon'}, \Theta_\varepsilon\rangle_{L^2(\vec{\mu})} \le 0.$$

But we have not yet used $\varepsilon \le \varepsilon'$, so we may exchange $\varepsilon, \varepsilon'$ to get

$$\langle \partial\Theta_{\varepsilon'} - \partial\Theta_\varepsilon, \partial\Theta_{\varepsilon'}\rangle_{L^2(\mu)} + \varepsilon'\langle \Theta_{\varepsilon'} - \Theta_\varepsilon, \Theta_{\varepsilon'}\rangle_{L^2(\vec{\mu})} \le 0.$$

Adding-up those inequalities and rearranging, we finally arrive at

$$(\varepsilon' - \varepsilon)\langle \Theta_\varepsilon - \Theta_{\varepsilon'}, \Theta_{\varepsilon'}\rangle_{L^2(\vec{\mu})} \ge \|\partial\Theta_\varepsilon - \partial\Theta_{\varepsilon'}\|_2^2 + \varepsilon\|\Theta_\varepsilon - \Theta_{\varepsilon'}\|_2^2.$$

In particular, $\langle \Theta_\varepsilon, \Theta_{\varepsilon'}\rangle_{L^2(\vec{\mu})} \ge \|\Theta_{\varepsilon'}\|_2^2$ and (6.2) follows since

$$\|\Theta_{\varepsilon'} - \Theta_{\varepsilon'}\|_2^2 = \|\Theta_{\varepsilon'}\|_2^2 + \|\Theta_\varepsilon\|_2^2 - 2\langle \Theta_\varepsilon, \Theta_{\varepsilon'}\rangle_{L^2(\vec{\mu})}.$$

Finally, if our extra assumption (6.3) is dropped, we may apply (6.2) with Θ_ε, $\Theta_{\varepsilon'}$ replaced by $\Theta_\varepsilon^\Delta$, $\Theta_{\varepsilon'}^\Delta$ and let then $\Delta \to \infty$. By construction, $\Theta_\varepsilon^\Delta \to \Theta_\varepsilon$ and $\Theta_{\varepsilon'}^\Delta \to \Theta_{\varepsilon'}$ pointwise, and (6.2) follows by dominated convergence. $\square$

PROPOSITION 6.2 (Continuity of balanced loads). *Let $\{G_n\}_{n\ge1}$ be a sequence of finite graphs with local weak limit μ. Then*

$$\mathcal{L}_{G_n} \xrightarrow[n\to\infty]{\mathcal{P}(\mathbb{R})} \mathcal{L},$$

where $\mathcal{L} = \mathcal{L}_\mu$ is the law of the random variable $\partial\Theta_0 \in L^1(\mu)$.

PROOF. For $n \ge 1$ we let $\widehat{G_n}$ denote the network obtained by encoding a balanced allocation θ_n as $[0,1]$-valued marks on the oriented edges of G_n. The sequence $\{\mathcal{U}(\widehat{G_n})\}_{n\ge1}$ is tight, because $\{\mathcal{U}(G_n)\}_{n\ge1}$ converges weakly and the marks are $[0,1]$-valued. Consider any subsequential weak limit $(\mathbb{G}, o, \theta)$. By construction, $(\mathbb{G}, o)$ has law μ and θ is a.s. a balanced allocation on $\mathbb{G}$. Our goal is to establish that $\partial\theta(o) = \partial\Theta_0(\mathbb{G}, o)$ a.s. Set $\theta'(i,j) := \Theta_0(\mathbb{G}, i, j)$. The random rooted network $(\mathbb{G}, o, \theta, \theta')$ is unimodular, since $(\mathbb{G}, o, \theta)$ is a weak limit of finite networks and Θ_0 is Borel. Now,

$$\begin{aligned}\mathbb{E}\big[(\partial\theta(o) - \partial\theta'(o))^+\big] &= \mathbb{E}\Big[\sum_{i\sim o}(\theta(i,o) - \theta'(i,o))\mathbf{1}_{\partial\theta(o)>\partial\theta'(o)}\Big]\\ &= \mathbb{E}\Big[\sum_{i\sim o}(\theta(o,i) - \theta'(o,i))\mathbf{1}_{\partial\theta(i)>\partial\theta'(i)}\Big]\\ &= \mathbb{E}\Big[\sum_{i\sim o}(\theta'(i,o) - \theta(i,o))\mathbf{1}_{\partial\theta(i)>\partial\theta'(i)}\Big],\end{aligned}$$

where the second equality follows from unimodularity and the third one from the identities $\theta(o,i) = 1 - \theta(i,o)$ and $\theta'(o,i) = 1 - \theta'(i,o)$. Combining the first and last lines, we see that $\mathbb{E}[(\partial\theta(o) - \partial\theta'(o))^+]$ equals

$$\frac{1}{2}\mathbb{E}\Big[\sum_{i\sim o}(\theta(i,o) - \theta'(i,o))(\mathbf{1}_{\partial\theta(o)>\partial\theta'(o)} - \mathbf{1}_{\partial\theta(i)>\partial\theta'(i)})\Big].$$

The fact that θ, θ' are balanced across $\{i, o\}$ easily implies that $\theta(i,o) - \theta'(i,o)$ and $\mathbf{1}_{\partial\theta(o)>\partial\theta'(o)} - \mathbf{1}_{\partial\theta(i)>\partial\theta'(i)}$ can neither be simultaneously positive, nor simultaneously negative. Therefore, $\mathbb{E}[(\partial\theta(o) - \partial\theta'(o))^+] \le 0$. Exchanging the roles of θ, θ' yields $\partial\theta(o) = \partial\theta'(o)$ a.s., as desired. $\square$

7. Proof of Theorem 1.

We now complete the proof of Theorem 1.

PROPOSITION 7.1. *Let Θ be a Borel allocation. Then for all $t \in \mathbb{R}$,*

$$\int_{\mathcal{G}_\star}(\partial\Theta - t)^+\,d\mu \ge \sup_{\substack{f:\mathcal{G}_\star\to[0,1]\\ \text{Borel}}}\Big\{\frac{1}{2}\int_{\mathcal{G}_{\star\star}}\widehat{f}\,d\vec{\mu} - t\int_{\mathcal{G}_\star} f\,d\mu\Big\},$$

with equality for all $t \in \mathbb{R}$ if and only if Θ is balanced on μ.

