# Load Balancing in Hypergraphs 

Payam Delgosha and Venkat Anantharam<br>Department of Electrical Engineering and Computer Sciences<br>University of California, Berkeley<br>\{pdelgosha, ananth\} @ berkeley.edu

October 3, 2017


#### Abstract

Consider a simple locally finite hypergraph on a countable vertex set, where each edge represents one unit of load which should be distributed among the vertices defining the edge. An allocation of load is called balanced if load cannot be moved from a vertex to another that is carrying less load. We analyze the properties of balanced allocations of load. We extend the concept of balancedness from finite hypergraphs to their local weak limits in the sense of Benjamini and Schramm (2001) and Aldous and Steele (2004). To do this, we define a notion of unimodularity for hypergraphs which could be considered an extension of unimodularity in graphs. We give a variational formula for the balanced load distribution and, in particular, we characterize it in the special case of unimodular hypergraph Galton Watson processes. Moreover, we prove the convergence of the maximum load under some conditions. Our work is an extension to hypergraphs of Anantharam and Salez (2016), which considered load balancing in graphs, and is aimed at more comprehensively resolving conjectures of Hajek (1990).


Keywords: Load balancing, sparse graph limits, hypergraphs, local weak convergence, objective method, unimodularity, configuration model

## 1 Introduction

The problem of load balancing is ubiquitous in networks. As examples, consider the problem of routing traffic through a communication network or of assigning tasks among the servers in a cloud computing framework. What is common in these scenarios is a number of servers and a number of tasks whose load should be distributed among the servers. Examples of servers are paths through the network from a given source to a given destination in the routing scenario, or processors in a cloud computing framework. Examples of tasks are the amount of traffic to be routed from the source to the destination or the computational work to be done at the servers, respectively. Further, there are typically restrictions as to which resources are available to a given task. Performance considerations require that the allocation of the load of a task among the resources available to it should be done in a way that optimizes a measure of performance, such as delay or queue length. When the problem size is large, it may be expensive to compute the detailed characteristics of an optimal or sufficiently good allocation of the load. Instead, it is interesting to focus on the statistical characteristics of the allocation, such as the empirical distribution of the load faced by a typical resource in the network.

[^0]

Figure 1: Load balancing with 3 tasks and 4 servers. (a) illustrates the bipartite representation together with an allocation. While the load of $e_{1}$ could be served by nodes in $\left\{v_{1}, v_{2}, v_{3}\right\}$, half of its load is being assigned to $v_{1}$ and the other half to $v_{3}$, i.e. $\theta\left(e_{1}, v_{1}\right)=1 / 2, \theta\left(e_{1}, v_{2}\right)=0$ and $\theta\left(e_{1}, v_{3}\right)=1 / 2$. (b) shows the hypergraph representation. The total load at a node $i$ is denoted by $\partial \theta(i)$. For instance, $\partial \theta\left(v_{3}\right)=\theta\left(e_{1}, v_{3}\right)+\theta\left(e_{3}, v_{3}\right)=1 / 2$.

This paper is concerned with developing such a viewpoint in the context of a specific kind of load balancing problem which has broad applicability.

### 1.1 Model and Prior Work

We model the load balancing problem by a bipartite graph in which every node on the right represents a task and every node on the left represents a server. Each server is accessible to a certain subset of the tasks. Equivalently each task has access to only a certain subset of the servers. We view the bipartite graph as a hypergraph, with each vertex of the hypergraph representing a server and each hyperedge representing a task. The vertices of a hyperedge are then the servers that are accessible to it. Let $\left\{v_{1}, \ldots, v_{n}\right\}$ and $\left\{e_{1}, \ldots, e_{m}\right\}$ denote the set of servers and tasks, or equivalently vertices and hyperedges, respectively. Therefore, $v_{i} \in e_{j}$ means that server $v_{i}$ can be used to contribute to the performance of task $e_{j}$. See Figure 1 for an example. In general, we might want to consider the scenario where task $e_{j}$ has an amount of load equal to $l_{j}$, which could be arbitrarily allocated among the servers $v_{i} \in e_{j}$. For simplicity, we will consider in this paper only the case where all the $l_{j}$ equal 1 , but we leave the discussion general for the moment. Let $\theta$ be an allocation of the load of tasks among the servers, i.e. $\theta\left(e_{j}, v_{i}\right)$ is the amount of load coming from task $e_{j}$ assigned to server $v_{i}$. Hence, $\theta\left(e_{j}, v_{i}\right) \geq 0$ and $\sum_{v_{i} \in e_{j}} \theta\left(e_{j}, v_{i}\right)=l_{j}$. For a server $v_{i}$, let $\partial \theta\left(v_{i}\right)$ be the total amount of load assigned to $v_{i}$, i.e. $\partial \theta\left(v_{i}\right)=\sum_{e_{j}: v_{i} \in e_{j}} \theta\left(e_{j}, v_{i}\right)$.

This formulation of load balancing was studied by Hajek [Haj90] who, in particular, formulated the notion of a balanced allocation. It is natural to expect that a task would be happier to use servers that are currently handling less load, if available, as opposed to those handling more load. An allocation $\theta$ is said to be balanced if no task desires to change the allocation of its load. For finite load balancing problems, this turns out to be equivalent to the statement that the allocation minimizes $\sum_{i} f\left(\partial \theta\left(v_{i}\right)\right)$ for any given fixed strictly convex function $f$. One can think of $\sum_{i} f\left(\partial \theta\left(v_{i}\right)\right)$ as the aggregate cost we need to pay to process all the tasks. Hajek showed the existence of balanced allocations and uniqueness of the total load at nodes under any balanced allocation, and suggested algorithms to find a balanced allocation. It is particularly remarkable that the notion of a balanced allocation does not depend on the specific choice of the strictly convex cost function $f$.

With the aim of understanding the statistical characteristics of balanced allocations in large load


Figure 2: Hajek's example to show non-uniqueness of the load for infinite graphs. Consider the rooted 3 -regular graph with infinite depth as shown. We send all of the unit load corresponding to each edge in the direction of the shown arrows. The red path goes to infinity in both pictures. The allocation in (a) makes the total load at every vertex equal to 2 while that in (b) makes the total load at all vertices equal to 1 . Therefore, both are balanced.
balancing problems, Hajek assumed that each task could be performed by only two servers - hence the underlying hypergraph reduces to a graph - and he assumed that each edge in this graph carries one unit of load. He then studied such a load balancing problem in large random graphs [Haj90]. In particular, Hajek considered the sparse Erdős-Rényi model to generate these graphs, where $\alpha n$ edges are distributed among $n$ vertices uniformly at random, with $\alpha$ being a fixed parameter. It is well known that the asymptotic structure of the local neighborhood of a typical vertex in a sparse ErdősRényi model is given by a Poisson Galton-Watson tree (see, for instance, [ $\mathrm{DM}^{+} 10$, Proposition 2.6] for a precise formulation of this statement). This suggests that the behavior of balanced allocations in Galton-Watson trees might be a good proxy for the load distribution in large Erdős-Rényi graphs. Hajek conjectured that the recursive nature of a Galton-Watson process helps one analyze the distribution of balanced allocations by studying fixed point equations. He was even able to suggest the form of the fixed point equation for the Poisson Galton-Watson tree. However, it turns out that this approach is more subtle than it looks. For one thing, Hajek realized that the notion of balanced allocation in an infinite graph as a proxy for large graphs is not well defined $\left[\mathrm{H}^{+} 96\right]$. See Figure 2 for an example.

Hajek's conjecture for the graph regime (i.e. when each task could only be distributed among two servers) was settled by Anantharam and Salez [AS16]. They employed the framework of local weak convergence, also known as the objective method [BS01], [AS04], [AL07]. This framework introduces a notion of convergence for a sequence of finite graphs by representing each finite graph as a probability distribution on rooted graphs and then discussing convergence in probabilistic terms. Roughly speaking, the operational meaning of this technique is that convergence holds when the distribution of the local neighborhood of a vertex chosen uniformly at random in the graph converges to that of the local neighborhood around the root in the limit. For instance, a sequence of sparse Erdős-Rényi graphs converges in this sense to a Poisson Galton-Watson tree, consistent with our established intuition in this case.

A random rooted graph that can appear as the limit of a sequence of finite graphs is called sofic. Not all random rooted graphs are sofic. The reason is that all the vertices in a finite graph have the same chance of being chosen as the root. This should manifest itself as some form of stationarity in the limit, i.e. the limiting object should be invariant under changing the root. This property is called unimodularity, which is a necessary condition for being sofic. Whether the converse is true is an open problem [AL07].

Anantharam and Salez settled Hajek's conjecture by first defining a notion of balancedness for unimodular random rooted graphs. Moreover, they showed that if a sequence of finite graphs $G_{n}$
converges to a random rooted graph in the above local weak sense, the total load associated to a balanced allocation at a vertex chosen uniformly at random in $G_{n}$ converges in distribution to the total load associated to the balanced allocation at the root of the limit. Additionally, they managed to express a certain functional of the distribution of the load at the root of the Galton-Watson tree in terms of a fixed point distributional equation, settling Hajek's conjecture in the graph regime. Beyond this, they also proved the convergence of the maximum load for a sequence of finite graphs resulting from a certain configuration model to that of their local weak limit, under some additional conditions.

### 1.2 Our Contributions

We study the above load balancing problem in the more general regime where each task could have access to more than two servers, i.e. the underlying network is a hypergraph instead of a graph.

Our machinery for deriving results analogous to those in the graph regime will be a generalized method of local weak convergence on hypergraphs. One novelty of our development is to introduce a notion analogous to unimodularity for processes on random rooted hypergraphs. We believe that this generalized framework could be of independent interest in a variety of problems in which the underlying model is best expressed in terms of hypergraphs rather than graphs.

In particular, we prove that for any unimodular probability distribution on the set of rooted hypergraphs with finite expected degree, there exists a balanced allocation which is consistent with the local weak limit theory, i.e. the load distribution of a sequence of hypergraphs converges to that of the limit. For a special class of branching process on rooted hypertrees which is a generalization of Galton-Watson processes, we show that the distribution of the load at the root can be specified via a fixed point distributional equation. Finally, we study the convergence of the maximum load for a sequence of random hypergraphs generated from a configuration model to that of the limit, under some additional conditions.

## 2 Prerequisites and Notations

Throughout this paper $\mathbb{R}$ and $\mathbb{R}_{\geq 0}$ denote the set of real numbers and nonnegative real numbers, respectively. Moreover, $\mathbb{Q}$ denotes the set of rational numbers. $\mathbb{N}$ denotes the set of positive integers and $\mathbb{N}_{0}:=\mathbb{N} \cup\{0\}$. For a Polish space $X$, let $\mathcal{P}(X)$ and $\mathcal{M}(X)$ respectively denote the set of probability measures and nonnegative finite measures on $X$, with respect to the Borel $\sigma$-field of $X$. We use the abbreviations "a.s." and "a.e." for the phrases "almost surely" and "almost everywhere", respectively. For a real number $x \in \mathbb{R}$, we denote $\max \{x, 0\}$ by $x^{+}$, and we denote $\min \left\{x^{+}, 1\right\}$ by $[x]_{0}^{1}$.

For a Polish space $X$, we say that a sequence of probability measures $\mu_{n}$ converges weakly to a probability measure $\mu \in \mathcal{P}(X)$, and write $\mu_{n} \Rightarrow \mu$, if for any bounded continuous function $f: X \rightarrow \mathbb{R}$ we have

$$
\lim _{n \rightarrow \infty} \int f d \mu_{n}=\int f d \mu
$$

See [Bil71] and [Bil99] for more details on weak convergence of probability measures.
Given two measurable spaces $\left(X_{1}, \mathcal{F}_{1}\right)$ and $\left(X_{2}, \mathcal{F}_{2}\right)$, a measurable mapping $f: X_{1} \rightarrow X_{2}$, and a nonnegative measure $\mu_{1}$ on $\mathcal{F}_{1}$, the pushforward measure $f_{*}\left(\mu_{1}\right)$ on $\mathcal{F}_{2}$ is defined by

$$
f_{*}\left(\mu_{1}\right)(A)=\mu_{1}\left(f^{-1}(A)\right)
$$

for $A \in \mathcal{F}_{2}$.
We write $:=$ for equality by definition. We use the terms "node" and "vertex" interchangeably.

### 2.1 Hypergraphs

We work with simple hypergraphs defined on a countable vertex set, where each edge is a finite subset of the vertex set. For a hypergraph $H$, the sets of vertices and edges are denoted by $V(H)$ and $E(H)$, respectively. We write $H$ as $\langle V, E\rangle$, where $V=V(H)$ and $E=E(H)$. We say a hypergraph is simple if $E(H) \subset 2^{V(H)}$. This means that in any edge each vertex can show up at most once, and that each subset of vertices can show up at most once as an edge. All hypergraphs appearing in this paper will be simple, unless otherwise stated. For a vertex $i \in V(H)$, denote its degree by $\operatorname{deg}_{H}(i):=|\{e \in E(H): e \ni i\}|$. For a given hypergraph $H$, let

$$
\begin{equation*}
\Psi(H):=\{(e, i): e \in E(H), i \in e\} \tag{1}
\end{equation*}
$$

denote the set of all edge-vertex pairs in the hypergraph.
Definition 1. A hypergraph $H$ is said to be locally finite if $\operatorname{deg}_{H}(i)$ is finite for all $i \in V(H)$ and $e$ is finite for all $e \in E(H)$.

Note that the above definition does not imply that there is a uniform bound on edge sizes or vertex degrees. Hence, a locally finite hypergraph can have arbitrarily large edges or vertex degrees. Throughout this paper, we assume that all the hypergraphs are locally finite, unless otherwise stated. Thus, by default, the term hypergraph in this paper means a simple, locally finite hypergraph on a countable vertex set.

For technical reasons, it is sometimes easier to work with bounded hypergraphs in proofs and then relax the boundedness condition.

Definition 2. A hypergraph $H$ is said to be bounded if $\operatorname{deg}_{H}(i) \leq \Delta$ for all $i \in V(H)$ and also $|e| \leq L$ for all $e \in E(H)$, for finite constants $\Delta$ and $L$.

A path from node $i$ to node $j$ is an alternating vertex-edge sequence $i_{0}, e_{1}, i_{1}, e_{2}, i_{2}, \ldots, e_{n}, i_{n}$ with $i_{0}=i$ and $i_{n}=j$, and $i_{k} \in e_{k+1}$ for $0 \leq k<n$. The length of such a path is defined to be $n$. The distance between vertices $i$ and $j$, denoted by $d_{H}(i, j)$, is defined to be the length of the shortest path between $i$ and $j$ if $i \neq j$, and 0 when $i=j$. A path is called closed if $i_{0}=i_{n}$.

Definition 3. A hypergraph $H$ is called a hypertree if there is no closed path $i_{0}, e_{1}, i_{1}, \ldots, e_{n-1}, i_{n-1}, e_{n}, i_{n}$ with $n \geq 2$ such that $i_{j} \neq i_{l}$ and $e_{j} \neq e_{l}$ for $1 \leq j \neq l \leq n$.

Remark 1. Note that a hypertree need not be connected. It is straightforward to prove that if there is a path between vertices $i$ and $j$ in a hypertree, then the shortest path between these vertices is unique.

Definition 4. For a hypergraph $H$ and a subset $W \subset V(H)$, define $E_{H}(W):=\{e \in E(H): e \subset W\}$. $E_{H}(W)$ is comprised of the edges of $H$ with all endpoints in the set $W$. For $i \in V(H)$ and $d \geq 0$, define $V_{i, d}^{H}:=\left\{j \in V(H): d_{H}(i, j) \leq d\right\}$ and $D_{i, d}^{H}:=\left\{j \in V(H): d_{H}(i, j)=d\right\}$. In particular, $V_{i, 0}^{H}=D_{i, 0}^{H}=\{i\}$.

Definition 5. A vertex rooted hypergraph is a hypergraph $H$ with a distinguished vertex $i \in V(H)$. We denote this by $(H, i)$. An edge-vertex rooted hypergraph is a hypergraph with a distinguished edge $e \in E(H)$ and a distinguished vertex $i \in V(H)$ such that $i \in e$. This is denoted by $(H, e, i)$.

Definition 6. We say that two hypergraphs $H$ and $H^{\prime}$ are isomorphic and write $H \equiv H^{\prime}$ when there is a bijection $\phi: V(H) \rightarrow V\left(H^{\prime}\right)$ such that $e \in E(H)$ if and only if $\phi(e):=\{\phi(j): j \in e\} \in E\left(H^{\prime}\right)$. Also we say two vertex rooted hypergraphs $(H, i)$ and $\left(H^{\prime}, i^{\prime}\right)$ are isomorphic and write $(H, i) \equiv\left(H^{\prime}, i^{\prime}\right)$ if the above bijection $\phi$ exists and we have $\phi(i)=i^{\prime}$ as well. Furthermore, we say two edge-vertex rooted hypergraphs $(H, e, i)$ and $\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$ are isomorphic and write $(H, e, i) \equiv\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$ if the above bijection exists, and we have $\phi(i)=i^{\prime}$ and $\phi(e)=e^{\prime}$.

Instead of working with global isomorphisms as above, we can consider local isomorphisms, i.e. comparing two rooted hypergraphs up to some given depth.

Definition 7. We say two vertex rooted hypergraphs $(H, i)$ and $\left(H^{\prime}, i^{\prime}\right)$ are isomorphic up to depth d and write $(H, i) \equiv_{d}\left(H^{\prime}, i^{\prime}\right)$ if their truncations up to depth d are isomorphic, i.e. $\left\langle V_{i, d}^{H}, E_{H}\left(V_{i, d}^{H}\right)\right\rangle \equiv$ $\left\langle V_{i^{\prime}, d}^{H^{\prime}}, E_{H^{\prime}}\left(V_{i^{\prime}, d}^{H^{\prime}}\right)\right\rangle$ and also $\phi(i)=i^{\prime}$, where $\phi: V_{i, d}^{H} \rightarrow V_{i^{\prime}, d}^{H^{\prime}}$ is the vertex bijection establishing this isomorphism. Also, for $d \geq 1$, we say two edge-vertex rooted hypergraphs $(H, e, i)$ and $\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$ are isomorphic up to depth $d$ and write $(H, e, i) \equiv{ }_{d}\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$ if $\left\langle V_{i, d}^{H}, E_{H}\left(V_{i, d}^{H}\right)\right\rangle \equiv\left\langle V_{i^{\prime}, d}^{H^{\prime}}, E_{H^{\prime}}\left(V_{i^{\prime}, d}^{H^{\prime}}\right)\right\rangle$, $\phi(i)=i^{\prime}$, and $\phi(e)=e^{\prime}$, where $\phi(e):=\{\phi(j): j \in e\}$. Here $\phi: V_{i, d}^{H} \rightarrow V_{i^{\prime}, d}^{H^{\prime}}$ is the vertex bijection establishing this isomorphism.

Definition 8. Given two hypergraphs $H$ and $H^{\prime}$, for $i \in V(H)$ and $i^{\prime} \in V\left(H^{\prime}\right)$ we say that $(H, i)$ has a local embedding up to depth $d \geq 1$ into $\left(H^{\prime}, i^{\prime}\right)$ and write $(H, i) \hookrightarrow_{d}\left(H^{\prime}, i^{\prime}\right)$ if there is an injective mapping $\phi: V_{i, d}^{H} \hookrightarrow V_{i^{\prime}, d}^{H^{\prime}}$ such that:

1. $\phi(i)=i^{\prime}$, and
2. for all $e \in E_{H}\left(V_{i, d}^{H}\right)$, we have $\phi(e) \in E\left(H^{\prime}\right)$ where $\phi(e):=\{\phi(j): j \in e\}$.

Definition 9. Given a hypergraph $H$, a node $i \in V(H)$ and $d \geq 1$, let $(H, i)_{d}$ denote the vertex rooted hypergraph $\left(\left\langle V_{i, d}^{H}, E_{H}\left(V_{i, d}^{H}\right)\right\rangle, i\right)$. In fact, $(H, i)_{d}$ is the d-neighborhood of vertex $i$.

### 2.2 Balanced allocations on a hypergraph

Definition 10. An allocation on hypergraph $H=\langle V, E\rangle$ is a mapping $\theta: \Psi(H) \rightarrow[0,1]$ such that $\theta(e, i)$ with $i \in e \in E$ tells us how much load from resource $e$ is being given to node $i$. More formally, it is characterized by the properties:

$$
\begin{gathered}
\theta(e, i) \geq 0, \quad \forall e \in E(H), i \in e, \text { and } \\
\sum_{j \in e} \theta(e, j)=1, \quad \forall e \in E(H)
\end{gathered}
$$

In any allocation, a given vertex $i \in V(H)$ receives a portion $\theta(e, i)$ of the total unit load of resource $e$. The total load at the vertex is then the sum of portions it receives from resources $e \ni i$. The following definition establishes the notation to discuss this load.

Definition 11. Given an allocation $\theta$ on a hypergraph $H=\langle V, E\rangle$, define the function $\partial \theta: V(H) \rightarrow$ $\mathbb{R}_{\geq 0} b y$

$$
\partial \theta(i):=\sum_{e: i \in e} \theta(e, i), \quad \text { for all } i \in V(H)
$$

Definition 12. For a hypergraph $H=\langle V, E\rangle$, an allocation $\theta$ is called balanced if for all $e \in E$ and $i, j \in e$ we have

$$
\partial \theta(i)>\partial \theta(j) \quad \Rightarrow \quad \theta(e, i)=0
$$

Much of the paper is concerned with understanding the structure of balanced allocations on hypergraphs, and of the load resulting from such allocations. As we will soon see via examples, balanced allocations can exhibit phenomena analogous to phase transitions in statistical mechanics models. This is because the hard constraint defining balancedness can be thought of as analogous to a zero temperature limit. Following this analogy further, it is therefore convenient to deal with what might be called a positive temperature notion of balancedness, and then to send the temperature to zero. This is captured in the concept of $\epsilon$-balance.

Definition 13. For a hypergraph $H=\langle V, E\rangle$, an allocation $\theta$ is called $\epsilon$-balanced, if for all $e \in E$ and $i \in e$ we have

$$
\theta(e, i)=\frac{\exp (-\partial \theta(i) / \epsilon)}{\sum_{j \in e} \exp (-\partial \theta(j) / \epsilon)}
$$

Remark 2. Let $\theta$ be an $\epsilon$-balanced allocation on a hypergraph $H$. Note that if $e \in E$ and $i, j \in e$ are such that $\partial \theta(i)>\partial \theta(j)$ then

$$
\theta(e, i)=\frac{\exp (-\partial \theta(i) / \epsilon)}{\sum_{j} \exp (-\partial \theta(j) / \epsilon)} \leq \frac{1}{1+\exp \left(\frac{\partial \theta(i)-\partial \theta(j)}{\epsilon}\right)}
$$

Roughly speaking, if $\partial \theta(i)>\partial \theta(j)$ and $\epsilon$ is small, then $\theta(e, i) \approx 0$ and hence $\theta$ is approximately balanced. Also, roughly speaking, the smaller $\epsilon$ is, the more balanced an $\epsilon$-balanced allocation is.

In the above, we defined balancedness when all the loads come from the edges of the hypergraph. We can generalize this to the case where, in addition to the internal load imposed by the edges, we have external load as well. External load is modeled by a function $b: V(H) \rightarrow \mathbb{R}$, called the baseload function. For a vertex $i \in V(H), b(i)$ denotes the external load applied to node $i$. Throughout this paper, we assume that each baseload function is bounded, but we do not assume a uniform bound on all baseload functions. More precisely, for each baseload function $b$ on a given hypergraph $H$, we assume that there exists $M<\infty$ such that $|b(i)|<M$ for all $i \in V(H)$, where the constant $M$ may depend on $H$ and $b$. The concept of balancedness can be extended to the scenario with baseloads as follows.

Definition 14. For a hypergraph $H=\langle V, E\rangle$, together with a baseload function $b: V(H) \rightarrow \mathbb{R}$, an allocation $\theta: \Psi(H) \rightarrow[0,1]$ is called balanced with respect to the baseload b, if for all $e \in E$ and $i, j \in e$ we have

$$
\partial \theta(i)+b(i)>\partial \theta(j)+b(j) \quad \Rightarrow \quad \theta(e, i)=0 .
$$

Note that $\partial \theta(i)+b(i)$ is the total load at node $i$ where $\partial \theta(i)$ is the internal load and $b(i)$ is the contribution from the external load. We use the notation $\partial_{b} \theta$ as a shorthand for $\partial \theta+b$.

The concept of an $\epsilon$-balanced allocations can be similarly extended to the scenario with baseloads.
Definition 15. For a given hypergraph $H=\langle V, E\rangle$, together with a baseload function $b: V(H) \rightarrow \mathbb{R}$, we say an allocation $\theta: \Psi(H) \rightarrow[0,1]$ is $\epsilon$-balanced with respect to the baseload $b$, if for all $e \in E(H)$ and $i \in e$ we have

$$
\begin{equation*}
\theta(e, i)=\frac{\exp \left(-\partial_{b} \theta(i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-\partial_{b} \theta(j) / \epsilon\right)} \tag{2}
\end{equation*}
$$

It is known that if the hypergraph is finite, then balanced allocations exist with respect to any baseload and the resulting load vector is the same for all balanced allocations for the given baseload (see Theorem 2 and Corollary 5 in [Haj90]). This result is stated in the following proposition.

Proposition 1. If $H=\langle V, E\rangle$ is a finite hypergraph, and $b: V(H) \rightarrow \mathbb{R}$ is a given baseload function, then there exists at least one balanced allocation $\theta$ on $H$ with respect to the baseload $b$. Moreover, if $\theta$ and $\theta^{\prime}$ are two balanced allocations on $H$ with respect to $b$, then $\partial_{b} \theta(i)=\partial_{b} \theta^{\prime}(i)$ for all $i \in V(H)$.

Later, in Section 4, we will study $\epsilon$-balanced allocations with baseload for hypergraphs that are not necessarily finite. In particular, we will show in Corollary 4 therein that for bounded hypergraphs, for any baseload function, the total load at any vertex corresponding to any $\epsilon$-balanced allocation is uniquely defined.


Figure 3: A graph with 3 vertices and three edges and two different balanced allocations with zero baseload. Note that the total load at each vertex is the same for the two allocations and is equal to 1 at each vertex.

Note that for finite hypergraphs, although the balanced allocations with respect to a given baseload might not be uniquely defined, the total loads at the vertices resulting from any two balanced allocations necessarily have to be the same. See Figure 3 for an example. The case for infinite hypergraphs is more complicated though; in this case, the loads at the vertices also may not be unique, see Figure 2. See $\left[\mathrm{H}^{+} 96\right]$ for more discussion on this. However, we can state a weak uniqueness result in this case. The proof is given in Appendix A.

Proposition 2. Given the hypergraph $H=\langle V, E\rangle$ with the baseload function $b: V(H) \rightarrow \mathbb{R}$, suppose $\theta$ and $\theta^{\prime}$ are two balanced allocations on $H$ with respect to the baseload b. If $\sum_{i \in V(H)}\left|\partial_{b} \theta(i)-\partial_{b} \theta^{\prime}(i)\right|<\infty$ then $\partial_{b} \theta(i)=\partial_{b} \theta^{\prime}(i)$ for all $i \in V(H)$.

## $2.3 \quad \mathcal{H}_{*}$ and $\mathcal{H}_{* *}$

It is easy to check that the isomorphism between vertex rooted hypergraphs defined in Definition 6 is an equivalence relation. For a vertex rooted hypergraph $(H, i)$, let $[H, i]$ denote the equivalence class corresponding to $(H, i)$. Also, the isomorphism between edge-vertex rooted hypergraphs defined in Definition 6 is an equivalence relation. Let $[H, e, i]$ be the equivalence class corresponding to $(H, e, i)$.

Definition 16. Let $\mathcal{H}_{*}$ be the set of all equivalence classes of connected vertex rooted hypergraphs and $\mathcal{H}_{* *}$ the set of all equivalence classes of connected edge-vertex rooted hypergraphs. Hence, each element of $\mathcal{H}_{*}$ is of the form $[H, i]$, where $[H, i]$ denotes the equivalence class of $(H, i)$, where $i \in V(H)$ for some connected hypergraph $H=\langle V, E\rangle$. Similarly, each element of $\mathcal{H}_{* *}$ is of the form $[H, e, i]$, where $e \in E(H)$ and $i \in V(H)$ such that $i \in e$ for some connected hypergraph $H=\langle V, E\rangle$.

Definition 17. For two vertex rooted hypergraphs $(H, i)$ and $\left(H^{\prime}, i^{\prime}\right)$, define

$$
d_{*}\left((H, i),\left(H^{\prime}, i^{\prime}\right)\right):=\frac{1}{1+m^{*}}
$$

where $m^{*}:=\sup \left\{m \geq 1:(H, i) \equiv_{m}\left(H^{\prime}, i^{\prime}\right)\right\}$, and $m^{*}:=0$ if there is no $m \geq 1$ satisfying this. For two equivalence classes $[H, i] \in \mathcal{H}_{*}$ and $\left[H^{\prime}, i^{\prime}\right] \in \mathcal{H}_{*}$, define $d_{\mathcal{H}_{*}}\left([H, i],\left[H^{\prime}, i^{\prime}\right]\right)$ to be $d_{*}\left((H, i),\left(H^{\prime}, i^{\prime}\right)\right)$ where $(H, i)$ and $\left(H^{\prime}, i^{\prime}\right)$ are arbitrary members of $[H, i]$ and $\left[H^{\prime}, i^{\prime}\right]$, respectively. For two edge-vertex rooted hypergraphs $(H, e, i)$ and $\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$, define

$$
d_{* *}\left((H, e, i),\left(H^{\prime}, e^{\prime}, i^{\prime}\right)\right):=\frac{1}{1+m^{*}}
$$

where $m^{*}:=\sup \left\{m \geq 1:(H, e, i) \equiv_{m}\left(H^{\prime}, e^{\prime}, i^{\prime}\right)\right\}$, and $m^{*}:=0$ if there is no $m \geq 1$ satisfying this. For two equivalence classes $[H, e, i] \in \mathcal{H}_{* *}$ and $\left[H^{\prime}, e^{\prime}, i^{\prime}\right] \in \mathcal{H}_{* *}$, define $d_{\mathcal{H}_{* *}}\left([H, e, i],\left[H^{\prime}, e^{\prime}, i^{\prime}\right]\right)$ to be $d_{* *}\left((H, e, i),\left(H^{\prime}, e^{\prime}, i^{\prime}\right)\right)$ where $(H, e, i)$ and $\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$ are arbitrary members of $[H, e, i]$ and $\left[H^{\prime}, e^{\prime}, i^{\prime}\right]$, respectively.

Since all members of $[H, i]$ are isomorphic, it is not difficult to see that $d_{\mathcal{H}_{*}}$ is well-defined. Note that $(H, i) \equiv_{m}\left(H^{\prime}, i^{\prime}\right)$ and $\left(H^{\prime}, i^{\prime}\right) \equiv_{m^{\prime}}\left(H^{\prime \prime}, i^{\prime \prime}\right)$ implies $(H, i) \equiv_{\min \left\{m, m^{\prime}\right\}}\left(H^{\prime \prime}, i^{\prime \prime}\right)$. Hence $d_{\mathcal{H}_{*}}$ satisfies the triangle inequality. Moreover, $d_{\mathcal{H}_{*}}\left([H, i],\left[H^{\prime}, i^{\prime}\right]\right)=0$ iff $(H, i) \equiv_{m}\left(H^{\prime}, i^{\prime}\right)$ for all $m$, i.e. $[H, i]=\left[H^{\prime}, i^{\prime}\right]$. Hence $d_{\mathcal{H}_{*}}$ defines a metric on $\mathcal{H}_{*}$. We will show in Appendix B that $\mathcal{H}_{*}$ with the metric $d_{\mathcal{H}_{*}}$ is a Polish space (see Corollary 3). Similarly, $d_{\mathcal{H}_{* *}}$ is well defined and gives a metric on $\mathcal{H}_{* *}$. In Appendix B we will also show that $\mathcal{H}_{* *}$ with $d_{\mathcal{H}_{* *}}$ is a Polish space.

Remark 3. One can think of a function $f$ on $\mathcal{H}_{*}$ as a function on vertex rooted hypergraphs which is loyal to the isomorphism relation, i.e. $f((H, i))=f\left(\left(H^{\prime}, i^{\prime}\right)\right)$ whenever $(H, i) \equiv\left(H^{\prime}, i^{\prime}\right)$. This allows us to abuse notation and write $f(H, i)$ instead of $f([H, i])$. We will follow a similar convention for functions on $\mathcal{H}_{* *}$.
Definition 18. $\mathcal{T}_{*}$ and $\mathcal{T}_{* *}$ denote the set of equivalence classes of connected vertex rooted hypertrees and connected edge-vertex rooted hypertrees, respectively. It can be checked that $\mathcal{T}_{*}$ (respectively $\mathcal{T}_{* *}$ ) is a closed subset of $\mathcal{H}_{*}$ (respectively $\mathcal{H}_{* *}$ ).

### 2.4 Operators for total load and average load: $\partial$ and $\nabla$

Definition 19. For a function $f: \mathcal{H}_{* *} \rightarrow \mathbb{R}$, define $\partial f: \mathcal{H}_{*} \rightarrow \mathbb{R}$ as follows. For an equivalence class $[H, i] \in \mathcal{H}_{*}$, pick an arbitrary $\left(H^{\prime}, i^{\prime}\right) \in[H, i]$ and define

$$
\partial f([H, i])=\sum_{e^{\prime} \in E\left(H^{\prime}\right), e^{\prime} \ni i^{\prime}} f\left(\left[H^{\prime}, e^{\prime}, i^{\prime}\right]\right)
$$

Remark 4. Note that in the above definition, $\left[H^{\prime}, e^{\prime}, i^{\prime}\right]$ denotes the equivalence class of edge-vertex rooted hypergraph $\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$. Also, since all the representatives of $[H, i]$ are isomorphic, it is easy to check that the above expression is not dependent on the specific choice of $\left(H^{\prime}, i^{\prime}\right)$. More precisely, if $\left(H_{1}, i_{1}\right)$ and $\left(H_{2}, i_{2}\right)$ are both members of $[H, i]$, then if $\phi: V\left(H_{1}\right) \rightarrow V\left(H_{2}\right)$ is the function establishing the isomorphism, $\phi$ gives a one to one mapping between the set $\left\{e \in E\left(H_{1}\right), e \ni i_{1}\right\}$ and $\left\{e \in E\left(H_{2}\right)\right.$, $e \ni$ $\left.i_{2}\right\}$ and also $\left[H_{1}, e, i_{1}\right] \equiv\left[H_{2}, \phi(e), i_{2}\right]$. Hence

$$
\sum_{e \in E\left(H_{1}\right), e \ni i_{1}} f\left(\left[H_{1}, e, i_{1}\right]\right)=\sum_{e \in E\left(H_{2}\right), e \ni i_{2}} f\left(\left[H_{2}, e, i_{2}\right]\right),
$$

which shows that $\partial f$ is well-defined.
Remark 5. By the above discussion, we may write $\partial f(H, i)=\sum_{e \ni i} f(H, e, i)$, where by $(H, i)$ we mean any arbitrary member of $[H, i]$.
Remark 6. By abuse of notation, we can think of $\partial f$ as a function on $\mathcal{H}_{* *}$ by identifying $\partial f(H, e, i):=$ $\partial f(H, i)$. This will be helpful when we have functions both on $\mathcal{H}_{*}$ and $\mathcal{H}_{* *}$ and want to unify the domain.
Definition 20. For a function $f: \mathcal{H}_{* *} \rightarrow \mathbb{R}$, define the function $\nabla f: \mathcal{H}_{* *} \rightarrow \mathbb{R}$ as follows. Given $[H, e, i] \in \mathcal{H}_{* *}$, take an arbitrary representative $\left(H^{\prime}, e^{\prime}, i^{\prime}\right) \in[H, e, i]$ and define

$$
\nabla f([H, e, i]):=\frac{1}{\left|e^{\prime}\right|} \sum_{j^{\prime} \in e^{\prime}} f\left(\left[H^{\prime}, e^{\prime}, j^{\prime}\right]\right)
$$

Remark 7. As in our discussion in Remark 4, it can be easily checked that the above expression does not depend on the specific choice $\left(H^{\prime}, e^{\prime}, i^{\prime}\right)$. We can therefore abuse notation and write $\nabla f(H, e, i)=$ $\frac{1}{|e|} \sum_{j \in e} f(H, e, j)$.
Definition 21. For a distribution $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$, define

$$
\operatorname{deg}(\mu):=\int_{\mathcal{H}_{*}} \operatorname{deg}_{H}(i) d \mu
$$

### 2.5 From $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ to its directed version $\vec{\mu} \in \mathcal{M}\left(\mathcal{H}_{* *}\right)$

Definition 22. Given $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ with $\operatorname{deg}(\mu)<\infty$, define the measure $\vec{\mu} \in \mathcal{M}\left(\mathcal{H}_{* *}\right)$ as the the one with the property that for any Borel function $f: \mathcal{H}_{* *} \rightarrow[0, \infty)$, we have

$$
\int_{\mathcal{H}_{* *}} f d \vec{\mu}=\int_{\mathcal{H}_{*}} \partial f d \mu .
$$

Note that $\operatorname{deg}(\mu)=\int_{\mathcal{H}_{* *}} 1 d \vec{\mu}=\vec{\mu}\left(\mathcal{H}_{* *}\right)$ is the total mass of $\vec{\mu}$. Hence the assumption $\operatorname{deg}(\mu)<\infty$ guarantees $\vec{\mu}\left(\mathcal{H}_{* *}\right)<\infty$ and so $\vec{\mu} \in \mathcal{M}\left(\mathcal{H}_{* *}\right)$.

This following useful lemma is proved in Appendix C.
Lemma 1. (i) Assume $A \subseteq \mathcal{H}_{*}$ happens $\mu$-almost surely Then $\tilde{A} \subseteq \mathcal{H}_{* *}$ defined as

$$
\tilde{A}:=\left\{[H, e, i] \in \mathcal{H}_{* *}:[H, i] \in A\right\}
$$

happens $\vec{\mu}$-almost everywhere, i.e. $\vec{\mu}(\tilde{A})=\vec{\mu}\left(\mathcal{H}_{* *}\right)$.
(ii) Assume $B \subseteq \mathcal{H}_{* *}$ happens $\vec{\mu}$-almost everywhere. Then, $\tilde{B}:=\left\{[H, i] \in \mathcal{H}_{*}:[H, e, i] \in B \forall e \ni\right.$ i) happens $\mu$-almost surely.

The following fact relating the convergence of a sequence of functions on $\mathcal{H}_{*}$ and that of their counterparts on $\mathcal{H}_{* *}$ will be useful later. This is proved in Appendix C.

Lemma 2. Let $f_{k}: \mathcal{H}_{*} \rightarrow \mathbb{R}, k \geq 0$, be measurable functions such that we have $\lim _{k \rightarrow \infty} f_{k}=f_{0}$, $\mu$-almost surely. Then, if we define their $\mathcal{H}_{* *}$ counterparts $\tilde{f}_{k}: \mathcal{H}_{* *} \rightarrow \mathbb{R}$ via

$$
\tilde{f}_{k}(H, e, i):=f_{k}(H, i)
$$

then we have $\tilde{f}_{k} \rightarrow \tilde{f}_{0}, \vec{\mu}$-almost everywhere.
Another useful lemma is the following one, which relates the convergence of a sequence of functions on $\mathcal{H}_{* *}$ to that of their $\partial$ on $\mathcal{H}_{*}$. This lemma is also proved in Appendix C.
Lemma 3. Given $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ and a sequence of functions $f_{k}: \mathcal{H}_{* *} \rightarrow \mathbb{R}, k \geq 0$, assume that we have $f_{k} \rightarrow f_{0} \vec{\mu}$-almost everywhere. Then we have $\partial f_{k} \rightarrow \partial f_{0} \mu$-almost surely.

### 2.6 Local Weak Convergence

For a finite hypergraph $H$, define $u_{H} \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ by choosing a vertex uniformly at random in $H$ as the root. More precisely, if $i$ is a vertex in $H$ and $H(i)$ denotes the connected component of $i$, define

$$
u_{H}:=\frac{1}{|V(H)|} \sum_{i \in V(H)} \delta_{[H(i), i]} .
$$

The reason why we take the connected component of $H$ is that $\mathcal{H}_{*}$ is the space of equivalence classes of connected vertex rooted hypergraphs.

If, for a sequence of finite hypergraphs $\left\{H_{n}\right\}, u_{H_{n}}$ converges weakly to some measure $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$, we say that $\mu$ is the "local weak limit" of the sequence $H_{n}$. The following lemma is useful in checking when local weak convergence occurs. See Appendix D for a proof.

Lemma 4. Given a sequence $\left\{\mu_{n}\right\}_{n \geq 1}$ in $\mathcal{P}\left(\mathcal{H}_{*}\right)$, and $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ such that $\operatorname{supp}(\mu) \subseteq \mathcal{T}_{*}, \mu_{n} \Rightarrow \mu$ iff the following condition is satisfied: for all $d \geq 1$ and for all rooted hypertrees $(T, i)$ with depth at most $d$, if

$$
A_{(T, i)}:=\left\{[H, j] \in \mathcal{H}_{*}:(H, j)_{d} \equiv(T, i)\right\}
$$

then $\mu_{n}\left(A_{(T, i)}\right) \rightarrow \mu\left(A_{(T, i)}\right)$.

The way to understand $u_{H_{n}} \Rightarrow \mu$ is that the local structure around a uniformly chosen vertex in $H_{n}$, where local means up to a fixed depth, looks more and more similar to that corresponding to $\mu$, hence the term "local" weak convergence. In particular, Lemma 4 says that if we have a sequence of finite hypergraphs $H_{n}$, then $u_{H_{n}} \Rightarrow \mu$, where $\mu \in \mathcal{P}\left(\mathcal{T}_{*}\right)$, if and only if for each $d \geq 1$ and all rooted hypertrees $(T, i)$, if we choose a vertex $v$ in $H_{n}$ uniformly at random, the probability that the local structure of $H_{n}$ rooted at $v$, i.e. $\left(H_{n}, v\right)_{d}$, is isomorphic to $(T, i)$ of depth at most $d$ converges to the probability that the rooted tree with law $\mu$ up to depth $d$ is isomorphic to $T$. See [BS01], [AS04], [AL07], [Bor14] for a review of the notion of local weak convergence in the graph regime.

### 2.7 Balanced allocations with respect to a distribution on $\mathcal{H}_{*}$

Definition 23. A function $\Theta: \mathcal{H}_{* *} \rightarrow[0,1]$ is called a Borel allocation, or just an allocation, if $\Theta$ is a Borel function and also

$$
\nabla \Theta(H, e, i)=\frac{1}{|e|}
$$

From Definition 20 and using our simplified notational conventions, we can equivalently say that $\Theta$ is an allocation precisely when it is a Borel function and $\sum_{i \in e} \Theta(H, e, i)=1$ for the edge-vertex rooted hypergraph $(H, e, i)$. It may seem strange that the condition is required only at the root edge. This will become more clear when we discuss unimodular measures in the next subsection, see especially Proposition 3.

Definition 24. Assume $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$. A Borel allocation $\Theta: \mathcal{H}_{* *} \rightarrow[0,1]$ is called balanced with respect to $\mu$ if for $\vec{\mu}$ almost all $[H, e, i] \in \mathcal{H}_{* *}$ and any $\left(H^{\prime}, e^{\prime}, i^{\prime}\right) \in[H, e, i]$ we have:

$$
\forall j_{1}, j_{2} \in e^{\prime}, \quad \partial \Theta\left(\left[H^{\prime}, e^{\prime}, j_{1}\right]\right)>\partial \Theta\left(\left[H^{\prime}, e^{\prime}, j_{2}\right]\right) \quad \Rightarrow \quad \Theta\left(\left[H^{\prime}, e^{\prime}, j_{1}\right]\right)=0
$$

Remark 8. As in our discussion in Remarks 4 and 7, the above predicate does not depend on the specific choice of $\left(H^{\prime}, e^{\prime}, i^{\prime}\right) \in[H, e, i]$. Hence, by abuse of notation, we may write the above predicate simply as "for $\vec{\mu}$ almost all $(H, e, i)$ and $j_{1}, j_{2} \in e, \partial \Theta\left(H, j_{1}\right)>\partial \Theta\left(H, j_{2}\right)$ implies $\Theta\left(H, e, j_{1}\right)=0$ ".

Similar to our notion of $\epsilon$-balanced allocation for a specific hypergraph, we can define $\epsilon$-balanced allocations with respect to a measure $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$. As in the discussion in Remark 8, we use a simplified language in describing the definition.
Definition 25. Assume $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$. A Borel allocation $\Theta_{\epsilon}: \mathcal{H}_{* *} \rightarrow[0,1]$ is called $\epsilon$-balanced with respect to $\mu$ if for $\vec{\mu}$ almost all $[H, e, i] \in \mathcal{H}_{* *}$ we have

$$
\Theta_{\epsilon}(H, e, i)=\frac{\exp \left(-\partial \Theta_{\epsilon}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-\partial \Theta_{\epsilon}(H, j) / \epsilon\right)}
$$

### 2.8 Unimodularity

Definition 26. A probability measure $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ is called unimodular if for every Borel function $f: \mathcal{H}_{* *} \rightarrow[0, \infty)$ we have

$$
\int f d \vec{\mu}=\int \nabla f d \vec{\mu}
$$

We denote the set of unimodular measures on $\mathcal{H}_{*}$ by $\mathcal{U}$.
See [AL07] for a definition of unimodularity for graphs. It can be easily checked that our definition of unimodularity reduces to the definition in [AL07] when we restrict to graphs, i.e. when we restrict $\mu$ to have support on hypergraphs with all edges having size two.

It can be shown that for a finite hypergraph, $u_{H}$ defined in Section 2.6 is unimodular. Moreover, if a sequence of finite hypergraphs has a local weak limit $\mu$, then $\mu$ is unimodular ${ }^{1}$. See Appendix E for a proof. Roughly speaking, unimodular measures are extensions of the kinds of measures on equivalence classes of vertex rooted hypergraphs that arise from choosing the root uniformly at random in finite hypergraphs.

The following property of unimodular measures is crucial in our analysis. It essentially says that "everything shows at the root" of a unimodular measure. See Appendix E for its proof.

Proposition 3. Assume $\tau: \mathcal{H}_{* *} \rightarrow \mathbb{R}$ and $\mu \in \mathcal{U}$ is a unimodular probability measure such that $\tau=1$ $\vec{\mu}$-almost everywhere. Then there exists some $A \subset \mathcal{H}_{* *}$ such that $\vec{\mu}\left(A^{c}\right)=0$ and

$$
\forall[H, e, i] \in A \quad \tau\left(\left[H, e^{\prime}, i^{\prime}\right]\right)=1, \quad \forall e^{\prime} \in E(H), i^{\prime} \in e^{\prime}, \forall(H, e, i) \in[H, e, i] .
$$

Note that this statement is consistent with our intuition regarding unimodular measures: when some property holds at the root, since the root is chosen "uniformly" and so all vertices have the same "weight", that property should hold everywhere. See Lemma 2.3 in [AL07] for a version of the above statement for graphs.

### 2.9 Unimodular Galton-Watson hypertrees

In this section, we introduce an analogue of Galton Watson processes on graphs for hypertrees. In the graph regime, a Galton-Watson process is defined by generating the degree of the root at random from a given distribution and then, iteratively, the degree of each child is generated at random and so forth. In order to generalize the notion of a Galton-Watson process to hypertrees, since there might exist edges of different sizes, one needs to make sense of the "degree" of edges of each possible size at each node.