PROOF. Fix a Borel $f:\mathcal{G}_\star \to [0,1]$. Since $(\partial\Theta - t)^+ \ge (\partial\Theta - t)f$, we have

$$\int_{\mathcal{G}_\star}(\partial\Theta - t)^+\,d\mu \ge \int_{\mathcal{G}_\star} f\,\partial\Theta\,d\mu - t\int_{\mathcal{G}_\star} f\,d\mu. \tag{7.1}$$

Using the unimodularity of μ and the identity $\Theta + \Theta^* = 1$, we also have

$$\begin{aligned}\int_{\mathcal{G}_\star} f\,\partial\Theta\,d\mu &= \frac{1}{2}\int_{\mathcal{G}_{\star\star}}\big(f(G,o)\Theta(G,i,o) + f(G,i)\Theta(G,o,i)\big)\,d\vec{\mu}(G,i,o)\\ &\ge \frac{1}{2}\int_{\mathcal{G}_{\star\star}}\big(f(G,o)\wedge f(G,i)\big)\,d\vec{\mu}(G,i,o).\end{aligned} \tag{7.2}$$

Combining (7.1) and (7.2) yields the inequality. Let us examine the equality case. First, equality holds in (7.1) if and only if for μ-a.e. $(G,o) \in \mathcal{G}_\star$,

$$\partial\Theta(G,o) > t \quad \Longrightarrow \quad f(G,o) = 1,$$

$$\partial\Theta(G,o) < t \quad \Longrightarrow \quad f(G,o) = 0.$$

Second, equality holds in (7.2) if and only if for $\vec{\mu}$-a.e. $(G,i,o) \in \mathcal{G}_{\star\star}$,

$$f(G,i) < f(G,o) \quad \Longrightarrow \quad \Theta(G,i,o) = 0.$$

If Θ is balanced on μ, then the choice $f = \mathbf{1}_{\{\partial\Theta > t\}}$ clearly satisfies all those requirements, so that equality holds for each $t \in \mathbb{R}$ in the proposition. This proves the *if* part and shows that the supremum in Proposition 7.1 is attained, because at least one balanced allocation exists by Proposition 6.1. Now, for the *only if* part, suppose that equality is achieved in Proposition 7.1. Then the above requirements imply that for $\vec{\mu}$-a.e. $(G,i,o) \in \mathcal{G}_{\star\star}$,

$$\partial\Theta(G,i) < t < \partial\Theta(G,o) \quad \Longrightarrow \quad \Theta(G,i,o) = 0.$$

Since this must be true for all $t \in \mathbb{Q}$, it follows that Θ is balanced on μ. $\square$

PROOF OF THEOREM 1. Existence, continuity and the variational characterization were established in Propositions 6.1, 6.2 and 7.1, respectively. Now, let Θ, Θ' be Borel allocations, and assume that Θ is balanced. Applying Proposition 7.1 to Θ and Θ' shows that for all $t \in \mathbb{R}$,

$$\int_{\mathcal{G}_\star} (\partial\Theta - t)^+ \, d\mu \le \int_{\mathcal{G}_\star} (\partial\Theta' - t)^+ \, d\mu.$$

Moreover, (6.1) guarantees that $\partial\Theta, \partial\Theta'$ have the same mean. As already mentioned below the statement of Theorem 1, those two conditions imply

$$\int_{\mathcal{G}_\star} (f \circ \partial\Theta) \, d\mu \le \int_{\mathcal{G}_\star} (f \circ \partial\Theta') \, d\mu,$$

for any convex function $f : [0,\infty) \to \mathbb{R}$. We have just proved (i) $\Longrightarrow$ (iii). On the other hand, (iii) $\Longrightarrow$ (ii) is obvious. In particular, Θ_0 satisfies (ii) and (iii). The *only if* part of Proposition 7.1 shows that (iii) $\Longrightarrow$ (i). The implication (iv) $\Longrightarrow$ (iii) is obvious given that Θ_0 satisfies (iii). Thus, it only remains to prove (ii) $\Longrightarrow$ (iv). Assume that Θ minimizes $\int (f \circ \partial\Theta) \, d\mu$ for some strictly convex function $f : [0,\infty) \to \mathbb{R}$, and let m denote the value of this minimum. Since Θ_0 satisfies (ii), we also have $\int (f \circ \partial\Theta_0) \, d\mu = m$. But then $\Theta' := (\Theta_0 + \Theta)/2$ is an allocation and by convexity,

$$\int_{\mathcal{G}_\star} (f \circ \partial\Theta') \, d\mu \le \int_{\mathcal{G}_\star} \frac{(f \circ \partial\Theta) + (f \circ \partial\Theta_0)}{2} \, d\mu = m.$$

This inequality contradicts the definition of m, unless it is an equality. This forces $\partial\Theta = \partial\Theta_0$ μ-a.e., since f is strictly convex. $\square$

8. Response functions. As many other graph-theoretical problems, load balancing has a simple recursive structure when specialized to trees. However, the exact formulation of this recursion requires the possibility to *condition* the allocation to take a certain value at a given edge, and we first need to give a proper meaning to this operation. Let $G=(V,E)$ be a locally finite graph and $b:V\to\mathbb{R}$ a function called the *baseload*. An allocation θ is *balanced with respect to b* if

$$b(i)+\partial\theta(i)<b(j)+\partial\theta(j)\quad\Longrightarrow\quad\theta(i,j)=0,$$

for all $(i,j)\in\vec{E}$. This is precisely the definition of balancing, except that the load *felt* by each vertex $i\in V$ is shifted by a certain amount $b(i)$. Similarly, θ is *ε-balanced with respect to b* if for all $(i,j)\in\vec{E}$,

$$\theta(i,j)=\left[\frac{1}{2}+\frac{b(i)+\partial\theta(i)-b(j)-\partial\theta(j)}{2\varepsilon}\right]_0^1. \tag{8.1}$$

The arguments used in Proposition 5.1 are easily extended to this situation.