To this end, we introduce the notion of type as a generalization of the notion of degree. The type of each node is a vector of integers specifying how many edges of each size a node is connected to. More precisely, since we want all the hypergraphs to be locally finite, we define the set of types, denoted by $\Lambda$, as

$$
\begin{equation*}
\Lambda:=\left\{\gamma \in \mathbb{N}_{0}^{\{2,3, \ldots\}}: \gamma(k)=0, k>k_{0} \text { for some } k_{0} \geq 2\right\} \tag{3}
\end{equation*}
$$

where $\mathbb{N}_{0}:=\mathbb{N} \cup\{0\}$. For a type $\gamma \in \Lambda, \gamma(k)$ determines the number of edges of size $k$ a node is connected to. For instance $(2,1,0,1)$ means a node is connected to 2 edges of size 2 , 1 edge of size 3 and 1 edge of size 5 (we haven't shown the rest of the sequence, which is zero).

For $\gamma \in \Lambda$ define

$$
\begin{equation*}
\|\gamma\|_{1}:=\sum_{k \geq 2} \gamma(k) \tag{4}
\end{equation*}
$$

and

$$
\begin{equation*}
h(\gamma):=\max \{k \geq 2: \gamma(k)>0\} \tag{5}
\end{equation*}
$$

For $k \geq 2$, define $\mathbb{e}_{k} \in \Lambda$ to be the vector with value 1 at coordinate $k$ and zero elsewhere.
Assume $P \in \mathcal{P}(\Lambda)$ is a probability distribution over the set of types, such that $\mathbb{E}[\Gamma(m)]<\infty$ for $m \geq 2$, where $\Gamma$ is a random variable with law $P$. For $m \geq 2$ such that $\mathbb{E}[\Gamma(m)]>0$, we define the size biased distributions $\hat{P}_{m}$ as

$$
\begin{equation*}
\hat{P}_{m}(\gamma)=\frac{(\gamma(m)+1) P\left(\gamma+\mathbb{e}_{m}\right)}{\mathbb{E}[\Gamma(m)]} \tag{6}
\end{equation*}
$$

[^1]
(3, 1, 2, 2, 1, 1)

Figure 4: The tree rooted at $\emptyset$ generated by the sequence $\left\{\gamma_{a}\right\}_{a \in \mathbb{N}_{\text {vertex }}}$ where $\gamma_{\emptyset}=(2,1), \gamma_{(3,1,2)}$, which determines the type of $(3,1,2)$ in the subtree below $(3,1,2)$, is equal to (2) (which results in the single edge of size 2 below $(3,1,2)$ ), and $\gamma_{a}$ is the zero vector for all other nodes $a$.
where it is easy to check that the normalizing term makes $\hat{P}_{m}$ a probability distribution. In case $\mathbb{E}[\Gamma(m)]=0$, or equivalently $\Gamma(m)=0$ with probability one, we define $\hat{P}_{m}$ to be an arbitrary distribution, e.g. $\hat{P}_{m}(\gamma)=1$ when $\gamma$ is the type with all coordinates being zero.

Let $\emptyset$ denote the root of the Galton-Watson hypertree and let $\mathbb{N}_{\text {vertex }}$ denote the set of vertices, so $\emptyset \in \mathbb{N}_{\text {vertex }}$. Let $\mathbb{N}_{\text {edge }}$ denote the set of hyperedges of the Galton-Watson hypertree. Each non-root element of $\mathbb{N}_{\text {vertex }}$ will be of the form $\left(s_{1}, e_{1}, i_{1}, \ldots, s_{k}, e_{k}, i_{k}\right)$ where $s_{j} \geq 2, e_{j} \geq 1$, and $1 \leq i_{j} \leq s_{j}-1$ for all $1 \leq j \leq k$. The semantics of $\left(s_{1}, e_{1}, i_{1}\right)$ is that it is the vertex numbered $i_{1}$ of the $e_{1}$-th copy of a hyperedge of size $s_{1}$ attached to the root, and so on. Thus, for example $(3,5,2,5,8,3)$ represents the vertex numbered 3 of the eight hyperedge of size 5 that attaches to the vertex labeled 2 of the fifth hyperedge of size 3 that attaches to the root. The elements of $\mathbb{N}_{\text {edge }}$ are thus of the form $\left(s_{1}, e_{1}, i_{1}, \ldots, s_{k}, e_{k}\right)$, where $s_{j} \geq 2, e_{j} \geq 1$, and $1 \leq i_{j} \leq s_{j}-1$ for all $1 \leq j \leq k-1$.

Given a sequence of types $\left\{\gamma_{a}\right\}_{a \in \mathbb{N}_{\text {vertex }}}$, we can construct a hypertree with vertex set and edge set $\mathbb{N}_{\text {vertex }}$ and $\mathbb{N}_{\text {edge }}$, respectively, where for $a \in \mathbb{N}_{\text {vertex }}, \gamma_{a}$ determines the type of the node $a$ in the subtree below node $a$. See Figure 4 for an example.

Definition 27. Let $P \in \mathcal{P}(\Lambda)$ such that $\mathbb{E}[\Gamma(m)]<\infty$ for $m \geq 2$, where $\Gamma$ has the distribution $P$. Construct a random rooted tree $(T, \emptyset)$ by generating $\left(\Gamma_{a}, a \in \mathbb{N}_{\text {vertex }}\right)$ independently such that $\Gamma_{\emptyset}$ has law $P$ and for any non-root node $a=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{k}, e_{k}, i_{k}\right), \Gamma_{a}$ has law $\hat{P}_{s_{k}}$. Then $U G W T(P)$ is the law of $[((T, \emptyset), \emptyset)]$ where $[((T, \emptyset), \emptyset)]$ denotes the equivalence class of $((T, \emptyset), \emptyset)$ in $\mathcal{H}_{*}$, with $(T, \emptyset)$ being the connected component of $\emptyset$ in $T$.

One important observation is that $\operatorname{UGWT}(P)$, as in Definition 27, is unimodular. See Appendix F for the proof.

Proposition 4. Assume $P \in \mathcal{P}(\Lambda)$ is a distribution over types such that $\mathbb{E}[\Gamma(k)]<\infty$ for $k \geq 2$, then $U G W T(P)$ is unimodular.

Note that if $(T, \emptyset)$ is generated as in Definition 27 and $\left(s_{1}, e_{1}, i_{1}\right)$ is a child of the root present in $T$, the subtree rooted at $i_{1}$ has a similar distribution to $(T, \emptyset)$ except that the type of its root has law $\hat{P}_{s_{1}}$. It is useful for our subsequent discussion to define a notation for this distribution.

Definition 28. Let $P \in \mathcal{P}(\Lambda)$ such that $\mathbb{E}[\Gamma(l)]<\infty$ for $l \geq 2$ and fix some $k \geq 2$. Construct a random rooted tree $(T, \emptyset)$ by generating $\left(\Gamma_{a}, a \in \mathbb{N}_{\text {vertex }}\right)$ independently such that $\Gamma_{\emptyset}$ has law $\hat{P}_{k}$ and for any non-root node $a=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{r}, e_{r}, i_{r}\right), \Gamma_{a}$ has law $\hat{P}_{s_{r}}$. Then $G W T_{k}(P)$ denotes the law of $[(T, \emptyset), \emptyset]$ where $[(T, \emptyset), \emptyset]$ denotes the equivalence class of $((T, \emptyset), \emptyset)$ in $\mathcal{H}_{*}$.

### 2.10 Equivalence classes of marked hypergraphs: $\overline{\mathcal{H}}_{*}(\Xi)$ and $\overline{\mathcal{H}}_{* *}(\Xi)$

Recall that $\mathcal{H}_{*}$ and $\mathcal{H}_{* *}$ are Polish spaces of isomorphism classes of vertex rooted hypergraphs and edge-vertex rooted hypergraphs respectively. We can extend the procedure by which these spaces were created to hypergraphs with marks on their edges. Hypergraphs with marks on their edges would be called "hypernetworks", following the terminology in Aldous and Lyons [AL07]. However, we prefer to call them marked hypergraphs.

Definition 29. Assume $H$ is a locally finite simple hypergraph on a countable vertex set and $\Xi$ is a complete separable metric space. $A \Xi$-valued edge mark on $H$ is a function

$$
\xi: \Psi(H) \rightarrow \Xi
$$

where we recall that $\Psi(H)$ denotes the set of edge-vertex pairs of $H$. A hypergraph carrying such a mark is called a marked hypergraph, and is denoted $(H, \zeta)$.

A vertex rooted marked hypergraph is a marked hypergraph $(H, \xi)$ together with a distinguished vertex $i \in V(H)$, and is denoted $((H, \xi), i)$. An edge-vertex rooted marked hypergraph is a marked hypergraph $(H, \xi)$ together with a distinguished edge $e \in E(H)$, and a distinguished vertex $i \in e$. It is denoted $((H, \xi), e, i)$.

Remark 9. To simplify the notation, we employ the notation $\bar{H}$ to denote a marked hypergraph, where its mark function is denoted by $\xi_{\bar{H}}$, and its underlying unmarked hypergraph is denoted by $H$. Moreover, we may use $V(\bar{H}), E(\bar{H})$ and $\Psi(\bar{H})$ instead of $V(H), E(H)$ and $\Psi(H)$.

Definition 30. We call two vertex rooted marked hypergraphs $\left(\bar{H}_{1}, i_{1}\right)$ and $\left(\bar{H}_{2}, i_{2}\right)$ isomorphic, and write $\left(\bar{H}_{1}, i_{1}\right) \equiv\left(\bar{H}_{2}, i_{2}\right)$ if there is a bijection $\phi: V\left(\bar{H}_{1}\right) \rightarrow V\left(\bar{H}_{2}\right)$, such that $\phi\left(i_{1}\right)=i_{2}$, and $e \in E\left(\bar{H}_{1}\right)$ iff $\phi(e) \in E\left(\bar{H}_{2}\right)$, where $\phi(e):=\{\phi(i): i \in e\}$. Moreover, for $(e, i) \in \Psi\left(\bar{H}_{1}\right)$, we require that $\xi_{\bar{H}_{1}}(e, i)=\xi_{\bar{H}_{2}}(\phi(e), \phi(i))$.

We say two edge-vertex rooted marked hypergraphs $\left(\bar{H}_{1}, e_{1}, i_{1}\right)$ and $\left(\bar{H}_{2}, e_{2}, i_{2}\right)$ are isomorphic, and write $\left(\bar{H}_{1}, e_{1}, i_{1}\right) \equiv\left(\bar{H}_{2}, e_{2}, i_{2}\right)$, if there is a bijection $\phi: V\left(\bar{H}_{1}\right) \rightarrow V\left(\bar{H}_{2}\right)$, such that $\phi\left(i_{1}\right)=i_{2}, \phi\left(e_{1}\right)=$ $e_{2}$, and such that $e \in E\left(\bar{H}_{1}\right)$ iff $\phi(e) \in E\left(\bar{H}_{2}\right)$. Moreover, we must have $\xi_{\bar{H}_{1}}(e, i)=\xi_{\bar{H}_{2}}(\phi(e), \phi(i))$ for $(e, i) \in \Psi\left(\bar{H}_{1}\right)$.
Definition 31. Let $\overline{\mathcal{H}}_{*}(\Xi)$ be the space of equivalence classes of connected vertex rooted marked hypergraphs with marks taking values in $\Xi$, with the equivalence class of $(\bar{H}, i)$ being denoted $[\bar{H}, i]$. Similarly, let $\overline{\mathcal{H}}_{* *}(\Xi)$ be the space of equivalence classes of edge-vertex rooted marked hypergraphs with marks taking values in $\Xi$, with the equivalence class of ( $\bar{H}, e, i$ ) being denoted $[\bar{H}, e, i]$.

We endow $\overline{\mathcal{H}}_{*}(\Xi)$ with the metric $\bar{d}_{*}$ where the distance between $\left[\bar{H}_{1}, i_{1}\right]$ and $\left[\bar{H}_{2}, i_{2}\right]$ is defined in the following way: take arbitrary representatives $\left(\bar{H}_{1}^{\prime}, i_{1}^{\prime}\right) \in\left[\bar{H}_{1}, i_{1}\right]$ and $\left(\bar{H}_{2}^{\prime}, i_{2}^{\prime}\right) \in\left[\bar{H}_{2}, i_{2}\right]$, then let $m^{*}$ be the supremum over all $m$ such that $\left(H_{1}^{\prime}, i_{1}^{\prime}\right) \equiv_{m}\left(H_{2}^{\prime}, i_{2}^{\prime}\right)$, and the $\Xi$-distance between the corresponding marks up to level $m$ is at most $1 / m$, i.e. if $\phi$ is the level $m$ isomorphism, then

$$
d_{\Xi}\left(\xi_{\bar{H}_{1}}(\tilde{e}, \tilde{i}), \xi_{\bar{H}_{2}}(\phi(\tilde{e}), \phi(\tilde{i}))\right) \leq \frac{1}{m}, \quad \forall \tilde{e} \in E_{H_{1}^{\prime}}\left(V_{i_{1}^{\prime}, m}^{H_{1}^{\prime}}\right)
$$

where $d_{\Xi}$ denotes the metric on $\Xi$. If there is no $m$ satisfying the above conditions, we set $m^{*}$ to be 0 . Then, $\bar{d}_{*}\left(\left[\bar{H}_{1}, i_{1}\right],\left[\bar{H}_{2}, i_{2}\right]\right)$ is defined to be $1 /\left(1+m^{*}\right)$. Since all the members in $[\bar{H}, i]$ are isomorphic as vertex rooted marked hypergraphs, $\bar{d}_{*}$ can be easily checked to be well defined. One can also check that it is a metric; in particular it satisfies the triangle inequality.

Similarly, we endow $\overline{\mathcal{H}}_{* *}(\Xi)$ with the metric $\bar{d}_{* *}$ where the distance between $\left[\bar{H}_{1}, e_{1}, i_{1}\right]$ and $\left[\bar{H}_{2}, e_{2}, i_{2}\right]$ is defined in the following way: take arbitrary representatives $\left(\bar{H}_{1}^{\prime}, e_{1}^{\prime}, i_{1}^{\prime}\right) \in\left[\bar{H}_{1}, e_{1}, i_{1}\right]$ and
$\left(\bar{H}_{2}^{\prime}, e_{2}^{\prime}, i_{2}^{\prime}\right) \in\left[\bar{H}_{2}, e_{2}, i_{2}\right]$, then let $m^{* *}$ be the supremum over all $m$ such that $\left(H_{1}^{\prime}, e_{1}^{\prime}, i_{1}^{\prime}\right) \equiv_{m}\left(H_{2}^{\prime}, i_{2}^{\prime}, e_{2}^{\prime}\right)$ and the $\Xi$-distance between the corresponding marks up to level $m$ is at most $1 / m$. If there is no $m$ satisfying these conditions, we set $m^{*}$ to be 0 . Finally, define $\bar{d}_{* *}\left(\left[\bar{H}_{1}, e_{1}, i_{1}\right],\left[\bar{H}_{2}, e_{2}, i_{2}\right]\right)$ to be $1 /\left(1+m^{* *}\right)$.

In Appendix B we prove that $\overline{\mathcal{H}}_{*}(\Xi)$ and $\overline{\mathcal{H}}_{* *}(\Xi)$ with their respective metrics are Polish spaces, see Proposition 19.

Similar to what we did for $\mathcal{H}_{*}$ and $\mathcal{H}_{* *}$, we can define $\partial$ and $\nabla$ operators as follows, where we use a simplified notation, whose validity can be justified as in Remarks 5 and 7:

$$
\partial f(\bar{H}, i):=\sum_{e \in E(\bar{H}), e \ni i} f(\bar{H}, e, i),
$$

and

$$
\nabla f(\bar{H}, e, i):=\frac{1}{|e|} \sum_{j \in e} f(\bar{H}, e, j)
$$

Remark 10. The preceding notation, strictly speaking, applies only to real valued functions on $\mathcal{H}_{* *}$. However, if we consider the function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow \Xi$ defined as $f(\bar{H}, e, i)=\xi_{\bar{H}}(e, i)$, then, when $\Xi$ has an additive structure, we may define $\partial f: \overline{\mathcal{H}}_{*}(\Xi) \rightarrow \Xi$ as

$$
\partial f(\bar{H}, i):=\sum_{e \in E(H), e \ni i} f(\bar{H}, e, i)=\sum_{e \in E(\bar{H}), e \ni i} \xi_{\bar{H}}(e, i) .
$$

By abuse of notation, we may use the notation $\partial \xi_{\bar{H}}(i)$ instead of $\partial f(\bar{H}, i)$ with $f$ defined above, and similarly for $\nabla \xi_{\bar{H}}(i)$. We will use such notation in this document because the marks we are interested in will be real valued.

For a probability measure $\mu \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}(\Xi)\right)$, we define $\vec{\mu} \in \mathcal{M}\left(\overline{\mathcal{H}}_{* *}(\Xi)\right)$ in a manner similar to what was done in Section 2.5. Namely, $\vec{\mu} \in \mathcal{M}\left(\overline{\mathcal{H}}_{* *}(\Xi)\right)$ is defined by requiring that for every nonnegative Borel function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow[0, \infty)$ we have

$$
\int f d \vec{\mu}=\int \partial f d \mu
$$

A probability measure $\mu \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}(\Xi)\right)$ is called unimodular if for every nonnegative Borel function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow[0, \infty)$ we have

$$
\int f d \vec{\mu}=\int \nabla f d \vec{\mu}
$$

By removing marks, we get a natural projection from $\Pi: \overline{\mathcal{H}}_{*}(\Xi) \rightarrow \mathcal{H}_{*}$ defined as

$$
\begin{equation*}
\Pi([\bar{H}, i])=[H, i] \tag{7}
\end{equation*}
$$

This can easily be checked to be continuous.
As was done in Section 2.8, choosing a vertex uniformly at random from a finite marked hypergraph results in a unimodular measure. More precisely, if $H$ is a finite hypergraph together with a mark $\xi$ taking values in $\Xi$,

$$
u_{\bar{H}}:=\frac{1}{|V(\bar{H})|} \sum_{i \in V(\bar{H})} \delta_{[\bar{H}(i), i]}
$$

is unimodular. Here, $\bar{H}(i)$ denotes the connected component of $i$ in $\bar{H}$. Moreover, the local weak limit of finite marked hypergraphs, i.e. the weak limit of the measures $u_{H_{n}}$, is unimodular, if it exists. See Appendix E for a proof.

## 3 Main Results

We now summarize the main results of the paper. We first prove some properties of balanced Borel allocations, as defined in Definition 24.

Theorem 1. Let $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ be a unimodular probability measure such that $\operatorname{deg}(\mu)<\infty$. The following are true.

1. There exists a Borel allocation $\Theta: \mathcal{H}_{* *} \rightarrow[0,1]$ which is balanced with respect to $\mu$.
2. Let $\Theta$ be a balanced Borel allocation with respect to $\mu$. Then we have the following variational characterization of the mean excess load under $\Theta$ above the load level $t$. For any $t \in \mathbb{R}$ :

$$
\int(\partial \Theta-t)^{+} d \mu=\max _{\substack{f: \mathcal{H} * \rightarrow[0,1] \\ \text { Borel }}} \int \tilde{f}_{\text {min }} d \vec{\mu}-t \int f d \mu
$$

where $\tilde{f}_{\text {min }}$ is defined as

$$
\tilde{f}_{\text {min }}(H, e, i):=\frac{1}{|e|} \min _{j \in e} f(H, j)
$$

3. The following are equivalent for a Borel allocation $\Theta: \mathcal{H}_{* *} \rightarrow[0,1]$ :
(a) $\Theta$ is balanced with respect to $\mu$.
(b) $\Theta$ minimizes $\int f \circ \partial \Theta d \mu$ among all Borel allocations, for some strictly convex function $f:[0, \infty) \rightarrow[0, \infty)$.
(c) $\Theta$ minimizes $\int f \circ \partial \Theta d \mu$ among all Borel allocations for every convex function $f:[0, \infty) \rightarrow$ $[0, \infty)$.
4. Assume $\Theta_{0}$ is a balanced allocation with respect to $\mu$ and $\Theta$ is any other allocation on $\mathcal{H}_{* *}$. Then $\Theta$ is balanced if and only if $\partial \Theta_{0}=\partial \Theta, \mu$-almost surely.
5. Let $\left\{H_{n}\right\}_{n \geq 1}$ be a sequence of finite hypergraphs with local weak limit $\mu$. Let $\mathcal{L}_{H_{n}}$ denote the distribution of the total load at a vertex in $H_{n}$ chosen uniformly at random. Namely, $\mathcal{L}_{H_{n}}=$ $\frac{1}{\left|V\left(H_{n}\right)\right|} \sum_{i \in V\left(H_{n}\right)} \delta_{\partial \theta_{n}(i)}$, where $\left(\partial \theta_{n}(i), i \in V\left(H_{n}\right)\right)$ denotes the load vector corresponding to any balanced allocation $\theta_{n}$, which we recall exists and is unique, due to Proposition 1. Let $\mathcal{L}$ denote the law of the total load at the root of the balanced allocation on $\mu$, i.e. the pushforward of $\mu$ under the mapping $\partial \Theta$, which we have just shown is well defined and unique, due to parts 1 and 4 of this theorem. Then $\mathcal{L}_{H_{n}}$ converges weakly to $\mathcal{L}$.

The proof of the above theorem is given in Section 6. Before that, we first investigate, in Section 4, the properties of $\epsilon$-balanced allocations for a specific hypergraph, as introduced in Definition 13. We then investigate, in Section 5, the properties of $\epsilon$-balanced allocations on $\mathcal{H}_{* *}$ for a given $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$, as defined in Definition 25. Then, by sending $\epsilon$ to zero, we prove part 1 of the above theorem in Section 6.1. The proofs of other parts of the theorem are given in Sections 6.2 through 6.5 respectively.

Note that, as a result of part 4 of the theorem, for a given $t \in \mathbb{R}$, the value of the integral $\int(\partial \Theta-t)^{+} d \mu$ for a balanced allocation $\Theta$ does not depend on the particular choice of $\Theta$. It only depends on $\mu$ and $t$, so it can be written as

$$
\begin{equation*}
\Phi_{\mu}(t):=\int(\partial \Theta-t)^{+} d \mu \tag{8}
\end{equation*}
$$

The function $\Phi_{\mu}$ is called the mean-excess function. Knowledge of $\Phi_{\mu}$ is equivalent to determining the distribution $\mathcal{L}$ of the load at the root associated to a balanced allocation $\Theta$. We now describe $\Phi_{\mu}$ for the class of unimodular Galton-Watson process defined in Section 2.9.

Recall the notation $\Lambda$ defined in (3). Assume $P \in \mathcal{P}(\Lambda)$ and $t \in \mathbb{R}$ are fixed. For a sequence of Borel probability measures $\left(Q_{l}, l \geq 2\right)$ on real numbers, let $F_{P, t}^{(k)}\left(\left\{Q_{l}\right\}_{l \geq 2}\right)$ be the distribution of the random variable

$$
\begin{equation*}
t-\sum_{k^{\prime} \geq 2} \sum_{i=1}^{\Gamma\left(k^{\prime}\right)}\left[1-X_{k^{\prime}, i, 1}^{+}-\cdots-X_{k^{\prime}, i, k^{\prime}-1}^{+}\right]_{0}^{1} \tag{9}
\end{equation*}
$$

where $\Gamma$ has law $\hat{P}_{k}$, and $X_{k^{\prime}, i, j}$ are random variables which are mutually independent and independent of $\Gamma$, with $X_{k^{\prime}, i, j}$ having law $Q_{k^{\prime}}$. Note that $\hat{P}_{k}$ is the size biased version of $P$ defined in (6). Also note that the first sum on the right hand side of (9) is a finite sum, because $\Gamma$ has finite support, pointwise.

Let $\mathcal{Q}$ be the set of sequences $\left\{Q_{l}\right\}_{l \geq 2}$ such that, for all $k \geq 2$, we have:

$$
\begin{equation*}
Q_{k}=F_{P, t}^{(k)}\left(\left\{Q_{l}\right\}_{l \geq 2}\right) \tag{10}
\end{equation*}
$$

Now, we are ready to provide a characterization of the mean excess function. We give the proof of the following result in Section 8.
Theorem 2. Let $P$ be a distribution on $\Lambda$ such that $\mathbb{E}\left[\|\Gamma\|_{1}\right]<\infty$ where $\Gamma$ has law $P$. Then, with $\mu:=U G W T(P)$, for any $t \in \mathbb{R}$, we have

$$
\begin{equation*}
\Phi_{\mu}(t)=\max _{\left\{Q_{k}\right\}_{k \geq 2} \in \mathcal{Q}}\left(\sum_{k=2}^{\infty} \frac{\mathbb{E}[\Gamma(k)]}{k} \mathbb{P}\left(\sum_{i=1}^{k} X_{k, i}^{+}<1\right)\right)-t \mathbb{P}\left(\sum_{k=2}^{h(\Gamma)} \sum_{i=1}^{\Gamma(k)} Y_{k, i}>t\right) \tag{11}
\end{equation*}
$$

where, in the first expression, $\Gamma$ is a random variable on $\Lambda$ with law $P$ and $\left\{X_{k, i}\right\}_{k, i}$ are i.i.d. such that $X_{k, i}$ has law $Q_{k}$. Also, in the second expression, $\Gamma$ has law $P$ and $\left\{Y_{k, i}\right\}_{k, i}$ are independent from each other and from $\Gamma$, with $Y_{k, i}$ having the law of the random variable $\left[1-\left(Z_{1}^{+}+\cdots+Z_{k-1}^{+}\right)\right]_{0}^{1}$, where $Z_{j}$ are i.i.d. with law $Q_{k}$.

For a finite hypergraph $H$, we define $\varrho(H)$ to be the maximum load corresponding to a balanced allocation on $H$, i.e. if $\theta$ is a balanced allocation on $H$,

$$
\begin{equation*}
\varrho(H):=\max _{v \in V(H)} \partial \theta(v) \tag{12}
\end{equation*}
$$

which is well defined due to Proposition 1. From [Haj90, Corollary 7] we know that there is a duality between this parameter and the subgraph of maximum edge density, i.e.

$$
\begin{equation*}
\varrho(H)=\max _{S \subseteq V(H), S \neq \emptyset} \frac{\left|E_{H}(S)\right|}{|S|} \tag{13}
\end{equation*}
$$

where $E_{H}(S)$ denotes the set of edges of $H$ with all endpoints in $S$.
For a unimodular probability distribution $\mu$ on $\mathcal{H}_{*}$ with finite $\operatorname{deg}(\mu)$, we define

$$
\begin{equation*}
\varrho(\mu):=\sup \left\{t \in \mathbb{R}: \Phi_{\mu}(t)>0\right\} \tag{14}
\end{equation*}
$$

where $\Phi_{\mu}($.$) is the mean excess function defined above. In other words, if \Theta$ is the balanced allocation corresponding to $\mu$ introduced in Theorem 1 and $\mathcal{L}_{\mu}$ is the law of $\partial \Theta$ under $\mu$, then

$$
\varrho(\mu)=\sup \left\{t \in \mathbb{R}: \mathcal{L}_{\mu}([t, \infty))>0\right\}
$$

One question is whether local weak convergence implies convergence of maximum load, i.e., if $H_{n}$ is a sequence of graphs with local weak limit $\mu$, does $\varrho\left(H_{n}\right)$ converge to $\varrho(\mu)$ ? Similar to the graph case, this is not true in general, since we can always add an arbitrary but bounded clique to boost $\varrho\left(H_{n}\right)$ without changing the local weak limit. We prove convergence, under some conditions, for the special case where the limit $\mu$ is the UGWT model defined in Section 2.9 and, for each $n, H_{n}$ is a random hypergraph obtained from a generalized hypergraph configuration model defined in Section 9.1.

Theorem 3. Let $P$ be a probability distribution on $\Lambda$ such that, if $\Gamma$ is a random variable with law $P$, $P(\Gamma(k)>0)>0$ for finitely many $k$ and $\mathbb{E}[\Gamma(k)]<\infty$ for all $k \geq 2$. Moreover, let $\mu:=U G W T(P)$. Then, if $\left\{H_{n}\right\}_{n=1}^{\infty}$ is a sequence of random hypergraphs obtained from a configuration model, under some conditions stated in Proposition 18, $\varrho\left(H_{n}\right)$ converges in probability to $\varrho(\mu)$.

This theorem is proved in Section 9.

## $4 \quad \epsilon$-balancing with baseloads

In this section, we analyze the properties of $\epsilon$-balanced allocations with respect to a baseload, which were introduced in Definition 15. Note that throughout this section, we are dealing with a given hypergraph, not a distribution on $\mathcal{H}_{*}$. By setting the baseload function $b$ to zero, our results here reduce to those for $\epsilon$-balanced allocations as introduced in Definition 13.

### 4.1 Existence

The existence of $\epsilon$-balanced allocations with respect to a baseload $b$ on hypergraphs is a consequence of the Schauder-Tychonoff fixed point theorem (see, for instance, [AMO09]). Here we give the details. Fix a hypergraph $H=\langle V, E\rangle$ and define the topological vector space $W$ to be:

$$
W:=\{\theta: \Psi(H) \rightarrow \mathbb{R}\}=\mathbb{R}^{\Psi(H)}
$$

with the product topology of $\mathbb{R}$. Here, we recall that $\Psi(H)$ is the set of all edge-vertex pairs of the hypergraph (1), and that this is a countable set. Define the following convex subset of functions with values in $[0,1]$ :

$$
A:=\{\theta: \Psi(H) \rightarrow[0,1]\} .
$$

Since we have employed the product topology, Tychonoff's theorem tells us that $A$ is a compact set (see, for instance, [Mun00]). Define the mapping $T: A \rightarrow A$ via:

$$
(T \theta)(e, i):=\frac{\exp \left(-\frac{\partial_{b} \theta(i)}{\epsilon}\right)}{\sum_{j \in e} \exp \left(-\frac{\partial_{b} \theta(j)}{\epsilon}\right)}
$$

We want to show that $T$ has a fixed point. In order to do so, we need to show that $T$ is continuous. Since we have employed the product topology, we need to show that, for all $(e, i) \in \Psi(H)$, the projected version $T_{e, i}$ defined as:

$$
T_{e, i}(\theta):=(T \theta)(e, i)
$$

which is a mapping from $A$ to $[0,1]$, is continuous. In order to show this, note that $T_{e, i}$ is the concatenation of a projection $\Pi_{e}: A \rightarrow \mathbb{R}^{\left|U_{e}\right|}$, where

$$
U_{e}:=\left\{\left(e^{\prime}, j\right): e^{\prime} \in E, e^{\prime} \cap e \neq \emptyset, j \in e \cap e^{\prime}\right\}
$$

and an addition function from $\mathbb{R}^{\left|U_{e}\right|}$ to $\mathbb{R}^{|e|}$, which gives us the vector $[\partial \theta(j)]_{j \in e}$, and then a function $f: \mathbb{R}^{|e|} \rightarrow \mathbb{R}$ defined as:

$$
f\left(\left[x_{j}\right]_{j \in e}\right):=\frac{e^{-\left(x_{i}+b(i)\right) / \epsilon}}{\sum e^{-\left(x_{j}+b(j)\right) / \epsilon}}
$$

Since all these three functions are continuous (note that $U_{e}$ is a finite set since we have assumed that the graph is locally finite and all edges have finite size), $T$ is also continuous. Therefore, since $W$ is Hausdorff and locally convex, $A$ is compact, and $T$ is continuous, the Schauder-Tychonoff fixed point theorem implies that $T$ has a fixed point (see, for instance, [AMO09, Theorem 8.2]). Note that $\theta^{\prime}:=T(\theta)$ satisfies $\sum_{i \in e} \theta^{\prime}(e, i)=1$ for any $\theta \in W$. Therefore this fixed point is an allocation in the sense of Definition 10, and is also $\epsilon$-balanced.

### 4.2 Monotony and Uniqueness

Intuitively, we expect that when we add more edges to a hypergraph and increase baseloads, the total load for an $\epsilon$-balanced allocation would increase. We also expect that the effect of an increase in baseload at any vertex tends to dissipate as one moves away from the vertex, when comparing the respective balanced allocations. Lemmas 5 and 7 below quantify these phenomena. They are formulated in the language of vertex rooted hypergraph embedding from Definition 8 .

Lemma 5 (depth 1 local contraction). Assume the vertex rooted hypergraph $(H, i)$ can be embedded up to depth 1 into the vertex rooted hypergraph $\left(H^{\prime}, i^{\prime}\right)$, i.e. $(H, i) \hookrightarrow_{1}\left(H^{\prime}, i^{\prime}\right)$, with embedding $\phi$ : $V_{i, 1}^{H} \hookrightarrow V_{i^{\prime}, 1}^{H^{\prime}}$. Let $\theta_{\epsilon}$ and $\theta_{\epsilon}^{\prime}$ be $\epsilon$-balanced allocations on $H$ and $H^{\prime}$ respectively, with respective baseload functions $b$ and $b^{\prime}$, with $b(i) \leq b^{\prime}\left(i^{\prime}\right)$. If

$$
M:=\max _{j: d_{H}(i, j)=1} \partial_{b} \theta_{\epsilon}(j)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(\phi(j)),
$$

then we have

$$
\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right) \geq \partial_{b} \theta_{\epsilon}(i)-\frac{\left|D_{i, 1}^{H}\right|}{\left|D_{i, 1}^{H}\right|+4 \epsilon} M^{+}
$$

where $D_{i, 1}^{H}$ is the set of nodes at distance one from node $i$ as was defined in Definition 4.
Note that, in this lemma, $\theta_{\epsilon}$ and $\theta_{\epsilon}^{\prime}$ are two arbitrary $\epsilon$-balanced allocations on $H$ and $H^{\prime}$ respectively, with respective baseload functions $b$ and $b^{\prime}$. We know from Section 4.1 that such allocations exist, but they might a priori not be unique. We will later prove uniqueness for the special case of bounded hypergraphs, which were introduced in Definition 2.

Before proving this result, we need the following tool, whose proof is given after the proof of Lemma 5.

Lemma 6. Assume that for $\epsilon>0$, the function $f_{\epsilon}: \mathbb{R}^{k} \rightarrow \mathbb{R}$ is defined in the following way:

$$
f_{\epsilon}\left(x_{1}, \ldots, x_{k}\right)=\frac{1}{1+\sum_{i=1}^{k} e^{-\frac{x_{i}}{\epsilon}}}
$$

Then, for arbitrary real valued sequences $\left(x_{1}, \ldots, x_{k}\right)$ and $\left(x_{1}^{\prime}, \ldots, x_{k}^{\prime}\right)$, we have,

$$
f_{\epsilon}\left(x_{1}, \ldots, x_{k}\right)-f_{\epsilon}\left(x_{1}^{\prime}, \ldots, x_{k}^{\prime}\right) \leq \frac{1}{4 \epsilon} \sum_{i=1}^{k}\left[x_{i}-x_{i}^{\prime}\right]^{+}
$$

Proof of Lemma 5. Since $\phi(e) \in E\left(H^{\prime}\right)$ for $e \ni i$, and $\theta_{\epsilon}^{\prime}$ is a nonnegative function, we have

$$
\begin{equation*}
\sum_{e^{\prime} \ni i^{\prime}} \theta_{\epsilon}^{\prime}\left(e^{\prime}, i^{\prime}\right) \geq \sum_{e \ni i} \theta_{\epsilon}^{\prime}\left(\phi(e), i^{\prime}\right) \tag{15}
\end{equation*}
$$

On the other hand, we have

$$
\begin{align*}
& \partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right) \stackrel{(a)}{\leq} \sum_{e \ni i} \theta_{\epsilon}(e, i)-\sum_{e^{\prime} \ni i^{\prime}} \theta_{\epsilon}^{\prime}\left(e^{\prime}, i^{\prime}\right) \\
& \stackrel{(b)}{\leq} \sum_{e \ni i} \theta_{\epsilon}(e, i)-\theta_{\epsilon}^{\prime}\left(\phi(e), i^{\prime}\right) \\
& \stackrel{(c)}{=} \sum_{e \ni i}\left(\frac{1}{1+\sum_{\substack{j \in e \\
j \neq i}} \exp \left(-\frac{\partial_{b} \theta_{\epsilon}(j)-\partial_{b} \theta_{\epsilon}(i)}{\epsilon}\right)}\right.  \tag{16}\\
&\left.-\frac{1}{1+\sum_{\substack{j \in e \\
j \neq i}} \exp \left(-\frac{\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(\phi(j))-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right)}{\epsilon}\right)}\right) \\
& \stackrel{(d)}{\leq} \frac{1}{4 \epsilon} \sum_{\substack{ }} \sum_{\substack{j \in e \\
j \neq i}}\left[\left(\partial_{b} \theta_{\epsilon}(j)-\partial_{b} \theta_{\epsilon}(i)\right)-\left(\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(\phi(j))-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right)\right)\right]^{+}
\end{align*}
$$

where $(a)$ results from $b(i) \leq b^{\prime}\left(i^{\prime}\right)$, (b) uses (15), (c) is a substitution from Definition 14, and (d) uses Lemma 6.

Now, let

$$
I:=\left\{(e, j): e \ni i, j \in e, j \neq i, \partial_{b} \theta_{\epsilon}(j)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(\phi(j)) \geq \partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right)\right\}
$$

Then the inequality in (16) together with the definition of $M$ implies

$$
\partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right) \leq \frac{1}{4 \epsilon}|I|\left(M-\left(\partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right)\right)\right) .
$$

Rearranging the terms, we get

$$
\partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right) \leq \frac{|I|}{|I|+4 \epsilon} M \leq \frac{\left|D_{i, 1}^{H}\right|}{\left|D_{i, 1}^{H}\right|+4 \epsilon} M^{+}
$$

which is exactly what we wanted to prove.
Proof of Lemma 6. First we prove the statement for $k=1$. In this case, the function $f_{\epsilon}(x)=\frac{1}{1+e^{-\frac{x}{\epsilon}}}$ is $\frac{1}{4 \epsilon}$-Lipschitz and increasing in $x$, hence the statement holds for $k=1$.

Now assume $k>1$ is arbitrary. We have

$$
\begin{equation*}
f_{\epsilon}\left(x_{1}, \ldots, x_{k}\right)-f_{\epsilon}\left(x_{1}^{\prime}, \ldots, x_{k}^{\prime}\right)=\sum_{i=1}^{k} \frac{e^{-\frac{x_{i}^{\prime}}{\epsilon}}-e^{-\frac{x_{i}}{\epsilon}}}{\left(1+\sum_{r=1}^{k} e^{-\frac{x_{r}}{\epsilon}}\right)\left(1+\sum_{s=1}^{k} e^{-\frac{x_{s}^{\prime}}{\epsilon}}\right)} \tag{17}
\end{equation*}
$$

Now, for each $1 \leq i \leq k$, if $x_{i}^{\prime} \geq x_{i}$ then

$$
\frac{e^{-\frac{x_{i}^{\prime}}{\epsilon}}-e^{-\frac{x_{i}}{\epsilon}}}{\left(1+\sum_{r=1}^{k} e^{-\frac{x_{r}}{\epsilon}}\right)\left(1+\sum_{s=1}^{k} e^{-\frac{x_{s}^{\prime}}{\epsilon}}\right)} \leq 0=\frac{1}{4 \epsilon}\left[x_{i}-x_{i}^{\prime}\right]^{+}
$$

On the other hand, if $x_{i}^{\prime}<x_{i}$, then we have

$$
\begin{aligned}
\frac{e^{-\frac{x_{i}^{\prime}}{\epsilon}}-e^{-\frac{x_{i}}{\epsilon}}}{\left(1+\sum_{r=1}^{k} e^{-\frac{x_{r}}{\epsilon}}\right)\left(1+\sum_{s=1}^{k} e^{-\frac{x_{s}^{\prime}}{\epsilon}}\right)} & \leq \frac{e^{-\frac{x_{i}^{\prime}}{\epsilon}}-e^{-\frac{x_{i}}{\epsilon}}}{\left(1+e^{-\frac{x_{i}}{\epsilon}}\right)\left(1+e^{-\frac{x_{i}^{\prime}}{\epsilon}}\right)} \\
& =\frac{1}{1+e^{-\frac{x_{i}}{\epsilon}}}-\frac{1}{1+e^{-\frac{x_{i}^{\prime}}{\epsilon}}} \\
& \leq \frac{1}{4 \epsilon}\left[x_{i}-x_{i}^{\prime}\right]^{+}
\end{aligned}
$$

where the last step uses the statement for $k=1$. Therefore, in either case, we have proved that for all $1 \leq i \leq k$ we have:

$$
\frac{e^{-\frac{x_{i}^{\prime}}{\epsilon}}-e^{-\frac{x_{i}}{\epsilon}}}{\left(1+\sum_{r=1}^{k} e^{-\frac{x_{r}}{\epsilon}}\right)\left(1+\sum_{s=1}^{k} e^{-\frac{x_{s}^{\prime}}{\epsilon}}\right)} \leq \frac{1}{4 \epsilon}\left[x_{i}-x_{i}^{\prime}\right]^{+}
$$

Substituting this into (17) we get the desired result.
Now we generalize Lemma 5 to depth $d$ local embeddings.
Lemma 7 (depth $d$ local contraction). Assume $(H, i) \hookrightarrow_{d}\left(H^{\prime}, i^{\prime}\right)$ with embedding $\phi: V_{i, d}^{H} \hookrightarrow V_{i^{\prime}, d}^{H^{\prime}}$. Also let $\theta_{\epsilon}$ and $\theta_{\epsilon}^{\prime}$ be $\epsilon$-balanced allocations on $H$ and $H^{\prime}$ respectively, with respective baseload functions $b$ and $b^{\prime}$, where $b(j) \leq b^{\prime}(\phi(j))$ for all $j$ such that $d_{H}(i, j) \leq d-1$. Then we have:

$$
\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\left(i^{\prime}\right) \geq \partial_{b} \theta_{\epsilon}(i)-\left(\frac{L \Delta}{4 \epsilon+L \Delta}\right)^{d} M_{d}^{+}
$$

where

$$
M_{d}:=\max _{j \in D_{i, d}^{H}} \partial_{b} \theta_{\epsilon}(j)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(\phi(j)),
$$

and

$$
\begin{aligned}
L & :=\max _{e \in E_{H}\left(V_{i, d}^{H}\right)}|e|, \\
\Delta & :=\max _{j \in V_{i, d-1}^{H}} \operatorname{deg}_{H}(j) .
\end{aligned}
$$

Proof. Take some $j \in V(H)$ with $d_{H}(i, j)=k$. If $j^{\prime} \in V(H)$ is such that $d_{H}\left(j, j^{\prime}\right)=1$, the triangle inequality implies that

$$
\left|d_{H}\left(i, j^{\prime}\right)-d_{H}(i, j)\right| \leq d_{H}\left(j, j^{\prime}\right)=1
$$

Thus

$$
\begin{equation*}
D_{j, 1}^{H} \subset D_{i, k-1}^{H} \cup D_{i, k}^{H} \cup D_{i, k+1}^{H} \tag{18}
\end{equation*}
$$

Hence, if $d_{H}(i, j) \leq d-1$, using the same embedding map $\phi$ we have $(H, j) \hookrightarrow_{1}\left(H^{\prime}, \phi(j)\right)$. Hence, if we define

$$
M_{k}:=\max _{j: d_{H}(i, j)=k} \partial_{b} \theta_{\epsilon}(j)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(\phi(j)) \quad 0 \leq k \leq d
$$

and

$$
\alpha:=\frac{L \Delta}{4 \epsilon+L \Delta}
$$

then using Lemma 5 and (18) we have

$$
\begin{equation*}
M_{k} \leq \alpha\left(M_{k-1}^{+} \vee M_{k}^{+} \vee M_{k+1}^{+}\right) \quad 1 \leq k \leq d-1 \tag{19}
\end{equation*}
$$

and

$$
\begin{equation*}
M_{0} \leq \alpha M_{1}^{+} \tag{20}
\end{equation*}
$$

We show by induction that $M_{k}^{+} \leq \alpha M_{k+1}^{+}$for $0 \leq k \leq d-1$. For $k=0$, this follows from (20) and the fact that $x \mapsto x^{+}$is increasing. For $k \geq 1$ we have from (19) that

$$
M_{k}^{+} \leq \alpha\left(M_{k-1}^{+} \vee M_{k}^{+} \vee M_{k+1}^{+}\right) \leq \alpha\left(\left(\alpha M_{k}^{+} \vee M_{k}^{+}\right) \vee M_{k+1}^{+}\right)=\alpha\left(M_{k}^{+} \vee M_{k+1}^{+}\right)
$$

We claim that $M_{k}^{+} \leq \alpha M_{k+1}^{+}$. The above inequality means that either $M_{k}^{+} \leq \alpha M_{k}^{+}$or $M_{k}^{+} \leq \alpha M_{k+1}^{+}$. The latter case is precisely what we have claimed, while in the former case, as $\alpha<1$, we have $M_{k}=0$ and the inequality $M_{k}^{+} \leq \alpha M_{k+1}^{+}$is automatic. This implies that $M_{0} \leq \alpha^{d} M_{d}^{+}$and completes the proof.

Using the above local results, we now show that if we add edges to a hypergraph and/or increase the baseload, the total load should increase.

Proposition 5. Let $H$ and $H^{\prime}$ be two hypergraphs defined on the same vertex set $V=V(H)=V\left(H^{\prime}\right)$, such that $E(H) \subset E\left(H^{\prime}\right)$, and $H$ is bounded (but $H^{\prime}$ is not necessarily bounded). Suppose two bounded baseload functions $b, b^{\prime}: V \rightarrow \mathbb{R}$ are given, such that $b(i) \leq b^{\prime}(i)$ for all $i \in V$. If $\theta_{\epsilon}$ and $\theta_{\epsilon}^{\prime}$ are two $\epsilon$-balanced allocations on $H$ and $H^{\prime}$, respectively, for the respective baseload functions $b$ and $b^{\prime}$, then we have $\partial_{b} \theta_{\epsilon}(i) \leq \partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i)$ for all $i \in V$.

Proof. Since $H$ is bounded, there are constants $\Delta$ and $L$ such that, for all $i \in V, \operatorname{deg}_{H}(i) \leq \Delta$ and $|e| \leq L$ for all $e \in E(H)$. Also, as $b$ is bounded, for some $K>0$ and all $i \in V$, we have $|b(i)| \leq K$.

Now, fix some $i \in V$. Since $E(H) \subset E\left(H^{\prime}\right)$, we have $(H, i) \hookrightarrow_{1}\left(H^{\prime}, i\right)$ with the identity map as the embedding function. With this, define

$$
M:=\sup _{j \in V} \partial_{b} \theta_{\epsilon}(j)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(j)
$$

Note that, since $\partial_{b} \theta_{\epsilon}^{b}$ is bounded to $\Delta+K$, the above quantity is finite and well defined. Now, using Lemma 5, we have

$$
\partial_{b} \theta_{\epsilon}^{b}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{b^{\prime}}(i) \leq \frac{\Delta L}{4 \epsilon+\Delta L} M^{+}
$$

Taking the supremum over $i$ on the left hand side, we get

$$
M \leq \frac{\Delta L}{4 \epsilon+\Delta L} M^{+}
$$

which, since $\frac{\Delta L}{4 \epsilon+\Delta L}<1$, implies $M \leq 0$ and completes the proof.
If we have a fixed bounded hypergraph $H$ with a baseload function, and $\theta_{\epsilon}, \theta_{\epsilon}^{\prime}$ are two $\epsilon$-balanced allocations on $H$, repeating the above proposition twice, we get $\partial \theta_{\epsilon} \leq \partial \theta_{\epsilon}^{\prime}$ and $\partial \theta_{\epsilon}^{\prime} \leq \partial \theta_{\epsilon}$, which implies uniqueness. To sum up, we have:

Corollary 1. If a hypergraph $H$ is bounded, there is a unique $\epsilon$-balanced allocation with respect to any given baseload on it.