PROPOSITION 8.1 (Existence, uniqueness and monotony). *If G has bounded degree and if b is bounded, then there is a unique ε-balanced allocation with baseload b. Moreover, if $b'\le b$ is bounded and if $E'\subseteq E$, then the ε-balanced allocation θ' on $G'=(V,E')$ with baseload b' satisfies $b'+\partial\theta'\le b+\partial\theta$.*

As in Section 5, we then define an ε-balanced allocation in the general case by considering the truncated graph G^Δ with baseload the truncation of b to $[-\Delta,\Delta]$, and let then $\Delta\to\infty$. Monotony guarantees the existence of a limiting ε-balanced allocation. We shall need the following property.

PROPOSITION 8.2 (Nonexpansion). *Let θ,θ' be the ε-balanced allocations with baseloads $b,b':V\to\mathbb{R}$. Set $f=\partial\theta+b$ and $f'=\partial\theta'+b'$. Then*

$$\|f'-f\|_{\ell^1(V)}\le\|b'-b\|_{\ell^1(V)}.$$

PROOF. By considering $b''=b\wedge b'$ and using the triangle inequality, we may assume that $b\le b'$. Note that this implies $f\le f'$, thanks to Proposition 8.1. When G is finite, the claim trivially follows from conservation of mass:

$$\sum_{o\in V}\bigl(f'(o)-f(o)\bigr)=\sum_{o\in V}\bigl(b'(o)-b(o)\bigr).$$

This then extends to the case where G has bounded degrees with b,b' bounded as follows: choose finite subsets $V_1\subseteq V_2\subseteq\cdots$ such that $\bigcup_{n\ge1}V_n=V$. For each $n\ge1$, let θ_n,θ_n' denote the ε-balanced allocations on the subgraph induced by V_n, with baseloads the restrictions of b,b' to V_n. Then $\theta_n\to\theta$ and $\theta_n'\to\theta'$ pointwise, by compactness and uniqueness. Now, any finite $K\subseteq V$ is contained in V_n for large enough n, and since V_n is finite we know that $f_n := \partial\theta_n + b$ and $f'_n := \partial\theta'_n + b'$ satisfy

$$\sum_{i \in K} |f'_n(i) - f_n(i)| \leq \sum_{i \in V_n} |b'(i) - b(i)|.$$

Letting $n \to \infty$ yields the desired result, since K is arbitrary. Finally, for the general case, we may apply the result to the truncated graph G^Δ with baseloads the truncation of b, b' to $[-\Delta, \Delta]$, and let then $\Delta \to \infty$. □

Although the uniqueness in Proposition 8.1 does not extend to the $\varepsilon = 0$ case, the following weaker result will be useful in the next section.

PROPOSITION 8.3 (Weak uniqueness). *Assume that θ, θ' are balanced with respect to b and that $\|\partial\theta - \partial\theta'\|_{\ell^1(V)} < \infty$. Then, $\partial\theta = \partial\theta'$.*

PROOF. Fix $\delta > 0$. Since $\|\partial\theta - \partial\theta'\|_{\ell^1(V)} < \infty$, the level set $S := \{j \in V : \partial\theta'(j) - \partial\theta(j) > \delta\}$ must be finite. Therefore, it satisfies the conservation of mass:

$$\sum_{j \in S} \partial\theta'(j) - \partial\theta(j) = \sum_{(i,j) \in E(V-S,S)} \theta'(i,j) - \theta(i,j). \tag{8.2}$$

Now, if $(i, j) \in E(V - S, S)$ then clearly, $\partial\theta'(i) - \partial\theta(i) < \partial\theta'(j) - \partial\theta(j)$. Consequently, at least one of the following inequalities must hold:

$$b(j) - b(i) < \partial\theta(i) - \partial\theta(j) \quad \text{or} \quad b(j) - b(i) > \partial\theta'(i) - \partial\theta'(j).$$

The first one implies $\theta(i, j) = 1$ and the second $\theta'(i, j) = 0$, since θ, θ' are balanced with respect to b. In either case, we have $\theta'(i, j) \leq \theta(i, j)$. Thus, the right-hand side of (8.2) is nonpositive, hence so must the left-hand side be. This contradicts the definition of S unless $S = \varnothing$, that is, $\partial\theta' \leq \partial\theta + \delta$. Since δ is arbitrary, we conclude that $\partial\theta' \leq \partial\theta$. Equality follows by symmetry. □

Given $o \in V$ and $x \in \mathbb{R}$, we set $\mathfrak{f}^\varepsilon_{(G,o)}(x) = x + \partial\theta(o)$ where θ is the ε-balanced allocation with baseload x at o and 0 elsewhere. We call $\mathfrak{f}^\varepsilon_{(G,o)} : \mathbb{R} \to \mathbb{R}$ the *response function* of the rooted graph (G, o). Propositions 8.1 and 8.2 guarantee that $\mathfrak{f}^\varepsilon_{(G,o)}$ is nondecreasing and nonexpansive, that is,

$$x \leq y \quad \Longrightarrow \quad 0 \leq \mathfrak{f}^\varepsilon_{(G,o)}(y) - \mathfrak{f}^\varepsilon_{(G,o)}(x) \leq y - x. \tag{8.3}$$

Note for future use that the definition of $\mathfrak{f}^\varepsilon_{(G,o)}(x)$ also implies

$$0 \leq \mathfrak{f}^\varepsilon_{(G,o)}(x) - x \leq \deg(G, o). \tag{8.4}$$

When G is a tree, response functions turn out to satisfy a simple recursion.

9. Recursion on trees. We are now ready to state the promised recursion. Fix a tree $T = (V, E)$. Deleting $\{i, j\} \in E$ creates two disjoint subtrees, viewed as rooted at i and j and denoted $T_{i \to j}$ and $T_{j \to i}$, respectively.

PROPOSITION 9.1. *The response function $\mathfrak{f}^{\varepsilon}_{(T,o)}$ is invertible and*

$$\{\mathfrak{f}^{\varepsilon}_{(T,o)}\}^{-1} = \mathrm{Id} - \sum_{i \sim o} [1 - \{\mathfrak{f}^{\varepsilon}_{T_{i \to o}} + \varepsilon(2\mathrm{Id} - 1)\}^{-1}]_0^1, \tag{9.1}$$

where Id *denotes the identity function on* $\mathbb{R}$.

PROOF. $\mathfrak{f}^{\varepsilon}_{T_{i \to o}} + \varepsilon(2\mathrm{Id} - 1)$ increases continuously from $\mathbb{R}$ onto $\mathbb{R}$, so $\{\mathfrak{f}^{\varepsilon}_{T_{i \to o}} + \varepsilon(2\mathrm{Id} - 1)\}^{-1}$ exists and increases continuously from $\mathbb{R}$ onto $\mathbb{R}$. Consequently, the function $g : \mathbb{R} \to \mathbb{R}$ appearing in the right-hand side of (9.1) is continuously increasing from $\mathbb{R}$ onto $\mathbb{R}$, hence invertible. Given $x \in \mathbb{R}$, it now remains to prove that $t := \mathfrak{f}^{\varepsilon}_{(T,o)}(x)$ satisfies $g(t) = x$. By definition,

$$t = x + \partial\theta(o), \tag{9.2}$$

where θ denotes the ε-balanced allocation on T with baseload x at o and 0 elsewhere. Now fix $i \sim o$. The restriction of θ to $T_{i \to o}$ is clearly an ε-balanced allocation on $T_{i \to o}$ with baseload $\theta(o, i)$ at i and 0 elsewhere. This is precisely the allocation appearing in the definition of $\mathfrak{f}^{\varepsilon}_{T_{i \to o}}(\theta(o, i))$, hence

$$\mathfrak{f}^{\varepsilon}_{T_{i \to o}}(\theta(o, i)) = \partial\theta(i).$$

Thus, the fact that θ is ε-balanced along (o, i) may now be rewritten as

$$\theta(o, i) = \left[\frac{1}{2} + \frac{t - \mathfrak{f}^{\varepsilon}_{T_{i \to o}}(\theta(o, i))}{2\varepsilon}\right]_0^1. \tag{9.3}$$

But by definition, $x_i := \{\mathfrak{f}^{\varepsilon}_{T_{i \to o}} + \varepsilon(2\mathrm{Id} - 1)\}^{-1}(t)$ is the unique solution to

$$x_i = \frac{1}{2} + \frac{t - \mathfrak{f}^{\varepsilon}_{T_{i \to o}}(x_i)}{2\varepsilon}. \tag{9.4}$$

Comparing (9.3) and (9.4), we see that $\theta(o, i) = [x_i]_0^1$, that is, $\theta(i, o) = [1 - x_i]_0^1$. Re-injecting this into (9.2), we arrive exactly at the desired $x = g(t)$. $\square$

In the remainder of this section, we fix a vanishing sequence $\{\varepsilon_n\}_{n \ge 1}$ and study the pointwise limit $\mathfrak{f} = \lim_{n \to \infty} \mathfrak{f}^{\varepsilon_n}_{(T,o)}$, when it exists. Note that $\mathfrak{f}$ needs not be invertible. However, (8.3) and (8.4) guarantee that $\mathfrak{f}$ is nondecreasing with $\mathfrak{f}(\pm\infty) = \pm\infty$, so that is admits a well-defined right-continuous inverse

$$\mathfrak{f}^{-1}(t) := \sup\{x \in \mathbb{R} : \mathfrak{f}(x) \le t\}, \qquad t \in \mathbb{R}.$$

PROPOSITION 9.2. *Assume that* $\ell_o := \lim_{n\to\infty} \partial\Theta_{\varepsilon_n}(T,o)$ *exists for each* $o \in V$. *Then* $\mathfrak{f}_{T_{i\to j}} := \lim_{n\to\infty} \mathfrak{f}^{\varepsilon_n}_{T_{i\to j}}$ *exists pointwise for each* $(i,j) \in \vec{E}$, *and*

$$\mathfrak{f}^{-1}_{T_{i\to j}}(t) = t - \sum_{k\sim i, k\neq j} \left[1 - \mathfrak{f}^{-1}_{T_{k\to i}}(t)\right]_0^1, \tag{9.5}$$

for every $t \in \mathbb{R}$. *Moreover, for every* $o \in V$,

$$\ell_o > t \quad \Longleftrightarrow \quad \sum_{i\sim o} \left[1 - \mathfrak{f}^{-1}_{T_{i\to o}}(t)\right]_0^1 > t. \tag{9.6}$$

PROOF. Fix $(i,j) \in \vec{E}$, $x \in \mathbb{R}$, and let us show that $\{\mathfrak{f}^{\varepsilon_n}_{T_{i\to j}}(x)\}_{n\geq 1}$ converges. By definition, $\mathfrak{f}^{\varepsilon}_{T_{i\to j}}(x) = x + \partial\theta_\varepsilon(i)$, where θ_ε is the ε-balanced allocation on $T_{i\to j}$ with baseload x at i and 0 elsewhere. Since the set of allocations on $T_{i\to j}$ is compact, it is enough to consider two subsequential limits θ, θ' of $\{\theta_{\varepsilon_n}\}_{n\geq 1}$ and prove that $\partial\theta = \partial\theta'$. Passing to the limit in (8.1), we know that θ, θ' are balanced with respect to the above baseload. Writing $V_{i\to j}$ for the vertex set of $T_{i\to j}$, Proposition 8.3 reduces our task to proving

$$\|\partial\theta - \partial\theta'\|_{\ell^1(V_{i\to j})} < \infty. \tag{9.7}$$

Let $\theta^\star_\varepsilon$ be the restriction of Θ_ε to $T_{i\to j}$. Thus, $\theta^\star_\varepsilon$ is an allocation on $T_{i\to j}$ and it is ε-balanced with baseload $\theta^\star_\varepsilon(j,i)$ at i and 0 elsewhere. Consequently, Proposition 8.2 guarantees that for any finite $K \subseteq V_{i\to j} \setminus \{i\}$,

$$\|\partial\theta_\varepsilon - \partial\theta^\star_\varepsilon\|_{\ell^1(K)} \leq |x| + 1.$$

Applying this to $\varepsilon, \varepsilon' > 0$ and using the triangle inequality, we obtain

$$\|\partial\theta_\varepsilon - \partial\theta_{\varepsilon'}\|_{\ell^1(K)} \leq 2|x| + 2 + \|\partial\theta^\star_\varepsilon - \partial\theta^\star_{\varepsilon'}\|_{\ell^1(K)}.$$