## $4.3 \quad \epsilon$-balanced allocations for unbounded hypergraphs with respect to a baseload: canonical allocations

We now produce $\epsilon$-balanced allocations on a hypergraph $H$ with respect to a given baseload even when the hypergraph is not necessarily bounded. Note that we do not make any claims about uniqueness.

For a given hypergraph $H$ and $\Delta \in \mathbb{N}$, define $H^{\Delta}$ to be the hypergraph with vertex set $V(H)$ and edge set $E^{\Delta}$, where

$$
\begin{equation*}
E^{\Delta}:=\left\{e \in E(H):|e| \leq \Delta, \operatorname{deg}_{H}(i) \leq \Delta \quad \forall i \in e\right\} \tag{21}
\end{equation*}
$$

Given the baseload function $b$, define $\theta_{\epsilon}^{\Delta}$ to be the unique $\epsilon$-balanced allocation on $H^{\Delta}$ with respect to the baseload $b$, where the existence and uniqueness of $\theta_{\epsilon}^{\Delta}$ is a consequence of Corollary 1 above. Since $E^{\Delta}$ increases to $E(H)$ as $\Delta$ increases, Proposition 5 implies that $\partial_{b} \theta_{\epsilon}^{\Delta}$ is pointwise increasing in $\Delta$. Since it is also pointwise bounded, it is convergent. For a node $i \in V(H)$, define

$$
l_{i}:=\lim _{\Delta \rightarrow \infty} \partial_{b} \theta_{\epsilon}^{\Delta}(i)
$$

Now, for $e \in E(H)$ and $i \in e$, define

$$
\begin{equation*}
\theta_{\epsilon}(e, i):=\frac{\exp \left(-\frac{l_{i}}{\epsilon}\right)}{\sum_{j \in e} \exp \left(-\frac{l_{j}}{\epsilon}\right)} \tag{22}
\end{equation*}
$$

Now we prove that $\theta_{\epsilon}$ is an $\epsilon$-balanced allocation with respect to the baseload $b$. First, observe that, because of the normalizing term in the denominator, $\sum_{i \in e} \theta_{\epsilon}(e, i)=1$ for all $e \in E(H)$. We now show that $l_{i}=\partial_{b} \theta_{\epsilon}(i)$. Note that

$$
\begin{aligned}
\sum_{e \ni i} \theta_{\epsilon}(e, i) & =\sum_{e \ni i} \frac{\exp \left(-\frac{l_{i}}{\epsilon}\right)}{\sum_{j \in e} \exp \left(-\frac{l_{j}}{\epsilon}\right)} \\
& =\lim _{\Delta \rightarrow \infty} \sum_{e \ni i} \frac{\exp \left(-\frac{\partial_{b} \theta_{\epsilon}^{\Delta}(i)}{\epsilon}\right)}{\sum_{j \in e} \exp \left(-\frac{\partial_{b} \theta_{\epsilon}^{\Delta}(j)}{\epsilon}\right)}
\end{aligned}
$$

Observe that, for $\Delta>\max _{e \ni i}\left(|e| \vee \max _{j \in e} \operatorname{deg}_{H}(j)\right)$, we have $e \in E^{\Delta}$ for all $e \ni i$. Hence, the term inside the summation is $\theta_{\epsilon}^{\Delta}(e, i)$, because $\theta_{\epsilon}^{\Delta}$ is the unique $\epsilon$-balanced allocation on $H^{\Delta}$ with respect to the baseload $b$. Since we are taking $\Delta \rightarrow \infty$, we can assume it is big enough to get

$$
\partial_{b} \theta_{\epsilon}(i)=b(i)+\lim _{\Delta \rightarrow \infty} \sum_{e \ni i} \theta_{\epsilon}^{\Delta}(e, i)=\lim _{\Delta \rightarrow \infty} \partial_{b} \theta_{\epsilon}^{\Delta}(i)=l_{i} .
$$

Substituting this into (22), we conclude that $\theta_{\epsilon}$ is $\epsilon$-balanced, with respect to the baseload $b$.
Remark 11. Note that the above procedure gives rise to an $\epsilon$-balanced allocation with respect to any given baseload, but we do not know if it is the only possible $\epsilon$-balanced allocation or not. In fact, we have proved uniqueness for bounded hypergraphs only. To emphasize this and avoid confusion, we call the $\epsilon$-balanced allocation resulting from the procedure above the "canonical" allocation with respect to the given baseload. A special case of the procedure yields the canonical $\epsilon$-balanced allocation when there is no baseload.

Now, we generalize the monotonicity property of Proposition 5 to the case of not necessarily bounded hypergraphs, for these canonical allocations.

Proposition 6. Given hypergraphs $H$ and $H^{\prime}$ on the same vertex set $V$, with $E(H) \subseteq E\left(H^{\prime}\right)$, and baseload functions $b, b^{\prime}: V \rightarrow \mathbb{R}$ such that $b(i) \leq b^{\prime}(i)$ for all $i \in V$, if $\theta_{\epsilon}$ and $\theta_{\epsilon}^{\prime}$ are the canonical $\epsilon-$ balanced allocations on $H$ and $H^{\prime}$, with respect to baseloads $b$ and $b^{\prime}$, respectively, then $\partial_{b} \theta_{\epsilon}(i) \leq \partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i)$ for all $i \in V$.

Proof. Set $E:=E(H)$ and $E^{\prime}:=E\left(H^{\prime}\right)$. Note that since $E \subseteq E^{\prime}$ and the vertex sets are the same for $H$ and $H^{\prime}, E^{\Delta} \subseteq E^{\prime}$ for any $\Delta$. Thus, if $\theta_{\epsilon}^{\Delta}$ is the unique $\epsilon$-balanced allocation on $\left\langle V, E^{\Delta}\right\rangle$ with respect to the baseload $b$ and $\theta_{\epsilon}^{\prime}$ is the canonical $\epsilon$-balanced allocation on $\left\langle V, E^{\prime}\right\rangle$ with respect to the baseload $b^{\prime}$, Proposition 5 implies that $\partial_{b} \theta_{\epsilon}^{\Delta}(i) \leq \partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i)$ for all $i \in V$ (note that in Proposition 5 only the smaller hypergraph needs to be bounded, hence we only need to truncate $E$ ). By sending $\Delta$ to infinity, we get the desired result.

### 4.4 Nonexpansivity

Proposition 7. Let $H$ be a given hypergraph and $b, b^{\prime}$ be two baseload functions. Let $\theta$ and $\theta^{\prime}$ be the canonical $\epsilon$-balanced allocations on $H$ with respect to $b$ and $b^{\prime}$, respectively. Then, we have

$$
\left\|\partial_{b} \theta_{\epsilon}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\right\|_{l^{1}(V(H))} \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))} .
$$

Proof. To start with, assume that $H$ is bounded. Later, we will relax this assumption.
Consider first the special case $b(i) \geq b^{\prime}(i)$ for all $i \in V(H)$. Then, using the monotonicity property of Proposition 5, we have $\partial_{b} \theta_{\epsilon}(i) \geq \partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i)$ for all $i \in V(H)$. In particular,

$$
\left\|\partial_{b} \theta_{\epsilon}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\right\|_{l^{1}(V(H))}=\sum_{i \in V(H)} \partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i)
$$

Now, let $\left\{V_{n}\right\}$ be a nested sequence of finite subsets of $V(H)$ converging to $V(H)$, i.e. $V_{n} \uparrow V(H)$. Let $\theta_{n, \epsilon}$ and $\theta_{n, \epsilon}^{\prime}$ be the $\epsilon$-balanced allocations on $\left\langle V_{n}, E_{H}\left(V_{n}\right)\right\rangle$ with respect to the restrictions of $b$ and $b^{\prime}$ to $V_{n}$, respectively. The monotonicity property of Proposition 5 and an argument similar to the one given in Section 4.3 above, implies that $\partial_{b} \theta_{n, \epsilon}$ and $\partial_{b^{\prime}} \theta_{n, \epsilon}^{\prime}$ converge to $\epsilon$-balanced allocations on $H$ and $H^{\prime}$ with respect to $b$ and $b^{\prime}$, respectively. Since $H$ is bounded, there is a unique $\epsilon$-balanced allocations with respect to any baseload function. Therefore, we have $\partial_{b} \theta_{n, \epsilon} \uparrow \partial_{b} \theta_{\epsilon}$ and $\partial_{b^{\prime}} \theta_{n, \epsilon}^{\prime} \uparrow \partial_{b^{\prime}} \theta_{\epsilon}^{\prime}$ as $n \rightarrow \infty$. Using the conservation of mass, we have

$$
\sum_{i \in V_{n}} \partial_{b} \theta_{n, \epsilon}(i)-\partial_{b^{\prime}} \theta_{n, \epsilon}^{\prime}(i)=\sum_{i \in V_{n}} b(i)-b^{\prime}(i) \leq \sum_{i \in V(H)} b(i)-b^{\prime}(i)=\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
$$

where we have used the assumption $b \geq b^{\prime}$. In fact, for any finite subset of vertices $K \subset V(H)$, we have $K \subset V_{n}$ for $n$ large enough and so using $\partial_{b} \theta_{n, \epsilon} \uparrow \partial_{b} \theta_{\epsilon}$ and $\partial_{b^{\prime}} \theta_{n, \epsilon}^{\prime} \uparrow \partial_{b^{\prime}} \theta_{\epsilon}^{\prime}$ and also the monotonicity property of Proposition 5, we have

$$
\begin{aligned}
\sum_{i \in K} \partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i) & =\lim _{n \rightarrow \infty} \sum_{i \in K} \partial_{b} \theta_{n, \epsilon}(i)-\partial_{b^{\prime}} \theta_{n, \epsilon}^{\prime}(i) \\
& \leq \lim _{n \rightarrow \infty} \sum_{i \in V_{n}} \partial_{b} \theta_{n, \epsilon}(i)-\partial_{b^{\prime}} \theta_{n, \epsilon}^{\prime}(i) \\
& \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
\end{aligned}
$$

Since this holds for every finite subset $K \subset V(H)$, we conclude that

$$
\left\|\partial_{b} \theta_{\epsilon}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\right\|_{l^{1}(V(H))} \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
$$

Now, continuing to assume that $H$ is bounded, suppose the condition $b \geq b^{\prime}$ does not necessarily hold. Define $b^{\prime \prime}:=b \wedge b^{\prime}$ with $\theta_{\epsilon}^{\prime \prime}$ being the unique $\epsilon$-balanced allocation on $H$ with respect to $b^{\prime \prime}$. Due to monotonicity property of Proposition 5 , we have $\partial_{b} \theta_{\epsilon} \geq \partial_{b^{\prime \prime}} \theta_{\epsilon}^{\prime \prime}$ and $\partial_{b^{\prime}} \theta_{\epsilon}^{\prime} \geq \partial_{b^{\prime \prime}} \theta_{\epsilon}^{\prime \prime}$. Thus, using the above argument and the triangle inequality, we have

$$
\begin{aligned}
\left\|\partial_{b} \theta_{\epsilon}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\right\|_{l^{\prime}(V(H))} & \leq\left\|\partial_{b} \theta_{\epsilon}-\partial_{b^{\prime \prime}} \theta_{\epsilon}^{\prime \prime}\right\|_{l^{1}(V(H))}+\left\|\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}-\partial_{b^{\prime \prime}} \theta_{\epsilon}^{\prime \prime}\right\|_{l^{1}(V(H))} \\
& \leq\left\|b-b^{\prime \prime}\right\|_{l^{1}(V(H))}+\left\|b^{\prime}-b^{\prime \prime}\right\|_{l^{1}(V(H))} \\
& =\sum_{i \in V(H)} b(i)-b^{\prime \prime}(i)+b^{\prime}(i)-b^{\prime \prime}(i) \\
& =\sum_{i \in V(H)}\left|b(i)-b^{\prime}(i)\right|=\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
\end{aligned}
$$

Finally, we relax the boundedness assumption on $H$.
We take a not necessarily bounded hypergraph $H$ and let $H^{\Delta}$ be the truncation of $H$, as defined in Section 4.3. Let $\theta_{\epsilon}^{\Delta}$ and $\theta_{\epsilon}^{\prime, \Delta}$ be the unique $\epsilon$-balanced allocations on $H^{\Delta}$, with respect to the baseloads $b$ and $b^{\prime}$, respectively. Since $H^{\Delta}$ is bounded, using the above argument, we have

$$
\left\|\partial_{b} \theta_{\epsilon}^{\Delta}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime, \Delta}\right\|_{l^{1}(V(H))} \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))} .
$$

Hence, for any finite subset $K \subseteq V(H)$, we have

$$
\sum_{i \in K}\left|\partial_{b} \theta_{\epsilon}^{\Delta}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime, \Delta}(i)\right| \leq\left\|\partial_{b} \theta_{\epsilon}^{\Delta}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime, \Delta}\right\|_{l^{1}(V(H))} \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
$$

Sending $\Delta$ to infinity and using the facts that $\partial \theta_{\epsilon}^{\Delta}$ and $\partial \theta_{\epsilon}^{\prime, \Delta}$ converge to $\partial \theta_{\epsilon}$ and $\partial \theta_{\epsilon}^{\prime}$, respectively, and also the fact that $K$ is finite, we have

$$
\sum_{i \in K}\left|\partial_{b} \theta_{\epsilon}(i)-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}(i)\right| \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
$$

Since the above is true for all finite $K$, we have

$$
\left\|\partial_{b} \theta_{\epsilon}-\partial_{b^{\prime}} \theta_{\epsilon}^{\prime}\right\|_{l^{1}(V(H))} \leq\left\|b-b^{\prime}\right\|_{l^{1}(V(H))}
$$

and the proof is complete.

### 4.5 Regularity property for canonical $\epsilon$-balanced allocations with respect to a baseload

In this section, we give a regularity property of canonical allocations which is crucial in our analysis.
Let $T$ be a hypertree. For a node $i \in V(T)$ and $e \in E(T), e \ni i$, define $T_{e \rightarrow i}$ to the connected subtree with root $i$ that does not contain the part of $T$ directed from $e$. To be more precise, the vertex set of $T_{e \rightarrow i}$ is the set of vertices $j \in V(T)$ such that the shortest path from $i$ to $j$ does not contain $e$. The edge set of $T_{e \rightarrow i}$ contains all the edges with all their end points in this subset, i.e. $E_{T}\left(V\left(T_{e \rightarrow i}\right)\right)$ in the notation of Section 2.1.

Proposition 8. Let $T$ be a hypertree, $b$ a baseload function on $T$, and $\theta$ the canonical $\epsilon$-balanced allocation on $T$ with respect to the baseload b. Let $e \in E(T)$ and $i \in e$, and let $\theta_{T_{e \rightarrow i}}$ denote the restriction of $\theta$ to $T_{e \rightarrow i}$, i.e.

$$
\theta_{T_{e \rightarrow i}}\left(e^{\prime}, i^{\prime}\right)=\theta\left(e^{\prime}, i^{\prime}\right), \quad e^{\prime} \in E\left(T_{e \rightarrow i}\right), i^{\prime} \in e^{\prime}
$$

Then, $\theta_{T_{e \rightarrow i}}$ is the canonical $\epsilon$-balanced allocation on $T_{e \rightarrow i}$ with respect to the baseload function $\tilde{b}$ defined as

$$
\tilde{b}(i)=b(i)+\theta(e, i),
$$

and $\tilde{b}(j)=b(j)$ for $j \in V\left(T_{e \rightarrow i}\right) \backslash\{i\}$.
Proof. It is straightforward to check that $\theta_{T_{e \rightarrow i}}$ is an $\epsilon$-balanced allocation on $T_{e \rightarrow i}$ with baseload function $\tilde{b}$. Thus, the content of the theorem is the statement about this $\epsilon$-balanced allocation being the canonical $\epsilon$-balanced allocation on $T_{e \rightarrow i}$ with the baseload function $\tilde{b}$.

Let $\tilde{\theta}$ denote the canonical $\epsilon$-balanced allocation on $T_{e \rightarrow i}$ with respect to the baseload $\tilde{b}$. Moreover, let $\tilde{\theta}^{\Delta}$ denote the unique $\epsilon$-balanced allocation on the bounded tree $\left(T_{e \rightarrow i}\right)^{\Delta}$ with respect to the baseload $\tilde{b}$. Throughout this proof, we assume that $\Delta>\Delta_{0}$ where

$$
\Delta_{0}:=|e| \vee \max \left\{\operatorname{deg}_{H}(j): j \in e\right\}
$$

If $\Delta$ satisfies this property, then $e \in E\left(T^{\Delta}\right)$, and one can also check that

$$
\begin{equation*}
\left(T_{e \rightarrow i}\right)^{\Delta} \equiv\left(T^{\Delta}\right)_{e \rightarrow i} \tag{23}
\end{equation*}
$$

so we can unambiguously write $T_{e \rightarrow i}^{\Delta}$ for this tree. (Note that $T^{\Delta}$ need not be connected, but this is not relevant.) If $\theta^{\Delta}$ denotes the unique $\epsilon$-balanced allocation on $T^{\Delta}$ with respect to the baseload $b$ and $\theta_{T_{e \rightarrow i}}^{\Delta}$ is its restriction to $T_{e \rightarrow i}^{\Delta}$, then, since $\Delta \geq \Delta_{0}, \theta_{T_{e \rightarrow i}}^{\Delta}$ is the unique $\epsilon$-balanced allocation on $T_{e \rightarrow i}^{\Delta}$ with respect to the baseload $b^{\Delta}$ defined as $b^{\Delta}(i)=b(i)+\theta^{\Delta}(e, i)$ and $b^{\Delta}(j)=b(j)$ for $j \neq i$. Now, $\tilde{\theta}^{\Delta}$ and $\theta_{T_{e \rightarrow i}}^{\Delta}$ are canonical $\epsilon$-balanced allocations on $T_{e \rightarrow i}^{\Delta}$ with respect to the baseloads $\tilde{b}$ and $b^{\Delta}$, respectively. Using Proposition 7, we conclude that for an arbitrary vertex $k \in V\left(T_{e \rightarrow i}\right)$ we have

$$
\left|\partial_{b} \Delta \theta_{T_{e \rightarrow i}}^{\Delta}(k)-\partial_{\tilde{b}} \tilde{\theta}^{\Delta}(k)\right| \leq\left\|\partial_{b} \Delta \theta_{T_{e \rightarrow i}}^{\Delta}-\partial_{\tilde{b}} \tilde{\theta}^{\Delta}\right\|_{l^{1}\left(V\left(T_{e \rightarrow i}\right)\right)} \leq\left|\theta^{\Delta}(e, i)-\theta(e, i)\right| .
$$

Now, sending $\Delta$ to infinity and noting the fact that $\theta^{\Delta}(e, i) \rightarrow \theta(e, i)$, we have $\mid \partial_{b} \theta_{T_{e \rightarrow i}}^{\Delta}(k)-$ $\partial_{\tilde{b}} \tilde{\theta}^{\Delta}(k) \mid \rightarrow 0$. Since $\theta_{T_{e \rightarrow i}}^{\Delta}$ is the restriction to $T_{e \rightarrow i}^{\Delta}$ of the $\epsilon$-balanced allocation $\theta^{\Delta}$ on $T^{\Delta}$ with respect to the baseload $b$, we have $\partial_{b} \Delta \theta_{T_{e \rightarrow i}}^{\Delta}(k)=\partial_{b} \theta^{\Delta}(k)$ for all $k \in V\left(T_{e \rightarrow i}^{\Delta}\right)$. Since $\partial_{b} \theta^{\Delta}(k) \rightarrow \partial_{b} \theta(k)$ and $\partial_{\tilde{b}} \tilde{\theta}^{\Delta}(k) \rightarrow \partial_{\tilde{b}} \tilde{\theta}(k)$ for all $k \in V(T)$ as $\Delta \rightarrow \infty$, we conclude that

$$
\partial_{\tilde{b}} \tilde{\theta}(k)=\partial_{b} \theta(k), \quad \forall k \in V\left(T_{e \rightarrow i}\right)
$$

But it is straightforward to check that $\partial_{b} \theta(k)=\partial_{\tilde{b}} \theta_{T_{e \rightarrow i}}(k)$ for all $k \in V\left(T_{e \rightarrow i}\right)$. From the definition of $\epsilon$-balanced allocations, we conclude that $\theta_{T_{e \rightarrow i}}\left(e^{\prime}, i^{\prime}\right)$ equals $\tilde{\theta}\left(e^{\prime}, i^{\prime}\right)$ for all $e^{\prime} \in E\left(T_{e \rightarrow i}\right)$ and all $i^{\prime} \in e^{\prime}$, i.e. that $\theta_{T_{e \rightarrow i}}$ is the canonical $\epsilon$-balanced allocation on $T_{e \rightarrow i}$ with respect to the baseload $\tilde{b}$, which was what was to be shown.

## $5 \epsilon$-balanced allocations on $\mathcal{H}_{* *}$

In this section, we discuss how to find an $\epsilon$-balanced allocation on $\mathcal{H}_{* *}$, in the sense of Definition 25, with respect to a unimodular measure $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$. Recall that this is a Borel allocation $\Theta_{\epsilon}: \mathcal{H}_{* *} \rightarrow$ $[0,1]$ such that

$$
\Theta_{\epsilon}(H, e, i)=\frac{\exp \left(-\partial \Theta_{\epsilon}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-\partial \Theta_{\epsilon}(H, j) / \epsilon\right)}, \quad \vec{\mu} \text {-a.e.. }
$$

What we do here is in fact stronger, in the sense that we introduce a Borel allocation $\Theta_{\epsilon}$ such that the above condition is satisfied pointwise, i.e.

$$
\begin{equation*}
\Theta_{\epsilon}(H, e, i)=\frac{\exp \left(-\partial \Theta_{\epsilon}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-\partial \Theta_{\epsilon}(H, j) / \epsilon\right)} \quad \forall[H, e, i] \in \mathcal{H}_{* *} \tag{24}
\end{equation*}
$$

In fact, we can define $\Theta_{\epsilon}(H, e, i)$ to be $\theta_{\epsilon}^{H}(e, i)$ where $\theta_{\epsilon}^{H}$ is the canonical $\epsilon$-balanced allocation for $H$, as was introduced in Section 4.3. Defining $\Theta_{\epsilon}$ in this way guarantees that (24) is satisfied. Therefore, it remains to show that $\Theta_{\epsilon}$ is a Borel allocation. In the following, we construct $\Theta_{\epsilon}$ differently, but as we will see later, the $\Theta_{\epsilon}(H, e, i)$ we construct will be equal to $\theta_{\epsilon}^{H}(e, i)$ (see Remark 12 below).

For the construction, given $\Delta \in \mathbb{N}$, define $F_{\epsilon}^{\Delta}: \mathcal{H}_{*} \rightarrow \mathbb{R}$ via $F_{\epsilon}^{\Delta}(H, i):=\partial \theta_{\epsilon}^{H^{\Delta}}(i)$, where $H^{\Delta}$ is the truncated hypergraph introduced in Section 4.3. The uniqueness property of $\epsilon$-balanced allocations for bounded hypergraphs (Corollary 1) implies that the above definition does not depend on the specific choice of $(H, i)$ in the equivalence class. Therefore, $F_{\epsilon}^{\Delta}$ is well defined. We claim that $F_{\epsilon}^{\Delta}$ is a continuous function on $\mathcal{H}_{*}$. Indeed, if $\left(H_{1}, i_{1}\right) \equiv_{d}\left(H_{2}, i_{2}\right)$, then $\left(H_{1}^{\Delta}, i_{1}\right) \equiv_{d-1}\left(H_{2}^{\Delta}, i_{2}\right)$, and using Lemma 7 we have

$$
\left|\partial \theta_{\epsilon}^{H_{1}^{\Delta}}\left(i_{1}\right)-\partial \theta_{\epsilon}^{H_{2}^{\Delta}}\left(i_{2}\right)\right| \leq\left(\frac{\Delta^{2}}{4 \epsilon+\Delta^{2}}\right)^{d-1} \Delta .
$$

This implies that $F_{\epsilon}^{\Delta}$ is (uniformly) continuous. Moreover, Proposition 5 implies that $F_{\epsilon}^{\Delta}$ is pointwise increasing in $\Delta$. On the other hand, $F_{\epsilon}^{\Delta}(H, i) \leq \operatorname{deg}_{H}(i)$. Hence, there is a pointwise limit

$$
\begin{equation*}
F_{\epsilon}(H, i):=\lim _{\Delta \rightarrow \infty} F_{\epsilon}^{\Delta}(H, i) . \tag{25}
\end{equation*}
$$

We now define $\Theta_{\epsilon}: \mathcal{H}_{* *} \rightarrow[0,1]$ in the following way:

$$
\begin{equation*}
\Theta_{\epsilon}(H, e, i):=\frac{\exp \left(-F_{\epsilon}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-F_{\epsilon}(H, j) / \epsilon\right)} \tag{26}
\end{equation*}
$$

We want to show that $\Theta_{\epsilon}$ is an $\epsilon$-balanced allocation on $\mathcal{H}_{* *}$. Since $F_{\epsilon}^{\Delta}$ is continuous, $F_{\epsilon}$ is Borel, and hence $\Theta_{\epsilon}$ is a Borel map. Also, the normalization factor in the denominator guarantees that $\sum_{i \in e} \Theta_{\epsilon}(H, e, i)=1$, and hence $\Theta_{\epsilon}$ is a Borel allocation. Comparing (24) with (26), it suffices to show that

$$
\begin{equation*}
\partial \Theta_{\epsilon}(H, i)=F_{\epsilon}(H, i) \tag{27}
\end{equation*}
$$

In order to show this, note that

$$
\begin{aligned}
\partial \Theta_{\epsilon}(H, i) & =\sum_{e \ni i} \Theta_{\epsilon}(H, e, i) \\
& =\sum_{e \ni i} \frac{\exp \left(-F_{\epsilon}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-F_{\epsilon}(H, j) / \epsilon\right)} \\
& =\lim _{\Delta \rightarrow \infty} \sum_{e \ni i} \frac{\exp \left(-F_{\epsilon}^{\Delta}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-F_{\epsilon}^{\Delta}(H, j) / \epsilon\right)}
\end{aligned}
$$

Now, when $\Delta \geq \max _{e \ni i}|e|$ and also $\left.\Delta \geq \max _{j \in e, e \ni i} \operatorname{deg}_{H}(j)\right)$, we have $e \in E\left(H^{\Delta}\right)$, and we have

$$
\frac{\exp \left(-F_{\epsilon}^{\Delta}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-F_{\epsilon}^{\Delta}(H, j) / \epsilon\right)}=\frac{\exp \left(-\partial \theta_{\epsilon}^{H^{\Delta}}(H, i) / \epsilon\right)}{\sum_{j \in e} \exp \left(-\partial \theta_{\epsilon}^{H^{\Delta}}(H, j) / \epsilon\right)}=\theta_{\epsilon}^{H^{\Delta}}(H, e, i)
$$

Consequently,

$$
\begin{aligned}
\partial \Theta_{\epsilon}(H, i) & =\lim _{\Delta \rightarrow \infty} \sum_{e \ni i} \theta_{\epsilon}^{H^{\Delta}}(H, e, i) \\
& =\lim _{\Delta \rightarrow \infty} \partial \theta_{\epsilon}^{H^{\Delta}}(H, i) \\
& =\lim _{\Delta \rightarrow \infty} F_{\epsilon}^{\Delta}(H, i) \\
& =F_{\epsilon}(H, i) .
\end{aligned}
$$

Therefore, we have shown (27), which shows that $\Theta_{\epsilon}$ satisfies (24) and hence is an $\epsilon$-balanced allocation (both pointwise and $\vec{\mu}$-almost everywhere).
Remark 12. Note that (25) means that $F_{\epsilon}(H, i)=\partial \theta_{\epsilon}^{H}(i)$, where $\theta_{\epsilon}^{H}$ is the canonical $\epsilon$-balanced allocation in $H$. Hence, (26) implies that $\Theta_{\epsilon}(H, e, i)$ is pointwise equal to $\theta_{\epsilon}^{H}(e, i)$.

The following Proposition proves an almost sure uniqueness property for Borel $\epsilon$-balanced allocations which is similar in flavor to part 4 of Theorem 1. The proof of this statement is given in Appendix H.
Proposition 9. Assume $\mu$ is a unimodular measure on $\mathcal{H}_{*}$ such that $\operatorname{deg}(\mu)<\infty$. Given $\epsilon>0$, let $\Theta_{\epsilon}$ be the $\epsilon$-balanced allocation defined in this section, and let $\Theta_{\epsilon}^{\prime}$ be any other $\epsilon$-balanced allocation, both with respect to $\mu$. Then, we have $\partial \Theta_{\epsilon}=\partial \Theta_{\epsilon}^{\prime}, \mu-a . s$. and $\Theta_{\epsilon}=\Theta_{\epsilon}^{\prime}, \vec{\mu}-a . e$. .

## 6 Properties of balanced allocations

### 6.1 Existence

In this section we prove the existence of balanced allocations (part 1 of Theorem 1):
Proposition 10. Assume $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ is unimodular with $\operatorname{deg}(\mu)<\infty$. Then, there is a sequence $\epsilon_{k} \downarrow 0$ such that $\Theta_{\epsilon_{k}}$ converges to a balanced allocation $\Theta_{0}$, with the convergence being both in $L^{2}(\vec{\mu})$ and $\vec{\mu}$-almost everywhere.

Proof. First, we show that $\Theta_{\epsilon}$ is Cauchy in $L^{2}(\vec{\mu})$. To do so, we take $\epsilon, \epsilon^{\prime}>0$ and try to bound $\left\|\Theta_{\epsilon}-\Theta_{\epsilon^{\prime}}\right\|_{L^{2}(\vec{\mu})}$. For an integer $\Delta>0$, define the function $\Theta_{\epsilon}^{\Delta}$ on $\mathcal{H}_{* *}$ as follows:

$$
\Theta_{\epsilon}^{\Delta}(H, e, i)=\left\{\begin{array}{lc}
\theta_{\epsilon}^{H^{\Delta}}(e, i) & \text { if } e \in E\left(H^{\Delta}\right)  \tag{28}\\
0 & \text { otherwise }
\end{array}\right.
$$

where $H^{\Delta}$ is the truncated hypergraph introduced in Section 4.3 and $\theta_{\epsilon}^{H^{\Delta}}$ is the unique $\epsilon$-balanced allocation associated to it.

Take a locally finite hypergraph $H$ and an edge $e$ in $E\left(H^{\Delta}\right)$. As $\theta_{\epsilon}^{H^{\Delta}}$ is $\epsilon$-balanced on $H^{\Delta}$ and $e \in E\left(H^{\Delta}\right)$, for $i, j \in e$ we have

$$
\frac{\theta_{\epsilon}^{H^{\Delta}}(e, i)}{\theta_{\epsilon}^{H^{\Delta}}(e, j)}=\frac{\exp \left(-\partial \theta_{\epsilon}^{H^{\Delta}}(i) / \epsilon\right)}{\exp \left(-\partial \theta_{\epsilon}^{H^{\Delta}}(j) / \epsilon\right)}
$$

By the definition of $\Theta_{\epsilon}^{\Delta}$, if $e \in E\left(H^{\Delta}\right)$, then for all $j^{\prime} \in e, \Theta_{\epsilon}^{\Delta}\left(H, e, j^{\prime}\right)=\theta_{\epsilon}^{H^{\Delta}}\left(e, j^{\prime}\right)$ and $\partial \Theta_{\epsilon}^{\Delta}\left(H, j^{\prime}\right)=$ $\partial \theta_{\epsilon}^{H^{\Delta}}\left(j^{\prime}\right)$. Taking logarithms on both sides of the above equation, we have

$$
\partial \Theta_{\epsilon}^{\Delta}(H, i)+\epsilon \log \Theta_{\epsilon}^{\Delta}(H, e, i)=\partial \Theta_{\epsilon}^{\Delta}(H, j)+\epsilon \log \Theta_{\epsilon}^{\Delta}(H, e, j), \quad \forall i, j \in e
$$

with this equation holding pointwise whenever $e \in E\left(H^{\Delta}\right)$. As this equality holds for all $i, j \in e$, any two convex combinations of the values of $\partial \Theta_{\epsilon}^{\Delta}+\epsilon \log \Theta_{\epsilon}^{\Delta}$ evaluated at nodes in $e$ are equal, whenever $e \in E\left(H^{\Delta}\right)$. In particular, if $\Theta_{\epsilon^{\prime}}^{\Delta}$ denotes the $\epsilon^{\prime}$-balanced allocation on $\mathcal{H}_{* *}$ defined similarly to the above, then, whenever $e \in E\left(H^{\Delta}\right)$, we have $\sum_{i \in e} \Theta_{\epsilon}^{\Delta}(H, e, i)=\sum_{i \in e} \Theta_{\epsilon^{\prime}}^{\Delta}(H, e, i)=1$. Hence we have:

$$
\begin{align*}
\sum_{i \in e} \Theta_{\epsilon}^{\Delta}(H, e, i) & \left(\partial \Theta_{\epsilon}^{\Delta}(H, i)+\epsilon \log \Theta_{\epsilon}^{\Delta}(H, e, i)\right) \\
& =\sum_{i \in e} \Theta_{\epsilon^{\prime}}^{\Delta}(H, e, i)\left(\partial \Theta_{\epsilon}^{\Delta}(H, i)+\epsilon \log \Theta_{\epsilon}^{\Delta}(H, e, i)\right) \tag{29}
\end{align*}
$$

which holds pointwise, whenever $e \in E\left(H^{\Delta}\right)$. On the other hand, if $e \notin E\left(H^{\Delta}\right)$, then, by definition, $\Theta_{\epsilon}^{\Delta}(H, e, j)$ as well as $\Theta_{\epsilon^{\prime}}^{\Delta}(H, e, j)$ are zero for all $j \in e$. In this case the above equality again holds, with $0 \log 0$ being interpreted as 0 . In other words, pointwise on $\mathcal{H}_{* *}$, we have

$$
\nabla\left(\Theta_{\epsilon}^{\Delta}\left(\partial \Theta_{\epsilon}^{\Delta}+\epsilon \log \Theta_{\epsilon}^{\Delta}\right)\right)=\nabla\left(\Theta_{\epsilon^{\prime}}^{\Delta}\left(\partial \Theta_{\epsilon}^{\Delta}+\epsilon \log \Theta_{\epsilon}^{\Delta}\right)\right)
$$

where, by abuse of notation, we have treated $\partial \Theta_{\epsilon}^{\Delta}$ as a function on $\mathcal{H}_{* *}$ rather than on $\mathcal{H}_{*}$, via $\partial \Theta_{\epsilon}^{\Delta}(H, e, i):=\partial \Theta_{\epsilon}^{\Delta}(H, i)$, and likewise for $\partial \Theta_{\epsilon^{\prime}}^{\Delta}$. Rewriting the above identity, we have:

$$
\begin{equation*}
\nabla\left(\Theta_{\epsilon}^{\Delta} \partial \Theta_{\epsilon}^{\Delta}+\epsilon \Theta_{\epsilon}^{\Delta} \log \Theta_{\epsilon}^{\Delta}-\Theta_{\epsilon^{\prime}}^{\Delta} \partial \Theta_{\epsilon}^{\Delta}\right)=\epsilon \nabla\left(\Theta_{\epsilon^{\prime}}^{\Delta} \log \Theta_{\epsilon}^{\Delta}\right) \tag{30}
\end{equation*}
$$

Note that $\partial \Theta_{\epsilon}^{\Delta}$ is pointwise bounded by $\Delta$ by definition. Moreover, $\operatorname{deg}(\mu)<\infty$, which implies that $\vec{\mu}$ has finite total measure. Hence, all terms in the above equation have finite integral. On the other hand, from the definition of $\vec{\mu}$, we have:

$$
\int\left(\Theta_{\epsilon}^{\Delta}-\Theta_{\epsilon^{\prime}}^{\Delta}\right) \partial \Theta_{\epsilon}^{\Delta} d \vec{\mu}=\int \partial\left(\left(\Theta_{\epsilon}^{\Delta}-\Theta_{\epsilon^{\prime}}^{\Delta}\right) \partial \Theta_{\epsilon}^{\Delta}\right) d \mu=\int\left(\partial \Theta_{\epsilon}^{\Delta}-\partial \Theta_{\epsilon^{\prime}}^{\Delta}\right) \partial \Theta_{\epsilon}^{\Delta} d \mu
$$

Substituting this into (30) and using unimodularity, we have

$$
\begin{equation*}
\left\langle\partial \Theta_{\epsilon}^{\Delta}-\partial \Theta_{\epsilon^{\prime}}^{\Delta}, \partial \Theta_{\epsilon}^{\Delta}\right\rangle+\epsilon \int \Theta_{\epsilon}^{\Delta} \log \Theta_{\epsilon}^{\Delta} d \vec{\mu}=\epsilon \int \Theta_{\epsilon^{\prime}}^{\Delta} \log \Theta_{\epsilon}^{\Delta} d \vec{\mu} \tag{31}
\end{equation*}
$$

where $\langle.,$.$\rangle denotes the inner product of two functions in L^{2}(\mu)$. Now, changing the order of $\epsilon$ and $\epsilon^{\prime}$, we have

$$
\begin{equation*}
\left\langle\partial \Theta_{\epsilon^{\prime}}^{\Delta}-\partial \Theta_{\epsilon}^{\Delta}, \partial \Theta_{\epsilon^{\prime}}^{\Delta}\right\rangle+\epsilon^{\prime} \int \Theta_{\epsilon^{\prime}}^{\Delta} \log \Theta_{\epsilon^{\prime}}^{\Delta} d \vec{\mu}=\epsilon^{\prime} \int \Theta_{\epsilon}^{\Delta} \log \Theta_{\epsilon^{\prime}}^{\Delta} d \vec{\mu} \tag{32}
\end{equation*}
$$

Summing up these two equalities, we have:

$$
\begin{align*}
\left\|\partial \Theta_{\epsilon}^{\Delta}-\partial \Theta_{\epsilon^{\prime}}^{\Delta}\right\|_{2}^{2} & =\epsilon \int \Theta_{\epsilon^{\prime}}^{\Delta} \log \Theta_{\epsilon}^{\Delta} d \vec{\mu}-\epsilon \int \Theta_{\epsilon}^{\Delta} \log \Theta_{\epsilon}^{\Delta} d \vec{\mu}  \tag{33}\\
& +\epsilon^{\prime} \int \Theta_{\epsilon}^{\Delta} \log \Theta_{\epsilon^{\prime}}^{\Delta} d \vec{\mu}-\epsilon^{\prime} \int \Theta_{\epsilon^{\prime}}^{\Delta} \log \Theta_{\epsilon^{\prime}}^{\Delta} d \vec{\mu}
\end{align*}
$$

We now use the following information theoretic notation. For functions $\Theta, \Theta^{\prime}: \mathcal{H}_{* *} \rightarrow[0,1]$, we define

$$
\begin{aligned}
H(\Theta) & :=-\int \Theta \log \Theta d \vec{\mu}, \text { and } \\
D\left(\Theta \| \Theta^{\prime}\right) & :=\int \Theta \log \frac{\Theta}{\Theta^{\prime}} d \vec{\mu},
\end{aligned}
$$

where $0 \log 0$ is interpreted as 0 , and $\Theta(H, e, i)$ is assumed to be zero whenever $\Theta^{\prime}(H, e, i)$ is zero (by definition, $\Theta_{\epsilon}^{\Delta}$ and $\Theta_{\epsilon^{\prime}}^{\Delta}$ have this property). Rearranging the terms in (33),

$$
\begin{equation*}
\left\|\partial \Theta_{\epsilon}^{\Delta}-\partial \Theta_{\epsilon^{\prime}}^{\Delta}\right\|_{2}^{2}+\epsilon D\left(\Theta_{\epsilon^{\prime}}^{\Delta} \| \Theta_{\epsilon}^{\Delta}\right)+\epsilon^{\prime} D\left(\Theta_{\epsilon}^{\Delta} \| \Theta_{\epsilon^{\prime}}^{\Delta}\right)=\left(\epsilon-\epsilon^{\prime}\right)\left(H\left(\Theta_{\epsilon}^{\Delta}\right)-H\left(\Theta_{\epsilon^{\prime}}^{\Delta}\right)\right) \tag{34}
\end{equation*}
$$

Note that, since $\operatorname{deg}(\mu)<\infty$ and $\partial \Theta_{\epsilon}^{\Delta}$ and $\partial \Theta_{\epsilon^{\prime}}^{\Delta}$ are pointwise bounded to $\Delta$, all the terms are finite. With, $\mathcal{A}:=\left\{[H, e, i] \in \mathcal{H}_{* *}: e \in E\left(H^{\Delta}\right)\right\}$, we have

$$
\begin{align*}
D\left(\Theta_{\epsilon^{\prime}}^{\Delta} \| \Theta_{\epsilon}^{\Delta}\right) & =\int_{\mathcal{A}} \Theta_{\epsilon^{\prime}}^{\Delta} \log \frac{\Theta_{\epsilon^{\prime}}^{\Delta}}{\Theta_{\epsilon}^{\Delta}} d \vec{\mu} \\
& =\int_{\mathcal{A}} \nabla\left(\Theta_{\epsilon^{\prime}}^{\Delta} \log \frac{\Theta_{\epsilon^{\prime}}^{\Delta}}{\Theta_{\epsilon}^{\Delta}}\right) d \vec{\mu}  \tag{35}\\
& =\int_{\mathcal{A}} \frac{1}{|e|} D_{\mathrm{KL}}\left(\left(\Theta_{\epsilon^{\prime}}^{\Delta}(H, e, j)\right)_{j \in e} \|\left(\Theta_{\epsilon}^{\Delta}(H, e, j)\right)_{j \in e}\right) d \vec{\mu}(H, e, i) \\
& \geq 0
\end{align*}
$$

where $D_{\mathrm{KL}}$ denotes the standard KL divergence, and for $[H, e, i] \in \mathcal{A}$, by definition, $\left(\Theta_{\epsilon^{\prime}}^{\Delta}(H, e, j)\right)_{j \in e}$ and $\left(\Theta_{\epsilon^{\prime}}^{\Delta}(H, e, j)\right)_{j \in e}$ are vectors of nonnegative values summing up to one. Similarly, one can show that $D\left(\Theta_{\epsilon}^{\Delta} \| \Theta_{\epsilon^{\prime}}^{\Delta}\right) \geq 0$. Therefore, all the terms on the left hand side of (34) are nonnegative. As a result

$$
\begin{equation*}
\epsilon>\epsilon^{\prime} \quad \Rightarrow \quad H\left(\Theta_{\epsilon}^{\Delta}\right) \geq H\left(\Theta_{\epsilon^{\prime}}^{\Delta}\right) \tag{36}
\end{equation*}
$$

As $H\left(\Theta_{\epsilon}^{\Delta}\right) \geq 0$ for all $\epsilon>0$, the above inequality implies that $H\left(\Theta_{\epsilon}^{\Delta}\right)$ converges to some value as $\epsilon \downarrow 0$. Another result of (34) is that for $\epsilon>\epsilon^{\prime}$,

$$
\epsilon D\left(\Theta_{\epsilon^{\prime}}^{\Delta} \| \Theta_{\epsilon}^{\Delta}\right) \leq\left(\epsilon-\epsilon^{\prime}\right)\left(H\left(\Theta_{\epsilon}^{\Delta}\right)-H\left(\Theta_{\epsilon^{\prime}}^{\Delta}\right)\right)
$$

Dividing by $\epsilon$ and using $\epsilon>\epsilon^{\prime}$ and $H\left(\Theta_{\epsilon}^{\Delta}\right) \geq H\left(\Theta_{\epsilon^{\prime}}^{\Delta}\right)$, we get

$$
\begin{equation*}
D\left(\Theta_{\epsilon^{\prime}}^{\Delta} \| \Theta_{\epsilon}^{\Delta}\right) \leq H\left(\Theta_{\epsilon}^{\Delta}\right)-H\left(\Theta_{\epsilon^{\prime}}^{\Delta}\right) \tag{37}
\end{equation*}
$$

Using (35) and Pinsker's inequality (see, for instance, [CK11]),

$$
\begin{aligned}
D\left(\Theta_{\epsilon^{\prime}}^{\Delta} \| \Theta_{\epsilon}^{\Delta}\right) & =\int_{\mathcal{A}} \frac{1}{|e|} D_{\mathrm{KL}}\left(\left(\Theta_{\epsilon^{\prime}}^{\Delta}(H, e, j)\right)_{j \in e} \|\left(\Theta_{\epsilon}^{\Delta}(H, e, j)\right)_{j \in e}\right) d \vec{\mu} \\
& \stackrel{(a)}{\geq} \int_{\mathcal{A}} \frac{1}{|e|} \frac{1}{2} \sum_{j \in e}\left|\Theta_{\epsilon^{\prime}}^{\Delta}(H, e, j)-\Theta_{\epsilon}^{\Delta}(H, e, j)\right|^{2} d \vec{\mu} \\
& \stackrel{(b)}{=} \frac{1}{2} \int_{\mathcal{A}}\left|\Theta_{\epsilon^{\prime}}^{\Delta}-\Theta_{\epsilon}^{\Delta}\right|^{2} d \vec{\mu} \\
& =\frac{1}{2}\left\|\Theta_{\epsilon}^{\Delta}-\Theta_{\epsilon^{\prime}}^{\Delta}\right\|_{2}^{2}
\end{aligned}
$$

where (a) uses the Pinsker's inequality, and (b) uses unimodularity of $\mu$. Combining this with (37), for $\epsilon, \epsilon^{\prime}>0$ we have

$$
\left\|\Theta_{\epsilon}^{\Delta}-\Theta_{\epsilon^{\prime}}^{\Delta}\right\|_{2}^{2} \leq 2\left|H\left(\Theta_{\epsilon}^{\Delta}\right)-H\left(\Theta_{\epsilon^{\prime}}^{\Delta}\right)\right|
$$

Now, as we send $\Delta$ to infinity, $\Theta_{\epsilon}^{\Delta}$ converges pointwise to the function $\Theta_{\epsilon}$ defined in Section 5 . Moreover, $\operatorname{deg}(\mu)<\infty$ and the function $x \log x$ is bounded for $x \in[0,1]$. Hence, using dominated convergence theorem, we have

$$
\begin{equation*}
\left\|\Theta_{\epsilon}-\Theta_{\epsilon^{\prime}}\right\|_{2}^{2} \leq 2\left|H\left(\Theta_{\epsilon}\right)-H\left(\Theta_{\epsilon^{\prime}}\right)\right| . \tag{38}
\end{equation*}
$$

Similarly, sending $\Delta \rightarrow \infty$, (36) implies that

$$
\begin{equation*}
\epsilon>\epsilon^{\prime}>0 \quad \Rightarrow \quad H\left(\Theta_{\epsilon}\right) \geq H\left(\Theta_{\epsilon^{\prime}}\right) . \tag{39}
\end{equation*}
$$

This means that, as $H\left(\Theta_{\epsilon}\right) \geq 0$ for all $\epsilon>0, H\left(\Theta_{\epsilon}\right)$ is convergent as $\epsilon \downarrow 0$. In particular, this together with (38) implies that there is a sequence of positive values $\epsilon_{k}$ converging to zero such that $\Theta_{\epsilon_{k}}$ converges to some $\Theta_{0}: \mathcal{H}_{* *} \rightarrow[0,1]$ in $L^{2}(\vec{\mu})$. Therefore, there is a subsequence of this sequence converging to $\Theta_{0} \vec{\mu}$-almost everywhere. Without loss of generality, we may assume this subsequence is the whole sequence. With this, $\Theta_{\epsilon_{k}}$ converges to $\Theta_{0}$ both in $L^{2}(\vec{\mu})$ and $\vec{\mu}$-almost everywhere. Using Lemma 3, we have $\partial \Theta_{\epsilon_{k}} \rightarrow \partial \Theta_{0} \mu$-almost surely. Lemma 3 then implies that for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}, \partial \Theta_{\epsilon_{k}}(H, j) \rightarrow \partial \Theta_{0}(H, j)$ for all $j \in e$. Following Remark 6, we can treat $\partial \Theta_{0}$ and the $\partial \Theta_{\epsilon_{k}}$ as functions on $\mathcal{H}_{* *}$ instead of $\mathcal{H}_{*}$. Then, Lemma 2 implies that

$$
\begin{equation*}
\partial \Theta_{\epsilon_{k}}(H, e, j) \rightarrow \partial \Theta_{0}(H, e, j) \forall j \in e, \quad \vec{\mu}-\text { a.e.. } \tag{40}
\end{equation*}
$$

Now, we are ready to show that $\Theta_{0}$ is actually $\vec{\mu}$-balanced. To do so, assume that $\partial \Theta_{0}(H, i)>$ $\partial \Theta_{0}(H, j)$ for some $i, j \in e$. Equivalently, $\partial \Theta_{0}(H, e, i)>\partial \Theta_{0}(H, e, j)$. Then (40)implies that, outside a measure zero set, for some fixed $\delta$, and for $k$ large enough,

$$
\partial \Theta_{\epsilon_{k}}(H, e, i)-\partial \Theta_{\epsilon_{k}}(H, e, j)>\delta \quad \vec{\mu} \text {-a.e.. }
$$

On the other hand, using the definition of an $\epsilon$-balanced allocation, we have

$$
\Theta_{\epsilon_{k}}(H, e, i) \leq \frac{1}{1+\exp \left(-\frac{\partial \Theta_{\epsilon_{k}}(H, e, j)-\partial \Theta_{\epsilon_{k}}(H, e, i)}{\epsilon_{k}}\right)} \leq \frac{1}{1+\exp \left(\delta / \epsilon_{k}\right)}
$$

Sending $k$ to infinity, since $\delta$ is fixed, the above inequality implies that $\Theta_{\epsilon_{k}}(H, e, i)$ converges to zero. Also, we know that $\Theta_{\epsilon_{k}} \rightarrow \Theta_{0} \vec{\mu}$-almost everywhere. Thus, we have shown that

$$
\partial \Theta_{0}(H, i)>\partial \Theta_{0}(H, j) \text { for } i, j \in e \quad \Rightarrow \quad \Theta_{0}(H, e, i)=0, \quad \vec{\mu} \text {-a.e., }
$$

which shows that $\Theta_{0}$ is balanced and the proof is complete.