Since $\{\partial\theta^\star_{\varepsilon_n}\}_{n\geq 1}$ converges by assumption, we may pass to the limit to obtain $\|\partial\theta - \partial\theta'\|_{\ell^1(K)} \leq 2|x| + 2$. But K is arbitrary, so (9.7) follows. This shows that $\mathfrak{f}_{T_{i\to j}} := \lim_{n\to\infty} \mathfrak{f}^{\varepsilon_n}_{T_{i\to j}}$ exists pointwise. We now recall two classical facts about nondecreasing functions $\mathfrak{f}: \mathbb{R} \to \mathbb{R}$ with $\mathfrak{f}(\pm\infty) = \pm\infty$. First, $\mathfrak{f}^{-1}$ is nondecreasing, so that its discontinuity set $\mathcal{D}(f^{-1})$ is countable. Second, the pointwise convergence $\mathfrak{f}_n \to \mathfrak{f}$ implies $\mathfrak{f}_n^{-1}(t) \to \mathfrak{f}^{-1}(t)$ for every $t \in \mathbb{R} \setminus \mathcal{D}(\mathfrak{f}^{-1})$. Consequently, letting $\varepsilon \to 0$ in (9.1) proves (9.5) for $t \notin \mathcal{D} := \mathcal{D}(\mathfrak{f}^{-1}_{T_{i\to j}}) \cup \bigcup_{k\sim i} \mathcal{D}(\mathfrak{f}^{-1}_{T_{k\to i}})$. The equality then extends to $\mathbb{R}$ since $\mathcal{D}$ is countable and both sides of (9.5) are right-continuous in t. Replacing $T_{i\to j}$ with (T,o) in the above argument shows that $\mathfrak{f}_{(T,o)} := \lim_{n\to\infty} \mathfrak{f}^{\varepsilon_n}_{(T,o)}$ exists and satisfies

$$\mathfrak{f}^{-1}_{(T,o)}(t) = t - \sum_{i\sim o} \left[1 - \mathfrak{f}^{-1}_{T_{i\to o}}(t)\right]_0^1, \qquad t \in \mathbb{R}.$$

Finally, recall that $\mathfrak{f}^{\varepsilon_n}_{(T,o)}(0) = \partial\Theta_{\varepsilon_n}(T,o)$ for all $n \geq 1$, so that $\mathfrak{f}_{(T,o)}(0) = \ell_o$. But $\mathfrak{f}_{(T,o)}(0) > t \Longleftrightarrow \mathfrak{f}^{-1}_{(T,o)}(t) < 0$ by definition of $\mathfrak{f}^{-1}_{(T,o)}$, so (9.6) follows. $\square$

10. Proof of Theorem 2. In all this section, $t \in \mathbb{R}$ is fixed. We manipulate networks rather than graphs, where each $(i, j) \in \vec{E}$ is equipped with a mark $\xi(i, j) \in [0, 1]$. The marks are assumed to satisfy the local recursion

$$\xi(i, j) = \left[1 - t + \sum_{k \sim i, k \neq j} \xi(k, i)\right]_0^1, \qquad (i, j) \in \vec{E}. \tag{10.1}$$

We start with a simple lemma.

LEMMA 10.1. $\partial\xi(i) \wedge \partial\xi(j) > t \iff \xi(i, j) + \xi(j, i) > 1$.

PROOF. We check the equivalence separately in each case. By assumption,

$$\xi(i, j) = [1 - t + \partial\xi(i) - \xi(j, i)]_0^1, \tag{10.2}$$

$$\xi(j, i) = [1 - t + \partial\xi(j) - \xi(i, j)]_0^1. \tag{10.3}$$

- If $0 < \xi(i, j), \xi(j, i) < 1$, then the equivalence trivially holds since we may safely remove the truncation $[\cdot]_0^1$ from (10.2)–(10.3) to obtain
$$\partial\xi(i) - t = \xi(i, j) + \xi(j, i) - 1 = \partial\xi(j) - t.$$
- If $\xi(j, i) = 0$, then we have $1 - t + \partial\xi(j) - \xi(i, j) \leq 0$ thanks to (10.3), and hence $\partial\xi(j) \leq t$. Thus, both sides of the equivalence are false.
- If $\xi(i, j) = 1, \xi(j, i) > 0$, then using $\xi(i, j) = 1$ in (10.2) gives $\partial\xi(i) - t \geq \xi(j, i)$ and since $\xi(j, i) > 0$ we obtain $\partial\xi(i) > t$. Similarly, using $\xi(j, i) > 0$ in (10.3) gives $\partial\xi(j) > t + \xi(i, j) - 1$ and since $\xi(i, j) = 1$ we obtain $\partial\xi(j) > t$. Thus, both sides of the equivalence are true.

The other possible cases follow by exchanging $\xi(i, j)$ and $\xi(j, i)$. $\square$

We are ready for the proof of Theorem 2, which we divide into two parts. The notation are those of Theorem 2, that is, $\mu := \mathrm{UGWT}(\pi)$, where π is a fixed probability distribution on $\mathbb{N}$ with finite, nonzero mean.

PROPOSITION 10.1. *If $Q \in \mathcal{P}([0, 1])$ satisfies $Q = F_{\pi,t}(Q)$, then*

$$\Psi_{\mathcal{L}_\mu}(t) \geq \frac{\mathbb{E}[D]}{2}\mathbb{P}(\xi_1 + \xi_2 > 1) - t\mathbb{P}(\xi_1 + \cdots + \xi_D > t),$$

where $D \sim \pi$ and where $\{\xi_k\}_{k \geq 1}$ are i.i.d. with law Q, independent of D.

PROOF. Kolmogorov's extension theorem allows us to convert the consistency equation $Q = F_{\pi,t}(Q)$ into a random rooted tree $\mathbb{T} \sim \mathrm{UGWT}(\pi)$ equipped with marks satisfying (10.1) a.s., such that conditionally on the structure of $[\mathbb{T}, o]_h$, the marks from generation h to $h - 1$ are i.i.d. with law Q. This random rooted network is easily checked to be unimodular. Thus, we may apply Proposition 7.1 with $f = \mathbf{1}_{\partial\xi > t}$. By Lemma 10.1, we have $\widehat{f} = \mathbf{1}_{\xi+\xi^* > 1}$, and hence

$$\Psi_{\mathcal{L}_\mu}(t) \geq \tfrac{1}{2}\vec{\mu}\big(\xi + \xi^* > 1\big) - t\mu(\partial\xi > t).$$

This is precisely the desired result, since we have by construction

$$\mu(\partial\xi > t) = \mathbb{P}(\xi_1 + \cdots + \xi_D > t), \qquad \vec{\mu}\big(\xi + \xi^* > 1\big) = \mathbb{E}[D]\mathbb{P}(\xi_1 + \xi_2 > 1),$$

where $D \sim \pi$ and $\xi_1, \xi_2, \ldots$ are i.i.d. with law Q, independent of D. □

PROPOSITION 10.2. *There exists $Q \in \mathcal{P}([0, 1])$ with $Q = F_{\pi,t}(Q)$ and*

$$\Psi_{\mathcal{L}_\mu}(t) = \frac{\mathbb{E}[D]}{2}\mathbb{P}(\xi_1 + \xi_2 > 1) - t\mathbb{P}(\xi_1 + \cdots + \xi_D > t),$$

where $D \sim \pi$ and where $\{\xi_k\}_{k \geq 1}$ are i.i.d. with law Q, independent of D.