### 6.2 Variational characterization

In this section, we prove the variational characterization (part 2) of Theorem 1.
Proposition 11. Assume $\Theta$ is a Borel allocation on $\mathcal{H}_{* *}$ and $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ is unimodular with $\operatorname{deg}(\mu)<$ $\infty$. Then, for all $t \in \mathbb{R}$ we have

$$
\begin{equation*}
\int(\partial \Theta-t)^{+} d \mu \geq \sup _{\substack{f: \mathcal{H}_{* \rightarrow[0,1]}^{\text {Borel }}}} \int \tilde{f}_{\text {min }} d \vec{\mu}-t \int f d \mu \tag{41}
\end{equation*}
$$

where $\tilde{f}_{\text {min }}$ is defined as

$$
\tilde{f}_{\min }(H, e, i)=\frac{1}{|e|} \min _{j \in e} f(H, j)
$$

Furthermore, equality happens for all $t \in \mathbb{R}$ if and only if $\Theta$ is balanced. Moreover, when $\Theta$ is balanced, the function $f=\mathbb{1}_{\partial \Theta>t}$ achieves the supremum.

Proof. Note that for any real number $x$, we have $x^{+} \geq x y$ for $y \in[0,1]$. Therefore, for any Borel function $f: \mathcal{H}_{*} \rightarrow[0,1]$, we have

$$
\begin{equation*}
\int(\partial \Theta-t)^{+} d \mu \geq \int(\partial \Theta-t) f d \mu=\int f \partial \Theta d \mu-t \int f d \mu \tag{42}
\end{equation*}
$$

The assumption that $\operatorname{deg}(\mu)<\infty$ guarantees that both integrals on the RHS are finite, and hence $\int(\partial \theta-t) f d \mu$ exists. It is easy to see that $f \partial \Theta=\partial(\tilde{f} \Theta)$, where $\tilde{f}: \mathcal{H}_{* *} \rightarrow \mathbb{R}$ is defined as $\tilde{f}(H, e, i):=$ $f(H, i)$. Therefore, using the unimodularity of $\mu$, we have

$$
\begin{equation*}
\int f \partial \Theta d \mu=\int \tilde{f} \Theta d \vec{\mu}=\int \nabla(\tilde{f} \Theta) d \vec{\mu} \tag{43}
\end{equation*}
$$

Now, we have

$$
\begin{align*}
\nabla(\tilde{f} \Theta)(H, e, i) & =\frac{1}{|e|} \sum_{j \in e} \tilde{f}(H, e, j) \Theta(H, e, j) \\
& =\frac{1}{|e|} \sum_{j \in e} f(H, j) \Theta(H, e, j)  \tag{44}\\
& \stackrel{(a)}{\geq} \frac{1}{|e|} \min _{j \in e} f(H, j) \\
& =\tilde{f}_{\min }(H, e, i)
\end{align*}
$$

where (a) holds since $\sum_{j \in e} \Theta(H, e, j)=1$ and $\Theta(H, e, j) \geq 0$ for all $j \in e$. This, together with (42) and (43), proves (41).

Now, we show that equality holds if and only if $\Theta$ is balanced. First, assume $\Theta$ is balanced. We show that equality holds in (41) for all $t \in \mathbb{R}$. Take $f=\mathbb{1}_{\partial \Theta>t}$. Then, (42) becomes an equality. Therefore, it remains to show that (44) also becomes an equality, $\vec{\mu}$-almost everywhere. Note that if $f(H, j)=0$ for all $j \in e$ or $f(H, j)=1$ for all $j \in e$, equality holds in (44). Thereby, it suffices to consider the case that for some $j \in e, f(H, j)=0$ and for some other $j^{\prime} \in e, f\left(H, j^{\prime}\right)=1$. This implies $\partial \Theta\left(H, j^{\prime}\right) \geq t>\Theta(H, j)$. As $\Theta$ is balanced, outside a measure zero set, we can conclude from the above that $\Theta\left(H, e, j^{\prime}\right)=0$, and so its contribution in (44) vanishes. Since this is true for any $j^{\prime}$ with $f\left(H, j^{\prime}\right)=1$, both sides of the inequality (a) in (44) become equal to zero. As this argument holds outside a measure zero set, the above discussion shows that (41) becomes an equality when $\Theta$ is balanced. Moreover, the function $f=\mathbb{1}_{\partial \Theta>t}$ achieves the supremum.

Now, we show that if (41) is an equality for all $t \in \mathbb{R}$, then $\Theta$ is balanced. Fix some $t \in \mathbb{R}$. First, we claim that the supremum on the RHS of (41) is a maximum. To see this, note that we have already shown in Section 6.1 that a balanced allocation $\Theta_{0}$ with respect to $\mu$ exists, and the above discussion then implies that $f=\mathbb{1}_{\partial \Theta_{0}>t}$ achieves the equality in (41) with $\Theta$ being replaced with $\Theta_{0}$. But the RHS of (41) has no dependence on $\Theta$. Thereby, the RHS of (41) always has $f:=\mathbb{1}_{\partial \Theta_{0}>t}$ as a maximizer, whenever $\mu$ is unimodular with bounded mean. Since $f \in[0,1]$, we have $(\partial \Theta-t)^{+} \geq(\partial \Theta-t) f$ pointwise. However, equality holds in (41), so by (44), we must have

$$
\begin{equation*}
(\partial \Theta-t) f=(\partial \Theta-t)^{+} \quad \mu-\text { a.s.. } \tag{45}
\end{equation*}
$$

Using Lemma 1 , this means that

$$
\begin{equation*}
(\partial \Theta-t) \tilde{f}=(\partial \Theta-t)^{+} \quad \vec{\mu}-\text { a.e. } \tag{46}
\end{equation*}
$$

where, by abuse of notation, we have treated $\partial \Theta$ as being defined on $\mathcal{H}_{* *}$ via $\partial \Theta(H, e, i)=\partial \Theta(H, i)$. On the other hand, (44) must be an equality and

$$
\begin{equation*}
\nabla(\tilde{f} \Theta)=\tilde{f}_{\min } \quad \vec{\mu}-\text { a.e.. } \tag{47}
\end{equation*}
$$

Now, we show that $\Theta$ must be balanced by checking the conditions in Defintion 24. For the above fixed $t$, if for some $[H, e, i] \in \mathcal{H}_{* *}$ and some $j \in e$ we have

$$
\partial \Theta(H, e, i)>t>\partial \Theta(H, e, j)
$$

then (46) implies that outside a measure zero set, we can conclude that $f(H, i)=1$ and $f(H, j)=0$. This can be seen by comparing the left hand side and right hand side of (46) in each of the two cases, and then recalling that $\tilde{f}(H, e, i)=f(H, i)$.

Hence, by definition, $\tilde{f}_{\text {min }}(H, e, i)=0$. Therefore, (47) implies that outside a measure zero set, we have

$$
\begin{aligned}
0=\nabla(\tilde{f} \Theta)(H, e, i) & =\frac{1}{|e|} \sum_{k \in e} f(H, k) \Theta(H, e, k) \\
& \geq \frac{1}{|e|} f(H, i) \Theta(H, e, i) \\
& =\frac{1}{|e|} \Theta(H, e, i)
\end{aligned}
$$

which implies that $\Theta(H, e, i)=0$.
So far we have shown that, for a fixed $t$, for almost all $[H, e, i] \in \mathcal{H}_{* *}$, if for some $j \in e$ we have $\Theta(H, e, i)>t>\Theta(H, e, j)$, then $\Theta(H, e, i)=0$. Since this holds for all $t \in \mathbb{Q}$, and $\mathbb{Q}$ is countable and dense in $\mathbb{R}$, we can conclude that for $\vec{\mu}$-almost every $[H, e, i] \in \mathcal{H}_{* *}, \partial \Theta(H, i)>\partial \Theta(H, j)$ for some $j \in e$ implies that $\Theta(H, e, i)=0$. Using Proposition 3, we can conclude that for $\vec{\mu}$-almost all $[H, e, i] \in \mathcal{H}_{* *}, \partial \Theta\left(H, e, j_{1}\right) \geq \partial \Theta\left(H, e, j_{2}\right)$ for some $j_{1}, j_{2} \in e$ implies that $\partial \Theta\left(H, j_{1}\right)=0$. Thus, $\Theta$ is balanced.

### 6.3 Optimality

In this section, we prove part 3 of Theorem 1. We do this in three steps:
$(a) \Rightarrow(c)$ Let $\Theta$ be a balanced allocation with respect to $\mu$. Let $\Theta^{\prime}$ be any other allocation, and $f:[0, \infty) \rightarrow[0, \infty)$ be convex. We will show that

$$
\int f \circ \partial \Theta d \mu \leq \int f \circ \partial \Theta^{\prime} d \mu
$$

Using Proposition 11, since $\Theta$ is balanced and $\Theta^{\prime}$ is a Borel allocation we have, for any $t \in \mathbb{R}$,

$$
\begin{equation*}
\int(\partial \Theta-t)^{+} d \mu=\sup _{\substack{f: \mathcal{H}_{*} \rightarrow[0,1] \\ \text { Borel }}} \int \tilde{f}_{\min } d \vec{\mu}-t \int f d \mu \leq \int\left(\partial \Theta^{\prime}-t\right)^{+} d \mu \tag{48}
\end{equation*}
$$

On the other hand,

$$
\int \partial \Theta d \mu=\int \Theta d \vec{\mu}=\int \nabla \Theta d \vec{\mu}=\int \frac{1}{|e|} d \vec{\mu}
$$

The same chain of equalities holds for $\Theta^{\prime}$. Hence, $\int \partial \Theta d \vec{\mu}=\int \partial \Theta^{\prime} d \vec{\mu}$. Standard results in the convex ordering of random variables (see Theorem 3.A. 1 of [SS07] for instance) show that (48) implies that
$\Theta \preceq_{c x} \Theta^{\prime}$, where $\preceq_{c x}$ denotes the partial order defined by the property that $\int f \circ \Theta d \mu \leq \int f \circ \Theta^{\prime} d \mu$ for any convex function $f$.
$(c) \Rightarrow(b)$ Just restrict to the given strictly convex function.
$\overline{(b) \Rightarrow(a)}$ Assume $\Theta_{0}$ is a balanced allocation w.r.t. $\mu$ (from Section 6.1 we know it exists). As we showed above, $\int f \circ \partial \Theta_{0} d \mu \leq \int f \circ \partial \Theta d \mu$ for the given strictly convex function $f:[0, \infty) \rightarrow[0, \infty)$. But $\Theta$ is a minimizer. Therefore,

$$
\int f \circ \partial \Theta d \mu=\int f \circ \partial \Theta_{0} d \mu
$$

Now, define $\Theta^{\prime}: \mathcal{H}_{* *} \rightarrow[0,1]$ as $\Theta^{\prime}=\left(\Theta+\Theta_{0}\right) / 2$. Note that $\Theta^{\prime}$ is a Borel allocation. Also, using the convexity of $f$, we have

$$
\begin{equation*}
\int f \circ \partial \Theta^{\prime} d \mu \leq \frac{1}{2}\left(\int f \circ \partial \Theta d \mu+\int f \circ \partial \Theta_{0} d \mu\right)=\int f \circ \partial \Theta_{0} d \mu \tag{49}
\end{equation*}
$$

On the other hand, since $\Theta^{\prime}$ is a Borel allocation and $\Theta_{0}$ is balanced, the $(a) \Rightarrow(c)$ part implies that

$$
\int f \circ \partial \Theta_{0} d \mu \leq \int f \circ \partial \Theta^{\prime} d \mu
$$

Hence, we have equality in (49). As $f$ is nonnegative, it must be the case that

$$
f \circ \partial \Theta^{\prime}=\frac{f \circ \partial \Theta_{0}+f \circ \partial \Theta}{2} \quad \mu-\text { a.s.. }
$$

Now, since $f$ is strictly convex, this implies that $\partial \Theta=\partial \Theta_{0}, \mu$-almost surely. Therefore, for any value of $t$, we have

$$
\int(\partial \Theta-t)^{+} d \mu=\int\left(\partial \Theta_{0}-t\right)^{+} d \mu
$$

Comparing this with (41), we realize that $\Theta$ achieves the equality therein, which means, from Proposition 11 , that $\Theta$ is balanced.

### 6.4 Uniqueness of $\partial \Theta$

In this section we prove part 4 of Theorem 1. If $\Theta$ is balanced, part 3 of the theorem, which was proved in the preceding section, implies that for a strictly convex function $f, \int f \circ \partial \Theta d \mu$ is minimized. Then, in the proof of the $(b) \Rightarrow(a)$ part of Section 6.3 it was shown that $\partial \Theta=\partial \Theta_{0} \mu$-almost surely. This is precisely what we want to show.

For the converse, assume that $\partial \Theta=\partial \Theta_{0}, \mu$-almost surely. Then, for all $t$, we have

$$
\int(\partial \Theta-t)^{+} d \mu=\int\left(\partial \Theta_{0}-t\right)^{+} d \mu
$$

so Proposition 11 implies that $\Theta$ is balanced, by the same logic that was used at the end of the preceding section.

### 6.5 Continuity with respect to the local weak limit

In this section we prove Part 5 of Theorem 1. Prior to that, we need some notation and tools.
Recall the notion of marked hypergraphs from Section 2.10. Let $\Pi: \overline{\mathcal{H}}_{*}(\Xi) \rightarrow \mathcal{H}_{*}$ be the function that removes marks, i.e.

$$
\begin{equation*}
\Pi([\bar{H}, i])=[H, i] \tag{50}
\end{equation*}
$$

where $H$ is the underlying hypergraph associated to $\bar{H}$. It can be easily checked that $\Pi$ is a continuous map.

Note that allocations could be considered as marks with values in $[0,1]$. Hence, we can capture the notion of a balanced allocation using the formalism in Section 2.10, via the following definition.

Consider the function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow \Xi$ defined as $f(\bar{H}, e, i)=\xi_{\bar{H}}(e, i)$. When $\Xi$ has an additive structure (e.g. $\Xi=[0,1]$ ) we may consider $\partial f: \overline{\mathcal{H}}_{*}(\Xi) \rightarrow \Xi$, defined as

$$
\begin{equation*}
\partial f(\bar{H}, i)=\sum_{e \in E(\bar{H}), e \ni i} f(\bar{H}, e, i)=\sum_{e \in E(\bar{H}), e \ni i} \xi_{\bar{H}}(e, i) . \tag{51}
\end{equation*}
$$

By abuse of notation, we may write $\partial \xi_{\bar{H}}(i)$ instead of $\partial f(\bar{H}, i)$ with the above $f$.
Definition 32. A measure $\bar{\mu} \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}([0,1])\right)$ is called balanced if for $\vec{\mu}$-almost every $[\bar{H}, e, i] \in$ $\overline{\mathcal{H}}_{*}([0,1])$, we have

$$
\begin{equation*}
\sum_{j \in e} \xi_{\bar{H}}(e, j)=1 \tag{52}
\end{equation*}
$$

and

$$
\begin{equation*}
\partial \xi_{\bar{H}}(j)>\partial \xi_{\bar{H}}\left(j^{\prime}\right) \quad \text { for some } j, j^{\prime} \in e \quad \Longrightarrow \quad \xi_{\bar{H}}(e, j)=0 \tag{53}
\end{equation*}
$$

Before proving the result of this section, we state the following lemmas and postpone their proofs until the end of this section.

Lemma 8. If $\bar{\mu}_{n}$ is a sequence of balanced probability measures on $\overline{\mathcal{H}}_{*}([0,1])$ with weak limit $\bar{\mu}$, then $\bar{\mu}$ is also balanced.

Lemma 9. Assume $\theta_{1}, \ldots, \theta_{n}$ and $\theta_{1}^{\prime}, \ldots, \theta_{n}^{\prime}$ are non-negative real numbers such that $\sum_{i} \theta_{i}=\sum_{i} \theta_{i}^{\prime}=$ 1. Also, assume that nonnegative real numbers $l_{1}, \ldots, l_{n}$ and $l_{1}^{\prime}, \ldots, l_{n}^{\prime}$ are given such that for $1 \leq i, j \leq$ $n, l_{i}>l_{j}$ implies $\theta_{i}=0$. Similarly, assume that $l_{i}^{\prime}>l_{j}^{\prime}$ implies $\theta_{i}^{\prime}=0$. Then, we have

$$
\sum_{i=1}^{n}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} \leq 0
$$

Lemma 10. Assume $K$ is a compact subset of $\mathcal{H}_{*}$ and $\Xi$ is a compact metric space. Define $\bar{K} \subset \overline{\mathcal{H}}_{*}(\Xi)$ as

$$
\bar{K}:=\left\{[\bar{H}, i] \in \overline{\mathcal{H}}_{*}(\Xi):[H, i] \in K\right\}
$$

Then, $\bar{K}$ is compact in $\overline{\mathcal{H}}_{*}(\Xi)$.
Proposition 12. Let $\left\{H_{n}\right\}_{n \geq 1}$ be a sequence of finite hypergraphs with local weak limit $\mu$. Then, if $\mathcal{L}_{H_{n}}$ denotes the distribution of balanced load on $H_{n}$ with a vertex chosen uniformly at random, and $\mathcal{L}$ is the law of $\partial \Theta$ corresponding to the balanced allocation $\Theta$ on $\mu$, we have

$$
\mathcal{L}_{H_{n}} \Rightarrow \mathcal{L}
$$

Proof. Define $\theta_{n}$ to be a balanced allocation on $H_{n}$ and $\bar{H}_{n}$ to be the marked hypergraph obtained by adding $\theta_{n}$ to $H_{n}$ as edge marks, i.e. $\xi_{\bar{H}_{n}}(e, i)=\theta_{n}(e, i)$ for $(e, i) \in \Psi(H)$. Note that $H_{n}$ is a finite hypergraph; hence, $\theta_{n}$ is well defined and $\partial \theta_{n}(i)$ is unique for $i \in V\left(H_{n}\right)$. Now, define $\bar{\mu}_{n} \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}([0,1])\right)$ to be the distribution of $\left[\bar{H}_{n}, v\right]$ where $v \in V\left(H_{n}\right)$ is chosen uniformly at random. We claim that $\left\{\bar{\mu}_{n}\right\}_{n=1}^{\infty}$ is a tight sequence, which means that it has a convergent subsequence. Since $\mu_{n}$ converges weakly to $\mu$, Prokhorov's theorem implies that $\mu_{n}$ is tight in $\mathcal{P}\left(\mathcal{H}_{*}\right)$ (see, for instance,
[Bil99, Theorems 5.1 and 5.2]). Consequently, for $\epsilon>0$, there is a compact set $K \subset \mathcal{H}_{*}$ such that $\mu_{n}\left(K^{c}\right) \leq \epsilon$ for all $n$. Define

$$
\bar{K}:=\left\{[\bar{H}, i] \in \overline{\mathcal{H}}_{*}([0,1]):[H, i] \in K\right\}
$$

From Lemma $10, \bar{K}$ is compact in $\overline{\mathcal{H}}_{*}([0,1])$. It is easy to see that $\bar{\mu}_{n}\left(\bar{K}^{c}\right)=\mu_{n}\left(K^{c}\right)$ which means that $\bar{\mu}_{n}$ is a tight sequence. Hence, it has a subsequence converging weakly to some $\bar{\mu} \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}\right)$. In order to simplify the notation, assume that this subsequence is the whole sequence.

With the projection map $\Pi$ defined in (7), define $\nu_{n}$ and $\nu$ to be the pushforward measures on $\mathcal{H}_{*}$ corresponding to $\bar{\mu}_{n}$ and $\bar{\mu}$, respectively. As $\Pi$ removes marks, we have $\nu_{n}=\mu_{n}$ and $\nu=\mu$. Now, note that $\bar{\mu}_{n} \Rightarrow \bar{\mu}$ implies $(\Pi)_{*} \bar{\mu}_{n} \Rightarrow(\Pi)_{*} \bar{\mu}$ which means that $\mu_{n} \Rightarrow(\Pi)_{*} \bar{\mu}$. On the other hand, we know $\mu_{n} \Rightarrow \mu$. Thereby, $(\Pi)_{*} \bar{\mu}=\mu$.

Note that, with the above construction, $\mathcal{L}_{H_{n}}$ is the pushforward of $\bar{\mu}_{n}$ under the mapping $[\bar{H}, i] \mapsto$ $\partial \xi_{\bar{H}}(i)$ defined in (51). Therefore, as $\bar{\mu}_{n} \Rightarrow \bar{\mu}, \mathcal{L}_{H_{n}}$ converges weakly to the law of $\partial \xi_{\bar{H}}(i)$ under $\bar{\mu}$. Consequently, to show that $\mathcal{L}_{H_{n}} \Rightarrow \mathcal{L}$, it suffices to show that

$$
\begin{equation*}
\partial \xi_{\bar{H}}(i)=\partial \Theta_{0}([H, i]) \quad \bar{\mu}-\text { a.s.. } \tag{54}
\end{equation*}
$$

From Proposition 20 in Appendix E we know that $\bar{\mu}$ is unimodular. Hence, we have

$$
\begin{aligned}
& \int\left(\partial \xi_{\bar{H}}(i)-\partial \Theta_{0}([H, i])\right)^{+} d \bar{\mu}(\bar{H}, i) \\
& \quad=\int\left(\partial_{\bar{H}} \xi(i)-\partial \Theta_{0}([H, i])\right) \mathbb{1}_{\partial \xi_{\bar{H}}(i)>\partial \Theta_{0}([H, i])} d \bar{\mu}([\bar{H}, i]) \\
& \quad=\int\left(\xi_{\bar{H}}(e, i)-\Theta_{0}([H, e, i])\right) \mathbb{1}_{\partial \xi_{\bar{H}}(i)>\partial \Theta_{0}([H, i])} d \vec{\mu}([\bar{H}, e, i]) \\
& \quad=\int \frac{1}{|e|} \sum_{j \in e}\left(\xi_{\bar{H}}(e, j)-\Theta_{0}([H, e, j])\right) \mathbb{1}_{\partial \xi_{\bar{H}}(j)>\partial \Theta_{0}([H, j])} d \overrightarrow{\vec{\mu}}([\bar{H}, e, i)
\end{aligned}
$$

where the last equality follows from unimodularity of $\bar{\mu}$. From Lemma $8, \bar{\mu}$ is balanced in the sense of Definition 32 above. Also, Lemma 9 together with the balancedness of $\bar{\mu}$ and $\Theta_{0}$, implies that the integrand is non-positive almost everywhere, which means that $\partial \xi_{\bar{H}}(i) \leq \partial \Theta_{0}([H, i]) \bar{\mu}$-almost surely. It could be proved in a similar fashion that $\partial \Theta_{0}(H, i) \leq \partial \xi_{\bar{H}}(i) \bar{\mu}$-almost surely. This proves (54).

So far, we have shown that $\mathcal{L}_{H_{n}}$ has a subsequence $\mathcal{L}_{H_{n_{k}}}$ such that $\mathcal{L}_{H_{n_{k}}} \Rightarrow \mathcal{L}$. This argument could be repeated for any subsequence of $H_{n}$, i.e. any subsequence $H_{n_{k}}$ has a further subsequence $H_{n_{k_{l}}}$ such that $\mathcal{L}_{H_{n_{k_{l}}}} \Rightarrow \mathcal{L}$. This implies that $\mathcal{L}_{H_{n}} \Rightarrow \mathcal{L}$ (see, for instance, Theorem 2.2 in [Bil71]).

Proof of Lemma 8. First, we show that $\bar{\mu}$ satisfies (53). Define

$$
A_{* *}:=\left\{[\bar{H}, e, i] \in \overline{\mathcal{H}}_{* *}([0,1]): \forall j, j^{\prime} \in e, \partial \xi_{\bar{H}}(j)>\partial \xi_{\bar{H}}\left(j^{\prime}\right) \Rightarrow \xi_{\bar{H}}(e, j)=0\right\}
$$

and

$$
A_{*}:=\left\{[\bar{H}, i] \in \overline{\mathcal{H}}_{*}([0,1]): \forall e \ni i, j, j^{\prime} \in e, \partial \xi_{\bar{H}}(j)>\partial \xi_{\bar{H}}\left(j^{\prime}\right) \Rightarrow \xi_{\bar{H}}(e, j)=0\right\}
$$

We claim balancedness of $\bar{\mu} \in \mathcal{P}\left(\mathcal{H}_{*}([0,1])\right)$ is equivalent to $\bar{\mu}\left(A_{*}\right)=1$. By definition, $\bar{\mu}$ being balanced means $\bar{\mu}\left(A_{* *}\right)=\bar{\mu}\left(\overline{\mathcal{H}}_{* *}([0,1])\right)$. This is equivalent to $\int \mathbb{1}_{A_{* *}} d \vec{\mu}=\int 1 d \vec{\mu}$, or $\int \partial \mathbb{1}_{A_{* *}} d \bar{\mu}=\int \operatorname{deg}_{H}(i) d \bar{\mu}$. But we have $\partial \mathbb{1}_{A_{* *}} \leq \operatorname{deg}_{H}(i)$ pointwise. Therefore, $\partial \mathbb{1}_{A_{* *}}(H, i, \xi)=\operatorname{deg}_{H}(i), \bar{\mu}$-almost surely. This is equivalent to $[\bar{H}, e, i] \in A_{* *}$ for all $e \ni i, \bar{\mu}$-almost surely, or $[\bar{H}, i] \in A_{*} \bar{\mu}$-almost surely which is in turn equivalent to $\bar{\mu}\left(A_{*}\right)=1$.

Now, we claim that $A_{*}$ is closed in $\overline{\mathcal{H}}_{*}([0,1])$ (with respect to the topology induced by the distance $\bar{d}_{*}$ defined in Section 2.10). Equivalently, we show that $A_{*}^{c}$ is open. Take an arbitrary $[\bar{H}, i] \in A_{*}^{c}$.

Since $[\bar{H}, i] \notin A_{*}$, there exists $e \ni i$ and $j, j^{\prime} \in e$ such that $\partial \xi_{\bar{H}}(j)>\partial \xi_{\bar{H}}\left(j^{\prime}\right)$ but $\xi_{\bar{H}}(e, j)>0$. If $\left[\bar{H}^{\prime}, i^{\prime}\right] \in \overline{\mathcal{H}}_{*}([0,1])$ is such that $d:=\bar{d}_{*}\left([\bar{H}, i],\left[\bar{H}^{\prime}, i^{\prime}\right]\right)$ satisfies $d<1 / 4$, then $[H, i] \equiv_{3}\left[H^{\prime}, i^{\prime}\right]$ (recall that $H$ and $H^{\prime}$ are the underlying unmarked hypergraphs associated to $\bar{H}$ and $\bar{H}^{\prime}$, respectively). Moreover, if we define

$$
\epsilon=\min \left\{\frac{\partial \xi_{\bar{H}}(j)-\partial \xi_{\bar{H}}\left(j^{\prime}\right)}{3}, \frac{\xi_{\bar{H}}(e, j)}{2}\right\},
$$

and

$$
\Delta=\max _{k: d_{H}(i, k) \leq 3} \operatorname{deg}_{H}(k),
$$

then if

$$
d<\frac{1}{1+\frac{\Delta}{\epsilon}}
$$

one can check that $\partial \xi_{\bar{H}^{\prime}}(\phi(j))>\partial \xi_{\bar{H}^{\prime}}\left(\phi\left(j^{\prime}\right)\right)$ while $\xi_{\bar{H}^{\prime}}(\phi(e), \phi(j))>0$, where $\phi$ is the local isomorphism between $H$ and $H^{\prime}$ following from $[H, i] \equiv \equiv_{3}\left[H^{\prime}, i^{\prime}\right]$. This in particular means that $\left[\bar{H}^{\prime}, i^{\prime}\right] \notin A_{*}$. Consequently, the ball with radius $\min \{1 / 4,1 /(1+(\Delta / \epsilon))\}$ around $[\bar{H}, i]$ is in $A_{*}^{c}$. This means that $A_{*}$ is closed.

Notice that since $\bar{\mu}_{n}$ is balanced, $\bar{\mu}_{n}\left(A_{*}\right)=1$. On the other hand, as $\bar{\mu}_{n} \Rightarrow \bar{\mu}$ and $A_{*}$ is closed, we have

$$
\bar{\mu}\left(A_{*}\right) \geq \limsup \bar{\mu}_{n}\left(A_{*}\right)=1
$$

which means that $\bar{\mu}\left(A_{*}\right)=1$ and $\bar{\mu}$ is balanced.
Now we turn to showing that $\bar{\mu}$ also satisfies (52). Similar to the above approach, if we define

$$
B_{* *}:=\left\{[\bar{H}, e, i] \in \mathcal{H}_{* *}([0,1]): \sum_{j \in e} \xi_{\bar{H}}(e, j)=1\right\},
$$

and

$$
B_{*}:=\left\{[\bar{H}, i] \in \mathcal{H}_{*}([0,1]):[\bar{H}, e, i] \in B_{* *} \forall e \ni i\right\},
$$

we can show that for $\overrightarrow{\vec{\mu}}$-almost every $[\bar{H}, e, i] \in \mathcal{H}_{* *}([0,1])$ we have $\sum_{j \in e} \xi_{\bar{H}}(e, j)=1$. Hence, all the conditions in Definition 32 are satisfied and $\bar{\mu}$ is balanced.

Proof of Lemma 9. Define $L:=\min _{i} l_{i}$ and $A:=\left\{1 \leq i \leq n: l_{i}=L\right\}$. Likewise, let $L^{\prime}:=\min _{i} l_{i}^{\prime}$ and $A^{\prime}:=\left\{1 \leq i \leq n: l_{i}^{\prime}=L^{\prime}\right\}$. The given condition implies that $\theta_{j}=0$ for $j \notin A$ and similarly $\theta_{j}^{\prime}=0$ for $j \notin A^{\prime}$.

First assume that $L>L^{\prime}$. In this case, we have

$$
\begin{aligned}
\sum_{i}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} & \stackrel{(a)}{=} \sum_{i \in A \cup A^{\prime}}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} \\
& =\sum_{i \in A \backslash A^{\prime}}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}}+\sum_{i \in A^{\prime}}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} \\
& \stackrel{(b)}{=} \sum_{i \in A \backslash A^{\prime}} \theta_{i} \mathbb{1}_{L>l_{i}^{\prime}}+\sum_{i \in A^{\prime}}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} \\
& \stackrel{(c)}{=} \sum_{i \in A \backslash A^{\prime}} \theta_{i} \mathbb{1}_{L>l_{i}^{\prime}}+\sum_{i \in A^{\prime}}\left(\theta_{i}-\theta_{i}^{\prime}\right) \\
& \leq \sum_{i \in A \backslash A^{\prime}} \theta_{i}+\sum_{i \in A^{\prime}} \theta_{i}-\theta_{i}^{\prime} \\
& =1-1=0
\end{aligned}
$$

where (a) holds since $\theta$ and $\theta^{\prime}$ are zero outside $A \cup A^{\prime},(b)$ uses $\theta_{i}^{\prime}=0$ for $i \notin A^{\prime}$, and $(c)$ uses the fact that for $i \in A^{\prime}, l_{i} \geq L>L^{\prime}=l_{i}^{\prime}$.

Now, for the case $L \leq L^{\prime}$ we have

$$
\begin{aligned}
\sum_{i=1}^{n}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} & =\sum_{i \in A}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}}+\sum_{i \in A^{\prime} \backslash A}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} \\
& \stackrel{(a)}{=} \sum_{i \in A^{\prime} \backslash A}\left(\theta_{i}-\theta_{i}^{\prime}\right) \mathbb{1}_{l_{i}>l_{i}^{\prime}} \\
& \stackrel{(b)}{=} \sum_{i \in A^{\prime} \backslash A}-\theta_{i}^{\prime} \mathbb{1}_{l_{i}>l_{i}^{\prime}} \\
& \leq 0
\end{aligned}
$$

where (a) follows from the fact that for $i \in A, l_{i}=L \leq L^{\prime} \leq l_{i}^{\prime}$ and (b) employs $\theta_{i}=0$ for $i \notin A$. This completes the proof.

Proof of Lemma 10. We take a sequence $\left[\bar{H}_{n}, i_{n}\right]$ in $\bar{K}$ and try to find a converging subsequence. Notice that $\left[H_{n}, i_{n}\right]$ is a sequence in $K$. Therefore, it has a convergent subsequence. Without loss of generality, assume this subsequence is the whole sequence, converging to some $[H, i] \in \mathcal{H}_{*}$. By reducing to a further subsequence, we may assume that $d_{\mathcal{H}_{*}}\left(\left[H_{n}, i_{n}\right],[H, i]\right) \leq 1 /(n+1)$. This means that, for all $m \geq n$, we have $\left(H_{m}, i_{m}\right) \equiv_{n}(H, i)$. Since $(H, i)$ is locally finite, there are finitely many marks up to level $n$ in $(H, i)$ which all have values in the compact space $\Xi$. Hence, there is a subsequence where the marks in the first $n$ levels in $\bar{H}_{n}$ are convergent. With this, we may associate marks to $(H, i)_{n}$ using the limiting values. Since this is true for all $n$, via a diagonalization argument we may construct a marked rooted hypergraph $(\bar{H}, i)$ together with a subsequence $\left(\bar{H}_{m_{l}}, i_{m_{l}}\right)$ such that the underlying unmarked hypergraph is identical to $H$ and also, for any integer $k$, the marks up to depth $k$ in $\left(\bar{H}_{m_{l}}, i_{m_{l}}\right)$ converge to those in $(\bar{H}, i)$ as $m \rightarrow \infty$. This means that $\left[\bar{H}_{m_{l}}, i_{m_{l}}\right] \rightarrow[\bar{H}, i]$ and completes the proof.

## 7 Response Functions

Let $H$ be a fixed hypergraph (not necessary bounded) and $i \in V(H)$ be a vertex in $H$. Fix $\epsilon>0$. The response function $\rho_{(H, i)}^{\epsilon}: \mathbb{R} \rightarrow \mathbb{R}$ is defined as follows: $\rho_{(H, i)}^{\epsilon}(t)$ is the total load at node $i$ corresponding to the canonical $\epsilon$-balanced allocation with respect to an external load with value $t$ at node $i$ (recall the definition of canonical $\epsilon$-balanced allocations from Section 4.3). More precisely, given $t \in \mathbb{R}$, let $b_{t, i}: V(H) \rightarrow \mathbb{R}$ be the baseload function such that $b_{t, i}(i)=t$ and $b_{t, i}(j)=0$ for $j \neq i$. Moreover, let $\theta_{\epsilon}^{b_{t, i}}$ be the canonical $\epsilon$-balanced allocation on $H$ with respect to the baseload $b_{t, i}$, as was defined in Section 4.3. We then define

$$
\rho_{(H, i)}^{\epsilon}(t):=\partial_{b_{t, i}} \theta_{\epsilon}^{b_{t, i}}(i) .
$$

It turns out that this function has the following properties:
Proposition 13. Given a vertex rooted hypergraph $(H, i)$, for any $\epsilon>0$ and $x<y$, we have

$$
\begin{equation*}
0 \leq \rho_{(H, i)}^{\epsilon}(y)-\rho_{(H, i)}^{\epsilon}(x) \leq y-x \tag{55}
\end{equation*}
$$

Also,

$$
\begin{equation*}
0 \leq \rho_{(H, i)}^{\epsilon}(x)-x \leq \operatorname{deg}_{H}(i) \tag{56}
\end{equation*}
$$

Proof. Let $\theta_{\epsilon}^{b_{x, i}}$ be the canonical $\epsilon$-balanced allocation with respect to the baseload $b_{x, i}$, as defined above. Then, by definition, $\rho_{(H, i)}^{\epsilon}(x)=\partial_{b_{x, i}} \theta_{\epsilon}^{b_{x, i}}(i)$. Let $b_{y, i}$ and $\theta_{\epsilon}^{b_{y, i}}$ be defined similarly, with $x$ being replaced by $y$. As $x<y$, Proposition 6 implies that $\partial_{b_{x, i}} \theta_{\epsilon}^{b_{x, i}}(i) \leq \partial_{b_{y, i}} \theta_{\epsilon}^{b_{y, i}}(i)$, which means $\rho_{(H, i)}^{\epsilon}(x) \leq \rho_{(H, i)}^{\epsilon}(y)$.

In order to show that $\rho_{(H, i)}^{\epsilon}(y)-\rho_{(H, i)}^{\epsilon}(x) \leq y-x$, let $d_{x, i}$ be the baseload function such that $d_{x, i}(i)=0$ and $d_{x, i}(j)=-x$ for $j \in V(H), j \neq i$. Moreover, let $d_{y, i}$ be the baseload function such that $d_{y, i}(i)=0$ and $d_{y, i}(j)=-y$ for $j \in V(H), j \neq i$. Let $\theta_{\epsilon}^{b_{x, i}, \Delta}$ and $\theta_{\epsilon}^{b_{y, i}, \Delta}$ be the canonical $\epsilon$-balanced allocations on $H^{\Delta}$ with respect to $b_{x, i}$ and $b_{y, i}$, respectively. It is easy to see that, on $H^{\Delta}$, $\theta_{\epsilon}^{b_{x, i}, \Delta}$ and $\theta_{\epsilon}^{b_{y, i}, \Delta}$ are $\epsilon$-balanced with respect to $d_{x, i}$ and $d_{y, i}$ as well, respectively. This is because $b_{x, i}(j)=d_{x, i}(j)+x$ and $b_{y, i}(j)=d_{y, i}(j)+y$, for all $j \in V(H)$. As $d_{x, i}(i)=0$ and $b_{x, i}(i)=x$, we have

$$
\partial_{d_{x, i}} \theta_{\epsilon}^{b_{x, i}, \Delta}(i)=\partial_{b_{x, i}} \theta_{\epsilon}^{b_{x, i}, \Delta}(i)-x .
$$

Likewise, $\partial_{d_{y, i}} \theta_{\epsilon}^{b_{y, i}, \Delta}(i)=\partial_{b_{y, i}} \theta_{\epsilon}^{b_{y, i}, \Delta}(i)-y$. On the other hand, as $d_{x, i}(j) \geq d_{y, i}(j)$ for all $j \in V(H)$, using Proposition 5, $\partial_{d_{x, i}} \theta_{\epsilon}^{b_{x, i}, \Delta}(i) \geq \partial_{d_{y, i}} \theta_{\epsilon}^{b_{y, i}, \Delta}(i)$. Comparing with the above, this means that

$$
\partial_{b_{x, i}} \theta_{\epsilon}^{b_{x, i}, \Delta}(i)-x \geq \partial_{b_{y, i}} \theta_{\epsilon}^{b_{y, i}, \Delta}(i)-y
$$

Sending $\Delta \rightarrow \infty$, we have $\partial_{b_{y, i}} \theta_{\epsilon}^{b_{y, i}}(i)-\partial_{b_{x, i}} \theta_{\epsilon}^{b_{x, i}} \leq y-x$. This means that $\rho_{(H, i)}^{\epsilon}(y)-\rho_{(H, i)}^{\epsilon}(x) \leq y-x$.
To show (56), simply note that $\rho_{(H, i)}^{\epsilon}(x)-x$ is $\partial \theta_{\epsilon}^{b_{x, i}}(i)$, which is the sum of $\operatorname{deg}_{H}(i)$ many numbers in the interval $[0,1]$, hence is in the interval $\left[0, \operatorname{deg}_{H}(i)\right]$.

In view of Proposition 13, it is convenient to define the following class of functions:
Definition 33. For $C>0$, let $\Upsilon(C)$ denote the class of functions $g: \mathbb{R} \rightarrow \mathbb{R}$ satisfying the following conditions:
(i) $g$ is nondecreasing.
(ii) $g$ is nonexpansive, i.e. for $x \leq y$ we have $g(y)-g(x) \leq y-x$.
(iii) $0 \leq g(x)-x \leq C$.

In this notation, Proposition 13 implies that for any vertex rooted hypergraph $(H, i)$ and $\epsilon>0$,

$$
\rho_{(H, i)}^{\epsilon}(.) \in \Upsilon\left(\operatorname{deg}_{H}(i)\right)
$$

It is easy to check the following properties of this class of functions:
Lemma 11. For any function $g \in \Upsilon(C)$ we have
(i) $g$ is continuous.
(ii) $\lim _{x \rightarrow \pm \infty} g(x)= \pm \infty$.
(iii) For any $t \in \mathbb{R}$, the set $\{x \in \mathbb{R}: g(x)=t\}$ is a nonempty bounded closed interval in $\mathbb{R}$.
(iv) The function $g^{-1}: \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$
\begin{equation*}
g^{-1}(t):=\max \{x \in \mathbb{R}: g(x)=t\} \tag{57}
\end{equation*}
$$

is well defined, nondecreasing and right-continuous. Moreover, if $D$ denotes its set of discontinuities of $g^{-1}, D$ is countable. Furthermore, the set $\{x: g(x)=t\}$ has exactly one element iff $t \notin D$.
(v) Let $D$ be the set of discontinuities of $g^{-1}$ and $t \notin D$. If for some $a, a<g^{-1}(t)$, then we have $g(a)<t$. Moreover, if for some a we have $a>g^{-1}(t)$, then we have $g(a)>t$.

### 7.1 Recursion for Response functions on Hypertrees

Recall from Section 4.5 that if $T$ is a hypertree, then for $(e, i) \in \Psi(T), T_{e \rightarrow i}$ is obtained by removing $e$ from $T$ and looking at the connected component of the resulting hypertree rooted at $i$.

Given $x \in \mathbb{R}$, let $b_{x}$ be the baseload function on $T_{e \rightarrow i}$ with $b_{x}(i)=x$ and $b_{x}(j)=0, j \neq i$. With this, let $\theta_{\epsilon}^{b_{x}}$ be the canonical $\epsilon$-balanced allocation on $T_{e \rightarrow i}$ with respect to $b_{x}$. As was discussed above, $\rho_{T_{e \rightarrow i}}^{\epsilon}(x)=x+\partial \theta_{\epsilon}^{b_{x}}(i)$.
$\vec{N}$ Now we attempt to find some recursive expressions for such response functions. Before that, we need the following general lemma, whose proof is postponed to the end of this section.

Lemma 12. Given $C>0$ and a collection of nondecreasing functions $g_{i}:[0,1] \rightarrow \mathbb{R}, 1 \leq i \leq n$, the set of fixed point equations

$$
\begin{equation*}
\theta_{i}=\frac{e^{-g_{i}\left(\theta_{i}\right)}}{\sum_{j=1}^{n} e^{-g_{j}\left(\theta_{j}\right)}+C} \quad 1 \leq i \leq n \tag{58}
\end{equation*}
$$

has a unique solution $\left(\theta_{1}, \ldots, \theta_{n}\right) \in[0,1]^{n}$.
Proposition 14. Assume $T$ is a locally finite hypertree (not necessarily bounded), $\epsilon>0$, and $(e, i) \in$ $\Psi(T)$. Then, $\rho_{T_{e \rightarrow i}}^{\epsilon}($.$) is an invertible function, and for t \in \mathbb{R}$, we have

$$
\begin{equation*}
\left(\rho_{T_{e \rightarrow i}}^{\epsilon}\right)^{-1}(t)=t-\sum_{e^{\prime} \ni i: e^{\prime} \neq e}\left(1-\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{\epsilon}\right) \tag{59}
\end{equation*}
$$

where $\left\{\zeta_{e^{\prime}, j}^{\epsilon}\right\}$ for $e^{\prime} \ni i, e^{\prime} \neq e, j \in e^{\prime}, j \neq i$ are the unique solutions to the set of equations

$$
\begin{equation*}
\zeta_{e^{\prime}, j}^{\epsilon}=\frac{\exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon}\left(\zeta_{e^{\prime}, j}^{\epsilon}\right)}{\epsilon}\right)}{e^{-t / \epsilon}+\sum_{l \in e^{\prime}, l \neq i} \exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow l}}^{\epsilon}\left(\zeta_{e^{\prime}, l}^{\epsilon}\right)}{\epsilon}\right)} \tag{60}
\end{equation*}
$$

Proof. Note that Lemma 11 implies that the set $A(t):=\left\{x \in \mathbb{R}: \rho_{T_{e \rightarrow i}}^{\epsilon}(x)=t\right\}$ is not empty. Hence, in order to show that $\rho_{T_{e \rightarrow i}}^{\epsilon}($.$) is invertible, we should show that A(t)$ is a singleton for all $t \in \mathbb{R}$. By definition, $x \in A(t)$ means that we have

$$
\begin{equation*}
t=x+\sum_{e^{\prime} \ni i, e^{\prime} \neq e}\left(1-\sum_{j \in e^{\prime}, j \neq i} \theta_{\epsilon}^{b_{x}}\left(e^{\prime}, j\right)\right) \tag{61}
\end{equation*}
$$

For each $\left(e^{\prime}, j\right)$ in the above summation, as $T_{e^{\prime} \rightarrow j}$ is a subtree of $T_{e \rightarrow i}$, we can treat $\theta_{\epsilon}^{b_{x}}$ as an allocation on $T_{e^{\prime} \rightarrow j}$ via the identity projection. With this, Proposition 8 implies that $\theta_{\epsilon}^{b_{x}}$ is the canonical $\epsilon-$ balanced allocation on $T_{e^{\prime} \rightarrow j}$ with respect to the baseload function that evaluates to $\theta^{\epsilon}\left(e^{\prime}, j\right)$ at $j$ and zero elsewhere. Thereby, by the definition of the response function, for each such pair $\left(e^{\prime}, j\right)$, we have $\partial \theta_{\epsilon}^{b_{x}}(j)=\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon}\left(\theta^{\epsilon}\left(e^{\prime}, j\right)\right)$. Consequently,

$$
\theta_{\epsilon}^{b_{x}}\left(e^{\prime}, j\right)=\frac{\exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon}\left(\theta_{\epsilon}^{b x}\left(e^{\prime}, j\right)\right)}{\epsilon}\right)}{e^{-t / \epsilon}+\sum_{l \in e^{\prime}, l \neq i} \exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow l}}^{\epsilon}\left(\theta_{\epsilon}^{b x}\left(e^{\prime}, l\right)\right)}{\epsilon}\right)}
$$

Lemma 12 guarantees that for each $e^{\prime} \ni i, e^{\prime} \neq e$, there is a unique set of solutions to these equations. From (61), we realize that for any $t \in \mathbb{R}$, there is a unique solution for $x$. This implies the invertibility of $\rho_{T_{e \rightarrow i}}^{\epsilon}$. Rearranging (61), we get (59), with $\zeta_{e^{\prime}, j}=\theta_{\epsilon}^{b_{x}}\left(e^{\prime}, j\right)$.

Now, we will send $\epsilon$ to zero in the above Proposition and show that, under some conditions, the sequence of response functions converges pointwise to a limit which satisfies a certain fixed point equation. To do so, we need the following lemma, whose proof is given at the end of this section.

Lemma 13. Assume $\left\{g_{n}\right\}_{n=1}^{\infty}$ is a sequence of functions in $\Upsilon(C)$ that converge pointwise to a function $g$, i.e. $g(x)=\lim _{n \rightarrow \infty} g_{n}(x)$ for all $x \in \mathbb{R}$. Then, $g$ is also in $\Upsilon(C)$. Furthermore, if $D_{n}$ denotes the set of discontinuities of $g_{n}$ and $D$ denotes the set of discontinuities of $g$, then for $t \notin\left(\bigcup_{n} D_{n}\right) \cup D$ we have

$$
\begin{equation*}
\lim _{n \rightarrow \infty} g_{n}^{-1}(t)=g^{-1}(t) \tag{62}
\end{equation*}
$$

Now, we can write recursive equations for the limit of $\rho_{T_{e \rightarrow i}}^{\epsilon}$, if it exists.
Proposition 15. Assume $T$ is a locally finite hypertree (not necessarily bounded) and $\epsilon_{n}$ is a sequence of positive numbers converging to zero, with $\theta_{\epsilon_{n}}$ being the canonical $\epsilon_{n}$-balanced allocation on $T$. If $l_{i}:=\lim _{n \rightarrow \infty} \partial \theta_{\epsilon_{n}}(i)$ exists for all $i \in V(T)$, then, for all $(e, i) \in \Psi(T), \rho_{T_{e \rightarrow i}}^{\epsilon_{n}}($.$) converges pointwise$ to some $\rho_{T_{e \rightarrow i}}(.) \in \Upsilon\left(\operatorname{deg}_{T}(i)-1\right)$. Moreover, for all $t \in \mathbb{R}$, we have

$$
\begin{equation*}
\rho_{T_{e \rightarrow i}}^{-1}(t)=t-\sum_{e^{\prime} \ni i: e^{\prime} \neq e}\left[1-\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}\right]_{0}^{1} \tag{63}
\end{equation*}
$$

where the inverse functions are defined as in (57). Furthermore, for a node $i \in V(T)$ we have

$$
\begin{equation*}
l_{i}>t \quad \Longleftrightarrow \quad \sum_{e \ni i}\left[1-\sum_{j \in e, j \neq i}\left(\rho_{T_{e \rightarrow j}}^{-1}(t)\right)^{+}\right]_{0}^{1}>t \tag{64}
\end{equation*}
$$

Proof. First we fix $x \in \mathbb{R}$ and $(e, i) \in \Psi(T)$ and show that, as $n \rightarrow \infty, \rho_{T_{e \rightarrow i}}^{\epsilon_{n}}(x)$ is convergent. We call the limit $\rho_{T_{e \rightarrow i}}(x)$. Let $\theta_{\epsilon_{n}}^{b_{x}}$ denote the canonical $\epsilon_{n}$-balanced allocation on $T_{e \rightarrow i}$ with baseload $b_{x}$, which equals $x$ at $i$ and is zero elsewhere. By definition, $\rho_{T_{e \rightarrow i}}^{\epsilon_{n}}(x)=x+\partial \theta_{\epsilon_{n}}^{b_{x}}(i)$. Thus, it suffices to show that $\partial \theta_{\epsilon_{n}}^{b_{x}}(i)$ is convergent. Note that $\theta_{\epsilon_{n}}^{b_{x}}$ is a sequence in the compact space $[0,1]^{\Psi\left(T_{e \rightarrow i}\right)}$ equipped with the product topology. On the other hand, $\partial \theta_{\epsilon_{n}}^{b_{x}}(i)$ is a bounded sequence depending only on a finite number of coordinates, namely $\operatorname{deg}_{T}(i)-1$. As a result, in order to show that $\partial \theta_{\epsilon_{n}}^{b_{x}}(i)$ is convergent, it suffices to show that if $\theta_{1}$ and $\theta_{2}$ are two subsequential limits of $\theta_{\epsilon_{n}}^{b_{x}}$ in $[0,1]^{\Psi\left(T_{e \rightarrow i}\right)}$, then $\partial \theta_{1}(i)=\partial \theta_{2}(i)$. Passing to the limit in (2), we realize that both $\theta_{1}$ and $\theta_{2}$ are balanced allocations on $T_{e \rightarrow i}$ with respect to the baseload $b_{x}$. Using Proposition 2, it suffices to show that $\left\|\partial \theta_{1}-\partial \theta_{2}\right\|_{l^{1}\left(V\left(T_{e \rightarrow i}\right)\right)}<\infty$. From Proposition 8, we know that the restriction of $\theta_{\epsilon_{n}}$ to $T_{e \rightarrow i}$ is the canonical $\epsilon_{n}$-balanced allocation with baseload $\theta_{\epsilon_{n}}(e, i)$ at $i$ and zero elsewhere. Hence, if $K$ is a finite subset of $V\left(T_{e \rightarrow i}\right) \backslash\{i\}$, using Proposition 7, we have

$$
\sum_{j \in K}\left|\partial \theta_{\epsilon_{n}}^{b_{x}}(j)-\partial \theta_{\epsilon_{n}}(j)\right| \leq|x|+\theta_{\epsilon_{n}}(e, i) \leq|x|+1
$$

Using the triangle inequality, for integers $n$ and $m$,

$$
\sum_{j \in K}\left|\partial \theta_{\epsilon_{n}}^{b_{x}}(j)-\partial \theta_{\epsilon_{m}}^{b_{x}}(j)\right| \leq 2(|x|+1)+\sum_{j \in K}\left|\partial \theta_{\epsilon_{n}}(j)-\partial \theta_{\epsilon_{m}}(j)\right|
$$

Now, send $n$ to infinity along the subsequence of $\theta_{\epsilon_{n}}^{b_{x}}$ that converges to $\theta_{1}$. Likewise, send $m$ to infinity along the subsequence of $\theta_{\epsilon_{n}}^{b_{x}}$ that converges to $\theta_{2}$. Using the assumption that $\partial \theta_{\epsilon_{n}}(j)$ is convergent, and since $K$ is finite, we get

$$
\sum_{j \in K}\left|\partial \theta_{1}(j)-\partial \theta_{2}(j)\right| \leq 2(|x|+1)
$$

Since $K$ is arbitrary, sending $K$ to $V\left(T_{e \rightarrow i}\right) \backslash\{i\}$ we get

$$
\begin{aligned}
\sum_{j \in V\left(T_{e \rightarrow i}\right)}\left|\partial \theta_{1}(j)-\partial \theta_{2}(j)\right|= & \left|\partial \theta_{1}(i)-\partial \theta_{2}(i)\right|+ \\
& \sum_{j \in V\left(T_{e \rightarrow i}\right) \backslash\{i\}}\left|\partial \theta_{1}(j)-\partial \theta_{2}(j)\right| \\
& \leq 2(|x|+1)+2 \operatorname{deg}_{T}(i)<\infty
\end{aligned}
$$

which means $\left\|\partial \theta_{1}-\partial \theta_{2}\right\|_{l^{1}\left(V\left(T_{e \rightarrow i}\right)\right)}<\infty$. This means that for all $x \in \mathbb{R}$ and $(e, i) \in \Psi(T), \rho_{T_{e \rightarrow i}}^{\epsilon_{n}}(x) \rightarrow$ $\rho_{T_{e \rightarrow i}}(x)$. As $\rho_{T_{e \rightarrow i}}^{\epsilon_{n}}(.) \in \Upsilon\left(\operatorname{deg}_{T}(i)-1\right)$, this in particular implies that $\rho_{T_{e \rightarrow i}}(.) \in \Upsilon\left(\operatorname{deg}_{T}(i)-1\right)$.

Now, we prove (63). Using part (iv) of Lemma 11, both sides of (63) are right continuous functions of $t$. Hence, if $D$ denotes the union of the discontinuity sets of $\rho_{T_{e \rightarrow i}}^{-1}$ and $\rho_{T_{e^{\prime} \rightarrow j}}^{-1}$ for $e^{\prime} \ni i$ and $j \in e^{\prime}$, $D$ is countable and hence $D^{c}$ is dense in $\mathbb{R}$. Thus, due to the right continuity, it suffices to show (63) for $t \notin D$.

Now, take some $t \in D^{c}$. Using Lemma 13 we have

$$
\rho_{T_{e \rightarrow i}}^{-1}(t)=\lim _{n \rightarrow \infty}\left(\rho_{T_{e \rightarrow i}}^{\epsilon_{n}}\right)^{-1}(t)
$$

for $t$ in a dense subset of $D^{c}$. By abuse of notation, we continue to denote this subset by $D^{c}$. Using Proposition 14, $\left(\rho_{T_{e \rightarrow i}}^{\epsilon_{n}}\right)^{-1}$ is expressed in terms of $\zeta_{e^{\prime}, j}^{\epsilon_{n}}$ which are solutions to (60). Comparing this to (63), it suffices to prove that for each $e^{\prime} \ni i, e^{\prime} \neq e$, we have

$$
\begin{equation*}
\lim _{n \rightarrow \infty} 1-\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{\epsilon_{n}}=\left[1-\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}\right]_{0}^{1} \tag{65}
\end{equation*}
$$

Fixing such an $e^{\prime}$, since for each $j \in e^{\prime}, j \neq i$, the sequence $\left\{\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right\}_{n=1}^{\infty}$ is in the compact set $[0,1]$, it suffices to show that if for all $j \in e^{\prime}, j \neq i$, there is a subsequence $\zeta_{e^{\prime}, j}^{\epsilon_{k}}$ converging to some $\zeta_{e^{\prime}, j}^{*}$, then

$$
\begin{equation*}
1-\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{*}=\left[1-\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}\right]_{0}^{1} \tag{66}
\end{equation*}
$$

Without loss of generality and in order to simplify the notation, we may assume that the subsequence is the whole sequence, i.e. for all $j \in e^{\prime}, j \neq i, \zeta_{e^{\prime}, j}^{\epsilon_{n}} \rightarrow \zeta_{e^{\prime}, j}^{*}$. We show (66) in different cases:

Case I: $\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+} \leq 1$ : since $\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}$is nonnegative, it suffices to show that $\zeta_{e^{\prime}, j}^{*}=\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}$for each $j \in e^{\prime}, j \neq i$. For such a $j$, we do this in two subcases:

Case Ia: First, assume $\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t) \leq 0$, in which case we should show $\zeta_{e^{\prime}, j}^{*}=0$. If this is not the case, as $\zeta_{e^{\prime}, j}^{*} \in[0,1]$, there must be the case that $\zeta_{e^{\prime}, j}^{*}>0 \geq \rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)$. Then, part $(v)$ of Lemma 11
implies that $\rho_{T_{e^{\prime} \rightarrow j}}\left(\zeta_{e^{\prime}, j}^{*}\right) \geq t+\delta$ for some $\delta>0$. With this,

$$
\begin{aligned}
\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right)-t= & \rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right)-\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{*}\right) \\
& +\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{*}\right)-\rho_{T_{e^{\prime} \rightarrow j}}^{*}\left(\zeta_{e^{\prime}, j}^{*}\right) \\
& +\rho_{T_{e^{\prime} \rightarrow j}}\left(\zeta_{e^{\prime}, j}^{*}\right)-t \\
\geq- & \left|\zeta_{e^{\prime}, j}^{\epsilon_{n}}-\zeta_{e^{\prime}, j}^{*}\right| \\
& +\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{*}\right)-\rho_{T_{e^{\prime} \rightarrow j}}\left(\zeta_{e^{\prime}, j}^{*}\right) \\
& +\rho_{T_{e^{\prime} \rightarrow j}}^{*}\left(\zeta_{e^{\prime}, j}^{*}\right)-t .
\end{aligned}
$$

Note that the first two terms converge to zero as $n \rightarrow \infty$ and hence they could be made smaller than $\delta / 3$ by choosing $n$ large enough. Thus $\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right)-t \geq \delta / 3$ for $n$ large enough. On the other hand

$$
\begin{aligned}
\zeta_{e^{\prime}, j}^{\epsilon_{n}} & \leq \frac{1}{1+\exp \left(-\frac{t-\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right)}{\epsilon_{n}}\right)} \\
& \leq \frac{1}{1+e^{\delta /\left(3 \epsilon_{n}\right)}}
\end{aligned}
$$

Sending $n$ to infinity, since $\delta$ is fixed, we realize that $\zeta_{e^{\prime}, j}^{\epsilon_{n}}$ converges to zero, which is a contradiction with the assumption that $\zeta_{e^{\prime}, j}^{*}>0$.

Case Ib: Now consider the case $\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)>0$. If $\zeta_{e^{\prime}, j}^{*}>\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)$, following a similar argument as in case Ia, we can conclude that $\zeta_{e^{\prime}, j}^{*}=0$, which is a contradiction. Hence, assume $\zeta_{e^{\prime}, j}^{*}<\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)$. Since $t \in D^{c}$ is a continuity point of $\rho_{T_{e^{\prime} \rightarrow j}}^{-1}$, using Lemma 11 part $(v)$, we have $\rho_{T_{e^{\prime} \rightarrow j}}\left(\zeta_{e^{\prime}, j}^{*}\right) \leq t-\delta$ for some $\delta>0$. An argument similar to that in case Ia implies $\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right) \leq t-\delta / 3$ for $n$ large enough. Now

$$
\begin{aligned}
1-\sum_{l \in e^{\prime}, l \neq i} \zeta_{e^{\prime}, l}^{*} & =\lim _{n \rightarrow \infty} 1-\sum_{l \in e^{\prime}, l \neq i} \zeta_{e^{\prime}, l}^{\epsilon_{n}} \\
& =\lim _{n \rightarrow \infty} \frac{1}{1+\sum_{l \in e^{\prime}, l \neq i} \exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow l}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, l}^{\epsilon_{n}}\right)-t}{\epsilon_{n}}\right)} \\
& \leq \lim _{n \rightarrow \infty} \frac{1}{1+\exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow j}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j}^{\epsilon_{n}}\right)-t}{\epsilon_{n}}\right)} \\
& \leq \frac{1}{1+e^{\delta /\left(3 \epsilon_{n}\right)}} .
\end{aligned}
$$

Since $\delta$ is fixed, sending $n$ to infinity we realize that $\sum_{l \in e^{\prime}, l \neq i} \zeta_{e^{\prime}, l}^{*}=1$. This, together with our earlier assumption of Case I, means that

$$
\sum_{l \in e^{\prime}, l \neq i} \zeta_{e^{\prime}, l}^{*}=1 \geq \sum_{l \in e^{\prime}, l \neq i}\left(\rho_{T_{e^{\prime} \rightarrow l}}^{-1}(t)\right)^{+}=\sum_{l \in e^{\prime}, l \neq i, \rho_{T_{e^{\prime} \rightarrow l}^{-1}}^{-1}(t)>0} \rho_{T_{e^{\prime} \rightarrow l}}^{-1}(t)
$$

Since we have assumed that $\zeta_{e^{\prime}, j}^{*}<\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)$, this means there exists some $j^{\prime} \in e^{\prime}, j^{\prime} \neq i$ such that $\rho_{T_{e^{\prime} \rightarrow j^{\prime}}}^{-1}(t)>0$ and $\zeta_{e^{\prime}, j^{\prime}}^{*}>\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)>0$. This, as we discussed above, results in $\zeta_{e^{\prime}, j^{\prime}}^{*}=0$, which is a contradiction. Hence, $\zeta_{e^{\prime}, j}^{*}$ should be equal to $\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)$ and the proof of this case is complete.

Case II: $\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}>1$, in which case we need to show that $1-\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{*}=0$ to conclude (66). Since

$$
\begin{aligned}
\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{*} & =\lim _{n \rightarrow \infty} \sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{\epsilon_{n}} \\
& =\lim _{n \rightarrow \infty} \frac{\sum_{l \in e^{\prime}, l \neq i} \exp \left(-\rho_{T_{e^{\prime} \rightarrow l}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, l}^{\epsilon_{n}}\right) / \epsilon\right)}{e^{-t / \epsilon}+\sum_{l \in e^{\prime}, l \neq i} \exp \left(-\rho_{T_{e^{\prime} \rightarrow l}^{n}}\left(\zeta_{e^{\prime}, l}^{\epsilon_{n}}\right) / \epsilon\right)} \\
& \leq 1<\sum_{j \in e^{\prime}, j \neq i}\left(\rho_{T_{e^{\prime} \rightarrow j}}^{-1}(t)\right)^{+}
\end{aligned}
$$

there should exist some $j^{*} \in e^{\prime}, j^{*} \neq i$ such that $\zeta_{e^{\prime}, j^{*}}^{*}<\rho_{T_{e^{\prime} \rightarrow j^{*}}}^{-1}(t)$ and $\rho_{T_{e^{\prime} \rightarrow j^{*}}}^{-1}>0$. Since $t \notin D$, using Lemma 11 we have $\rho_{T_{e^{\prime} \rightarrow j^{*}}}\left(\zeta_{e^{\prime}, j^{*}}^{*}\right) \leq t-\delta$ for some $\delta>0$. Using calculations similar to those in case Ia above, for $n$ large enough we have $\rho_{T_{e^{\prime} \rightarrow j^{*}}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j^{*}}^{\epsilon_{n}}\right) \leq t-\delta / 3$. Now,

$$
\begin{aligned}
1-\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{\epsilon_{n}} & =\frac{1}{1+\sum_{j \in e^{\prime}, j \neq i} \exp \left(-\frac{\rho_{e_{e^{\prime} \rightarrow j}}^{\varepsilon_{n}}\left(\zeta_{e^{\prime}, j}^{e e_{n}}\right)-t}{\epsilon_{n}}\right)} \\
& \leq \frac{1}{1+\exp \left(-\frac{\rho_{T_{e^{\prime} \rightarrow j^{*}}}^{\epsilon_{n}}\left(\zeta_{e^{\prime}, j^{*}}^{\epsilon_{n}}\right)-t}{\epsilon_{n}}\right)} \\
& \leq \frac{1}{1+e^{\delta /\left(3 \epsilon_{n}\right)}} .
\end{aligned}
$$

Since $\delta$ is fixed, sending $n$ to infinity we conclude that $1-\sum_{j \in e^{\prime}, j \neq i} \zeta_{e^{\prime}, j}^{*}=0$, which completes the argument of this case.

Having verified (66) in all cases, we conclude (63). Now, we prove (64). Note that in the above discussion we started with the rooted tree $T_{e \rightarrow i}$. However, it can be verified that all the arguments are valid if we start with the tree $T$ rooted at an arbitrary vertex $i \in V(T)$, i.e. $(T, i)$. Convergence of $\rho_{(T, i)}$ is also similar. In this case, repeating the above argument, (63) becomes

$$
\begin{equation*}
\rho_{(T, i)}^{-1}(t)=t-\sum_{e \ni i}\left[1-\sum_{j \in e, j \neq i}\left(\rho_{T_{e \rightarrow j}}^{-1}(t)\right)^{+}\right]_{0}^{1} . \tag{67}
\end{equation*}
$$

On the other hand, note that

$$
l_{i}=\lim _{n \rightarrow \infty} \partial \theta_{n}(i)=\lim _{n \rightarrow \infty} \rho_{(T, i)}^{\epsilon_{n}}(0)=\rho_{(T, i)}(0) .
$$

Hence, $l_{i}>t$ iff $\rho_{(T, i)}(0)>t$, or equivalently, using part $(v)$ of Lemma 11, $\rho_{(T, i)}^{-1}(t)<0$. Substituting into (67), we conclude (64), and the proof is complete.

Proof of Lemma 12. Assume $\left(\theta_{1}, \ldots, \theta_{n}\right)$ and $\left(\theta_{1}^{\prime}, \ldots, \theta_{n}^{\prime}\right)$ are two distinct solutions to this set of equations. We claim that it can not be the case that $\theta_{i}^{\prime}>\theta_{i}$ for some $i$ and $\theta_{j}^{\prime} \leq \theta_{j}$ for some other $j \neq i$. Assume this holds. Since the right hand side of (58) is positive, all $\theta_{i}$ 's and $\theta_{i}^{\prime}$ 's are positive. On the other hand, we have

$$
\frac{e^{-g_{i}\left(\theta_{i}^{\prime}\right)}}{e^{-g_{j}\left(\theta_{j}^{\prime}\right)}}=\frac{\theta_{i}^{\prime}}{\theta_{j}^{\prime}}>\frac{\theta_{i}}{\theta_{j}}=\frac{e^{-g_{i}\left(\theta_{i}\right)}}{e^{-g_{j}\left(\theta_{j}\right)}},
$$

which means

$$
g_{i}\left(\theta_{i}^{\prime}\right)+g_{j}\left(\theta_{j}\right)<g_{j}\left(\theta_{j}^{\prime}\right)+g_{i}\left(\theta_{i}\right)
$$

But $\theta_{i}^{\prime}>\theta_{i}$ implies $g_{i}\left(\theta_{i}^{\prime}\right) \geq g_{i}\left(\theta_{i}\right)$ since $g_{i}$ is nondecreasing. On the other hand, $\theta_{j}^{\prime} \leq \theta_{j}$ implies $g_{j}\left(\theta_{j}^{\prime}\right) \geq g_{j}\left(\theta_{j}\right)$ which is a contradiction with the above inequality.

Hence, without loss of generality, we may assume that $\theta_{i}^{\prime} \geq \theta_{i}$ for all $1 \leq i \leq n$. If it is not the case that $\theta_{i}=\theta_{i}^{\prime}$ for all $1 \leq i \leq n$, then

$$
\frac{\sum_{i=1}^{n} e^{-g_{i}\left(\theta_{i}\right)}}{\sum_{i=1}^{n} e^{-g_{i}\left(\theta_{i}\right)}+C}=\sum \theta_{i}<\sum \theta_{i}^{\prime}=\frac{\sum_{i=1}^{n} e^{-g_{i}\left(\theta_{i}^{\prime}\right)}}{\sum_{i=1}^{n} e^{-g_{i}\left(\theta_{i}^{\prime}\right)}+C}
$$

On the other hand, $\theta_{i}^{\prime} \geq \theta_{i}$ and $g_{i}$ being nondecreasing implies that $e^{-g_{i}\left(\theta_{i}\right)} \geq e^{-g_{i}\left(\theta_{i}^{\prime}\right)}$ for all $1 \leq i \leq n$, which is in contradiction with the above inequality.

Proof of Lemma 13. Sending $n$ to infinity in the three conditions of Definition 33 and using the fact that $g_{n}$ converges pointwise to $g$ implies that $g$ is in $\Upsilon(C)$. To show (62), given $t \notin\left(\bigcup_{n} D_{n}\right) \cup D$, define $x_{n}:=g_{n}^{-1}(t)$ and $x:=g^{-1}(t)$. Since $g_{n} \in \Upsilon(C)$, we have $x_{n} \leq g_{n}\left(x_{n}\right) \leq x_{n}+C$ or $x_{n} \in[t-C, t]$, which is a compact set. Hence, it suffices to show that any subsequential limit of $x_{n}$ is equal to $x$. Thus, without loss of generality, we may assume that $x_{n} \rightarrow x^{\prime}$, and we show that $x^{\prime}=x$.

If $x^{\prime}<x$, since $t$ is a continuity point for $g$, Lemma 11 part $(i v)$ implies that $g\left(x^{\prime}\right)<t$. Thereby, $g\left(x^{\prime}\right) \leq t-\delta$ for some $\delta>0$. Now,

$$
\begin{aligned}
g_{n}\left(x_{n}\right) & =g_{n}\left(x_{n}\right)-g_{n}\left(x^{\prime}\right)+g_{n}\left(x^{\prime}\right)-g\left(x^{\prime}\right)+g\left(x^{\prime}\right) \\
& \leq\left|x_{n}-x^{\prime}\right|+\left|g_{n}\left(x^{\prime}\right)-g\left(x^{\prime}\right)\right|+g\left(x^{\prime}\right),
\end{aligned}
$$

where the last inequality employs the fact that $g_{n} \in \Upsilon(C)$. Since $x_{n} \rightarrow x^{\prime}$ and $g_{n}$ converges pointwise to $g$, for large $n$ the first two terms could be made smaller than $\delta / 3$. Thus, for large $n, g_{n}\left(x_{n}\right) \leq t-\delta / 3$, which is a contradiction with $g_{n}\left(x_{n}\right)=t$. The assumption $x^{\prime}>x$ similarly results in contradiction. As a result, $x^{\prime}=x$, and the proof is complete.

## 8 Characterization of the Mean Excess Function for Galton Watson Processes

In this section, we prove Theorem 2. This is done in two steps. First, in Section 8.1, we prove that for any set of fixed points $\left\{Q_{l}\right\}_{l \geq 2}$, the LHS of (11) is no less than the RHS. This is proved in Proposition 16. Later, in Section 8.2, we show that there exists a set of fixed points achieving the maximum in (11). This is done in Proposition 17.

### 8.1 Lower Bound

In this section, we use the indexing notation $\mathbb{N}_{\text {vertex }}$ and $\mathbb{N}_{\text {edge }}$, which was introduced in Section 2.9. The level of a vertex $\left(s_{1}, e_{1}, i_{1}, \ldots, s_{k}, e_{k}, i_{k}\right) \in \mathbb{N}_{\text {vertex }}$ is defined to be $k$, and the level of $\emptyset$ is defined to be zero. Likewise, the level of an edge $\left(s_{1}, e_{1}, i_{1}, \ldots, s_{k}, e_{k}\right) \in \mathbb{N}_{\text {edge }}$ is defined to be $k$. Also, recall that for $v \in \mathbb{N}_{\text {vertex }}, s \geq 2$ and $e \geq 1,(v, s, e)$ is an element in $\mathbb{N}_{\text {edge }}$ obtained by concatenating $(s, e)$ to the end of the string representing $v$ in $\mathbb{N}_{\text {vertex }}$. Likewise, for $s \geq 2, e \geq 1$ and $1 \leq i \leq s-1$, $(v, s, e, i) \in \mathbb{N}_{\text {vertex }}$ is defined similarly.

We need the following tool before proving our lower bound. In the following lemma, $\Pi: \overline{\mathcal{H}}_{*}(\mathbb{R}) \rightarrow$ $\mathcal{H}_{*}$ is the projection defined in (7).

Lemma 14. Let $t \in \mathbb{R}$ together with distributions $P$ and $\left\{Q_{k}\right\}_{k \geq 2} \in \mathcal{Q}$ be given as in Theorem 2. Given a probability distribution $W$ on $\Lambda$, there is a random marked rooted tree $\left(\overline{\mathbb{T}}_{W}, \emptyset\right)$, with marks taking values in $\mathbb{R}$, and with vertex set and edge set $\mathbb{N}_{\text {vertex }}$ and $\mathbb{N}_{\text {edge }}$, respectively, such that the underlying unmarked rooted tree is a Galton Watson tree such that the type of the root is distributed according to $W$, and the type of a non-root vertex $v=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{r}, e_{r}, i_{r}\right)$ in the subtree below $v$ is distributed according to $\hat{P}_{s_{r}}$. Moreover, the marks of $\mathbb{T}_{W}$ satisfy

$$
\begin{equation*}
\xi_{\mathbb{T}_{W}}(e, i)=t-\sum_{e^{\prime} \ni i, e^{\prime} \neq e}\left[1-\sum_{j \in e^{\prime}, j \neq i} \xi_{\bar{T}_{W}}\left(e^{\prime}, j\right)^{+}\right]_{0}^{1}, \tag{68}
\end{equation*}
$$

for all $(e, i) \in \Psi\left(\overline{\mathbb{T}}_{W}\right)$. Furthermore, for any $L \geq 1$, conditioned on the structure of the tree up to depth $L$, the set of marks from edges in level $L$ towards vertices in that level are independent, and for any edge-vertex pair $(e, i)$ both in level $L, \xi_{\overline{\mathbb{T}}_{W}}(e, i)$ is distributed according to $Q_{k}$ where $k$ is the size of $e$.

In particular, when $W=P$, the measure $\nu \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}(\mathbb{R})\right)$, which is defined to be the law of $\left[\overline{\mathbb{T}}_{P}, \emptyset\right]$, is unimodular and $(\Pi)_{*} \nu=U G W T(P)$.

Proof. We first generate the collection of random variables $\Gamma_{\emptyset},\left(\Gamma_{v}, X_{v}\right)_{v \in \mathbb{N}_{\mathrm{vertex}} \backslash\{\emptyset\}}$, such that $\Gamma_{v}$ for $v \in \mathbb{N}_{\text {vertex }}$ takes value in $\Lambda$ and $X_{v}$ for $v \in \mathbb{N}_{\text {vertex }} \backslash\{\emptyset\}$ takes values in $\mathbb{R}$, with the following properties: $(i)\left(\Gamma_{v}\right)_{v \in \mathbb{N}_{\text {vertex }}}$ are independent from each other such that $\Gamma_{\emptyset}$ has law $W$ and for $v=$ $\left(s_{1}, e_{1}, i_{1}, \ldots, s_{r}, e_{r}, i_{r}\right) \in \mathbb{N}_{\text {vertex }}, \Gamma_{v}$ has law $\hat{P}_{s_{r}} ;(i i)$ For any $v \in \mathbb{N}_{\text {vertex }} \backslash\{\emptyset\}$, we have

$$
\begin{equation*}
X_{v}=t-\sum_{k=2}^{h\left(\Gamma_{v}\right)} \sum_{l=1}^{\Gamma_{v}(k)}\left[1-X_{(v, k, l, 1)}^{+}-\cdots-X_{(v, k, l, k-1)}^{+}\right]_{0}^{1} \tag{69}
\end{equation*}
$$

(iii) For any $L \geq 1, X_{v}$ for nodes $v$ at level $L$ are independent and for $v=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{L}, e_{L}, i_{L}\right)$ at level $L, X_{v}$ is distributed according to $Q_{s_{L}}$.

We construct the law of the above random variables satisfying the above conditions using Kolmogorov's extension theorem (see, for instance, [Tao11]). For an integer $L \geq 1$, define $A_{L}$ to be the set of nodes in $\mathbb{N}_{\text {vertex }}$ with level at most $L$. For each $L \geq 1$, we introduce the law of a subset of the above family of random variables, namely $\Gamma_{\emptyset},\left(X_{v}, \Gamma_{v}\right)_{v \in A_{L} \backslash\{\emptyset\}}$, and denote this law by $\nu_{L}$. To start with, we generate $\Gamma_{v}, v \in A_{L}$ independently such that $\Gamma_{\emptyset}$ has law $W$ and $\Gamma_{v}$ for $v=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{r}, e_{r}, i_{r}\right)$, $X_{v}$ has law $\hat{P}_{s_{r}}$. In the next step, we generate $X_{v}$ for nodes $v$ with depth equal to $L$ independently such that $X_{v}$ for $v=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{L}, e_{L}, i_{L}\right)$ has law $Q_{s_{L}}$. Next, we define $X_{v}$ for nodes at levels 1 through $L-1$ using the relation (69) inductively starting from level $L-1$ all the way up to level 1. Using the fact that $Q_{k}=F_{P, t}^{(k)}\left(\left\{Q_{l}\right\}_{l \geq 2}\right)$, see (9), and also that $\Gamma_{v}$ for $v=\left(s_{1}, e_{1}, i_{1}, \ldots, s_{r}, e_{r}, i_{r}\right)$ has law $\hat{P}_{s_{r}}$, it is evident that the set of measures $\left\{\nu_{L}\right\}_{L \geq 1}$ are consistent. Therefore, Kolmogorov's extension theorem implies that the set of random variables with the conditions stated above exist.

Now, we turn these random variables into a marked random rooted tree $\overline{\mathbb{T}}_{W}$ having vertex set and edge set $\mathbb{N}_{\text {vertex }}$ and $\mathbb{N}_{\text {edge }}$, respectively. To do so, we first construct the underlying unmarked tree $\overline{\mathbb{T}}_{W}$ given the types $\Gamma_{v}, v \in \mathbb{N}_{\text {vertex }}$. In the next step, for any edge $e$ and vertex $v$ being at the same level in the tree, we set $\xi_{\overline{\mathbb{T}}_{W}}(e, v):=X_{v}$. It can be easily seen that the "upward" marks, i.e. marks from edges towards nodes above them, are immediately unambiguously defined. To see this, for an edge $e$ at level 1 , we define $\xi_{\overline{\mathbb{T}}_{W}}(e, \emptyset)$ using (68). We then inductively go down one level at a time, to define all the other upward marks.

Now, we show that the measure $\nu$, which is defined to be the law of $\left[\overline{\mathbb{T}}_{P}, \emptyset\right]$, is unimodular. Our proof technique is similar to that of Lemma 20 in Appendix F. Take a Borel function $f: \overline{\mathcal{H}}_{* *}(\mathbb{R}) \rightarrow[0, \infty)$
and note that due to our above construction,

$$
\int f d \vec{\nu}=\int \partial f d \nu=\mathbb{E}\left[\sum_{k=2}^{h\left(\Gamma_{\emptyset}\right)} \sum_{l=1}^{\Gamma_{\emptyset}(k)} f\left(\overline{\mathbb{T}}_{P},(k, l), \emptyset\right)\right]
$$

Using the symmetry in the construction, we have

$$
\int f d \vec{\nu}=\sum_{\gamma \in \Lambda} P(\gamma) \sum_{k=2}^{h(\gamma)} \gamma(k) \mathbb{E}\left[f\left(\overline{\mathbb{T}}_{P},(k, 1), \emptyset\right) \mid \Gamma_{\emptyset}=\gamma\right] .
$$

As all the terms are nonnegative, we may change the order of summation. Also using the definition of $\hat{P}$, if $\Gamma$ is a random variable with law $P$, we get

$$
\int f d \vec{\nu}=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \sum_{\gamma \in \Lambda} \hat{P}_{k}(\gamma) \mathbb{E}\left[f\left(\overline{\mathbb{T}}_{P},(k, 1), \emptyset\right) \mid \Gamma_{\emptyset}=\gamma+\mathbb{e}_{k}\right]
$$

Now, for each $k \geq 2$, define $\widetilde{P}_{k}$ to be the law of the random variable $\Gamma_{k}+\mathbb{C}_{k}$ where $\Gamma_{k}$ has law $\hat{P}_{k}$. With this, for each $k \geq 2$, the inner summation over $\gamma$ in the above expression could be interpreted as an expectation with respect to a tree with root type distribution $\widetilde{P}_{k}$, i.e. $\overline{\mathbb{T}}_{\widetilde{P}_{k}}$. In fact, this shows that

$$
\begin{equation*}
\int f d \vec{\nu}=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \mathbb{E}\left[f\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}},(k, 1), \emptyset\right)\right] \tag{70}
\end{equation*}
$$

Now, note that for each $k \geq 2$, due to the definition of $\widetilde{P}_{k}$, the tree $\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow \emptyset}$ rooted at $\emptyset$ (which we recall is obtained by removing $(k, 1)$ and then taking the subtree rooted at $\emptyset$ ) has an underlying unmarked structure which is precisely $\mathrm{GWT}_{k}(P)$. This, in particular, implies that $\xi_{\overline{\mathbb{T}}_{\tilde{P}_{k}}}((k, 1), \emptyset)$ has law $Q_{k}$. Also, by construction, for $1 \leq i \leq k-1$, $\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow(k, 1, i)}$ have independent underlying unmarked structures, all with law $\operatorname{GWT}_{k}(P)$. Moreover, the downward marks in $\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow \emptyset}$ and $\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow(k, 1, i)}$ for $1 \leq i \leq k-1$ are independent from each other and have the same distribution. Thereby, $\xi_{\bar{T}_{\tilde{P}_{k}}}((k, 1), v)$ for $v \in(k, 1)$ are independent and all have distribution $Q_{k}$. Now, we claim that for all $v \in(k, 1),\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow v}$ are equal in distribution. To see this, note that for all $v \in(k, 1)$, the unmarked structure of $\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow v}$ is $\operatorname{GWT}_{k}(P)$ and the downward marks are constructed following the same recipes. As was discussed above, to construct upward marks, we start from the root and go down inductively. The fact that $\left.\xi_{\overline{\mathbb{T}}_{\widetilde{P}_{k}}}(k, 1), v\right)$ for $v \in(k, 1)$ are i.i.d. with law $Q_{k}$ guarantees that the downward marks in $\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}\right)_{(k, 1) \rightarrow v}$ are equally distributed for all $v \in(k, 1)$. This in particular implies that for $1 \leq i \leq k-1$,

$$
\mathbb{E}\left[f\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}},(k, 1), \emptyset\right)\right]=\mathbb{E}\left[f\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}},(k, 1),(k, 1, i)\right)\right] .
$$

Using this and writing (70) for $\int \nabla f d \mu$, we conclude that $\int f d \vec{\nu}=\int \nabla f d \vec{\nu}$ which completes the proof of the unimodularity of $\nu$.

Before proving our lower bound, we state a modified version of our variational representation in Proposition 11 and a general lemma. The proof of the following lemmas are given at the end of this section.

Lemma 15. Assume $\mu$ is a distribution on $\mathcal{H}_{*}$ with $\operatorname{deg}(\mu)<\infty$ and $\nu$ is a unimodular distribution on $\overline{\mathcal{H}}_{*}(\mathbb{R})$ such that $(\Pi)_{*} \nu=\mu$. Then, for any Borel allocation $\Theta: \mathcal{H}_{* *} \rightarrow[0,1]$ and any function $f: \overline{\mathcal{H}}_{*}(\mathbb{R}) \rightarrow[0,1]$, we have

$$
\int(\partial \Theta-t)^{+} d \mu \geq \int \tilde{f}_{m i n} d \vec{\nu}-t \int f d \nu
$$

where

$$
\tilde{f}_{m i n}([\bar{H}, e, i]):=\frac{1}{|e|} \min _{j \in e} f([\bar{H}, j])
$$

Lemma 16. Assume $x_{1}, \ldots, x_{n}$ are real numbers. Then $x_{i}<\left[1-\sum_{j \neq i} x_{j}^{+}\right]_{0}^{1}$ for all $1 \leq i \leq n$ if and only if $\sum x_{i}^{+}<1$.

Proposition 16. Assume $P$ is a distribution on $\Lambda$ such that $\mathbb{E}\left[\|\Gamma\|_{1}\right]<\infty$ where $\Gamma$ has law $P$. Then, with $\mu=U G W T(P)$, for any $t \in \mathbb{R}$, and any set of probability distributions on real numbers $\left\{Q_{k}\right\}_{k \geq 2}$ such that for all $k \geq 2$ we have $Q_{k}=F_{P, t}^{(k)}\left(\left\{Q_{l}\right\}_{l \geq 2}\right)$, it holds that

$$
\Phi_{\mu}(t) \geq\left(\sum_{k=2}^{\infty} \frac{\mathbb{E}[\Gamma(k)]}{k} \mathbb{P}\left(\sum_{i=1}^{k} X_{k, i}^{+}<1\right)\right)-t \mathbb{P}\left(\sum_{k=2}^{h(\Gamma)} \sum_{i=1}^{\Gamma(k)} Y_{k, i}>t\right)
$$

Here, in the first expression, $\Gamma$ is a random variable on $\Lambda$ with law $P$ and $\left\{X_{k, i}\right\}_{k, i}$ are independent such that $X_{k, i}$ has law $Q_{k}$. Also, in the second expression, $\Gamma$ has law $P$ and $\left\{Y_{k, i}\right\}_{k, i}$ are independent from each other and from $\Gamma$ such that $Y_{k, i}$ has the law of the random variable $\left[1-\left(Z_{1}^{+}+\cdots+Z_{k-1}^{+}\right)\right]_{0}^{1}$ where $Z_{j}$ are i.i.d. with law $Q_{k}$.

Proof. Note that the condition $\mathbb{E}\left[\|\Gamma\|_{1}\right]<\infty$ guarantees that $\operatorname{deg}(\mu)<\infty$. Define the functions $F: \overline{\mathcal{H}}_{* *}(\mathbb{R}) \rightarrow[0,1]$ and $f: \overline{\mathcal{H}}_{*} \rightarrow[0,1]$ as

$$
F([\bar{H}, e, i])=\left[1-\sum_{j \in e, j \neq i} \xi_{\bar{H}}(e, j)^{+}\right]_{0}^{1}
$$

and $f([\bar{H}, i]):=\mathbb{1}_{\partial F([\bar{H}, i])>t}$. Using the unimodular measure $\nu$ constructed in Lemma 14 and the variational characterization in Lemma 15, we have

$$
\begin{equation*}
\Phi_{\mu}(t) \geq \int \tilde{f}_{\min } d \vec{\nu}-t \int f d \nu \tag{71}
\end{equation*}
$$

where

$$
\tilde{f}_{\min }(\bar{H}, e, j)=\frac{1}{|e|} \min _{j \in e} f(\bar{H}, j)=\frac{1}{|e|} \mathbb{1}[\partial F(\bar{H}, j)>t \forall j \in e] .
$$

Following the proof of Lemma 14, and in particular Equation (70) therein, we have

$$
\begin{equation*}
\int \tilde{f}_{\min } d \vec{\nu}=\sum_{k \geq 2} \mathbb{E}[\Gamma(k)] \mathbb{E}\left[\tilde{f}_{\min }\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}},(k, 1), \emptyset\right)\right] \tag{72}
\end{equation*}
$$

Due to the definition of $f, \tilde{f}_{\min }\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}},(k, 1), \emptyset\right)=\frac{1}{k} \mathbb{1}\left[\partial F\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}, v\right)>t \forall v \in(k, 1)\right]$. For $v \in(k, 1)$,

$$
\begin{aligned}
\partial F\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}, v\right)= & \sum_{e^{\prime} \ni v} F\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}, e^{\prime}, v\right) \\
= & {\left[1-\sum_{w \in(k, 1), w \neq v} \xi_{\overline{\mathbb{P}}_{\widetilde{P}_{k}}}((k, 1), w)^{+}\right]_{0}^{1}+} \\
& \sum_{e^{\prime} \ni v, e^{\prime} \neq(k, 1)}\left[1-\sum_{w \in e^{\prime}, w \neq v} \xi_{\overline{\mathbb{T}}_{\widetilde{P}_{k}}}\left(e^{\prime}, w\right)^{+}\right]_{0}^{1} .
\end{aligned}
$$

Using (68), this yields

$$
\partial F\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}, v\right)=\left[1-\sum_{w \in(k, 1), w \neq v} \xi_{\overline{\mathbb{P}}_{\widetilde{P}_{k}}}((k, 1), w)^{+}\right]_{0}^{1}+t-\xi_{\overline{\mathbb{P}}_{\widetilde{P}_{k}}}((k, 1), v)
$$

Therefore, $\partial F\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}}, v\right)>t$ for all $v \in(k, 1)$ if and only if for all $v \in(k, 1)$,

$$
\xi_{\overline{\mathbb{T}}_{\tilde{P}_{k}}}((k, 1), v)<\left[1-\sum_{w \in(k, 1), w \neq v} \xi_{\overline{\mathbb{T}}_{\tilde{P}_{k}}}((k, 1), w)^{+}\right]_{0}^{1}
$$

Using Lemma 16 , this is equivalent to $\sum_{v \in(k, 1)} \xi_{\overline{\mathbb{T}}_{\widetilde{P}_{k}}}((k, 1), v)^{+}<1$. Therefore,

$$
\mathbb{E}\left[\tilde{f}_{\min }\left(\overline{\mathbb{T}}_{\widetilde{P}_{k}},(k, 1), \emptyset\right)\right]=\frac{1}{k} \mathbb{P}\left(\sum_{v \in(k, 1)} \xi_{\overline{\mathbb{T}}_{\tilde{P}_{k}}}((k, 1), v)^{+}<1\right)
$$

But as was shown in Lemma $14, \xi_{\overline{\mathbb{P}}_{\tilde{P}_{k}}}((k, 1), v)$ for $v \in(k, 1)$ are i.i.d. with law $Q_{k}$. Consequently, substituting in (72) we get

$$
\begin{equation*}
\int \tilde{f}_{\min } d \vec{\nu}=\sum_{k=2}^{\infty} \frac{\mathbb{E}[\Gamma(k)]}{k} \mathbb{P}\left(\sum_{i=1}^{k} X_{k, i}^{+}<1\right) \tag{73}
\end{equation*}
$$

where $X_{k, i}, k \geq 2,1 \leq i \leq k$ are independent such that $X_{k, i}$ has law $Q_{k}$.
On the other hand,

$$
\begin{aligned}
\int f d \nu=\mathbb{P}\left(\partial F\left(\overline{\mathbb{T}}_{P}, \emptyset\right)>t\right) & =\mathbb{P}\left(\sum_{k=2}^{h\left(\Gamma_{\emptyset}\right)} \sum_{i=1}^{\Gamma_{\emptyset}(k)} F\left(\overline{\mathbb{T}}_{P},(k, i), \emptyset\right)>t\right) \\
& =\mathbb{P}\left(\sum_{k=2}^{h\left(\Gamma_{\emptyset}\right)} \sum_{i=1}^{\Gamma_{\emptyset}(k)}\left[1-\sum_{j=1}^{k-1} \xi_{\overline{\mathbb{T}}_{P}}((k, i),(k, i, j))^{+}\right]_{0}^{1}>t\right) .
\end{aligned}
$$

But, as we saw in Lemma $14, \xi_{\overline{\mathbb{T}}_{P}}((k, i),(k, i, j))$ for $k \geq 2, i \leq \Gamma_{\emptyset}(k), 1 \leq j \leq k-1$, are independent, and $\xi_{\bar{T}_{P}}((k, i),(k, i, j))$ has law $Q_{k}$. Thereby,

$$
\begin{equation*}
\int f d \nu=\mathbb{P}\left(\sum_{k=2}^{h(\Gamma)} \sum_{i=1}^{\Gamma(k)} Y_{k, i}>t\right) \tag{74}
\end{equation*}
$$

with $Y_{k, i}$ being independent from each other such that $Y_{k, i}$ has the law of the random variable [1-$\left.\left(Z_{1}^{+}+\cdots+Z_{k-1}^{+}\right)\right]_{0}^{1}$ with $Z_{j}$ 's being i.i.d. with law $Q_{k}$. The proof is complete by substituting (73) and (74) into (71).