PROOF. Let $\mathbb{T} \sim \mathrm{UGWT}(\pi)$. Thanks to Proposition 6.1, we have

$$\partial\Theta_\varepsilon(\mathbb{T}, o) \xrightarrow[\varepsilon \to 0]{L^2} \partial\Theta_0(\mathbb{T}, o).$$

In particular, there is a deterministic vanishing sequence $\varepsilon_1, \varepsilon_2, \ldots$ along which the convergence holds almost surely. This almost-sure convergence automatically extends from the root to all vertices, since under a unimodular measure μ, *everything shows at the root* [2], Lemma 2.3. More precisely,

$$\mu(A) = 1 \quad \Longrightarrow \quad \mu(\widetilde{A}) = 1,$$

for any Borel set $A \subseteq \mathcal{G}_\star$, where $\widetilde{A}$ consists of those $(G, o) \in \mathcal{G}_\star$ such that $(G, i) \in A$ for all vertices i of G. Here, we apply it to $\mu = \mathrm{UGWT}(\pi)$ and

$$A = \big\{(G, o) \in \mathcal{G}_\star : \partial\Theta_{\varepsilon_n}(G, o) \xrightarrow[n\to\infty]{} \partial\Theta_0(G, o)\big\}.$$

Thus, $\mathbb{T}$ satisfies almost surely the assumption of Proposition 9.2. Consequently, the marks $\xi(i, j) := [1 - \mathfrak{f}^{-1}_{\mathbb{T}_{i\to j}}(t)]_0^1$ satisfy (10.1) almost surely, and

$$\partial\Theta_0(\mathbb{T}, o) > t \quad \Longleftrightarrow \quad \partial\xi(o) > t.$$

This ensures that $f = \mathbf{1}_{\partial\xi > t}$ satisfies the requirements for equality in Proposition 7.1, and we may then use Lemma 10.1 to rewrite the conclusion as

$$\Psi_{\mathcal{L}_\mu}(t) = \tfrac{1}{2}\vec{\mu}\big(\xi + \xi^* > 1\big) - t\mu(\partial\xi > t).$$

Now, $D = \deg(\mathbb{T}, o)$ has law π and conditionally on D, the subtrees $\{\mathbb{T}_{i\to o}\}_{i\sim o}$ are i.i.d. copies of a homogenous Galton–Watson tree $\widehat{\mathbb{T}}$ with offspring distribution $\widehat{\pi}$. Since $\xi(i, o)$ depends only on the subtree $\mathbb{T}_{i\to o}$, we obtain

$$\mu(\partial\xi > t) = \mathbb{P}(\xi_1 + \cdots + \xi_D > t), \qquad \vec{\mu}\big(\xi + \xi^* > 1\big) = \mathbb{E}[D]\mathbb{P}(\xi_1 + \xi_2 > 1),$$

where $\xi_1, \xi_2, \ldots$ are i.i.d. copies of $[1 - \mathfrak{f}^{-1}_{\widehat{\mathbb{T}}}(t)]_0^1$, independent of D. In turn, removing the root of $\widehat{\mathbb{T}}$ splits it into a $\widehat{\pi}$-distributed number of i.i.d. copies of $\widehat{\mathbb{T}}$, so that the law Q of $[1 - \mathfrak{f}^{-1}_{\widehat{\mathbb{T}}}(t)]_0^1$ satisfies $Q = F_{\pi,t}(Q)$. □

11. Proof of Theorem 3. In this final section, we prove Theorem 3. This main ingredient is Proposition 11.1, which states that dense subgraphs must be large under the pairing model. Fix a degree sequence $\mathbf{d} = \{d(i)\}_{1\le i\le n}$ and set $2m = \sum_{i=1}^n d(i)$. We need two preparatory lemmas.

LEMMA 11.1. *Fix a subset of vertices $S \subseteq \{1,\dots,n\}$. Then the number of edges of $\mathbb{G}[\mathbf{d}]$ with both end-points in S is stochastically dominated by a binomial random variable with mean $\frac{1}{m}(\sum_{i\in S} d_i)^2$.*

PROOF. We assume that $s := \sum_{i\in S} d_i < m$, otherwise the claim is trivial. It is classical that $\mathbb{G}[\mathbf{d}]$ can be generated sequentially: at each step $1 \le t \le m$, a half-edge is selected and paired with a uniformly chosen other half-edge. The selection rule is arbitrary, and we choose to give priority to half-edges whose end-point lies in S. Let X_t be the number of edges with both end-points in S after t steps. Then $\{X_t\}_{0\le t\le m}$ is a Markov chain with $X_0 = 0$ and transitions

$$X_{t+1} := \begin{cases} X_t + 1, & \text{with conditional probability } \dfrac{(s - X_t - t - 1)^+}{2m - 2t - 1}, \\ X_t, & \text{otherwise.} \end{cases}$$

For every $0 \le t < m$, the fact that $X_t \ge 0$ ensures that

$$\frac{(s - X_t - t - 1)^+}{2m - 2t - 1} \le \frac{s - t - 1}{2m - 2t - 1}\mathbf{1}_{(t<s)} \le \frac{s}{2m}\mathbf{1}_{(t<s)},$$

where the second inequality uses the condition $s < m$. This shows that X_m is in fact stochastically dominated by a binomial $(s, \frac{s}{2m})$, which is enough. $\square$