Proof of Lemma 15. By interpreting $\Theta$ as a function on $\mathcal{H}_{* *}(\mathbb{R})$ via $\Theta(\bar{H}, e, i):=\Theta(H, e, i)$ (where we recall that $H$ is the underlying unmarked hypergraph associated to $\bar{H}$ ), and also using the inequality $x^{+} \geq x y$ which holds for $y \in[0,1]$, we have

$$
\begin{equation*}
\int(\partial \Theta-t)^{+} d \mu \int(\partial \Theta-t)^{+} d \nu \geq \int f \partial \Theta d \nu-t \int f d \nu \tag{75}
\end{equation*}
$$

Since $\operatorname{deg}(\mu)<\infty$, all the integrals are finite and well defined. Due to the definition of $\vec{\nu}$ and the unimodularity of $\nu$ we have $\int f \partial \Theta d \nu=\int f \Theta d \vec{\nu}=\int \nabla(f \Theta) d \nu$. On the other hand,

$$
\nabla(f \Theta)(\bar{H}, e, i)=\frac{1}{|e|} \sum_{j \in e} f(\bar{H}, j) \Theta(H, e, j) \geq \frac{1}{|e|} \min _{j \in e} f(\bar{H}, j)=\tilde{f}_{\min }(\bar{H}, i)
$$

where the inequality holds since $\sum_{j \in e} \Theta(H, e, j)=1$ and $\Theta([H, e, j]) \geq 0$ for all $j \in e$. Substituting this into (75) completes the proof.

Proof of Lemma 16. First, assume that $x_{i}<\left[1-\sum_{j \neq i} x_{j}^{+}\right]_{0}^{1}$ for all $i$. If $x_{i} \leq 0$ for all $i$ then nothing remains to be proved. Hence, assume that $x_{i}>0$ for some $i$. Since $x_{i}<\left[1-\sum_{j \neq i} x_{j}^{+}\right]_{0}^{1}$, we have $0<\sum_{j \neq i} x_{j}^{+}<1$, which means that $1-\sum_{j \neq i} x_{j}^{+} \in[0,1]$. Therefore, $\left[1-\sum_{j \neq i} x_{j}^{+}\right]_{0}^{1}=1-\sum_{j \neq i} x_{j}^{+}$. On the other hand,

$$
x_{i}^{+}=x_{i}<\left[1-\sum_{j \neq i} x_{j}^{+}\right]_{0}^{1}=1-\sum_{j \neq i} x_{j}^{+},
$$

which implies $\sum x_{i}^{+}<1$.
In order to prove the other direction, take $1 \leq i \leq n$ and note that

$$
x_{i} \leq x_{i}^{+}=\sum_{k=1}^{n} x_{k}^{+}-\sum_{j \neq i} x_{j}^{+}<1-\sum_{j \neq i} x_{j}^{+} .
$$

Moreover, $\sum_{j \neq i} x_{j}^{+} \leq \sum x_{i}^{+}<1$. Thereby, $1-\sum_{j \neq i} x_{j}^{+} \in[0,1]$ and $1-\sum_{j \neq i} x_{j}^{+}=\left[1-\sum_{j \neq i} x_{j}^{+}\right]_{0}^{1}$. Substituting this into the above inequality completes the proof.

### 8.2 Upper Bound

In this section, we show that there exists a family of probability distributions $\left\{Q_{k}\right\}_{k \geq 2}$ satisfying the fixed points equations (10) and achieving the maximum on the RHS of (11).

Proposition 17. Assume $P$ is a distribution on $\Lambda$ such that $\mathbb{E}\left[\|\Gamma\|_{1}\right]<\infty$, where $\Gamma$ has law $P$. Let $\mu=U G W T(P)$. Given $t \in \mathbb{R}$, there exists a family of probability distributions $\left\{Q_{l}\right\}_{l \geq 2}$ on the set of real numbers such that, for each $k \geq 2, Q_{k}=F_{P, t}^{(k)}\left(\left\{Q_{l}\right\}_{l \geq 2}\right)$, and such that we have

$$
\Phi_{\mu}(t)=\left(\sum_{k=2}^{\infty} \frac{\mathbb{E}[\Gamma(k)]}{k} \mathbb{P}\left(\sum_{i=1}^{k} X_{k, i}^{+}<1\right)\right)-t \mathbb{P}\left(\sum_{k=2}^{h(\Gamma)} \sum_{i=1}^{\Gamma(k)} Y_{k, i}>t\right)
$$

Here, in the first expression, $\Gamma$ is a random variable on $\Lambda$ with law $P$ and $\left\{X_{k, i}\right\}_{k, i}$ are i.i.d. such that $X_{k, i}$ has law $Q_{k}$. Also, in the second expression, $\Gamma$ has law $P$ and $\left\{Y_{k, i}\right\}_{k, i}$ are independent from each other and from $\Gamma$ such that $Y_{k, i}$ has the law of the random variable $\left[1-\left(Z_{1}^{+}+\cdots+Z_{k-1}^{+}\right)\right]_{0}^{1}$ where $Z_{j}$ are i.i.d. with law $Q_{k}$.

Proof. The condition $\mathbb{E}\left[\|\Gamma\|_{1}\right]<\infty$ guarantees that $\operatorname{deg}(\mu)<\infty$. Therefore, Proposition 10 implies that there exists a sequence $\epsilon_{n} \downarrow 0$ such that the sequence of $\epsilon_{n}$-balanced allocations $\Theta_{\epsilon_{n}}$ converges $\vec{\mu}$-a.e. to a Borel allocation $\Theta_{0}$ which is balanced with respect to $\mu$. As $\vec{\mu}$ is supported on $\mathcal{T}_{* *}$, Proposition 3 then implies that for $\vec{\mu}$-almost all $[T, e, i] \in \mathcal{T}_{* *}$, we have that for all $\left(e^{\prime}, i^{\prime}\right) \in \Psi(T)$, $\Theta_{\epsilon_{n}}\left(T, e^{\prime}, i^{\prime}\right) \rightarrow \Theta_{0}\left(T, e^{\prime}, i^{\prime}\right)$. Since all hypertrees in $\mathcal{T}_{* *}$ are locally finite, this means that for $\vec{\mu}$-almost all $[T, e, i] \in \mathcal{T}_{* *}$, we have that for all $i^{\prime} \in V(T), \partial \Theta_{\epsilon_{n}}(T, i) \rightarrow \partial \Theta_{0}(T, i)$. Recall from Remark 12 that $\Theta_{\epsilon_{n}}(T, e, i)=\theta_{\epsilon_{n}}^{T}(e, i)$ where $\theta_{\epsilon_{n}}^{T}$ is the canonical $\epsilon_{n}$-balanced allocation for $T$. Thereby, for $\vec{\mu}$-almost $[T, e, i] \in \mathcal{T}_{* *}$, the conditions of Proposition 15 are satisfied. Thereby, for $\vec{\mu}$-almost all $[T, e, i]$, for all $\left(e^{\prime}, i^{\prime}\right) \in \Psi(T), \rho_{T_{e^{\prime} \rightarrow i^{\prime}}}^{\epsilon_{n}}($.$) converges pointwise to some \rho_{T_{e^{\prime} \rightarrow i^{\prime}}}($.$) . Moreover, for \vec{\mu}$-almost all $[T, e, i]$, we have that for all $\left(e^{\prime}, i^{\prime}\right) \in \Psi(T)$,

$$
\begin{equation*}
\rho_{T_{e^{\prime} \rightarrow i^{\prime}}}^{-1}(t)=t-\sum_{e^{\prime \prime} \ni i^{\prime}: e^{\prime \prime} \neq e^{\prime}}\left[1-\sum_{j \in e^{\prime \prime}, j \neq i^{\prime}}\left(\rho_{T_{e^{\prime \prime} \rightarrow j}^{-1}}^{-1}(t)\right)^{+}\right]_{0}^{1} \tag{76}
\end{equation*}
$$

and

$$
\begin{equation*}
\partial \Theta_{0}\left(T, i^{\prime}\right)>t \quad \Longleftrightarrow \quad \sum_{e^{\prime \prime} \ni i^{\prime}}\left[1-\sum_{j \in e^{\prime \prime}, j \neq i^{\prime}}\left(\rho_{e_{e^{\prime \prime} \rightarrow j}^{-1}}^{-1}(t)\right)^{+}\right]_{0}^{1}>t \tag{77}
\end{equation*}
$$

For $[T, e, i]$ such that $\Theta_{\epsilon_{n}}\left(e^{\prime}, i^{\prime}\right)$ is not convergent for some $\left(e^{\prime}, i^{\prime}\right) \in \Psi(T)$, we may define $\rho_{T_{e^{\prime \prime} \rightarrow i^{\prime \prime}}}^{-1}($. arbitrarily for $\left(e^{\prime \prime}, i^{\prime \prime}\right) \in \Psi(T)$. This will not impact our argument, as this happens only on a measure zero set. Using Lemma 1 part (ii), we conclude that for $\mu$-almost all $[T, i] \in \mathcal{T}_{*}$, we have that for all $\left(e^{\prime}, i^{\prime}\right) \in \Psi(T),(76)$ and (77) hold.

With this, we define the functions $F$ and $G$ on $\mathcal{T}_{* *}$ as follows:

$$
G(T, e, i):=\rho_{T_{e \rightarrow i}}^{-1}(t)
$$

and

$$
F(T, e, i):=\left[1-\sum_{j \in e, j \neq i}\left(\rho_{T_{e \rightarrow j}}^{-1}(t)\right)^{+}\right]_{0}^{1}=\left[1-\sum_{j \in e, j \neq i} G(T, e, j)^{+}\right]_{0}^{1}
$$

Moreover, define the function $f: \mathcal{T}_{*} \rightarrow\{0,1\}$ as $f(T, i)=\mathbb{1}_{\partial F(T, i)>t}$. From (77), $\mu$-a.s. we have $f=\mathbb{1}_{\partial \Theta_{0}>t}$. Hence, using the variational characterization in Proposition 11, we have

$$
\begin{equation*}
\int\left(\partial \Theta_{0}-t\right)^{+} d \mu=\int \tilde{f}_{\min } d \vec{\mu}-t \int f d \mu \tag{78}
\end{equation*}
$$

where $\tilde{f}_{\text {min }}: \mathcal{H}_{* *} \rightarrow[0,1]$ is defined as

$$
\tilde{f}_{\min }(T, e, i)=\frac{1}{|e|} \min _{j \in e} f(T, j)
$$

With this and the definition of $f$, we have

$$
\tilde{f}_{\min }(T, e, i)= \begin{cases}\frac{1}{|e|} & \partial F(T, j)>t \forall j \in e \\ 0 & \text { otherwise }\end{cases}
$$

From (76), for $\vec{\mu}$-almost every $[T, e, i]$ we have

$$
G(T, e, j)=t-\sum_{e^{\prime} \ni j, e^{\prime} \neq e} F\left(T, e^{\prime}, j\right)=t-\partial F(T, j)+F(T, e, j)
$$

Therefore, $\vec{\mu}$-almost everywhere, $\partial F(T, j)>t$ iff $F(T, e, j)>G(T, e, j)$. Consequently, $\vec{\mu}$-a.e., we have

$$
\tilde{f}_{\min }(T, e, i)= \begin{cases}\frac{1}{|e|} & F(T, e, j)>G(T, e, j) \forall j \in e  \tag{79}\\ 0 & \text { otherwise }\end{cases}
$$

Note that, by definition, we have $F(T, e, j)=\left[1-\sum_{l \in e, l \neq j} G(T, e, l)^{+}\right]_{0}^{1}$. Thereby, using Lemma 16 in Section 8.1, $\vec{\mu}$-a.e. we have

$$
\tilde{f}_{\min }(T, e, i)= \begin{cases}\frac{1}{|e|} & \sum_{j \in e} G(T, e, j)^{+}<1 \\ 0 & \text { otherwise }\end{cases}
$$

Note that the condition $\mathbb{E}\left[\|\Gamma\|_{1}\right]<\infty$ in particular implies that for all $k \geq 2$ we have $\mathbb{E}[\Gamma(k)]<\infty$. With this, using Lemma 20 in Appendix F, we have

$$
\int \tilde{f}_{\min } d \vec{\mu}=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \sum_{\gamma \in \Lambda} \hat{P}_{k}(\gamma) \mathbb{E}\left[f(\mathbb{T},(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma+\mathbb{e}_{k}\right]
$$

where $\mathbb{T}$ is the random rooted hypertree of Definition 28 . Following the argument in the proof of Proposition 4 in Appendix F, this can be written as

$$
\int \tilde{f}_{\min } d \vec{\mu}=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \mathbb{E}\left[\tilde{f}_{\min }\left(\tilde{\mathbb{T}}_{k},(k, 1), \emptyset\right)\right]
$$

where for $k \geq 2, \tilde{\mathbb{T}}_{k}$ is a tree with root $\emptyset$ that has an edge $(k, 1)$ of size $k$ connected to the root, with the type of the other edges connected to the root being $\hat{P}_{k}$, and with the subtrees at the other vertices of all the edges (including the edge $(k, 1)$ ) generated according to the rules of $\operatorname{UGWT}(P)$. Now, using the definition of $\tilde{f}_{\text {min }}$, we have

$$
\int \tilde{f}_{\min } d \vec{\mu}=\sum_{k=2}^{\infty} \frac{\mathbb{E}[\Gamma(k)]}{k} \mathbb{P}\left(\sum_{v \in(k, 1)} G\left(\tilde{\mathbb{T}}_{k},(k, 1), v\right)^{+}<1\right)
$$

Now, for every $k$ and $1 \leq i \leq k-1$, let $\tilde{\mathbb{T}}_{k, i}$ be the hypertree below vertex $(k, 1, i)$ rooted at $(k, 1, i)$. Moreover, let $\tilde{\mathbb{T}}_{k, k}$ be the hypertree rooted at $\emptyset$ obtained from $\tilde{\mathbb{T}}_{k}$ by removing the edge $(k, 1)$ and all its subtree. Now, due to the construction of $\tilde{\mathbb{T}}_{k}, \tilde{\mathbb{T}}_{k, i}$, for $1 \leq i \leq k$ are i.i.d. $\mathrm{GWT}_{k}(P)$ hypertrees. Hence, $G\left(\tilde{\mathbb{T}}_{k},(k, 1), v\right)$ for $v \in(k, 1)$ are independent and identically distributed. Let $Q_{k}$ be the common distribution. This means that

$$
\begin{equation*}
\int \tilde{f}_{\min } d \vec{\mu}=\sum_{k=2}^{\infty} \frac{\mathbb{E}[\Gamma(k)]}{k} \mathbb{P}\left(\sum_{i=1}^{k} X_{k, i}^{+}<1\right) \tag{80}
\end{equation*}
$$

where for each $k \geq 2, X_{k, i}, 1 \leq i \leq k$ are i.i.d. with law $Q_{k}$.

On the other hand, with $\mathbb{T}$ being the random rooted hypertree of Definition 28 , we have

$$
\begin{aligned}
\int f d \mu & =\mathbb{P}(\partial F(\mathbb{T}, \emptyset)>t)=\mathbb{P}\left(\sum_{k=2}^{h\left(\Gamma_{\emptyset}\right)} \sum_{i=1}^{\Gamma_{\emptyset}(k)} F(\mathbb{T},(k, i), \emptyset)>t\right) \\
& =\mathbb{P}\left(\sum_{k=2}^{h\left(\Gamma_{\emptyset}\right)} \sum_{i=1}^{\Gamma_{\emptyset}(k)}\left[1-\sum_{j=1}^{k-1} G(\mathbb{T},(k, i),(k, i, j))^{+}\right]_{0}^{1}>t\right),
\end{aligned}
$$

where $\Gamma_{\emptyset}$ is the type of the root in $\mathbb{T}$. But by definition, $G(\mathbb{T},(k, i),(k, i, j))=\rho_{\mathbb{T}_{(k, i) \rightarrow(k, i, j)}}^{-1}(t)$. But $\mathbb{T}_{(k, i) \rightarrow(k, i, j)}$ for $2 \leq k \leq \Gamma_{\emptyset}, 1 \leq i \leq \Gamma_{\emptyset}(k), 1 \leq j \leq k-1$ are independent and $\mathbb{T}_{(k, i) \rightarrow(k, i, j)}$ has law $\mathrm{GWT}_{k}(P)$. Comparing this to the above definition of the distributions $Q_{k}, k \geq 2$, we realize that $G(\mathbb{T},(k, i),(k, i, j))$ for $2 \leq k \leq \Gamma_{\emptyset}, 1 \leq i \leq \Gamma_{\emptyset}(k), 1 \leq j \leq k-1$ are independent and $G(\mathbb{T},(k, i),(k, i, j))$ has law $Q_{k}$. Consequently, as $\Gamma_{\emptyset}$ in the above expression has law $P$, we have

$$
\begin{equation*}
\int f d \mu=\mathbb{P}\left(\sum_{k=2}^{h(\Gamma)} \sum_{i=1}^{\Gamma(k)} Y_{k, i}>t\right) \tag{81}
\end{equation*}
$$

where $\Gamma$ has law $P$ and $\left\{Y_{k, i}\right\}_{k, i}$ are independent from each other and from $\Gamma$ such that $Y_{k, i}$ has the law of the random variable $\left[1-\left(Z_{1}^{+}+\cdots+Z_{k-1}^{+}\right)\right]_{0}^{1}$ where $Z_{j}$ are i.i.d. with law $Q_{k}$. This together with (80) and the variational expression (78), completes the proof.

## 9 Convergence of Maximum Load

In this section, we first introduce our configuration model and conditions under which it converges to the unimodular Galton-Watson hypertree model defined in Section 2.9. This is done in Section 9.1 below, specifically Theorem 4. We then state the conditions under which we prove Theorem 3 in Section 9.2, Proposition 18 and give the proof.

Before introducing our configuration model, we need to formally define multihypergraphs. Here, we only work with finite multihypergraphs.
Definition 34. A finite multihypergraph $H=\left\langle V, E=\left(e_{j}, j \in J\right)\right\rangle$ consists of a finite vertex set $V$ together with a finite edge index set $J$ such that each hyperedge $e_{j}$ is a multiset of vertices in $V$.

Here, the assumption of $e_{j}$ being a multiset allows for vertices to appear more than once in each edge. In this case, we call such an edge a "self loop". Moreover, it might be the case that $e_{j}=e_{j^{\prime}}$ for $j \neq j^{\prime} \in I$, in which case we call $e_{j}$ and $e_{j^{\prime}}$ "multiple edges".

### 9.1 Configuration model on Hypergraphs

We proposed a generalized Galton Watson process for hypertrees in Section 2.9 and showed it is unimodular. In this section, we propose a configuration model which converges to it in the local weak sense under certain conditions.

Assume that, for each integer $n$, a type sequence $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ is given such that $\gamma_{i}^{(n)} \in \Lambda$ and

$$
\begin{gather*}
\gamma_{i}^{(n)}(k)=0 \quad \forall 1 \leq i \leq n, k>n, \text { and }  \tag{82a}\\
k \mid \sum_{i=1}^{n} \gamma_{i}^{(n)}(k) \quad \forall 2 \leq k \leq n, \tag{82b}
\end{gather*}
$$

where the latter means that $k$ divides $\sum_{i=1}^{n} \gamma_{i}^{(n)}(k)$. In what follows, we generate a random multihypergraph on the vertex set $\{1, \ldots, n\}$ such that the type of node $i$ is $\gamma_{i}^{(n)}$. For each $2 \leq k \leq n$ and $1 \leq i \leq n$, we attach $\gamma_{i}^{(n)}(k)$ many objects $e_{i, 1}^{k}, \ldots, e_{i, \gamma_{i}^{(n)}(k)}^{k}$ called " $1 / k$-edges" to the node $i$. For each $k$, let $\Delta^{(n)}(k)$ be defined as the set of all $1 / k$-edges, i.e.

$$
\Delta^{(n)}(k):=\bigcup_{i=1}^{n}\left\{e_{i, 1}^{k}, \ldots, e_{i, \gamma_{i}^{(n)}(k)}^{k}\right\},
$$

and let $\Delta^{(n)}:=\cup_{k} \Delta^{(n)}(k)$ be the set of all "partial edges", where by a partial edge we mean a $1 / k$-edge for some $k$. Also, let $\Delta_{i}^{(n)}$ be the set of all partial edges connected to a node $i \in\{1, \ldots, n\}$, i.e.

$$
\Delta_{i}^{(n)}=\bigcup_{k: \gamma_{i}^{(n)}(k)>0}\left\{e_{i, 1}^{k}, \ldots, e_{i, \gamma_{i}^{(n)}(k)}^{k}\right\}
$$

For a partial edge $e \in \Delta^{(n)}$, define $\nu(e)$ to be the node it corresponds to, i.e. $\nu\left(e_{i, j}^{k}\right)=i$. Also, define $|e|$ to be the size of $e$, i.e. $\left|e_{i, j}^{k}\right|=k$.

We say that a permutation $\sigma_{k}$ is a $k$-matching on the set $\Delta^{(n)}(k)$ if it is a permutation on $\Delta^{(n)}(k)$ with no fixed points and also with all the cycles having size exactly equal to $k$. Due to the condition $k \mid \sum_{i=1}^{n} \gamma_{i}^{(n)}(k)$, such $k$-matchings exist for all $k$ such that $\Delta^{(n)}(k) \neq \emptyset$. In fact, if for a finite nonempty set $A$, whose cardinality is divisible by $k$, we denote the set of $k$-matchings on $A$ with $\mathcal{M}_{k}(A)$, it could be easily checked that $\left|\mathcal{M}_{k}(A)\right|=\frac{|A|!}{k^{|A| / k}(|A| / k)!}$.

Given this, for each $k$ such that $\Delta^{(n)}(k) \neq \emptyset$, we pick $\sigma_{k}$ uniformly at random from $\mathcal{M}_{k}\left(\Delta^{(n)}(k)\right)$, independently over $k$. With this, we generate a random multihypergraph $H_{n}$ on the set of vertices $\{1, \ldots, n\}$. This is done by identifying each cycle of $\sigma_{k}$ such that $\Delta^{(n)}(k) \neq \emptyset$ of the form $\left(e, \sigma_{k}(e), \ldots, \sigma_{k}^{(k-1)}(e)\right)$ with the edge $\left\{\nu(e), \nu\left(\sigma_{k}(e)\right), \ldots, \nu\left(\sigma_{k}^{(k-1)}(e)\right)\right\}$ in $H_{n}$. Here, $e \in \Delta^{(n)}(k)$ and $\sigma_{k}^{(l)}$ denotes the permutation $\sigma_{k}$ begin applied $l$ times. Note that it is possible that two realizations of permutations as above result in the same multihypergraph $H_{n}$. As an example, if $n=3, \gamma_{1}^{(n)}(2)=2$, $\gamma_{2}^{(n)}(2)=\gamma_{3}^{(n)}(2)=1$, and $\gamma_{i}^{(n)}(k)=0$ for $k \neq 2,1 \leq i \leq 3$, then the two permutations $\sigma_{2}$ and $\sigma_{2}^{\prime}$ on $\Delta^{(n)}(2)$ presented in the cycle notation as $\sigma_{2}=\left(\left(e_{1,1}^{2}, e_{2,1}^{2}\right)\left(e_{1,2}^{2}, e_{3,1}^{2}\right)\right.$ and $\sigma_{2}^{\prime}=\left(e_{1,2}^{2}, e_{2,1}^{2}\right)\left(e_{1,1}^{2}, e_{3,1}^{2}\right)$ would result in the same multihypergraph (which turns out to be a simple graph in this example).

Note that, in general, $H_{n}$ might not be simple. In particular, it might contain edges which contain a vertex more than once (we call such an edge a "self loop"), or it might be the case that two edges exist in $H_{n}$ with the exact same multiset of vertices (we call such an edge a "multiple edge"), or these two can happen simultaneously. Having generated $H_{n}$, we may generate a simple hypergraph $H_{n}^{e}$ by deleting all such self loops and multiple edges in $H_{n}$, i.e. we first remove all self loops and, subsequently, we delete all edges with the same set of endpoints. See Figure 5 for an example of an instance of the configuration model.

If one fixes a type sequence $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ for each $n$, then the above configuration model generates a sequence of random hypergraphs $H_{n}^{e}$. Since $H_{n}^{e}$ is random, $u_{H_{n}^{e}}$ is also a random probability distribution over $\mathcal{H}_{*}$. Let $\mathbb{E}\left[u_{H_{n}^{e}}\right]$ be the expectation with respect to the randomness of $H_{n}^{e}$, i.e. a weighted average over all the possible configurations of $H_{n}^{e}$ (the number of such configurations is indeed finite for each $n$ ). Hence, $\mathbb{E}\left[u_{H_{n}^{e}}\right]$ is a sequence of probability distributions over $\mathcal{H}_{*}$ and one can ask whether it has a weak limit. On the other hand, one can build all $H_{n}^{e}$ on a common probability space under which $H_{n}^{e}$ are independent for different $n$. Then, for some $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$, we say that $u_{H_{n}^{e}} \Rightarrow \mu$ almost surely when outside a measure zero set in this common probability space, this convergence holds. The following theorem shows that under some conditions on the type sequence, $H_{n}^{e}$ has a local weak limit. See Appendix G for a proof.


Figure 5: An example of an instance of the configuration model on $n=5$ vertices with vertex types $\gamma_{1}^{(n)}=(0,2), \gamma_{2}^{(n)}=(1,2), \gamma_{3}^{(n)}=(0,3), \gamma_{4}^{(n)}=(1,1), \gamma_{5}^{(n)}=(1,2)$. Here, a type $\gamma$ is represented by a vector where the first coordinate is $\gamma(2)$, the second is $\gamma(3)$ and so on. (a) illustrates the set of partial edges connected to vertices. (b) depicts the multihypergraph $H_{n}$ formed by matching the partial edges using the permutations $\sigma_{2}$ and $\sigma_{3}$ which are represented in the cycle notation as $\sigma_{2}=\left(e_{2,1}^{2}, e_{4,1}^{2}\right)\left(e_{5,1}^{2}, e_{5,2}^{2}\right)$ and $\sigma_{3}=\left(e_{1,1}^{3}, e_{2,1}^{3}, e_{3,1}^{3}\right)\left(e_{1,2}^{3}, e_{2,2}^{3}, e_{3,2}^{3}\right)\left(e_{3,3}^{3}, e_{4,1}^{3}, e_{5,1}^{3}\right)$. (c) illustrates the simple hypergraph $H_{n}^{e}$ formed by removing the multiple edges of size 3 on vertices $1,2,3$ and also the self loop on vertex 5 .

Theorem 4. Assume that a probability distribution $P$ on $\Lambda$ is given and define $I:=\{k \geq 2: P(\{\gamma \in$ $\Lambda: \gamma(k)>0\})>0\}$. Furthermore, let $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ be a type sequence satisfying (82a) and (82b) such that

$$
\begin{gather*}
\gamma_{i}^{(n)}(k)=0 \quad \forall k \notin I \quad \forall n, i,  \tag{83a}\\
\lim _{n \rightarrow} \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\gamma_{i}^{(n)}=\gamma}=P(\gamma) \quad \forall \gamma \in \Lambda,  \tag{83b}\\
\limsup _{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n}\left\|\gamma_{i}^{(n)}\right\|_{1}^{2}<\infty \tag{83c}
\end{gather*}
$$

Additionally, assume that there are constants $c_{1}, c_{2}, c_{3}, \epsilon>0$ such that for $n$ large enough,

$$
\begin{gather*}
\max _{1 \leq i \leq n}\left\|\gamma_{i}^{(n)}\right\|_{1} \leq c_{1}(\log n)^{c_{2}},  \tag{84a}\\
\max _{1 \leq i \leq n} h\left(\gamma_{i}^{(n)}\right) \leq c_{1}(\log n)^{c_{2}},  \tag{84b}\\
\forall 2 \leq k \leq n \quad \Delta^{(n)}(k)=\emptyset \quad \text { or } \quad\left|\Delta^{(n)}(k)\right| \geq c_{3} n^{\epsilon} . \tag{84c}
\end{gather*}
$$

Then, if $H_{n}^{e}$ is the random hypergraph generated from the configuration model corresponding to $\gamma^{(n)}$ described above, we have $u_{H_{n}^{e}} \Rightarrow U G W T(P)$ almost surely.

Remark 13. Note the the index set I would let us include only certain edge sizes in our model. For instance, when $I=\{2\}$, this theorem reduces to a statement for the graph pairing model, and when $I=\{k: k \geq 2\}$ it allows for all edge sizes to be present.

Remark 14. For a fixed $k$ and for each $n$, define $X_{n}$ to be an integer valued random variable taking value $\gamma_{i}^{(n)}(k)$ with probability $1 / n$ for $1 \leq i \leq n$. Then, the condition (83c) implies that the sequence $\left\{X_{n}\right\}$ is uniformly integrable. Also, (83b) implies that $X_{n} \xrightarrow{d} \Gamma(k)$ where $\Gamma$ has law $P$. Thus, $\mathbb{E}\left[X_{n}\right] \rightarrow$ $\mathbb{E}[\Gamma(k)]$, i.e.

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n} \gamma_{i}^{(n)}(k)=\mathbb{E}[\Gamma(k)] \quad \forall k>1 \tag{85}
\end{equation*}
$$

This identity is useful in our future analysis.
Remark 15. It can be proved that, under some regularity conditions for $P$, when the type sequence is generated i.i.d., that the conditions of Theorem 4 are satisfied with probability one. However, we omit the proof of such a claim here, since it is not central to our discussion.

We will use the following simplified version of the above theorem in this section.
Corollary 2. Assume that a probability distribution $P$ over $\Lambda$ is given such that for a finite set $I \subset\{2,3, \ldots$,$\} , if \Gamma$ is a random variable with law $P$, we have $\mathbb{P}(\Gamma(k)>0)>0$ for $k \in I$ and $\mathbb{P}(\Gamma(k)>0)=0$ for $k \notin I$. Furthermore, assume $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ is a type sequence satisfying (82a), (82b), (83a) and (83b). Additionally, assume that for some $\theta>0$,

$$
\begin{equation*}
\limsup _{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^{n} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}}<\infty \tag{86}
\end{equation*}
$$

Then, $u_{H_{n}^{e}} \Rightarrow U G W T(P)$ almost surely.
Proof. We check that in this regime, all the conditions of Theorem 4 are satisfied. Note that

$$
\frac{1}{n} \sum_{i=1}^{n} \gamma_{i}^{(n)}(k)^{2} \leq \frac{2!}{\theta^{2}} \frac{1}{n} \sum_{i=1}^{n} e^{\theta \gamma_{i}^{(n)}(k)} \leq \frac{2!}{\theta^{2}} \frac{1}{n} \sum_{i=1}^{n} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}}
$$

which, together with (86), implies (83c). In order to show (84a), note that (86) implies there is a constant $\lambda<\infty$ such that $\frac{1}{n} \sum_{i=1}^{n} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}}<\lambda$ for all $n$. Now, $\exp \left(\theta \max _{1 \leq i \leq n}\left\|\gamma_{i}^{(n)}\right\|_{1}\right) \leq$ $\sum_{i=1}^{n} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}} \leq n \lambda$. Hence, $\max _{1 \leq i \leq n}\left\|\gamma_{i}^{(n)}\right\|_{1} \leq(\log n+\log \lambda) / \theta$, which shows (84a). On the other hand, $h\left(\gamma_{i}^{(n)}\right) \leq \max \{k: k \in I\}$ for all $i$ and $n$. Hence, (84b) also holds.

For a fixed $k \in I$, using (85), which follows from (83c) which was proved above, for $n$ large enough we have $\left|\Delta^{(n)}(k)\right|=\sum_{i=1}^{n} \gamma_{i}^{(n)}(k) \geq n \mathbb{E}[\Gamma(k)] / 2$. Since there are finitely many $k \in I$, we have that for $n$ large enough $\left|\Delta^{(n)}(k)\right| \geq n \min _{k^{\prime} \in I} \mathbb{E}\left[\Gamma\left(k^{\prime}\right)\right] / 2$. But for $k^{\prime} \in I$, we have $\mathbb{E}\left[\Gamma\left(k^{\prime}\right)\right] \geq \mathbb{P}\left(\Gamma\left(k^{\prime}\right)>0\right)>0$ by assumption. Consequently, (84c) holds with $\epsilon=1$ and $c_{3}:=\min _{k^{\prime} \in I} \mathbb{E}\left[\Gamma\left(k^{\prime}\right)\right] / 2>0$. This means that all the conditions of Theorem 4 are satisfied and $u_{H_{n}^{e}} \Rightarrow \operatorname{UGWT}(P)$ almost surely.

### 9.2 Statement of the result

Here, we state the conditions under which we prove Theorem 3.
Proposition 18. Assume that a probability distribution $P$ over $\Lambda$ is given such that for a finite set $I \subset\{2,3, \ldots$,$\} , if \Gamma$ is a random variable with law $P$, we have $\mathbb{P}(\Gamma(k)>0)>0$ for $k \in I$ and $\mathbb{P}(\Gamma(k)>0)=0$ for $k \notin I$. Moreover, assume that for all $k \in I, \mathbb{E}[\Gamma(k)]<\infty$. Also, with $k_{\text {min }}$ being the minimum element in $I$, assume that $\mathbb{P}\left(\Gamma\left(k_{\min }\right)=0\right)+\mathbb{P}\left(\Gamma\left(k_{\min }\right)=1\right)<1$. Moreover, assume that a sequence of types $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ is given satisfying (82a), (82b), (83a) (83b) and (86). Then, if $H_{n}^{e}$ is the simple hypergraph generated by the configuration model in Section 9.1, $\varrho\left(H_{n}^{e}\right)$ converges in probability to $\varrho(\mu)$, where $\mu=U G W T(P)$.

Before proving this proposition, we need the following two lemmas, whose proofs are given at the end of this section.

Lemma 17. With the assumptions of Proposition 18, there is a constant $c_{4}>0$ such that for $n$ large enough, for any subset $S \subset\{1, \ldots, n\}$, the number of edges in $H_{n}^{e}$ with all endpoints in $S$ is stochastically dominated by the sum of s independent Bernoulli random variables, each having mean $c_{4} s^{k_{\text {min }}-1} / n^{k_{\text {min }}-1}$, where $s:=\sum_{i \in S}\left\|\gamma_{i}^{(n)}\right\|_{1}$.
Lemma 18. With the assumptions of Proposition 18, if $t>\frac{1}{k_{\min }-1}$, there exists $\delta>0$ such that if $Z_{\delta, t}^{(n)}$ denotes the number of subsets $S$ of $\{1, \ldots, n\}$ with size at most $n \delta$ such that $H_{n}^{e}$ has at least $t|S|$ many edges with all endpoints in $S$, we have $\mathbb{P}\left(Z_{\delta, t}^{(n)}>0\right) \rightarrow 0$ as $n \rightarrow \infty$.
Proof of Proposition 18. Let $\mu_{n}$ denote $u_{H_{n}^{e}}$, which is a random probability distribution on $\mathcal{H}_{*}$. Then, Corollary 2 guarantees that $\mu_{n} \Rightarrow \mu$ almost surely. As a result, if $\mathcal{L}_{n}$ is the law of the balanced load for $H_{n}^{e}$ and $\mathcal{L}$ is the law of $\partial \Theta$, where $\Theta$ is the balanced allocation corresponding to $\mu$, then, using Theorem 1, we have $\mathcal{L}_{n} \Rightarrow \mathcal{L}$ almost surely. Now, let $t:=\varrho(\mu)$, and fix $\epsilon>0$. Due to the definition of $\varrho(\mu), \mathcal{L}((t-\epsilon, \infty))>0$. As a result, using the portmanteau theorem (see, for instance, [Bil99, Theorem 2.1]), since $\mu_{n} \Rightarrow \mu$ almost surely, we have

$$
\liminf _{n} \mathcal{L}_{n}((t-\epsilon, \infty))>0 \quad \text { a.s.. }
$$

This means that

$$
\begin{equation*}
\mathbb{P}\left(\varrho\left(H_{n}^{e}\right) \leq t-\epsilon\right)=\mathbb{P}\left(\mathcal{L}_{n}((t-\epsilon, \infty))=0\right) \rightarrow 0 \tag{87}
\end{equation*}
$$

Now we show that $\mathbb{P}\left(\varrho\left(H_{n}^{e}\right) \geq t+\epsilon\right)$ also converges to zero. To do so, fix some $\delta>0$ and note that

$$
\begin{align*}
\mathbb{P}\left(\varrho\left(H_{n}^{e}\right) \geq t+\epsilon\right) & =\mathbb{P}\left(\mathcal{L}_{n}([t+\epsilon, \infty))>0\right) \\
& =\mathbb{P}\left(\mathcal{L}_{n}([t+\epsilon, \infty))>\delta\right)+\mathbb{P}\left(0<\mathcal{L}_{n}([t+\epsilon, \infty)) \leq \delta\right) \tag{88}
\end{align*}
$$

The portmanteau theorem implies that $\mathbb{P}\left(\mathcal{L}_{n}([t+\epsilon, \infty))>\delta\right)$ converges to zero. Now, we argue that the second term also converges to zero. If $\theta_{n}$ denotes the balanced allocation on $H_{n}^{e}$, then, by the definition of $\mathcal{L}_{n}$, on the event $0<\mathcal{L}_{n}([t+\epsilon, \infty)) \leq \delta$, the set $S:=\left\{1 \leq i \leq n: \partial \theta_{n}(i) \geq t+\epsilon\right\}$ is non-empty and $|S| \leq \delta n$. Now if $e \in E\left(H_{n}^{e}\right)$ with $i, j \in e$ such that $i \in S$ and $j \in S^{c}$, then, since $\partial \theta_{n}(j)<t+\epsilon \leq \partial \theta_{n}(i)$, we have $\theta_{n}(e, i)=0$. Hence, for all $i \in S, \partial \theta_{n}(i)=\sum_{e \subseteq S, e \in E\left(H_{n}^{e}\right)} \theta_{n}(e, i)$. Thus,

$$
\sum_{i \in S} \partial \theta_{n}(i)=\left|\left\{e \in E\left(H_{n}^{e}\right): e \subseteq S\right\}\right|=\left|E_{H_{n}^{e}}(S)\right|
$$

On the other hand, we have $\sum_{i \in S} \partial \theta_{n}(i) \geq|S|(t+\epsilon)$. Hence, $\left|E_{H_{n}^{e}}(S)\right| \geq(t+\epsilon)|S|$, where $E_{H_{n}^{e}}(S)$ denotes the set of edges in $H_{n}^{e}$ with all endpoints in $S$. As a result, if $Z_{\delta, t+\epsilon}^{(n)}$ denotes the number of subsets $S \subset\{1, \ldots, n\}$ with size at most $n \delta$ such that $H_{n}^{e}$ has at least $(t+\epsilon)|S|$ many edges with all endpoints in $S$, we have

$$
\begin{aligned}
& \mathbb{P}\left(0<\mathcal{L}_{n}([t+\epsilon, \infty))<\delta\right) \leq \mathbb{P}(\exists S \subseteq\{1, \ldots, n\}, 0<|S| \leq \delta n \\
&\left.\left|E_{H_{n}^{e}}(S)\right| \geq(t+\epsilon)|S|\right) \\
& \leq \mathbb{P}\left(Z_{\delta, t+\epsilon}^{(n)}>0\right)
\end{aligned}
$$

If we have $t \geq 1 /\left(k_{\min }-1\right)$, Lemma 18 above implies that $\mathbb{P}\left(Z_{\delta, t+\epsilon}^{(n)}>0\right)$ goes to zero as $n \rightarrow \infty$ and we are done. We now argue why this is the case. For this, note that Proposition 10 together with Proposition 3 imply that there exists a sequence of $\epsilon_{m}$-balanced allocations $\Theta_{\epsilon_{m}}$ such that for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}$, for all vertices $j \in V(H), \partial \Theta_{\epsilon_{m}}(H, j) \rightarrow \partial \Theta_{0}(H, j)$ where $\Theta_{0}$ is a balanced allocation with respect to $\mu$. Moreover, using Remark $12, \partial \Theta_{\epsilon_{m}}(H, j)=\partial \theta_{\epsilon_{m}}^{H}(j)$ where $\theta_{\epsilon_{n}}^{H}$ is the canonical $\epsilon$-balanced allocation on $H$.

On the other hand, the assumption $\mathbb{P}\left(\Gamma\left(k_{\min }\right)=0\right)+\mathbb{P}\left(\Gamma\left(k_{\min }\right)=1\right)<1$ guarantees that for all integer $M \geq 2$, there is a nonzero probability under $\mu$ that the Galton-Watson process has a finite sub-hypertree containing the root, and having $M$ edges with all these edge having size $k_{\min }$. It can be easily seen that a finite hypertree with $M$ edges all having size $c$ has $1+M(c-1)$ vertices and there is a balanced allocation on such a hypertree such that all the vertices get the same amount of load, which is equal to $\frac{M}{1+M(c-1)}$. Motivated by the discussion in the previous paragraph, for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}$, $\partial \Theta_{0}(H, i)=\lim _{m \rightarrow \infty} \partial \theta_{\epsilon_{m}}^{H}(i)$ where $\theta_{\epsilon_{m}}^{H}$ is the canonical $\epsilon_{m}$-balanced allocation on $H$. This, together with Proposition 6, implies that for each integer $M$, there is a nonzero probability under $\mu$ that $\partial \Theta_{0}(H, i)$ is at least $M /\left(1+M\left(k_{\min }-1\right)\right)$. Sending $M \rightarrow \infty$, this means that $t=\varrho(\mu) \geq 1 /\left(k_{\min }-1\right)$. As was discussed above, using Lemma 18 above, $\mathbb{P}\left(Z_{\delta, t+\epsilon}^{(n)}>0\right) \rightarrow 0$. Thus, $\mathbb{P}\left(\varrho\left(H_{n}^{e}\right) \geq t+\epsilon\right)$ goes to zero as $n$ goes to infinity. This together with (87) proves that $\varrho\left(H_{n}^{e}\right) \xrightarrow{p} \varrho(\mu)$.

Proof of Lemma 17. For $k \in I$, let $s_{k}:=\sum_{i \in S} \gamma_{i}^{(n)}(k)$ and $m_{k}:=\left|\Delta^{(n)}(k)\right|$. As the set of edges in $H_{n}^{e}$ is a subset of that of $H_{n}$, we may prove the result for $H_{n}$ instead. This allows us to directly analyze the configuration model. Let $A$ be the set of the partial edges connected to vertices in $S$. Note that $|A|=s$. Arbitrarily order the elements in $A$. At time $t=1$, we pick the smallest element in $A$. Let $k_{1}$ be the size of this edge. Then, we choose $k_{1}-1$ other partial edges in $\Delta^{(n)}\left(k_{1}\right)$ uniformly at random to form an edge in $H_{n}$. We continue this process until all the elements in $A$ are used up. More precisely, at time $t$, we pick the smallest available partial edge in $A$, namely $e_{t}$, and if $k_{t}$ is the size of $e_{t}$, we match $e_{t}$ with $k_{t}-1$ other elements in the available partial edges in $\Delta^{(n)}\left(k_{t}\right)$ uniformly at random to form an edge in $H_{n}$. At time $t$, let $s_{k_{t}}(t)$ and $m_{k_{t}}(t)$ be the number of available partial edges of size $k_{t}$ in $A$ and $\Delta^{(n)}\left(k_{t}\right)$, respectively. With this, if $p_{t}$ denotes the probability that $e_{t}$ is matched with partial edges all inside $A$,

$$
p_{t}=\frac{s_{k_{t}}(t)-1}{m_{k_{t}}(t)-1} \times \frac{s_{k_{t}}(t)-2}{m_{k_{t}}(t)-2} \times \cdots \times \frac{s_{k_{t}}(t)-\left(k_{t}-1\right)}{m_{k_{t}}(t)-\left(k_{t}-1\right)} \leq\left(\frac{s_{k_{t}}(t)}{m_{k_{t}}(t)}\right)^{k_{t}-1}
$$

Note that if $l_{t}$ denotes the number of partial edges of size $k_{t}$ among $e_{1}, \ldots, e_{t-1}, m_{k_{t}}(t)$ is precisely $m_{k_{t}}-l_{t} k_{t}$. On the other hand, $s_{k_{t}}(t) \leq s_{k_{t}}-l_{t}$, where equality holds only if all the partial edges of size $k$ among $e_{1}, \ldots, e_{t-1}$ are matched with partial edges outside $A$. With this,

$$
p_{t} \leq\left(\frac{k_{t} s_{k_{t}}}{m_{k_{t}}}\right)^{k_{t}-1}
$$

This is because, if $k_{t} s_{k_{t}} \leq m_{k_{t}}$, the even stronger inequality $p_{t} \leq\left(\frac{s_{k_{t}}}{m_{k_{t}}}\right)^{k_{t}-1}$ holds, while otherwise the RHS is at least 1 , and the inequality is trivial. Furthermore, as we saw in the proof of Corollary 2, there exists $\alpha>0$ such that for $n$ large enough and all $k \in I, m_{k} \geq n \alpha$. This together with the fact that $s_{k} \leq s$ for all $k \in I$ implies

$$
p_{t} \leq k_{\max }^{k_{\max }-1}\left(\frac{s}{n \alpha}\right)^{k_{\min }-1}
$$

where $k_{\text {max }}$ denotes the maximum element in $I$. As this upper bound is a constant, and the above process can continue for at most $s$ steps until we match all partial edges in $A$, the number of edges
with all endpoints in $S$ is stochastically dominated by $\operatorname{Binomial}\left(s, k_{\max }^{k_{\max }-1}\left(\frac{s}{n \alpha}\right)^{k_{\min }-1}\right)$. The proof is complete by setting $c_{4}:=k_{\max }^{k_{\max }-1} / \alpha^{k_{\min }-1}$.
Proof of Lemma 18. For positive integers $L$ and $r$, let $X_{L, r}^{(n)}$ denote the number of subsets $S \subset$ $\{1, \ldots, n\}$ with size $L$ such that $H_{n}^{e}$ has at least $r$ many edges with all endpoints in $S$. With this, $\mathbb{E}\left[Z_{\delta, t}^{(n)}\right]=\sum_{L=1}^{\lfloor n \delta\rfloor} \mathbb{E}\left[X_{L,\lceil L t\rceil}^{(n)}\right]$. Now, fix integers $L$ and $r$ such that $L \leq n \delta$, with $\delta>0$ sufficiently small. Let $\mathcal{S}_{L}$ denote the set of $S \subset\{1, \ldots, n\}$ with size equal to $L$. Using Lemma 17 and the fact that for a binomial random variable $Z, \mathbb{P}(Z \geq r) \leq(\mathbb{E}[Z])^{r} / r$ !, we have

$$
\mathbb{E}\left[X_{L, r}^{(n)}\right] \leq \sum_{S \in \mathcal{S}_{L}} \frac{1}{r!}\left(c_{4} \frac{\left(\sum_{i \in S}\left\|\gamma_{i}^{(n)}\right\|_{1}\right)^{k_{\min }}}{n^{k_{\min }-1}}\right)^{r}
$$

Using the inequality $x^{m} \leq m!e^{x}$ which holds for $x \geq 0$ and integer $m$, this yields

$$
\begin{aligned}
\mathbb{E}\left[X_{L, r}^{(n)}\right] & \leq \sum_{S \in \mathcal{S}_{L}} \frac{\left(r k_{\min }\right)!c_{4}^{r}}{r!\theta^{r k_{\min }} n^{r\left(k_{\min }-1\right)}} \prod_{i \in S} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}} \\
& \leq \frac{\left(r k_{\min }\right)!c_{4}^{r}}{r!\theta^{r k_{\min }} n^{r\left(k_{\min }-1\right)}} \frac{1}{L!}\left(\sum_{i=1}^{n} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}}\right)^{L}
\end{aligned}
$$

where $\theta>0$ is as in the statement of Corollary 2. Using the assumption (86), there exists $\lambda>0$ such that $\sum_{i=1}^{n} e^{\theta\left\|\gamma_{i}^{(n)}\right\|_{1}}<n \lambda$ for all $n$. Using this, together with the inequalities $L!\geq(L / e)^{L}$ and $\left(r k_{\min }\right)!/ r!\leq\left(r k_{\min }\right)^{r\left(k_{\min }-1\right)}$, and rearranging the terms, this yields

$$
\mathbb{E}\left[X_{L, r}^{(n)}\right] \leq\left(\frac{c_{5} r}{n}\right)^{r\left(k_{\min }-1\right)}\left(\frac{e n \lambda}{L}\right)^{L}
$$

where $c_{5}:=k_{\min } c_{4}^{\frac{1}{k_{\min }-1}} / \theta^{\frac{k_{\min }}{k_{\min }-1}}$. Note that we may assume $c_{5}>1$; otherwise, we may replace it with $c_{5} \vee 1$ which makes this quantity even bigger; for, the exponent of $c_{5}$ is positive.

Using this bound for $r=\lceil L t\rceil$, we have

$$
\begin{aligned}
\mathbb{E}\left[X_{L,\lceil L t\rceil}^{(n)}\right] & \leq\left(\frac{c_{5}\lceil L t\rceil}{n}\right)^{\lceil L t\rceil\left(k_{\min }-1\right)}\left(\frac{e n \lambda}{L}\right)^{L} \\
& =\left(\frac{c_{5}\lceil L t\rceil}{L}\right)^{\lceil L t\rceil\left(k_{\min }-1\right)}(e \lambda)^{L}\left(\frac{L}{n}\right)^{\lceil L t\rceil\left(k_{\min }-1\right)-L}
\end{aligned}
$$

Using $c_{5}>1, L / n<1, L t \leq\lceil L t\rceil \leq L(t+1)$ and the assumption $t>1 /\left(k_{\min }-1\right)$, we get the upper bound

$$
\mathbb{E}\left[X_{L,\lceil L t\rceil}^{(n)}\right] \leq f\left(\frac{L}{n}\right)^{L}
$$

where

$$
f(x):=c_{6} x^{t\left(k_{\min }-1\right)-1}
$$

with $c_{6}:=e \lambda\left(c_{5}(t+1)\right)^{(t+1)\left(k_{\min }-1\right)}$. Note that $f(L / n)^{L}=\exp (-n g(L / n))$, where

$$
g(x):=-x \log c_{6}-\left(t\left(k_{\min }-1\right)-1\right) x \log x
$$

As $t>1 /\left(k_{\min }-1\right)$, there exists $a>0$ such that $g(x)$ is strictly increasing and strictly positive in $(0, a)$. Now, we choose $\delta>0$ such that $\delta<a$ and also $f(\delta)<1$. This is possible since the assumption $t>1 /\left(k_{\min }-1\right)$ guarantees $f(\delta) \rightarrow 0$ as $\delta \downarrow 0$. With this, for any $0<\zeta<\delta$, we have

$$
\begin{aligned}
\mathbb{E}\left[Z_{\delta, t}^{(n)}\right] & \leq \sum_{L=1}^{\lfloor n \delta\rfloor} f(L / n)^{L} \\
& =\sum_{L=1}^{\lfloor n \zeta\rfloor} f(L / n)^{L}+\sum_{L=\lfloor n \zeta\rfloor+1}^{\lfloor n \delta\rfloor} \exp (-n g(L / n)) .
\end{aligned}
$$

Using the facts that $f$ is increasing in $(0, \infty), g$ is increasing in $(0, a)$ and $0<L / n \leq \delta<a$, we have

$$
\mathbb{E}\left[Z_{\delta, t}^{(n)}\right] \leq\left(\sum_{L=1}^{\lfloor n \zeta\rfloor} f(\zeta)^{L}\right)+n \delta \exp (-n g(\zeta))
$$

But $f(\zeta)<f(\delta)<1$. Therefore,

$$
\mathbb{E}\left[Z_{\delta, t}^{(n)}\right] \leq \frac{f(\zeta)}{1-f(\zeta)}+n \delta \exp (-n g(\zeta))
$$

Now, by sending $n$ to infinity, the second term vanishes, because $g(\zeta)>0$, and we get $\lim \sup \mathbb{E}\left[Z_{\delta, t}^{(n)}\right] \leq$ $f(\zeta) /(1-f(\zeta))$. Furthermore, by sending $\zeta \rightarrow 0$, we get $\mathbb{E}\left[Z_{\delta, t}^{(n)}\right] \rightarrow 0$ which means $\mathbb{P}\left(Z_{\delta, t}^{(n)}>0\right) \rightarrow 0$, as $Z_{\delta, t}^{(n)}$ is integer valued.

## A Weak uniqueness of balanced allocations

Proof of Proposition 2. For a fixed $\delta>0$, define the set

$$
A_{\delta}:=\left\{i \in V(H): \partial_{b} \theta(i)-\partial_{b} \theta^{\prime}(i)>\delta\right\} .
$$

By assumption, we have $\sum_{i \in V(H)}\left|\partial_{b} \theta(i)-\partial_{b} \theta^{\prime}(i)\right|<\infty$. Hence, $A_{\delta}$ is a finite set. Moreover

$$
\begin{equation*}
\sum_{i \in A_{\delta}} \partial_{b} \theta(i)-\partial_{b} \theta^{\prime}(i)=\sum_{i \in A_{\delta}} \sum_{e \ni i, e \nsubseteq A_{\delta}} \theta(e, i)-\theta^{\prime}(e, i) . \tag{89}
\end{equation*}
$$

Now, fix some $e \in E(H)$ such that $e \cap A_{\delta} \neq \emptyset$ and $e \nsubseteq A_{\delta}$. For $i \in e \cap A_{\delta}$ and $j \in e \backslash A_{\delta}$, we have

$$
\partial_{b} \theta(j)-\partial_{b} \theta^{\prime}(j) \leq \delta<\partial_{b} \theta(i)-\partial_{b} \theta\left(i^{\prime}\right)
$$

which means that

$$
\partial_{b} \theta(j)-\partial_{b} \theta(i)<\partial_{b} \theta^{\prime}(j)-\partial_{b} \theta^{\prime}(i)
$$

Hence it is either the case that $\partial_{b} \theta^{\prime}(j)>\partial_{b} \theta^{\prime}(i)$ or $\partial_{b} \theta(j)<\partial_{b} \theta(i)$.
If $\theta(e, i)=0$ for all $i \in e \cap A_{\delta}$, then $\sum_{i \in e \cap A_{\delta}} \theta(e, i)-\theta^{\prime}(e, i) \leq 0$. If $\theta\left(e, i^{*}\right) \neq 0$ for some $i^{*} \in e \cap A_{\delta}$, then $\partial_{b} \theta\left(i^{*}\right) \leq \partial_{b} \theta(j)$ for all $j \in e \backslash A_{\delta}$. Consequently, $\partial_{b} \theta^{\prime}(j)>\partial_{b} \theta^{\prime}\left(i^{*}\right)$ for all $j \in e \backslash A_{\delta}$; thereby, $\theta^{\prime}(e, j)=0$ for all $j \in e \backslash A_{\delta}$. This means that $\sum_{i \in e \cap A_{\delta}} \theta^{\prime}(e, i)=1 \geq \sum_{i \in e \cap A_{\delta}} \theta(e, i)$. Hence, we have
observed that, in either case, we have $\sum_{i \in e \cap A_{\delta}} \theta(e, i)-\theta^{\prime}(e, i) \leq 0$. Since this is true for all $e$ such that $e \cap A_{\delta} \neq \emptyset$ and $e \nsubseteq A_{\delta}$, substituting into (89) we realize that

$$
\sum_{i \in A_{\delta}} \partial_{b} \theta(i)-\partial_{b} \theta^{\prime}(i) \leq 0
$$

On the other hand, $\sum_{i \in A_{\delta}} \partial_{b} \theta(i)-\partial_{b} \theta^{\prime}(i) \geq \delta\left|A_{\delta}\right|$. Combining these two we conclude that $A_{\delta}=\emptyset$. Symmetrically, the set $B_{\delta}:=\left\{i \in V(H): \partial_{b} \theta^{\prime}(i)-\partial_{b} \theta(i)>\delta\right\}$ should be empty. Since $\delta$ is arbitrary, we conclude that $\partial \theta \equiv \partial \theta^{\prime}$, i.e. $\partial_{b} \theta \equiv \partial_{b} \theta^{\prime}$, which completes the proof.

## B $\quad \overline{\mathcal{H}}_{*}(\Xi)$ and $\overline{\mathcal{H}}_{* *}(\Xi)$ are Polish spaces

In this section, we prove that $\overline{\mathcal{H}}_{*}(\Xi)$ and $\overline{\mathcal{H}}_{* *}(\Xi)$ are Polish spaces when $\Xi$ is a Polish space. In particular, by setting $\Xi=\{\emptyset\}$, this means that $\mathcal{H}_{*}$ and $\mathcal{H}_{* *}$ are Polish spaces.

Proposition 19. Assume $\Xi$ is a Polish space. Then, $\overline{\mathcal{H}}_{*}(\Xi)$ and $\overline{\mathcal{H}}_{* *}(\Xi)$ are Polish spaces.
Proof. We give the proof for $\overline{\mathcal{H}}_{*}(\Xi)$ here. The proof for $\overline{\mathcal{H}}_{* *}(\Xi)$ is similar, and is therefore omitted.
First, we show $\overline{\mathcal{H}}_{*}(\Xi)$ is separable. Since $\Xi$ is separable, it has a countable dense subset $X=$ $\left\{\zeta_{1}, \zeta_{2}, \ldots\right\} \subseteq \Xi$. Define $A_{n}$ to be the set of all hypergraphs with $n$ vertices with marks taking values in $X$, i.e.

$$
A_{n}:=\left\{[\bar{H}, i] \in \overline{\mathcal{H}}_{*}(\Xi):|V(\bar{H})|=n, \xi_{\bar{H}}(e, i) \in X \forall(e, i) \in \Psi(\bar{H})\right\}
$$

Since there are finitely many hypergraphs on $n$ vertices and $X$ is countable, we see that $A_{n}$ is countable. Now, define $A=\cup_{n} A_{n}$, which is countable. We claim that $A$ is dense in $\overline{\mathcal{H}}_{*}(\Xi)$. To see this, for $[\bar{H}, i] \in \overline{\mathcal{H}}_{*}(\Xi)$ and $\epsilon>0$ given, pick $(\bar{H}, i) \in[\bar{H}, i]$. Then take $n$ large enough such that $\frac{1}{1+n}<\epsilon$. With $H$ being the underlying unmarked hypergraph associated to $\bar{H}$, we now define a marked rooted hypergraph $\left(\bar{H}^{\prime}, i^{\prime}\right)$ where the underlying unmarked hypergraph $H^{\prime}$ has the property that $\left(H_{\tilde{i}}{ }^{\prime}, i^{\prime}\right)$ is the truncation of $(H, i)$ up to depth $n$ and the mark function $\xi_{\bar{H}^{\prime}}$ is defined as follows. For $(\tilde{e}, \tilde{i}) \in \Psi\left(H^{\prime}\right)$, define $\xi_{\bar{H}^{\prime}}(\tilde{e}, \tilde{i}) \in X$ such that $d_{\Xi}\left(\xi_{H^{\prime}}(\tilde{e}, \tilde{i}), \xi_{\bar{H}}(\tilde{e}, \tilde{i})\right)<1 /(n+1)$. In this way, we have

$$
\bar{d}_{*}\left([\bar{H}, i],\left[\bar{H}^{\prime}, i^{\prime}\right]\right) \leq \frac{1}{1+n}<\epsilon
$$

But, $\left(\bar{H}^{\prime}, i^{\prime}\right)$ is finite, and the edge marks are in $X$; hence, $\left[\bar{H}^{\prime}, i^{\prime}\right] \in A$. Since $\epsilon$ was arbitrary, this shows that $A$ is dense in $\overline{\mathcal{H}}_{*}(\Xi)$. Thus, $\overline{\mathcal{H}}_{*}(\Xi)$ is separable.

Now, we turn to showing that $\overline{\mathcal{H}}_{*}(\Xi)$ is complete. Take a Cauchy sequence $\left[\bar{H}_{n}, i_{n}\right]$ in $\overline{\mathcal{H}}_{*}(\Xi)$ and let $\left(\bar{H}_{n}, i_{n}\right)$ be an arbitrary member of $\left[\bar{H}_{n}, i_{n}\right]$. Without loss of generality, by taking a subsequence if needed, we can assume that for $m>n$ we have

$$
\bar{d}_{*}\left(\left[\bar{H}_{n}, i_{n}\right],\left[\bar{H}_{m}, i_{m}\right]\right)<\frac{1}{1+n} .
$$

This means that, with $H_{k}$ being the underlying unmarked hypergraph associated to $\bar{H}_{k}$ for $k \geq 1$, we have

$$
\begin{equation*}
\left(H_{n}, i_{n}\right) \equiv_{n}\left(H_{m}, i_{m}\right) \quad \forall m>n \tag{90}
\end{equation*}
$$

and

$$
\begin{equation*}
d_{\Xi}\left(\xi_{\bar{H}_{n}}\left(e^{\prime}, i^{\prime}\right), \xi_{\bar{H}_{m}}\left(\phi_{n, m}\left(e^{\prime}\right), \phi_{n, m}\left(i^{\prime}\right)\right)\right)<\frac{1}{1+n} \tag{91}
\end{equation*}
$$

for all $e^{\prime} \in E_{H_{n}}\left(V_{i_{n}, n}^{H_{n}}\right)$ and $i^{\prime} \in e^{\prime}$, where $\phi_{n, m}$ is the depth $n$ isomorphism between $H_{n}$ and $H_{m}$. Note that we can choose $\phi_{n, m}$ for $n>m$ so that

$$
\begin{equation*}
\phi_{n, m}=\phi_{m-1, m} \circ \cdots \circ \phi_{n, n+1} \tag{92}
\end{equation*}
$$

In fact, since the RHS of (92) defines a depth $n$ isomorphism from $\left(H_{n}, i_{n}\right)$ to $\left(H_{m}, i_{m}\right)$, one can define $\phi_{n, m}$ in this way.

In view of (90), we can construct a rooted hypergraph $(H, i)$ so that $(H, i) \equiv_{n}\left(H_{n}, i_{n}\right)$ for all $n$. Further, there are depth $n$ isomorphisms, $\phi_{n}$, from $(H, i)$ to $\left(H_{n}, i_{n}\right)$, which satisfy the consistency condition

$$
\begin{equation*}
\phi_{m}=\phi_{n, m} \circ \phi_{n} \quad \forall m>n . \tag{93}
\end{equation*}
$$

So far we have constructed a rooted hypergraph $(H, i)$ such that $\left[H_{n}, i_{n}\right] \rightarrow[H, i]$. Now, we construct a marked rooted hypergraph $(\bar{H}, i)$, where its underlying unmarked rooted hypergraph is $(H, i)$, and the mark function $\xi_{\bar{H}}: \Psi(H) \rightarrow \Xi$ is defined as follows. Take $\left(e^{\prime}, i^{\prime}\right) \in \Psi(H)$ and choose $d$ such that $e^{\prime} \in E_{H}\left(V_{i, d}^{H}\right)$. We claim that the sequence

$$
\left\{\xi_{\bar{H}_{n}}\left(\phi_{n}\left(e^{\prime}\right), \phi_{n}\left(i^{\prime}\right)\right)\right\}_{n \geq d}
$$

is Cauchy in $\Xi$. Indeed, for $m>n$, using (91), we have

$$
d_{\Xi}\left(\xi_{\bar{H}_{n}}\left(\phi_{n}\left(e^{\prime}\right), \phi_{n}\left(i^{\prime}\right)\right), \xi_{\bar{H}_{m}}\left(\phi_{n, m} \circ \phi_{n}\left(e^{\prime}\right), \phi_{n, m} \circ \phi_{n}\left(i^{\prime}\right)\right)\right)<\frac{1}{1+n} .
$$

Using (93), this means

$$
d_{\Xi}\left(\xi_{\bar{H}_{n}}\left(\phi_{n}\left(e^{\prime}\right), \phi_{n}\left(i^{\prime}\right)\right), \xi_{\bar{H}_{m}}\left(\phi_{m}\left(e^{\prime}\right), \phi_{m}\left(i^{\prime}\right)\right)\right)<\frac{1}{1+n}
$$

which means that the sequence is Cauchy in $\Xi$. Since $\Xi$ is complete, we can define $\xi_{\bar{H}}\left(e^{\prime}, i^{\prime}\right)$ to be the limit of this sequence.

Now, we show that $\left[\bar{H}_{n}, i_{n}\right] \rightarrow[\bar{H}, i]$. For a given $d \in \mathbb{N}$, define

$$
A_{d}:=\left\{\left(e^{\prime}, i^{\prime}\right): e^{\prime} \in E_{H}\left(V_{i, d}^{H}\right), i^{\prime} \in e^{\prime}\right\}
$$

Since $H_{n}$ are locally finite, $H$ is also locally finite, and thus $A_{d}$ is finite. On the other hand, since $\xi_{\bar{H}_{n}}\left(\phi_{n}\left(e^{\prime}\right), \phi_{n}\left(i^{\prime}\right)\right) \rightarrow \xi_{\bar{H}}\left(e^{\prime}, i^{\prime}\right)$ for all $\left(e^{\prime}, i^{\prime}\right) \in A_{d}$, there exists a $N_{d}>d$ such that for all $n>N$, we have

$$
d_{\Xi}\left(\xi_{\bar{H}_{n}}\left(\phi_{n}\left(e^{\prime}\right), \phi_{n}\left(i^{\prime}\right)\right), \xi_{\bar{H}}\left(e^{\prime}, i^{\prime}\right)\right)<\frac{1}{1+d} \quad \forall\left(e^{\prime}, i^{\prime}\right) \in A_{d}
$$

Moreover, since $n>N_{d}>d,\left[H_{n}, i_{n}\right] \equiv_{d}[H, i]$. Hence,

$$
\bar{d}_{*}\left(\left[\bar{H}_{n}, i_{n}\right],[\bar{H}, i]\right)<\frac{1}{1+d} \quad \forall n>N_{d}
$$

Since $d$ was arbitrary, this means that $\left[\bar{H}_{n}, i_{n}\right] \rightarrow[\bar{H}, i]$ and $\overline{\mathcal{H}}_{*}(\Xi)$ is complete.
As was mentioned earlier, by setting $\Xi=\{\emptyset\}$, we conclude that $\mathcal{H}_{*}$ and $\mathcal{H}_{* *}$ are Polish spaces. This is explicitly stated below as a corollary.

Corollary 3. The spaces $\mathcal{H}_{*}$ and $\mathcal{H}_{* *}$ are Polish spaces.

## C Some properties of measures on $\mathcal{H}_{*}$

Proof of Lemma 1. Part (i): We have

$$
\vec{\mu}(\tilde{A})=\int \mathbb{1}_{\tilde{A}} d \vec{\mu}=\int \partial \mathbb{1}_{\tilde{A}} d \mu .
$$

But,

$$
\partial \mathbb{1}_{\tilde{A}}(H, i)=\sum_{e \ni i} \mathbb{1}_{\tilde{A}}(H, e, i)=\sum_{e \ni i} \mathbb{1}_{A}(H, i)=\operatorname{deg}_{H}(i) \mathbb{1}_{A}(H, i) .
$$

Hence

$$
\vec{\mu}(\tilde{A})=\int \mathbb{1}_{A}(H, i) \operatorname{deg}_{H}(i) d \mu=\int \operatorname{deg}_{H}(i) d \mu=\int d \vec{\mu}=\vec{\mu}\left(\mathcal{H}_{* *}\right)
$$

which shows that $\tilde{A}$ happens $\vec{\mu}$ almost everywhere and the proof is complete.
Part (ii): Note that

$$
\vec{\mu}(B)=\int_{\mathcal{H}_{* *}} \mathbb{1}[[H, e, i] \in B] d \vec{\mu}([H, e, i])=\int_{\mathcal{H}_{*}} \sum_{e \ni i} \mathbb{1}[[H, e, i] \in B] d \mu([H, i]) .
$$

On the other hand, $\vec{\mu}(B)=\vec{\mu}\left(\mathcal{H}_{* *}\right)=\operatorname{deg}(\mu)=\int \operatorname{deg}_{H}(i) d \mu([H, i])$. Moreover, for all $[H, i] \in \mathcal{H}_{*}$, $\sum_{e \ni i} \mathbb{1}[[H, e, i] \in B] \leq \operatorname{deg}_{H}(i)$. Consequently, it must be the case that for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}$, $\sum_{e \ni i} \mathbb{1}[[H, e, i] \in B]=\operatorname{deg}_{H}(i)$, or equivalently $[H, e, i] \in B$ for all $e \ni i$.

Proof of Lemma 2. If we define

$$
A:=\left\{[H, i] \in \mathcal{H}_{*}: f_{k}([H, i]) \rightarrow f_{0}([H, i])\right\}
$$

and

$$
\tilde{A}:=\left\{[H, e, i] \in \mathcal{H}_{* *}: \tilde{f}_{k}([H, e, i]) \rightarrow \tilde{f}_{0}([H, e, i])\right\}
$$

then we have

$$
\tilde{A}=\left\{[H, e, i] \in \mathcal{H}_{* *}:[H, i] \in A\right\}
$$

Then the proof is an immediate conseqeunce of Lemma 1.
Proof of Lemma 3. Define $B:=\left\{[H, e, i] \in \mathcal{H}_{* *}: f_{k}(H, e, i) \rightarrow f_{0}(H, e, i)\right\}$. As $\vec{\mu}(B)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$, from part (ii) of Lemma 1, for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}$, for all $e \ni i, f_{k}(H, e, i) \rightarrow f_{0}(H, e, i)$. This in particular implies that for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}, \partial f_{k}(H, i) \rightarrow \partial f_{0}(H, i)$.

## D Proof of Lemma 4

We first prove that if the condition mentioned in Lemma 4 is satisfied, then $\mu_{n} \Rightarrow \mu$. Fix $\epsilon>0$. Let $f: \mathcal{H}_{*} \rightarrow \mathbb{R}$ be a uniformly continuous and bounded function. There is some $\delta>0$ such that $\left|f([H, i])-f\left(\left[H^{\prime}, i^{\prime}\right]\right)\right|<\epsilon$ when $d_{\mathcal{H}_{*}}\left([H, i],\left[H^{\prime}, i^{\prime}\right]\right)<\delta$. Now choose $d$ such that $1 /(1+d)<\delta$. For all rooted trees $[H, i] \in \mathcal{T}_{*},[H, i] \in A_{(H, i)_{d}}$. Hence, one can find countably many rooted trees $\left\{T_{j}, i_{j}\right\}_{j=1}^{\infty}$ with depth at most $d$ such that $A_{\left(T_{j}, i_{j}\right)}$ partitions $\mathcal{T}_{*}$; hence, one can find finitely many $\left(T_{j}, i_{j}\right), 1 \leq j \leq m$ such that $\sum_{j=1}^{m} \mu\left(A_{\left(T_{j}, i_{j}\right)}\right) \geq 1-\epsilon$. If $\mathcal{A}$ denotes $\cup_{j=1}^{m} A_{\left(T_{j}, i_{j}\right)}$, then we have

$$
\begin{aligned}
&\left|\int f d \mu-\sum_{j=1}^{m} f\left(\left[T_{j}, i_{j}\right]\right) \mu\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| \leq \sum_{j=1}^{m}\left|\int_{A_{\left(T_{j}, i_{j}\right)}} f d \mu-f\left(\left[T_{j}, i_{j}\right]\right) \mu\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| \\
&+\|f\|_{\infty} \mu\left(\mathcal{A}^{c}\right) \\
& \leq \epsilon\left(1+\|f\|_{\infty}\right),
\end{aligned}
$$

where the last inequality uses the facts that $\mu\left(\mathcal{A}^{c}\right)<\epsilon$ and $\left|f([H, i])-f\left(\left[T_{j}, i_{j}\right]\right)\right|<\epsilon$ for $[H, i] \in A_{\left(T_{j}, i_{j}\right)}$, $1 \leq j \leq m$ since $1 /(1+d)<\epsilon$. Similarly, we have

$$
\begin{aligned}
\left|\int f d \mu_{n}-\sum_{j=1}^{m} f\left(\left[T_{j}, i_{j}\right]\right) \mu\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| \leq & \left|f f d \mu_{n}-\sum_{j=1}^{m} f\left(\left[T_{j}, i_{j}\right]\right) \mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| \\
& +\sum_{j=1}^{m}\left|f\left(T_{j}, i_{j}\right) \| \mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right)-\mu\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| \\
\leq & \|f\|_{\infty}\left(1-\sum_{j=1}^{m} \mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right)\right)+\epsilon \\
& +\|f\|_{\infty} \sum_{j=1}^{m}\left|\mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right)-\mu\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| .
\end{aligned}
$$

Combining these two, we have

$$
\begin{gathered}
\left|\int f d \mu_{n}-\int f d \mu\right| \leq\|f\|_{\infty}\left(1-\sum_{j=1}^{m} \mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right)\right)+\|f\|_{\infty}\left|\sum_{j=1}^{m} \mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right)-\mu\left(A_{\left(T_{j}, i_{j}\right)}\right)\right| \\
+\epsilon\left(2+\|f\|_{\infty}\right)
\end{gathered}
$$

Now, as $n$ goes to infinity, $\mu_{n}\left(A_{\left(T_{j}, i_{j}\right)}\right) \rightarrow \mu\left(A_{\left(T_{j}, i_{j}\right)}\right)$ by assumption. Also, $\mu\left(\mathcal{A}^{c}\right)<\epsilon$. Thus, we have

$$
\limsup _{n \rightarrow \infty}\left|\int f d \mu_{n}-\int f d \mu\right| \leq \epsilon\left(2+2\|f\|_{\infty}\right)
$$

Since $\|f\|_{\infty}<\infty$ and this is true for any $\epsilon>0$, we get $\int f d \mu_{n} \rightarrow \int f d \mu$; hence, $\mu_{n} \Rightarrow \mu$.
For the converse, note that $\mathbb{1}_{A_{(T, j)}}$ is a continuous function since $\mathbb{1}_{A_{(T, j)}}([H, i])=\mathbb{1}_{A_{(T, i)}}\left(\left[H^{\prime}, j^{\prime}\right]\right)$ for $d_{\mathcal{H}_{*}}\left([H, i],\left[H^{\prime}, i^{\prime}\right]\right)<1 /(1+d)$.

## E Some properties of Unimodular Measures

First, we prove Proposition 3. Our proof depends on the following lemma:
Lemma 19. Assume $\tau: \mathcal{H}_{* *} \rightarrow \mathbb{R}$ is a measurable function and $\mu \in \mathcal{P}\left(\mathcal{H}_{*}\right)$ is a unimodular measure such that $\tau=1, \vec{\mu}$-almost everywhere. Then, we have

1. With $\tau_{1}(H, e, i):=\mathbb{1}\left[\tau\left(H, e^{\prime}, i\right)=1, \forall e^{\prime} \ni i\right]$, it holds that $\tau_{1}=1 \vec{\mu}$-almost everywhere.
2. With $\tau_{2}(H, e, i):=\mathbb{1}\left[\tau\left(H, e, i^{\prime}\right)=1, \forall i^{\prime} \in e\right]$, it holds that $\tau_{2}=1 \vec{\mu}$-almost everywhere.

Proof. In order to prove the first part, note that from Lemma 1 part (ii), we have that for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}, \tau\left(H, e^{\prime}, i\right)=1$ for all $e^{\prime} \ni i$. Now, part $(i)$ of Lemma 1 implies that for $\vec{\mu}$-almost all $[H, e, i] \in \mathcal{H}_{* *}, \tau(H, e, i)=1$ for all $e^{\prime} \ni i$, which is precisely what we need to prove.

For the second part, since $\mu$ is unimodular, we have

$$
\int \mathbb{1}_{\tau=1} d \vec{\mu}=\int \nabla \mathbb{1}_{\tau=1} d \vec{\mu}=\int \frac{1}{|e|} \sum_{i^{\prime} \in e} \mathbb{1}_{\tau\left(H, e, i^{\prime}\right)=1} d \vec{\mu}(H, e, i)
$$

But since $\mathbb{1}_{\tau=1}=1$ holds $\vec{\mu}$-almost everywhere and $0 \leq \frac{1}{|e|} \sum_{i^{\prime} \in e} \mathbb{1}_{\tau\left(H, e, i^{\prime}\right)=1} \leq 1$, we conclude that

$$
\frac{1}{|e|} \sum_{i^{\prime} \in e} \mathbb{1}_{\tau\left(H, e, i^{\prime}\right)=1}=1, \quad \vec{\mu} \text {-a.e.. }
$$

Since the summands are either zero or one, this means that $\tau\left(H, e, i^{\prime}\right)=1$ for all $i^{\prime} \in e, \vec{\mu}$-almost everywhere, which is what we wanted to show.

Proof of Proposition 3. Define

$$
\begin{equation*}
A_{k}:=\left\{[H, e, i] \in \mathcal{H}_{* *}: \tau\left(H, e^{\prime}, i^{\prime}\right)=1 \forall i^{\prime} \in V(H): d_{H}\left(i, i^{\prime}\right) \leq k, \forall e^{\prime} \ni i^{\prime}\right\} \tag{94}
\end{equation*}
$$

Note that $A_{0}=\{[H, e, i]: \tau(H, e, i)=1\}$, for which it is known from the assumption that $\vec{\mu}\left(A_{0}\right)=$ $\vec{\mu}\left(\mathcal{H}_{* *}\right)$. Now, we want to show that $\vec{\mu}\left(A_{k}\right)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$ for all $k \geq 0$. We will do this by induction on $k$. Assume that $\vec{\mu}\left(A_{k}\right)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$. Hence, with $\phi_{k}(H, e, i):=\mathbb{1}\left[(H, e, i) \in A_{k}\right]$, we have $\phi_{k}=1 \vec{\mu}$-almost everywhere. Now, we will use Lemma 19 to prove that $\phi_{k+1}=1 \vec{\mu}$-almost everywhere.

To do so, using part 2 of Lemma 19, if we define

$$
B_{1}^{k}:=\left\{[H, e, i]: \forall i^{\prime} \in e, \phi_{k}\left(H, e, i^{\prime}\right)=1\right\}
$$

then we know $\vec{\mu}\left(B_{1}^{k}\right)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$. Then, applying part 1 of Lemma 19 for the function $\mathbb{1}_{B_{1}^{k}}$, we get that

$$
B_{2}^{k}:=\left\{[H, e, i]: \forall e^{\prime} \ni i,\left[H, e^{\prime}, i\right] \in B_{1}^{k}\right\}
$$

has the property that $\vec{\mu}\left(B_{2}^{k}\right)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$. On the other hand,

$$
B_{2}^{k}=\left\{[H, e, i]: \forall e^{\prime} \ni i, \forall i^{\prime} \in e^{\prime},\left[H, e^{\prime}, i^{\prime}\right] \in A_{k}\right\}=A_{k+1}
$$

Hence, we have proved that $\vec{\mu}\left(A_{k+1}\right)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$, which is the inductive step.
Thus, $\vec{\mu}\left(\cap_{k \in \mathbb{N}} A_{k}\right)=\vec{\mu}\left(\mathcal{H}_{* *}\right)$. Using the fact that the vertex set is countable, and that the hypergraphs corresponding to the elements of $\mathcal{H}_{*}$ are connected, the property $\tau$ holds for all the directed edges $\vec{\mu}$-almost everywhere, in $A:=\cap_{k \in \mathbb{N}} A_{k}$. Hence, the proof is complete.

Now, we prove that the local weak limit of finite marked hypergraphs is a unimodular probability distribution on $\overline{\mathcal{H}}_{*}(\Xi)$. By setting $\Xi=\{\emptyset\}$ in the following proposition, we can conclude that the local weak limit of finite simple hypergraphs is unimodular, as claimed in Section 2.8.

Proposition 20. Assume $\left\{\bar{H}_{n}\right\}$ is a sequence of finite marked hypergraphs, with $\xi_{\bar{H}_{n}}=\xi_{n}$ the associated edge mark functions, taking values in some metric space $\Xi$. Now, if

$$
\bar{\mu}_{n}:=u_{\bar{H}_{n}}=\frac{1}{\left|V\left(\bar{H}_{n}\right)\right|} \sum_{i \in V\left(\bar{H}_{n}\right)} \delta_{\left[\left(\bar{H}_{n}, i\right), i\right]}
$$

then $\bar{\mu}_{n} \in \mathcal{P}\left(\overline{\mathcal{H}}_{*}(\Xi)\right)$ is unimodular for each $n$. Moreover, if $\bar{\mu}_{n}$ converge weakly to some limit $\bar{\mu} \in$ $\mathcal{P}\left(\overline{\mathcal{H}}_{*}(\Xi)\right)$ such that $\operatorname{deg}(\bar{\mu})<\infty, \bar{\mu}$ is also unimodular.

Proof. First, we show that $\bar{\mu}_{n}$ is unimodular for each $n$. For this, take a Borel function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow$
$[0, \infty)$ and note that

$$
\begin{aligned}
\int f(\bar{H}, e, i) d \vec{\mu}_{n} & =\int \partial f(\bar{H}, i) d \bar{\mu}_{n} \\
& =\frac{1}{\left|V\left(\bar{H}_{n}\right)\right|} \sum_{i \in V\left(\bar{H}_{n}\right)} \partial f\left(\bar{H}_{n}, i\right) \\
& =\frac{1}{\left|V\left(\bar{H}_{n}\right)\right|} \sum_{i \in V\left(\bar{H}_{n}\right)} \sum_{e \ni i} f\left(\bar{H}_{n}, e, i\right) \\
& =\frac{1}{\left|V\left(\bar{H}_{n}\right)\right|} \sum_{e \in E\left(\bar{H}_{n}\right)} \sum_{i \in e} f\left(\bar{H}_{n}, e, i\right) \\
& =\frac{1}{\left|V\left(\bar{H}_{n}\right)\right|} \sum_{e \in E\left(\bar{H}_{n}\right)} \sum_{i \in e} \nabla f\left(\bar{H}_{n}, e, i\right) \\
& =\int \nabla f(\bar{H}, e, i) d \vec{\mu}_{n}
\end{aligned}
$$

Since this holds for all nonnegative Borel functions $f, \bar{\mu}_{n}$ is unimodular, by definition.
Now, for each $n$, define measures $\vec{\mu}_{n}^{(1)}$ and $\vec{\mu}_{n}^{(2)}$ on $\overline{\mathcal{H}}_{* *}(\Xi)$ so that for any Borel function $f$ : $\overline{\mathcal{H}}_{* *}(\Xi) \rightarrow[0, \infty)$, we have

$$
\begin{equation*}
\int f d \overrightarrow{\bar{\mu}}_{n}^{(1)}:=\int \sum_{e \ni i} f(\bar{H}, e, i) d \bar{\mu}_{n}([\bar{H}, i]), \tag{95}
\end{equation*}
$$

and

$$
\begin{equation*}
\int f d \vec{\mu}_{n}^{(2)}:=\int \sum_{e \ni i} \frac{1}{|e|} \sum_{j \in e} f(\bar{H}, e, j) d \bar{\mu}_{n}([\bar{H}, i]) \tag{96}
\end{equation*}
$$

We also define $\vec{\mu}^{(1)}$ and $\vec{\mu}^{(2)}$ for $\bar{\mu}$ in a similar fashion. Note that the the RHS of (95) is $\int \partial f d \bar{\mu}_{n}=$ $\int f d \vec{\mu}_{n}$. Therefore, $\vec{\mu}_{n}^{(1)}=\vec{\mu}_{n}$. Also, the RHS of (96) is $\int \partial \nabla f d \bar{\mu}_{n}=\int \nabla f d \vec{\mu}_{n}$. Hence, $\int f d \vec{\mu}_{n}^{(2)}=$ $\int \nabla f d \vec{\mu}_{n}$. Since we have shown that $\bar{\mu}_{n}$ is unimodular, this implies that $\vec{\mu}_{n}^{(1)}=\vec{\mu}_{n}^{(2)}$.

Now, we claim that $\overrightarrow{\vec{\mu}}^{(1)}=\overrightarrow{\vec{\mu}}^{(2)}$. To show this, take a bounded continuous function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow \mathbb{R}$. For $k>0$, define

$$
f_{k}(\bar{H}, e, i):= \begin{cases}f(\bar{H}, e, i) & \text { if } \operatorname{deg}_{H}(i) \leq k \\ 0 & \text { otherwise }\end{cases}
$$

It is easy to see that $f_{k}$ is continuous for each $k$, as $f$ is continuous. Moreover, $\partial f_{k}$ and $\partial \nabla f_{k}$ are bounded for each $k$. This, together with $\bar{\mu}_{n} \Rightarrow \bar{\mu}$, implies that

$$
\int f_{k} d \vec{\mu}_{n}^{(1)}=\int \partial f_{k} d \bar{\mu}_{n} \rightarrow \int \partial f_{k} d \bar{\mu}=\int f_{k} d \overrightarrow{\bar{\mu}}^{(1)}
$$

and

$$
\int f_{k} d \overrightarrow{\bar{\mu}}_{n}^{(2)}=\int \partial \nabla f_{k} d \bar{\mu}_{n} \rightarrow \int \partial \nabla f_{k} d \bar{\mu}=\int f_{k} d \overrightarrow{\vec{\mu}}^{(2)}
$$

This, together with the fact that $\vec{\mu}_{n}^{(1)}=\vec{\mu}_{n}^{(2)}$, implies that for all $k$

$$
\int f_{k} d \vec{\mu}^{(1)}=\int f_{k} d \overrightarrow{\bar{\mu}}^{(2)}
$$

Note that as all hypergraphs are locally finite, $f_{k} \rightarrow f$ pointwise. Thus, sending $k$ to infinity and using the dominated convergence theorem, we have that for any bounded continuous function $f: \overline{\mathcal{H}}_{* *}(\Xi) \rightarrow$ $[0, \infty)$,

$$
\int f d \overrightarrow{\vec{\mu}}^{(1)}=\int f d \overrightarrow{\vec{\mu}}^{(2)}
$$

Since $f$ can be an arbitrary bounded continuous function, we have $\overrightarrow{\vec{\mu}}^{(1)}=\overrightarrow{\vec{\mu}}^{(2)}$. But the definition of $\overrightarrow{\vec{\mu}}^{(1)}$ and $\overrightarrow{\vec{\mu}}^{(2)}$ then implies that for any nonnegative Borel function $f$ we have $\int f d \overrightarrow{\vec{\mu}}=\int \nabla f d \overrightarrow{\vec{\mu}}$, which means that $\bar{\mu}$ is unimodular.

## F Proof of unimodularity of UGWT $(P)$

Here we show that $\operatorname{UGWT}(P)$ is unimodular. Before that we prove the following lemma, which is useful in calculating $\int f d \mu$ when $\mu=\operatorname{UGWT}(P)$.

Lemma 20. Let $P \in \mathcal{P}(\Lambda)$, and $\Gamma$ a random variable with law $P$, where $\mathbb{E}[\Gamma(k)]<\infty$ for all $k \geq 2$. Moreover, let $\mu=U G W T(P)$. Then, for any Borel function $f: \mathcal{H}_{* *} \rightarrow[0, \infty)$, we have

$$
\int f d \vec{\mu}=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \sum_{\gamma \in \Lambda} \hat{P}_{k}(\gamma) \mathbb{E}\left[f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma+\mathbb{E}_{k}\right]
$$

where the expectation is with respect to the random rooted hypertree of Definition 28. Here, $\mathbb{e}_{k} \in \Lambda$ is such that $\mathbb{e}_{k}(k)=1$ and $\mathbb{e}_{k}(j)=0$ for $j \neq k$.

Proof. Due to the definition of $\vec{\mu}$, we have

$$
\int f d \vec{\mu}=\int \partial f d \mu=\mathbb{E}\left[\sum_{k=2}^{h\left(\Gamma_{\emptyset}\right)} \sum_{i=1}^{\Gamma_{\emptyset}(k)} f(T,(k, i), \emptyset)\right]
$$

Since $\Gamma_{\emptyset}$ has distribution $P$, we have

$$
\mathbb{E}\left[\sum_{k=1}^{h\left(\Gamma_{\emptyset}\right)} \sum_{i=1}^{\Gamma_{\emptyset}(k)} f(T,(k, i), \emptyset)\right]=\sum_{\gamma \in \Lambda} P(\gamma) \mathbb{E}\left[\sum_{k=2}^{h(\gamma)} \sum_{i=1}^{\gamma(k)} f(T,(k, i), \emptyset) \mid \Gamma_{\emptyset}=\gamma\right] .
$$

Now, due to symmetry, conditioned on $\Gamma_{\emptyset}=\gamma$, for a given $k \leq h(\gamma)$, all $f(T,(k, i), \emptyset)$ for $1 \leq i \leq \gamma(k)$ have the same distribution, hence

$$
\begin{aligned}
\int f d \vec{\mu} & =\sum_{\gamma \in \Lambda} P(\gamma) \sum_{k=2}^{h(\gamma)} \gamma(k) \mathbb{E}\left[f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma\right] \\
& =\sum_{\gamma \in \Lambda} P(\gamma) \sum_{k=2}^{\infty} \gamma(k) \mathbb{E}\left[f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma\right]
\end{aligned}
$$

where the second equality uses the fact that $\gamma(k)=0$ for $k>h(\gamma)$. Now, since all the terms are nonnegative, employing Tonelli's theorem to switch the order of integrals we have

$$
\int f d \vec{\mu}=\sum_{k=2}^{\infty} \sum_{\gamma \in \Lambda: \gamma(k)>0} P(\gamma) \gamma(k) \mathbb{E}\left[f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma\right] .
$$

Using the definition of $\hat{P}_{k}, P(\gamma) \gamma(k)$ is equal to $\mathbb{E}[\Gamma(k)] \hat{P}_{k}\left(\gamma-\mathbb{C}_{k}\right)$ for $\gamma \in \Lambda$ such that $\gamma(k)>0$, where $\mathbb{e}_{k} \in \Lambda$ is such that $\mathbb{e}_{k}(k)=1$ and $\mathbb{e}_{k}(j)=0$ for $j \neq k$. Hence, we have

$$
\begin{aligned}
\int f d \vec{\mu} & =\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \sum_{\gamma \in \Lambda: \gamma(k)>0} \hat{P}_{k}\left(\gamma-\mathbb{E}_{k}\right) \mathbb{E}\left[f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma\right] \\
& =\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \sum_{\gamma \in \Lambda} \hat{P}_{k}(\gamma) \mathbb{E}\left[f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma+\mathbb{C}_{k}\right]
\end{aligned}
$$

We can interpret the last expression above as follows. In computing $\int f d \vec{\mu}$, when we consider $f([T, e, \emptyset])$, the edge $e$ attached to the root $\emptyset$ of the tree $T$ is of size $k$ with probability $\Gamma(k)$. This explains the outer summation. Since the contribution to the integral is the same whichever edge of size $k$ connected to the root is picked, suppose the edge picked is the edge $(k, 1)$. Then the type of the rest of the edges connected to the root is given by $\hat{P}_{k}$, and so $\Gamma_{\emptyset}$ will equal $\gamma+\mathbb{e}_{k}$ with probability $\hat{P}_{k}(\gamma)$. This explains the inner summation.

Now we are ready to prove the unimodularity of $\operatorname{UGWT}(P)$ :
Proof of Proposition 4. We need to prove that for a nonnegative measurable function $f: \mathcal{H}_{* *} \rightarrow[0, \infty)$ we have $\int f d \vec{\mu}=\int \nabla f d \vec{\mu}$. Using Lemma 20, we have

$$
\begin{aligned}
\int \nabla f d \vec{\mu} & \stackrel{(a)}{=} \sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \mathbb{E}\left[\nabla f\left(\tilde{T}_{k},(k, 1), \emptyset\right)\right] \\
& =\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \frac{1}{k}\left(\mathbb{E}\left[f\left(\tilde{T}_{k},(k, 1), \emptyset\right)\right]+\sum_{i=1}^{k-1} \mathbb{E}\left[f\left(\tilde{T}_{k},(k, 1),(k, 1, i)\right)\right]\right)
\end{aligned}
$$

Here, for $k \geq 2, \tilde{T}_{k}$ denotes a tree with root $\emptyset$ that has an edge $(k, 1)$ of size $k$ connected to the root, with the type of the other edges connected to the root being $\hat{P}_{k}$, and with the subtrees at the other vertices of all the edges (including the edge $(k, 1)$ ) generated according to the rules of UGWT $(P)$. Step (a) is justified because, for each $k \geq 2$, we have

$$
\mathbb{E}\left[\nabla f\left(\tilde{T}_{k},(k, 1), \emptyset\right)\right]=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \sum_{\gamma \in \Lambda} \hat{P}_{k}(\gamma) \mathbb{E}\left[\nabla f(T,(k, 1), \emptyset) \mid \Gamma_{\emptyset}=\gamma+\mathbb{E}_{k}\right]
$$

Now, because of the symmetry in the construction of $\tilde{T}_{k}$, for each $i=1, \ldots, k-1,\left(\tilde{T}_{k},(k, 1),(k, 1, i)\right)$ has the same distribution as $\left(\tilde{T}_{k},(k, 1), \emptyset\right)$. Hence, $\mathbb{E}\left[f\left(\tilde{T}_{k},(k, 1),(k, 1, i)\right)\right]=\mathbb{E}\left[f\left(\tilde{T}_{k},(k, 1), \emptyset\right)\right]$. Substituting and simplifying we get

$$
\int \nabla f d \vec{\mu}=\sum_{k=2}^{\infty} \mathbb{E}[\Gamma(k)] \mathbb{E}[f(\tilde{T},(k, 1), \emptyset)]=\int f d \vec{\mu}
$$

where the last equality again uses Lemma 20. This completes the proof of the unimodularity of $\mu$.

## G Configuration model on hypergraphs and their local weak limit: Proof of Theorem 4

In this section, we prove Theorem 4. We prove this in two steps: first, in Section G.2, we prove that $\mathbb{E}\left[u_{H_{n}^{e}}\right]$ converges weakly to $\operatorname{UGWT}(P)$, where the expectation in $\mathbb{E}\left[u_{H_{n}^{e}}\right]$ is taken with respect
to the randomness in the construction of $H_{n}^{e}$. Later, in Section G.3, we conclude the almost sure convergence by a concentration argument. See [Bor14] for an argument on the local weak convergence of the configuration model in the graph regime.