LEMMA 11.2. *Let $X_{k,r}$ be the number of induced subgraphs with k vertices and at least r edges in $\mathbb{G}[\mathbf{d}]$. Then, for any $\theta > 0$,*

$$\mathbb{E}[X_{k,r}] \le \left(\frac{2r}{\theta^2 m}\right)^r \left(\frac{e}{k}\sum_{i=1}^n e^{\theta d_i}\right)^k.$$

PROOF. First observe that if $Z \sim \mathrm{Bin}(n, p)$ then by a simple union-bound,

$$\mathbb{P}(Z \ge r) \le \binom{n}{r} p^r \le \frac{n^r p^r}{r!} = \frac{\mathbb{E}[Z]^r}{r!}.$$

Thus, Lemma 11.1 ensures that the number Z_S of edges with both end-points in S satisfies

$$\mathbb{P}(Z_S \ge r) \le \frac{1}{r! m^r}\left(\sum_{i\in S} d_i\right)^{2r} \le \left(\frac{2r}{\theta^2 m}\right)^r \prod_{i\in S} e^{\theta d_i},$$

where we have used the crude bounds $x^{2r} \leq (2r)!e^x$ and $(2r)!/r! \leq (2r)^r$. The result follows by summing over all S with $|S| = k$ and observing that

$$\sum_{|S|=k} \prod_{i \in S} e^{\theta d_i} \leq \frac{1}{k!} \left(\sum_{i=1}^{n} e^{\theta d_i} \right)^k \leq \left(\frac{k}{e} \sum_{i=1}^{n} e^{\theta d_i} \right)^k.$$

The second inequality follows from the classical lower-bound $k! \geq (\frac{k}{e})^k$. □

We now fix $\{\mathbf{d}_n\}_{n\geq 1}$ as in Theorem 3. Let $Z_{\delta,t}^{(n)}$ be the number of subsets $\varnothing \subsetneq S \subseteq \{1, \dots, n\}$ such that $|S| \leq \delta n$ and $|E(S)| \geq t|S|$ in $\mathbb{G}_n := \mathbb{G}[\mathbf{d}_n]$.

PROPOSITION 11.1. *For each $t > 1$, there is $\delta > 0$ and $\kappa < \infty$ such that*

$$\mathbb{E}[Z_{\delta,t}^{(n)}] \leq \kappa \left(\frac{\ln n}{n} \right)^{t-1},$$

uniformly in $n \geq 1$. In particular, $Z_{\delta,t}^{(n)} = 0$ w.h.p. as $n \to \infty$.

PROOF. The assumptions of Theorem 3 guarantee that for some $\theta > 0$,

$$\alpha := \inf_{n\geq 1} \left\{ \frac{1}{n} \sum_{i=1}^{n} d_n(i) \right\} > 0 \quad \text{and} \quad \lambda := \sup_{n\geq 1} \left\{ \frac{1}{n} \sum_{i \in V} e^{\theta d_n(i)} \right\} < \infty.$$

Now, fix $t > 1$ and choose $\delta > 0$ small enough so that $f(\delta) < 1$, where

$$f(\delta) := \left(1 \vee \frac{2(1+t)}{\alpha \theta^2} \right)^{t+1} e\lambda \delta^{t-1}.$$

Using Lemma 11.2 and the trivial inequality $kt \leq \lceil kt \rceil \leq k(t+1)$, we have

$$\mathbb{E}[X_{k,\lceil kt \rceil}^{(n)}] \leq \left(\frac{2\lceil kt \rceil}{\theta^2 k \alpha} \right)^{\lceil kt \rceil} (e\lambda)^k \left(\frac{k}{n} \right)^{\lceil kt \rceil - k} \leq f^k \left(\frac{k}{n} \right).$$

Since f is increasing, we see that for any $1 \leq m \leq \delta n$,

$$\begin{aligned}
\mathbb{E}[Z_{\delta,t}^{(n)}] = \sum_{k=1}^{\lfloor \delta n \rfloor} \mathbb{E}[X_{k,\lceil kt \rceil}^{(n)}] &\leq \sum_{k=1}^{m-1} f^k \left(\frac{m}{n} \right) + \sum_{k=m}^{\lfloor \delta n \rfloor} f^k(\delta) \\
&\leq \frac{f(m/n)}{1 - f(m/n)} + \frac{f(\delta)^m}{1 - f(\delta)}.
\end{aligned}$$

Choose $m \sim c \ln n$ with c fixed. As $n \to \infty$, the first term is of order $(\frac{\ln n}{n})^{t-1}$ while the second is of order $f(\delta)^{c \ln n} \ll (\frac{\ln n}{n})^{t-1}$, if c is large enough. □

PROOF OF THEOREM 3. The assumptions on $\{\mathbf{d}_n\}_{n\geq 1}$ are more than sufficient to guarantee that a.s., the local weak limit of $\{\mathbb{G}_n\}_{n\geq 1}$ is $\mu := \mathrm{UGWT}(\pi)$

(see, e.g., [9]). Thus, the weak convergence $\mathcal{L}_{\mathbb{G}_n} \to \mathcal{L}_\mu$ holds a.s., by Theorem 1. Now, if $t < \varrho(\mu)$ then $\mathcal{L}_\mu((t,\infty)) > 0$, so the Portmanteau theorem ensures that $\liminf_n \mathcal{L}_{\mathbb{G}_n}((t,\infty)) > 0$ a.s. Consequently,

$$\mathbb{P}(\varrho(\mathbb{G}_n) \le t) = \mathbb{P}(\mathcal{L}_{\mathbb{G}_n}((t,\infty)) = 0) \underset{n\to\infty}{\longrightarrow} 0.$$

On the other-hand, if $t > \varrho(\mu)$ then $\mathcal{L}_\mu([t,\infty)) = 0$, so the Portmanteau theorem gives $\mathcal{L}_{\mathbb{G}_n}((t,\infty)) \to 0$ a.s. Thus, with δ as in Proposition 11.1,

$$\mathbb{P}(\varrho(\mathbb{G}_n) > t) \le \mathbb{P}(\mathcal{L}_{\mathbb{G}_n}([t,\infty)) > \delta) + \mathbb{P}(Z_{\delta,t}^{(n)} > 0) \underset{n\to\infty}{\longrightarrow} 0.$$

Note that the requirement $t > 1$ is fulfilled, since $\varrho(\mu) \ge 1$. Indeed, every node in a tree of size n has load $1 - \frac{1}{n}$, and the assumption $\pi_0 + \pi_1 < 1$ guarantees that the size of the random tree $\mathbb{T} \sim \mathrm{UGWT}(\pi)$ is unbounded. $\square$

Acknowledgements. The work was initiated at the follow-up meeting of the Newton Institute programme N°86: *Stochastic Processes in Communication Sciences*. The authors thank M. Lelarge, R. Sundaresan and an anonymous referee for fruitful discussions.

REFERENCES

[1] AGARWAL, R. P., MEEHAN, M. and O'REGAN, D. (2001). *Fixed Point Theory and Applications*. *Cambridge Tracts in Mathematics* **141**. Cambridge Univ. Press, Cambridge. MR1825411

[2] ALDOUS, D. and LYONS, R. (2007). Processes on unimodular random networks. *Electron. J. Probab.* **12** 1454–1508. MR2354165

[3] ALDOUS, D. and STEELE, J. M. (2004). The objective method: Probabilistic combinatorial optimization and local weak convergence. In *Probability on Discrete Structures*. *Encyclopaedia Math. Sci.* **110** 1–72. Springer, Berlin. MR2023650

[4] ALDOUS, D. J. (2001). The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms* **18** 381–418. MR1839499

[5] ALDOUS, D. J. and BANDYOPADHYAY, A. (2005). A survey of max-type recursive distributional equations. *Ann. Appl. Probab.* **15** 1047–1110. MR2134098

[6] BENJAMINI, I. and SCHRAMM, O. (2001). Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.* **6** no. 23, 13 pp. (electronic). MR1873300

[7] BILLINGSLEY, P. (1999). *Convergence of Probability Measures*, 2nd ed. Wiley, New York. MR1700749

[8] BOLLOBÁS, B. (1980). A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European J. Combin.* **1** 311–316. MR0595929

[9] BORDENAVE, C. (2012). Lecture notes on random graphs and probabilistic combinatorial optimization. Available at http://www.math.univ-toulouse.fr/~bordenave/coursRG.pdf.

[10] BORDENAVE, C., LELARGE, M. and SALEZ, J. (2013). Matchings on infinite graphs. *Probab. Theory Related Fields* **157** 183–208. MR3101844

[11] CAIN, J. A., SANDERS, P. and WORMALD, N. (2007). The random graph threshold for k-orientability and a fast algorithm for optimal multiple-choice allocation. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* 469–476. ACM, New York. MR2482873

[12] FERNHOLZ, D. and RAMACHANDRAN, V. (2007). The k-orientability thresholds for $G_{n,p}$. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms* 459–468. ACM, New York. MR2482872

[13] FOUNTOULAKIS, N., KHOSLA, M. and PANAGIOTOU, K. (2011). The multiple-orientability thresholds for random hypergraphs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms* 1222–1236. SIAM, Philadelphia, PA. MR2858395

[14] GAMARNIK, D., NOWICKI, T. and SWIRSZCZ, G. (2006). Maximum weight independent sets and matchings in sparse random graphs. Exact results using the local weak convergence method. *Random Structures Algorithms* **28** 76–106. MR2187483

[15] GAO, P. and WORMALD, N. C. (2010). Load balancing and orientability thresholds for random hypergraphs [extended abstract]. In *STOC'*10*—Proceedings of the* 2010 *ACM International Symposium on Theory of Computing* 97–103. ACM, New York. MR2743258

[16] HAJEK, B. (1990). Performance of global load balancing by local adjustment. *IEEE Trans. Inform. Theory* **36** 1398–1414. MR1080823

[17] HAJEK, B. (1996). Balanced loads in infinite networks. *Ann. Appl. Probab.* **6** 48–75. MR1389831

[18] JANSON, S. (2009). The probability that a random multigraph is simple. *Combin. Probab. Comput.* **18** 205–225. MR2497380

[19] KHANDWAWALA, M. and SUNDARESAN, R. (2014). Belief propagation for optimal edge cover in the random complete graph. *Ann. Appl. Probab.* **24** 2414–2454. MR3262507

[20] LECONTE, M., LELARGE, M. and MASSOULIÉ, L. (2012). Convergence of multivariate belief propagation, with applications to cuckoo hashing and load balancing. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms* 35–46. SIAM, Philadelphia, PA. MR3185378

[21] LELARGE, M. (2012). A new approach to the orientation of random hypergraphs. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms* 251–264. ACM, New York. MR3205213

[22] ŁUCZAK, T. (1992). Sparse random graphs with a given degree sequence. In *Random Graphs, Vol.* 2 (*Poznań*, 1989). *Wiley-Intersci. Publ.* 165–182. Wiley, New York. MR1166614

[23] LYONS, R. (2005). Asymptotic enumeration of spanning trees. *Combin. Probab. Comput.* **14** 491–522. MR2160416

[24] SALEZ, J. (2013). Weighted enumeration of spanning subgraphs in locally tree-like graphs. *Random Structures Algorithms* **43** 377–397. MR3094425

[25] SALEZ, J. and SHAH, D. (2009). Belief propagation: An asymptotically optimal algorithm for the random assignment problem. *Math. Oper. Res.* **34** 468–480. MR2554069

[26] SHAKED, M. and SHANTHIKUMAR, J. G. (1994). *Stochastic Orders and Their Applications*. Academic Press, Inc., Boston, MA. MR1278322

[27] STEELE, J. M. (2002). Minimal spanning trees for graphs with random edge lengths. In *Mathematics and Computer Science, II* (*Versailles*, 2002). *Trends Math.* 223–245. Birkhäuser, Basel. MR1940139

[28] VAN DER HOFSTAD, R. (2013). Random graphs and complex networks. Available at http://www.win.tue.nl/~rhofstad/NotesRGCN.html.

UNIVERSITY OF CALIFORNIA, BERKELEY
DEPARTMENT OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCES
253 CORY HALL
BERKELEY, CALIFORNIA 94720-1770
USA
E-MAIL: ananth@berkeley.edu

UNIVERSITÉ PARIS DIDEROT
UFR DE MATHÉMATIQUES
5 RUE THOMAS MANN
75205 PARIS CEDEX 13
FRANCE
E-MAIL: justin.salez@univ-paris-diderot.fr