Throughout this section, we employ the vertex and edge indexing notations $\mathbb{N}_{\text {vertex }}$ and $\mathbb{N}_{\text {edge }}$ defined in Section 2.9. Recall that for $a \in \mathbb{N}_{\text {vertex }}$ and integers $s \geq 2, e \geq 1$ and $1 \leq r \leq s-1,(a, s, e, r)$ is the element in $\mathbb{N}_{\text {vertex }}$ obtained from $a$ by concatenating $s, e, r$ at the end of the string representing $a$. For an edge $e \in \mathbb{N}_{\text {edge }}$ and an integer $r,(e, r)$ is defined similarly. Furthermore, $\mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$ in this section denotes the set of hypertrees with vertex set and edge set being subsets of $\mathbb{N}_{\text {vertex }}$ and $\mathbb{N}_{\text {edge }}$, respectively. Such a hypertree is treated to be rooted at $\emptyset$, unless otherwise stated. Moreover, for a sequence of types $\left\{\gamma_{a}\right\}_{a \in \mathbb{N}_{\text {vertex }}}, \mathcal{T}\left(\left\{\gamma_{a}\right\}_{a \in \mathbb{N}_{\text {vertex }}}\right)$ denotes the tree in $\mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$ in which the type of each node $a \in \mathbb{N}_{\text {vertex }}$ in the hypertree below that node is equal to $\gamma_{a}$ (recall Figure 4 from Section 2.9 as an example).

Before going through the proof, we need to define a procedure called the "exploration process" in Section G. 1 below.

## G. 1 Exploration process

Assume that a random hypergraph $H_{n}$ on the vertex set $\{1, \ldots, n\}$ is obtained from a given type sequence $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ satisfying (82a) and (82b). Note that, following our discussion in Section $9.1, H_{n}$ is identified by a set of random matchings $\sigma_{2}, \ldots, \sigma_{n}$. Here, we introduce a procedure that choses a node $v_{0}$ uniformly at random in $\{1, \ldots, n\}$ and explores the local neighborhood of that node in a breadth-first manner. This process at each step produces a hypertree in $\mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$, which turns out to be locally isomorphic to the neighborhood of $v_{0}$ given that the local neighborhood of $v_{0}$ in $H_{n}$ is tree-like. A similar process in the graph regime is introduced in [Bor14].

Formally speaking, the exploration process starts at time $t=0$ with choosing a vertex $v_{0} \in$ $\{1, \ldots, n\}$ uniformly at random as the root. At each time $t \geq 1$, we have a node indexing function $\phi_{t}: \mathbb{N}_{\text {vertex }} \rightarrow\{1, \ldots, n\} \cup\{\times\}$. To begin with, we define $\phi_{1}(\emptyset)=v_{0}$ and $\phi_{1}(\mathbf{i})=\times$ for $\emptyset \neq \mathbf{i} \in \mathbb{N}_{\text {vertex }}$. Also, at each time step $t$, we partition $\Delta^{(n)}$ into three sets: an "active set" $A_{t}$, a "connected set" $C_{t}$, and an "undiscovered set" $U_{t}$. These are initialized by setting $A_{1}=\Delta_{v_{0}}^{(n)}, C_{1}=\emptyset$ and $U_{1}=\Delta^{(n)} \backslash A_{1}$. Moreover, at time $t, N_{t} \subset\{1, \ldots, n\}$ contains the explored nodes at time $t$, and is initialized as $N_{1}=\left\{v_{0}\right\}$.

At time $t$, given the sets $A_{t}, U_{t}, C_{t}$ and $\phi_{t}$, we first form the set $W_{t}:=\left\{(\mathbf{i}, k, j): \phi_{t}(\mathbf{i}) \neq \times, e_{\phi_{t}(\mathbf{i}), j}^{k} \in\right.$ $\left.A_{t}\right\}$. If $W_{t}$ is nonempty, we define $\left(\mathbf{i}_{t}, k_{t}, j_{t}\right)$ to be the lexicographically smallest element in $W_{t}$ and let $e_{t}:=e_{\phi_{t}\left(\mathbf{i}_{t}\right), j_{t}}^{k_{t}}$. In fact, $e_{t}$ is the partial edge chosen at time $t$ to be matched with other partial edges to form an edge. Now, define

$$
e_{t, j}:=\sigma_{k_{t}}^{(j-1)}\left(e_{t}\right) \quad 1 \leq j \leq k_{t}-1
$$

which are the $k_{t}-1$ other partial edges matched with $e_{t}$. Also, for $1 \leq j \leq k_{t}-1$, let $u_{t, j}:=\nu\left(e_{t, j}\right)$ be the node associated to $e_{t, j}$. Moreover, we update the sets $C_{t}, A_{t}, U_{t}$ and $N_{t}$ as follows:

$$
\begin{align*}
& C_{t+1}=C_{t} \cup\left\{e_{t}, e_{t, 1}, \ldots, e_{t, k_{t}-1}\right\},  \tag{97a}\\
& A_{t+1}=A_{t} \backslash\left\{e_{t}, e_{t, 1}, \ldots, e_{t, k_{t}-1}\right\} \bigcup_{j=1}^{k_{t}-1}\left(\Delta_{u_{t, j}}^{(n)} \cap U_{t}\right),  \tag{97b}\\
& U_{t+1}=U_{t} \backslash \bigcup_{j=1}^{k_{t}-1} \Delta_{u_{t, j}}^{(n)},  \tag{97c}\\
& N_{t+1}=N_{t} \cup\left\{u_{t, 1}, \ldots, u_{t, k_{t}-1}\right\} . \tag{97d}
\end{align*}
$$

In order to update $\phi_{t}$, define $\tilde{j}_{t}$ to be the minimum $j$ such that

$$
\phi_{t}\left(\left(\mathbf{i}_{t}, k_{t}, j, 1\right)\right)=\times .
$$

With this, set $\phi_{t+1}$ to be equal to $\phi_{t}$ except for the following values:

$$
\phi_{t+1}\left(\left(\mathbf{i}_{t}, k_{t}, \tilde{j}_{t}, l\right)\right)=u_{t, l} \quad 1 \leq l \leq k_{t}-1
$$

This in particular means that the set of nodes in $\{1, \ldots, n\}$ that appear in the range of $\phi_{t+1}$ is precisely $N_{t+1}$.

At each time $t$, we define a rooted hypertree formed by the exploration process, which we denote by $R_{t}$, which is a member of $\mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$. $R_{t}$ is identified through the mapping $\phi_{t}$, i.e. its vertex set is $\left\{\mathbf{i} \in \mathbb{N}_{\text {vertex }}: \phi_{t}(\mathbf{i}) \neq \times\right\}$, and its edge set is $\left\{\left(\mathbf{i}_{r}, k_{r}, \tilde{j}_{r}\right), 1 \leq r \leq t\right\}$. This process continues until $A_{t}=\emptyset$, which results in exploring the connected component of $v_{0}$. Note that since the permutations determining the configuration process are random, the exploration process is in fact a random process. Let $\mathcal{F}_{t}$ be the sigma field generated by all the random variables defined above up to time $t$. Let $\tau$ be the stopping time corresponding to $A_{\tau}=\emptyset$. For the sake of simplicity, for $t>\tau$, we define $R_{t}=R_{\tau}$.

## G. 2 Convergence of expectation

In this section, we provide the proof of the convergence of $\mathbb{E}\left[u_{H_{n}^{e}}\right]$. This is done in two parts. Loosely speaking, we first show that, for any integer $d \geq 1$, with high probability, $H_{n}$ rooted at a vertex chosen uniformly at random up to depth $d$ is simple. Then, we prove that the multihypergraph $H_{n}$ obtained from the configuration model is locally and asymptotically similar to the Galton-Watson process. These two together complete the proof of Theorem 4. The former is proved in Proposition 21 below and the latter in Proposition 22.

Let $\bar{\mu}_{n}^{e}$ and $\mu$ denote $\mathbb{E}\left[u_{H_{n}^{e}}\right]$ and $\operatorname{UGWT}(P)$, respectively. In order to show that $\bar{\mu}_{n}^{e} \Rightarrow \mu$, using Lemma 4, it suffices to prove that for any $d \geq 1$, and with $(T, o)$ being a rooted hypertree of depth at $\operatorname{most} d, \bar{\mu}_{n}^{e}\left(A_{(T, o)}\right) \rightarrow \mu\left(A_{(T, o)}\right)$, where we recall that $A_{(T, o)}=\left\{[H, j] \in \mathcal{H}_{*}:(H, j)_{d} \equiv(T, o)\right\}$. Note that, with $v_{0}$ being chosen uniformly at random in $\{1, \ldots, n\}$, we have

$$
\begin{align*}
\bar{\mu}_{n}^{e}\left(A_{(T, o)}\right) & =\mathbb{E}\left[\frac{1}{n} \sum_{i=1}^{n} \delta_{\left[H_{n}^{e}(i), i\right]}\left(A_{(T, o)}\right)\right] \\
& =\mathbb{E}\left[\frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\left[H_{n}^{e}(i), i\right] \in A_{(T, o)}}\right]  \tag{98}\\
& =\frac{1}{n} \sum_{i=1}^{n} \mathbb{P}\left(\left(H_{n}^{e}, i\right)_{d} \equiv(T, o)\right) \\
& =\mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o)\right) .
\end{align*}
$$

Thus, motivated by the above discussion, we need to show that for all $d \geq 1$ and $(T, o)$ with depth at most $d, \mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o)\right) \rightarrow \mu\left(A_{(T, o)}\right)$. We prove this in two steps. First, we prove in Proposition 21 that the probability of $\left(H_{n}, v_{0}\right)_{d}$ being a simple hypertree goes to one as $n$ goes to infinity. Subsequently, in Proposition 22, we show that $\mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, i)\right)$ converges to $\mu\left(A_{(T, i)}\right)$.

In the following statement, $\gamma^{(n)}=\left(\gamma_{1}^{(n)}, \ldots, \gamma_{n}^{(n)}\right)$ is a fixed type sequence and $H_{n}$ is the random multihypergraph resulting from the configuration model. Moreover, for an integer $d \geq 1$ and vertex $v \in\{1, \ldots, n\},\left(H_{n}, v\right)_{d}$ is the multihypergraph rooted at vertex $v$ containing nodes with distance at most $d$ from $v$ and edges in $H_{n}$ with all endpoints among this set of vertices.

Proposition 21. Assume conditions (84a), (84b) and (84c) are satisfied with constants $c_{1}, c_{2}, c_{3}, \epsilon>$ 0 . Then, if $v_{0}$ is chosen uniformly at random from $\{1, \ldots, n\}$,

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left(\left(H_{n}, v_{0}\right)_{d} \text { is a simple hypertree }\right)=1
$$

Proof. Note that (84a) and (84b) imply that the degrees of all vertices and the sizes of all edges in $H_{n}$ are bounded by $\alpha_{n}:=c_{1}(\log n)^{c_{2}}$. Therefore, there are at most $\beta_{n}:=\left(\alpha_{n}\right)^{2 d+5}$ edges in $\left(H_{n}, v_{0}\right)_{d+1}$. Since at each step of the exploration process we form one edge, $\left(H_{n}, v_{0}\right)_{d+1}$ is completely observed up to step $\beta_{n}$. On the other hand, if at step $t$, the partial edge to be matched at that step, which is $e_{t}$ in our notation, is matched to partial edges outside $A_{t}$, all of the endpoints of the newly formed edge at step $t$ are in $N_{t}^{c}$, i.e. they are not observed so far. If in addition to this property, all the $k_{t}-1$ vertices of these partial edges that are matched with $e_{t}$ are distinct, no loops or multiple edges are formed at step $t$. If these properties hold for all $1 \leq t \leq \beta_{n},\left(H_{n}, v_{0}\right)_{d}$ will be a simple hypertree. Note that, in order to make sure that $\left(H_{n}, v_{0}\right)_{d}$ is a simple hypertree, we need to make sure that there is no loop in the exploration process up to depth $d+1$. This guarantees that even vertices at depth do not get connected to each other.

To formalize this, fix $1 \leq t \leq \beta_{n}$ and assume that the exploration process is not finished up to step $t$, and at step $t$, we need to match the partial edge $e_{t}$ of size $k_{t}$. Let $E_{t, 1}$ denote the event that $e_{t, 1}=\sigma_{k_{t}}^{(1)}\left(e_{t}\right) \in A_{t} \cap \Delta_{k_{t}}^{(n)}$. Moreover, for $2 \leq i \leq k_{t}-1$, let $E_{t, i}$ be te event that

$$
e_{t, i}=\sigma_{k_{t}}^{(i)}\left(e_{t}\right) \in\left(A_{t} \cap \Delta_{k_{t}}^{(n)}\right) \cup\left(\bigcup_{j=1}^{i-1} \Delta_{k_{t}}^{(n)}\left(v\left(e_{t, j}\right)\right)\right)
$$

where we recall that $\Delta_{k_{t}}^{(n)}\left(v\left(e_{t, j}\right)\right)$ denotes the set of partial edges of size $k_{t}$ connected to the vertex associated to $e_{t, j}$. Note that having chosen $e_{t, 1}, \ldots, e_{t, i-1}$, there are $\left|\Delta_{k_{t}}^{(n)}\right|-\left|\Delta_{k_{t}}^{(n)} \cap C_{t}\right|-i$ many candidates for $\sigma_{k_{t}}^{(i)}\left(e_{t}\right)$, each having the same chance of being chosen. We claim that $\left|A_{t} \cap \Delta_{k_{t}}^{(n)}\right| \leq$ $t \alpha_{n}^{2}$. The reason is that at each step in the exploration process, at most $\alpha_{n}$ many new vertices are added to $N_{t}$, each of whcih having at most $\alpha_{n}$ many partial edges of size $k_{t}$. On the other hand, $\left|\Delta_{k_{t}}^{(n)}\left(v\left(e_{t, j}\right)\right)\right| \leq \alpha_{n}$. Consequently,

$$
\left|\left(A_{t} \cap \Delta_{k_{t}}^{(n)}\right) \cup\left(\bigcup_{j=1}^{i-1} \Delta_{k_{t}}^{(n)}\left(v\left(e_{t, j}\right)\right)\right)\right| \leq(t+1) \alpha_{n}^{2}
$$

Additionally, at each step, at most $\alpha_{n}$ partial edges are added to $C_{t}$ to form a edge; therefore, $\left|C_{t}\right| \leq$ $t \alpha_{n}$. This, together with (84c), implies that for $1 \leq i \leq k_{t} \leq \alpha_{n}$, we have $\left|\Delta_{k_{t}}^{(n)}\right|-\left|\Delta_{k_{t}}^{(n)} \cap C_{t}\right|-i \geq$ $c_{3} n^{\epsilon}-(t+1) \alpha_{n}$. Since each of these candidates have the same probability of being chosen, we have

$$
\mathbb{P}\left(E_{t, i}\right) \leq \frac{(t+1) \alpha_{n}^{2}}{c_{3} n^{\epsilon}-(t+1) \alpha_{n}} \quad 1 \leq i \leq k_{t}-1
$$

If $E_{t}$ denotes $\cup_{i=1}^{k_{t}-1} E_{t, i}$, using $k_{t} \leq \alpha_{n}$ and the union bound, we have

$$
\mathbb{P}\left(\cup_{t=1}^{\beta_{n}} E_{t}\right) \leq \frac{\left(\beta_{n}+1\right)^{2} \alpha_{n}^{3}}{c_{3} n^{\epsilon}-\left(\beta_{n}+1\right) \alpha_{n}} .
$$

Note that, $\alpha_{n}$ and $\beta_{n}$ scale logarithmically in $n$, and $d$ is fixed. Hence, due to the $c_{3} n^{\epsilon}$ term in the denominator, the above probability goes to zero as $n$ goes to infinity. As was discussed above, outside the event $\cup_{t=1}^{\beta_{n}} E_{t}$, the rooted hypergraph $\left(H_{n}, v_{0}\right)_{d}$ is a simple hypertree.

Proposition 22. With the assumptions of Theorem 4, for an integer $d \geq 1$ and a rooted hypertree $(T, o)$ with depth at most $d$, if $v_{0}$ is a node chosen uniformly at random from $\{1, \ldots, n\}$, we have

$$
\lim _{n \rightarrow \infty} \mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o)\right)=\mu\left(A_{(T, o)}\right)
$$

Note that, the above proposition, together with the discussion in (98), implies that $\mathbb{E}\left[u_{H_{n}^{e}}\right] \Rightarrow \mu$. In Section G. 3 below, we show that $u_{H_{n}^{e}}$ is concentrated around its mean, which completes the proof of Theorem 4.

Proof. Given the rooted hypertree $(T, o)$ with depth at most $d$, let $\mathcal{C}_{[T, o]}$ denote the set of hypertrees $\tilde{T} \in \mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$ such that $(\tilde{T}, \emptyset) \equiv(T, o)$. For $\tilde{T} \in \mathcal{C}_{[T, o]}$, let $e_{1}^{\tilde{T}}, \ldots, e_{r}^{\tilde{T}}$ be the edges in $\tilde{T}$ with depth at most $d-1$, ordered lexicographically in $\mathbb{N}_{\text {edge }}$. Moreover, let $e_{r+1}^{\tilde{T}}, \ldots, e_{r+l}^{\tilde{T}}$ denote the edges in $\tilde{T}$ with depth precisely $d$, ordered lexicographically (if there is no such edge, $l=0$ ). For a node $\mathbf{i}$ in $\tilde{T}$, let $\gamma_{\mathbf{i}}^{\tilde{T}}$ denote the type of the vertex $\mathbf{i}$ in the subtree below $\mathbf{i}$. Furthermore, define

$$
\pi_{\tilde{T}}:=\mathbb{P}\left(\left(\mathcal{T}\left(\left\{\Gamma_{a}\right\}_{a \in \mathbb{N}_{\mathrm{vertex}}}\right), \emptyset\right)_{d}=(\tilde{T}, \emptyset)\right)
$$

under the probability in Definition 27. With the above notation, we have

$$
\begin{equation*}
\pi_{\tilde{T}}:=P\left(\gamma_{\emptyset}^{\tilde{T}}\right) \prod_{t=1}^{r} \prod_{j=1}^{\left|e_{t}^{\tilde{T}}\right|-1} \hat{P}_{\left|e_{t}^{\tilde{T}}\right|}\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}\right) \tag{99}
\end{equation*}
$$

By the definition of the distribution $\operatorname{UGWT}(P)$, we have

$$
\mu\left(A_{[T, o]}\right)=\sum_{\tilde{T} \in \mathcal{C}_{[T, o]}} \pi_{\tilde{T}} .
$$

Recall that $R_{t}$ denotes the hypertree in $\mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$ which results from the exploration process at step $t$. Note that if $\left(H_{n}, v_{0}\right)_{d}$ is a simple hypertree, we have $\left(H_{n}, v_{0}\right)_{d} \equiv\left(R_{\beta_{n}}, \emptyset\right)_{d}$ where $\beta_{n}$ is defined in Proposition 21 above. With this, we have

$$
\begin{aligned}
\mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o)\right)= & \mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o) \text { and }\left(H_{n}, v_{0}\right)_{d} \text { is a simple hypertree }\right) \\
& +\mathbb{P}\left(\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o) \text { and }\left(H_{n}, v_{0}\right)_{d} \text { is not a simple hypertree }\right)
\end{aligned}
$$

As is shown in Proposition 21, the second term converges to zero; therefore, we need to study only the first term. But, if $\left(H_{n}, v_{0}\right)_{d}$ is a simple hypertree, $\left(H_{n}^{e}, v_{0}\right)_{d} \equiv(T, o)$ if and only if $R_{r+l+1}=\tilde{T}$ for some $\tilde{T} \in \mathcal{C}_{[T, o]}$. Consequently, it suffices for us to show that

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \mathbb{P}\left(R_{r+l+1}=\tilde{T} \text { and }\left(H_{n}, v_{0}\right)_{d} \text { is a simple hypertree }\right)=\pi_{\tilde{T}} \quad \forall \tilde{T} \in \mathcal{C}_{[T, o]} \tag{100}
\end{equation*}
$$

For $1 \leq t \leq r$, let $E_{t}$ be the event defined in Proposition 21. Recall that $E_{t}^{c}$ is the event that $u_{t, 1}, \ldots, u_{t, k_{t}-1}$ are all distinct and are not in $N_{t}$. From Proposition 21, we know that the probabilities of both the events " $\left(H_{n}, v_{0}\right)_{d}$ is a simple hypertree" and $\cap_{t=1}^{r} E_{t}^{c}$ converge to 1 as $n \rightarrow \infty$. Therefore, it suffices to show that

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \mathbb{P}\left(\left(R_{r+l+1}=\tilde{T}\right) \cap\left(\cap_{t=1}^{r} E_{t}^{c}\right)\right)=\pi_{\tilde{T}} \quad \forall \tilde{T} \in \mathcal{C}_{[T, o]} \tag{101}
\end{equation*}
$$

To prove this, fix some $\tilde{T} \in \mathcal{T}\left(\mathbb{N}_{\text {vertex }}, \mathbb{N}_{\text {edge }}\right)$ with depth at most $d$ and define $S_{0}$ to be the event that $\gamma_{v_{0}}^{(n)}=\gamma_{\emptyset}^{\tilde{T}}$. From (83b), $\mathbb{P}\left(S_{0}\right) \rightarrow P\left(\gamma_{\emptyset}^{\tilde{T}}\right)$ as $n \rightarrow \infty$. Moreover, for $1 \leq t \leq r$, define $\tilde{S}_{t}$ to be the event that

$$
\gamma_{u_{t, s}}^{(n)}=\gamma_{\left(\mathbf{i}_{t}, k_{t}, \tilde{j}_{t}, s\right)}^{\tilde{T}}+\mathbb{E}_{k_{t}} \quad 1 \leq s \leq k_{t}-1
$$

and let $S_{t}=\tilde{S}_{t} \cap E_{t}^{c}$, which is in fact the intersection of $\tilde{S}_{t}$ and the event

$$
u_{t, j}, 1 \leq j \leq k_{t}-1 \text { are distinct and are not in } N_{t}
$$

With this, let $S^{t}=\cap_{i=0}^{t} S_{i}$. We claim that the event $\left(R_{r+l+1}=\tilde{T}\right) \cap\left(\cap_{t=1}^{r} E_{t}^{c}\right)$ coinsides with $S^{r}$. The reason is that on the event $S^{r}$, for each $1 \leq t \leq r$, the type of all the $k_{t}-1$ subnodes of edge formed at step $t$ matches with that of $e_{t}^{\tilde{T}}$. Moreover, on the event $\cap_{t=1}^{r} E_{t}^{c}$, there is no loops or cycles formed during the exploration process. In particular, those partial edges connected to the vertex $u_{t, s}$ which are added to the active set are not used until the process goes to vertex $u_{t, s}$ itself. Also, note that as $\tilde{T}$ has depth at most $d$, its structure is determined by the type of the vertices of depth at most $d-1$, which are subnodes of edges of depth at most $d-1$ in $\tilde{T}$, which are precisely $e_{1}^{\tilde{T}}, \ldots, e_{r}^{\tilde{T}}$.

Now, we prove by induction that for $0 \leq t \leq r$,

$$
\mathbb{P}\left(S^{t}\right) \rightarrow \pi_{\tilde{T}}(t):=P\left(\gamma_{\emptyset}^{\tilde{T}}\right) \prod_{t^{\prime}=1}^{t} \prod_{j=1}^{\left|e_{t^{\prime}}^{\tilde{T}}\right|-1} \hat{P}_{\left|e_{t^{\prime}}^{\tilde{T}}\right|}\left(\gamma_{\left(e_{t^{\prime}}^{\tilde{T}}, j\right)}^{\tilde{T}}\right)
$$

If $\pi_{\tilde{T}}(t-1)=0$, then $\mathbb{P}\left(S^{t}\right) \leq \mathbb{P}\left(S^{t-1}\right) \rightarrow 0=\pi_{\tilde{T}}(t)$. If $\pi_{\tilde{T}}(t-1) \neq 0$, we have $\mathbb{P}\left(S^{t-1}\right)>0$ for $n$ large enough. Note that $\mathbb{P}\left(S^{t}\right)=\mathbb{P}\left(S^{t-1} \cap \tilde{S}_{t} \cap E_{t}^{c}\right)$. Thereby,

$$
\mathbb{P}\left(S^{t-1}\right) \mathbb{P}\left(\tilde{S}_{t} \mid S^{t-1}\right)-\mathbb{P}\left(E_{t}\right) \leq \mathbb{P}\left(S^{t}\right) \leq \mathbb{P}\left(S^{t-1}\right) \mathbb{P}\left(\tilde{S}_{t} \mid S^{t-1}\right)
$$

But, we know that $\mathbb{P}\left(E_{t}\right) \rightarrow 0$. Consequently, it suffices to prove that

$$
\begin{equation*}
\mathbb{P}\left(\tilde{S}_{t} \mid S^{t-1}\right) \rightarrow \prod_{j=1}^{\left|e_{t}^{\tilde{T}}\right|-1} \hat{P}_{\left|e_{t}^{\tilde{T}}\right|}\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}\right) \tag{102}
\end{equation*}
$$

Since we construct one edge at a time in the exploration process, conditioned on $S^{t-1}$, the first $t-1$ edges are constructed in a way consistent with $\tilde{T}$. Therefore, it is easy to see that

$$
\mathbb{P}\left(\tilde{S}_{t} \mid S^{t-1}\right)=\mathbb{P}\left(\gamma_{u_{t, j}}^{(n)}=\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}+\mathbb{e}_{\left|e_{t}^{\tilde{T}}\right|} \text { for } 1 \leq j \leq\left|e_{t}^{\tilde{T}}\right|-1\right)
$$

For $1 \leq j \leq\left|e_{t}^{\tilde{T}}\right|-1$, let $\tilde{S}_{t, j}$ denote the event that $\gamma_{u_{t, j}}^{(n)}=\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}+\mathbb{e}_{\left|e_{t}^{\tilde{T}}\right|}$. Now, we study the probability of $\tilde{S}_{t, j}$ conditioned on $S^{t-1}$ and $\tilde{S}_{t, 1}, \ldots, \tilde{S}_{t, j-1}$. Note that, having chosen $e_{t, 1}, \ldots, e_{t, j-1}$, there are $\left|\Delta^{(n)}\left(k_{t}\right)\right|-\left|\Delta^{(n)}\left(k_{t}\right) \cap C_{t}\right|-j$ many candidates for $e_{t, j}$, each having the same chance. With $B_{t, j}:=\left\{e \in \Delta^{(n)}\left(\left|e_{t}^{\tilde{T}}\right|\right): \gamma_{v(e)}^{(n)}=\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}+\mathbb{E}_{\left|e_{t}^{\tilde{T}}\right|}\right\}$, the event $\tilde{S}_{t, j}$ happens iff $e_{t, j}$ is chosen among the set $B_{t, j} \backslash\left(C_{t} \cup\left\{e_{t, 1}, \ldots, e_{t, j-1}\right\}\right)$. Therefore,

$$
\mathbb{P}\left(\tilde{S}_{t, j} \mid S^{t-1}, \tilde{S}_{t, 1}, \ldots, \tilde{S}_{t, j-1}\right)=\frac{\left|B_{t, j} \backslash\left(C_{t} \cup\left\{e_{t, 1}, \ldots, e_{t, j-1}\right\}\right)\right|}{\left|\Delta^{(n)}\left(\left|e_{t}^{\tilde{T}}\right|\right)\right|-\left|\Delta^{(n)}\left(\left|e_{t}^{\tilde{T}}\right|\right) \cap C_{t}\right|-j}
$$

Note that

$$
\frac{1}{n}\left|B_{t, j}\right|=\frac{1}{n} \sum_{i=1}^{n}\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}\left(\left|e_{t}^{\tilde{T}}\right|\right)+1\right) \mathbb{1}\left[\gamma_{i}^{(n)}=\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{\tilde{T}}}+\mathbb{e}_{\left|e_{t}^{\tilde{T}}\right|}\right]
$$

Using (83b), we have

$$
\frac{1}{n}\left|B_{t, j}\right| \rightarrow\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}\left(\left|e_{t}^{\tilde{T}}\right|\right)+1\right) P\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}+\mathbb{e}_{\left|e_{t}^{\tilde{T}}\right|}\right)
$$

On the other hand, conditioned on $S^{t-1},\left|C_{t}\right|=\sum_{j=1}^{t-1}\left|e_{j}^{\tilde{T}}\right|$, which is a constant. Consequently,

$$
\begin{equation*}
\frac{1}{n}\left|B_{t, j} \backslash\left(C_{t} \cup\left\{e_{t, 1}, \ldots, e_{t, j-1}\right\}\right)\right| \rightarrow\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{\tilde{T}}}\left(\left|e_{t}^{\tilde{T}}\right|\right)+1\right) P\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{\tilde{T}}}+\mathbb{e}_{\left|e_{t}^{\tilde{T}}\right|}\right) \tag{103}
\end{equation*}
$$

Moreover, using (85),

$$
\begin{equation*}
\frac{1}{n}\left|\Delta^{(n)}\left(\left|e_{t}^{\tilde{T}}\right|\right)\right|=\frac{1}{n} \sum_{i=1}^{n} \gamma_{i}^{(n)}\left(\left|e_{t}^{\tilde{T}}\right|\right) \rightarrow \mathbb{E}\left[\Gamma\left(\left|e_{t}^{\tilde{T}}\right|\right)\right] \tag{104}
\end{equation*}
$$

Note that we are conditioning on $S^{t-1}$ and assuming that $\mathbb{P}\left(S^{t-1}\right) \neq 0$. On the other hand, $\left|e_{t}^{\tilde{T}}\right|$ is equal to the size of the partial edge $e_{t}$ which is a member of $\Delta^{(n)}$. Using the assumption (83a), we have $\left|e_{t}^{\tilde{T}}\right| \in I$ and hence $\mathbb{E}\left[\Gamma\left(\left|e_{t}^{\tilde{T}}\right|\right)\right]>0$. Putting (103) and (104) together, we have

$$
\mathbb{P}\left(\tilde{S}_{t, j} \mid S^{t-1}, \tilde{S}_{t, 1}, \ldots, \tilde{S}_{t, j-1}\right) \rightarrow \frac{\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}\left(\left|e_{t}^{\tilde{T}}\right|\right)+1\right) P\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}+\mathbb{e}_{\left|e_{t}^{\tilde{T}}\right|}\right)}{\mathbb{E}\left[\Gamma\left(\left|e_{t}^{\tilde{T}}\right|\right)\right]}=\hat{P}_{\left|e_{t}^{\tilde{T}}\right|}\left(\gamma_{\left(e_{t}^{\tilde{T}}, j\right)}^{\tilde{T}}\right)
$$

Multiplying for $1 \leq j \leq\left|e_{t}^{\tilde{T}}\right|-1$, we get (102) which completes the proof.

## G. 3 Almost sure convergence

In this section we prove that, with the assumptions of Theorem $4, u_{H_{n}^{e}} \Rightarrow \mathrm{UGWT}(P)$ almost surely.
For a fixed $n$, Let $\Delta^{(n)}\left(k_{1}\right) \ldots \Delta^{(n)}\left(k_{L}\right)$ be the nonempty sets among $\Delta^{(n)}(2), \ldots, \Delta^{(n)}(n)$. From (84b) we know that $L \leq c_{1}(\log n)^{c_{2}}$ and also $k_{i} \leq c_{1}(\log n)^{c_{2}}$ for $1 \leq i \leq L$. For the sake of simplicity, write $\sigma$ for $\left(\sigma_{k_{1}}, \ldots, \sigma_{k_{L}}\right)$ and $M_{i}$ for $\mathcal{M}_{k_{i}}\left(\Delta^{(n)}\left(k_{i}\right)\right), 1 \leq i \leq L$. From our construction, we know that $\sigma_{k_{i}}$ is drawn uniformly at random from $M_{i}$ and is independent from $\sigma_{k_{j}}, j \neq i$. With $H_{n}^{e}$ being the simple hypergraph constructed by the configuration model, for a fixed $d>0$ and a rooted tree ( $T, o$ ) of depth at most $d$, define

$$
F(\sigma):=\frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\left(H_{n}^{e}, i\right)_{d} \equiv(T, i)}
$$

From Proposition 22 we know that $\lim _{n \rightarrow \infty} \mathbb{E}[F(\sigma)]=\mu\left(A_{(T, o)}\right)$. We will show that $F$ is concentrated around its mean via a bounded difference argument.

Now, for each $1 \leq j \leq L$, fix a permutation $\pi_{k_{j}} \in M_{j}$ and define $\pi=\left(\pi_{k_{1}}, \ldots, \pi_{k_{L}}\right)$. Moreover, fix $1 \leq i \leq L$ and $e, e^{\prime} \in \Delta^{(n)}\left(k_{i}\right)$. With this, define $\pi_{k_{i}}^{\prime}:=\operatorname{swap}_{e, e^{\prime}} \circ \pi_{k_{i}} \circ \operatorname{swap}_{e, e^{\prime}}$, which is the conjugation of $\pi_{k_{i}}$ with the permutation that swaps $e$ and $e^{\prime}$. In fact, the cycle representation of $\pi_{k_{i}}^{\prime}$ is obtained by swapping $e$ and $e^{\prime}$ in the cycle representation of $\pi_{k_{i}}$. Moreover, let $\pi^{\prime}=\left(\pi_{k_{1}}, \ldots, \pi_{k_{i}}^{\prime}, \ldots, \pi_{k_{L}}\right)$ which differs from $\pi$ only on the $i^{\text {th }}$ coordinate. With this, the hypergraph obtained from $\pi$ and the hypergraph obtained from $\pi^{\prime}$ differ only in at most two edges. Since all edge sizes and degrees in the
graph are bounded to $c_{1}(\log n)^{c_{2}}$, there are at most $2\left(c_{1}(\log n)^{c_{2}}\right)^{2 d+1}$ many vertices in the hypergraph which have distance at most $d$ to a vertex in any of these two edges. Consequently,

$$
\begin{equation*}
\left|F(\pi)-F\left(\pi^{\prime}\right)\right| \leq \frac{2\left(c_{1}(\log n)^{c_{2}}\right)^{2 d+1}}{n} \tag{105}
\end{equation*}
$$

Now, fix $1 \leq i \leq L$ and $\pi_{k_{j}} \in M_{j}$ for $j \neq i$. Let $\sigma_{k_{i}}$ being chosen uniformly at random in $M_{i}$ and define $F_{i}\left(\sigma_{k_{i}}\right)=F\left(\pi_{k_{1}}, \ldots, \sigma_{k_{i}}, \ldots, \pi_{k_{L}}\right)$. Since $\Delta^{(n)}\left(k_{i}\right)$ is finite, we can equip it with an arbitrary total order. Let $X_{1}$ be the smallest element in $\Delta^{(n)}\left(k_{i}\right)$ and define $Y_{1}=\left(X_{1}, \sigma_{k_{i}}\left(X_{1}\right), \ldots, \sigma_{k_{i}}^{\left(k_{i}-1\right)}\left(X_{1}\right)\right)$, which is in fact the orbit of $X_{1}$, or in the configuration model language, the edge containing the partial edge $X_{1}$. Let $X_{2}$ be the smallest element that does not appear in $Y_{1}$ and let $Y_{2}=\left(X_{2}, \sigma_{k_{i}}\left(X_{2}\right), \ldots, \sigma_{k_{i}}^{\left(k_{i}-1\right)}\left(X_{2}\right)\right)$. We continue this process inductively, i.e. let $X_{j}$ be the smallest element that has not appeared in $Y_{1}, \ldots, Y_{j-1}$ and let $Y_{j}=\left(X_{j}, \sigma_{k_{i}}\left(X_{j}\right), \ldots, \sigma_{k_{i}}^{\left(k_{i}-1\right)}\left(X_{j}\right)\right)$. This process yields $Y_{1}, \ldots, Y_{\left|\Delta{ }^{(n)}\left(k_{i}\right)\right| / k_{i}}$. For $1 \leq j \leq\left|\Delta^{(n)}\left(k_{i}\right)\right| / k_{i}$, let $\mathcal{F}_{j}$ be the sigma field generated by $Y_{1}, \ldots, Y_{j}$. Moreover, let $Z_{j}=$ $\mathbb{E}\left[F_{i}\left(\sigma_{k_{i}}\right) \mid \mathcal{F}_{j}\right]$ for $1 \leq j \leq\left|\Delta^{(n)}\left(k_{i}\right)\right| / k_{i}$ and let $Z_{0}=\mathbb{E}\left[F_{i}\left(\sigma_{k_{i}}\right)\right]$. Note that $\pi_{k_{j}}, j \neq i$ are fixed; therefore, the randomness in the expression is with respect to $\sigma_{k_{i}}$ only. Indeed, $\left(Z_{j}, 0 \leq j \leq\left|\Delta^{(n)}\left(k_{i}\right)\right| / k_{i}\right)$ is a martingale. We claim that, almost surely, we have

$$
\left|Z_{j+1}-Z_{j}\right| \leq k_{i} \frac{2\left(c_{1}(\log n)^{c_{2}}\right)^{2 d+1}}{n}
$$

The reason is that changing the value of the $k_{i}$ variables in $Y_{j+1}$ can change the value of $F_{i}$ by at most $k_{i} \frac{2\left(c_{1}(\log n)^{c_{2}}\right)^{2 d+1}}{n}$ and the above inequality results from (105). Using Azuma's inequality and the fact that $k_{i} \leq c_{1}(\log n)^{c_{2}}$, we have

$$
\begin{equation*}
\mathbb{P}\left(\left|F_{i}\left(\sigma_{k_{i}}\right)-\mathbb{E}\left[F_{i}\left(\sigma_{k_{i}}\right)\right]\right|>\delta\right)<2 \exp \left(-\frac{\delta^{2} n^{2}}{4\left|\Delta^{(n)}\left(k_{i}\right)\right|\left(c_{1}(\log n)^{c_{2}}\right)^{4 d+3}}\right) \tag{106}
\end{equation*}
$$

To obtain an upper bound for $\left|\Delta^{(n)}\left(k_{i}\right)\right|$, note that

$$
\left|\Delta^{(n)}\left(k_{i}\right)\right|=\left(\sum_{j=1}^{n} \gamma_{j}^{(n)}\left(k_{i}\right)\right) \leq n\left(\frac{1}{n} \sum_{j=1}^{n}\left\|\gamma_{j}^{(n)}\right\|_{1}^{2}\right) .
$$

From (83c), there is a constant $\alpha$ independent of $n$ and $i$ that $\sum\left\|\gamma_{j}^{(n)}\right\|_{1}^{2}<\alpha n$. Hence, $\left|\Delta^{(n)}\left(k_{i}\right)\right|<\alpha n$. Incorporating this into (106), we have, for $1 \leq i \leq L$,

$$
\begin{equation*}
\mathbb{P}\left(\left|F_{i}\left(\sigma_{k_{i}}\right)-\mathbb{E}\left[F_{i}\left(\sigma_{k_{i}}\right)\right]\right|>\delta\right)<2 \exp \left(-\frac{\delta^{2} n}{4 \alpha\left(c_{1}(\log n)^{c_{2}}\right)^{4 d+3}}\right) \tag{107}
\end{equation*}
$$

Since this is true for all $i$ and $\pi_{k_{j}}, j \neq i$ and also the $\sigma_{k_{j}}$ are independent, using the above inequality $L$ times and using the fact that $L \leq c_{1}(\log n)^{c_{2}}$, we have

$$
\mathbb{P}(|F(\sigma)-\mathbb{E}[F(\sigma)]|>\delta) \leq 2 c_{1}(\log n)^{c_{2}} \exp \left(-\frac{\delta^{2} n}{4 \alpha\left(c_{1}(\log n)^{c_{2}}\right)^{4 d+3}}\right)
$$

As the sum of the RHS over $n$ is finite, using the Borel-Cantelli lemma and the fact that $\mathbb{E}[F(\sigma)] \rightarrow$ $\mu\left(A_{(T, o)}\right)$, we have $F(\sigma) \rightarrow \mu\left(A_{(T, o)}\right)$ almost surely. But there are countably many choices for $d$ and the rooted hypertree $(T, o)$. Thus, outside a measure zero set, $u_{H_{n}^{e}}\left(A_{(T, o)}\right) \rightarrow \mu\left(A_{(T, o)}\right)$ for all rooted tree $(T, o)$ with finite depth. The proof is complete, using Lemma 4.

## H Proof of Proposition 9

Proof of Proposition 9. Since $\Theta_{\epsilon}^{\prime}$ is $\epsilon$-balanced, from Definition 25, for $\vec{\mu}$-almost every $[H, e, i] \in \mathcal{H}_{* *}$, we have

$$
\begin{equation*}
\Theta_{\epsilon}^{\prime}(H, e, i)=\frac{\exp \left(-\frac{\partial \Theta_{\epsilon}^{\prime}(H, i)}{\epsilon}\right)}{\sum_{j \in e} \exp \left(-\frac{\partial \Theta_{\epsilon}^{\prime}(H, j)}{\epsilon}\right)} \tag{108}
\end{equation*}
$$

Using Proposition 3, there exists a $A \subset \mathcal{H}_{* *}$ such that $\vec{\mu}\left(A^{c}\right)=0$ and, for all $[H, e, i] \in A$, we have

$$
\Theta_{\epsilon}^{\prime}\left(H, e^{\prime}, i^{\prime}\right)=\frac{\exp \left(-\frac{\partial \Theta_{\epsilon}^{\prime}\left(H, i^{\prime}\right)}{\epsilon}\right)}{\sum_{j \in e^{\prime}} \exp \left(-\frac{\partial \Theta_{\epsilon}^{\prime}(H, j)}{\epsilon}\right)} \quad \forall\left(e^{\prime}, i^{\prime}\right) \in \Psi(H)
$$

Now, fix some $[H, e, i] \in A$ and take an arbitrary element of this equivalence class $(H, e, i) \in[H, e, i]$. The above equation guarantees that if we define the allocation $\theta_{\epsilon}^{\prime H}$ on $H$ as $\theta_{\epsilon}^{\prime H}\left(e^{\prime}, i^{\prime}\right):=\Theta_{\epsilon}^{\prime}\left(H, e^{\prime}, i^{\prime}\right)$ for $\left(e^{\prime}, i^{\prime}\right) \in \Psi(H)$, then $\theta_{\epsilon}^{\prime H}$ is an $\epsilon$-balanced allocation on $H$. Now, assume that $\theta_{\epsilon}^{H^{\Delta}}$ is the (unique) $\epsilon$-balanced allocation on the truncated hypergraph $H^{\Delta}$ defined in Section 4.3 (uniqueness comes from boundedness of $H^{\Delta}$ ). Proposition 5 then implies that $\partial \theta_{\epsilon}^{H^{\Delta}}\left(i^{\prime}\right) \leq \partial \theta_{\epsilon}^{\prime H}\left(i^{\prime}\right)$ for all $i^{\prime} \in V(H)$. Sending $\Delta$ to infinity, this means that $\partial \theta_{\epsilon}^{H}\left(i^{\prime}\right) \leq \partial \theta_{\epsilon}^{\prime H}\left(i^{\prime}\right)$ for all $i^{\prime} \in V(H)$, with $\theta_{\epsilon}^{H}$ being the canonical $\epsilon$-balanced allocation on $H$. Using Remark 12 and the definition of $\theta_{\epsilon}^{\prime H}$, this means that for $\vec{\mu}$-almost all $[H, e, i] \in \mathcal{H}_{* *}, \partial \Theta_{\epsilon}(H, i) \leq \partial \Theta_{\epsilon}^{\prime}(H, i)$. From part (ii) of Lemma 1, $\mu$-almost surely we have

$$
\begin{equation*}
\partial \Theta_{\epsilon} \leq \partial \Theta_{\epsilon}^{\prime} \tag{109}
\end{equation*}
$$

On the other hand, using unimodularity of $\mu$ and the fact that $\Theta_{\epsilon}$ is a Borel allocation, we have

$$
\int \partial \Theta_{\epsilon} d \mu=\int \Theta_{\epsilon} d \vec{\mu}=\int \nabla \Theta_{\epsilon} d \vec{\mu}=\int \frac{1}{|e|} d \vec{\mu}([H, e, i])
$$

Using the same logic, $\int \partial \Theta_{\epsilon}^{\prime} d \mu=\int \frac{1}{|e|} d \vec{\mu}([H, e, i])$. This means that $\int \partial \Theta_{\epsilon} d \mu=\int \partial \Theta_{\epsilon}^{\prime} d \mu$. As $\operatorname{deg}(\mu)<$ $\infty$, this common value is finite. This, together with (109), implies that $\partial \Theta_{\epsilon}=\partial \Theta_{\epsilon}^{\prime}, \mu$-almost surely. Therefore, Proposition 3 implies that for $\mu$-almost all $[H, i] \in \mathcal{H}_{*}$, we have $\partial \Theta_{\epsilon}(H, j)=\partial \Theta_{\epsilon}^{\prime}(H, j)$ for all $j \in V(H)$. Then, part ( $i$ ) of Lemma 1 implies that for $\vec{\mu}$-almost all $[H, e, i] \in \mathcal{H}_{* *}, \partial \Theta_{\epsilon}(H, j)=$ $\partial \Theta_{\epsilon}^{\prime}(H, j)$ for all $j \in V(H)$. Thereby, using (108) for $\Theta_{\epsilon}$ and $\Theta_{\epsilon}^{\prime}$, we have $\Theta_{\epsilon}(H, e, i)=\Theta_{\epsilon}^{\prime}(H, e, i)$ for $\vec{\mu}$-almost all $[H, e, i] \in \mathcal{H}_{* *}$, which completes the proof.

## Acknowledgements

The authors acknowledge support from the NSF grants ECCS-1343398, CNS- 1527846, CCF-1618145, the NSF Science \& Technology Center grant CCF- 0939370 (Science of Information), and the William and Flora Hewlett Foundation supported Center for Long Term Cybersecurity at Berkeley.

## References

[AL07] David Aldous and Russell Lyons. Processes on unimodular random networks. Electron. J. Probab, 12(54):1454-1508, 2007.
[AMO09] R.P. Agarwal, M. Meehan, and D. O’Regan. Fixed Point Theory and Applications. Cambridge Tracts in Mathematics. Cambridge University Press, 2009.
[AS04] David Aldous and J Michael Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. In Probability on discrete structures, pages 1-72. Springer, 2004.
[AS16] Venkat Anantharam and Justin Salez. The densest subgraph problem in sparse random graphs. The Annals of Applied Probability, 26(1):305-327, 2016.
[Bil71] Patrick Billingsley. Weak convergence of measures. Philadelphia: SIAM, 1971.
[Bil99] Patrick Billingsley. Convergence of probability measures. John Wiley \& Sons, 1999.
[Bor14] Charles Bordenave. Lecture notes on random graphs and probabilistic combinatorial optimization. 2014. http://www.math.univ-toulouse.fr/~bordenave/coursRG.pdf.
[BS01] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. Electron. J. Probab., 6:no. 23, 13 pp. (electronic), 2001.
[CK11] Imre Csiszar and János Körner. Information theory: coding theorems for discrete memoryless systems. Cambridge University Press, 2011.
$\left[\mathrm{DM}^{+} 10\right]$ Amir Dembo, Andrea Montanari, et al. Gibbs measures and phase transitions on sparse random graphs. Brazilian Journal of Probability and Statistics, 24(2):137-211, 2010.
[ $\left.\mathrm{H}^{+} 96\right]$ Bruce Hajek et al. Balanced loads in infinite networks. The Annals of Applied Probability, 6(1):48-75, 1996.
[Haj90] Bruce Hajek. Performance of global load balancing by local adjustment. Information Theory, IEEE Transactions on, 36(6):1398-1414, 1990.
[Mun00] James R Munkres. Topology. Prentice Hall, 2000.
[SS07] Moshe Shaked and J George Shanthikumar. Stochastic orders. Springer Science \& Business Media, 2007.
[Tao11] T. Tao. An Introduction to Measure Theory. Graduate studies in mathematics. American Mathematical Society, 2011.


[^0]:    *Parts of this work have been presented in the $55^{\text {th }}$ Annual Allerton Conference on Communication, Control, and Computing

[^1]:    ${ }^{1}$ Whether the converse is true is an open question, even in the graph regime

