Sperner's Colorings, Hypergraph Labeling Problems and Fair Division

Maryam Mirzakhani* Jan Vondrák†

Abstract

We prove three results about colorings of the simplex reminiscent of Sperner's Lemma, with applications in hardness of approximation and fair division.

First, we prove a coloring lemma conjectured by [5]: Let $V_{k,q} = \{\mathbf{v} \in \mathbb{Z}_+^k : \sum_{i=1}^k v_i = q\}$ and $E_{k,q} = \{\{\mathbf{a}+\mathbf{e}_1, \mathbf{a}+\mathbf{e}_2, \ldots, \mathbf{a}+\mathbf{e}_k\} : \mathbf{a} \in \mathbb{Z}_+^k, \sum_{i=1}^k a_i = q-1\}$. Then for every Sperner-admissible labeling ($\ell : V_{k,q} \to [k]$ such that $v_{\ell(\mathbf{v})} > 0$ for each $\mathbf{v} \in V_{k,q}$), there are at least $\binom{q+k-3}{k-2}$ non-monochromatic hyperedges in $E_{k,q}$. This implies an optimal Unique-Games hardness of $(k-1-\epsilon)$-approximation for the Hypergraph Labeling with Color Lists problem [2]: Given a k-uniform hypergraph $H = (V, E)$ with color lists $L(v) \subseteq [k]$ $\forall v \in V$, find a labeling $\ell(v) \in L(v)$ that minimizes the number of non-monochromatic hyperedges. We also show that a $(k-1)$-approximation can be achieved.

Second, we show that in contrast to Sperner's Lemma, there is a Sperner-admissible labeling of $V_{k,q}$ such that every hyperedge in $E_{k,q}$ contains at most 4 colors. We present an interpretation of this statement in the context of fair division: There is a preference function on $\Delta_{k,q} = \{\mathbf{x} \in \mathbb{R}_+^k : \sum_{i=1}^k x_i = q\}$ such that for any division of q units of a resource, $(x_1, x_2, \ldots, x_k) \in \Delta_{k,q}$ such that $\sum_{i=1}^k \lfloor x_i \rfloor = q-1$, at most 4 players out of k are satisfied.

Third, we prove that there are subdivisions of the simplex with a fractional labeling (analogous to a fractional solution for Min-CSP problems) such that every hyperedge in the subdivision uses only labelings with 1 or 2 colors. This means that a natural LP cannot distinguish instances of Hypergraph Labeling with Color Lists that can be labeled so that every hyperedge uses at most 2 colors, and instances that must have a rainbow hyperedge. We prove that this problem is indeed NP-hard for $k = 3$.

1 Introduction

We investigate hypergraph labeling problems of the following kind.

Hypergraph Labeling with Color Lists: *Given a hypergraph $H = (V, E)$ with weights $w(e), e \in E$ and color lists $L(v) \subseteq [k], v \in V$, find a labeling $\ell(v) \in L(v)$ for each $v \in V$ that minimizes the total weight of non-monochromatic hyperedges.*

This problem (in an equivalent form with assignment costs, referred to as Hypergraph Labeling) was introduced in [2] as a generalization of Uniform Metric Labeling [7], to incorporate relationships between multiple elements. (See Section 5 for a discussion of these problems and their equivalence in more detail.) Hypergraph Labeling with Color Lists can be cast in a more general framework involving submodular functions, as follows.

Submodular Labeling with Color Lists: *Given a submodular function $f : 2^V \to \mathbb{R}_+$ and color lists $L(v) \subseteq [k], v \in V$, find a labeling $\ell(v) \in L(v)$ that minimizes $\sum_{i=1}^k f(\ell^{-1}(i))$.*

Partitioning problems of this type have been investigated recently in [2, 3, 6, 5]. The main result of [3] is a 2-approximation for Submodular Multiway Partition, a special case of Submodular Labeling with Color Lists where the color lists are either singletons ("terminals") or equal to $[k]$ (unrestricted). This captures problems such as Hypergraph Multiway Cut and Node-weighted Multiway Cut (see [3]), where the color lists are similarly restricted to be singletons or $[k]$. Without this restriction, Submodular Labeling with Color Lists does not admit factors better than $\log n$, by a simple reduction from Set Cover [2]. An $O(k \log n)$-approximation for Submodular Labeling with Color Lists (in fact for a somewhat more general problem) has been given in [5].

For Hypergraph Labeling, [2] gave a Δ-approximation when all hyperedges have size at most Δ. This generalizes a 2-approximation for Uniform Metric Labeling [7] which corresponds to the $\Delta = 2$ case. On the hardness side, the strongest negative result was a hardness of $(2-\epsilon)$-approximation assuming the Unique Games Conjecture (for the special case of Uniform Metric Labeling [8]).

In [5], a statement somewhat reminiscent of Sperner's Lemma was conjectured, which would imply an integrality gap and also a hardness of $(k-1-\epsilon)$-approximation under the UGC (using [6]), for Hypergraph Labeling with Color Lists on k-uniform hypergraphs with label set $[k]$. This conjecture gives a lower bound

*Dept. of Mathematics, Stanford University, Stanford, CA; mmirzakh@stanford.edu.

†IBM Almaden Research Center, San Jose, CA; jvondrak@us.ibm.com.

on the number of non-monochromatic hyperedges for any feasible labeling of a certain hypergraph $H_{k,q}$ embedded in the simplex. We give the precise statement in Section 2. This statement was proved for $k = 3$ in [5]; it implies a Unique-Games hardness of $(2 - \epsilon)$-approximation for Hypergraph Labeling with Color Lists on 3-uniform hypergraphs with label set [3].

Our contribution. Our first result is a proof of the simplex coloring lemma conjectured by [5]. This lemma implies that assuming the Unique Games Conjecture, there is no $(k-1-\epsilon)$-approximation for Hypergraph Labeling with Color Lists on k-uniform hypergraphs with label set $[k]$. A Δ-approximation algorithm, where Δ is the maximum size of a hyperedge, was known for this problem [2]; we show that the same algorithm also gives a $(k-1)$-approximation (for label set $[k]$, independent of the hypergraph). Thus, we determine the optimal approximability of Hypergraph Labeling with Color Lists in terms of the label set size (modulo the UGC).

This result motivates us to consider other coloring questions related to Sperner's Lemma and the conjecture of [5]. We prove that in contrast to Sperner's Lemma, the hypergraph $H_{k,q}$ defined in [5] can be labeled in such a way that each hyperedge uses at most 4 colors. This implies in particular that the hardness result for Hypergraph Labeling with Color Lists holds even for hypergraphs with a feasible labeling such that each hyperedge uses at most 4 colors.

We also give an interpretation of this result in the setting of *fair division* (a well-known application of Sperner's Lemma [12]). Our result shows that for a certain restricted variant of fair division, not only is it impossible to satisfy all players but in fact all players except four will be always unsatisfied. We discuss this in Section 7.

Further, we consider simplicial subdivisions and ask what *fractional labelings* (in the sense of [6]) are possible for subdivisions of the simplex. We show that for sufficiently fine subdivisions, there is a fractional (Sperner-admisible) labeling that uses local labelings using at most 2 colors for each hyperedge (see Section 6 for a precise statement). In contrast, by Sperner's Lemma, for any admissible labeling there is a hyperedge with all k colors. This has consequences for the following problem.

Hypergraph j-Colors-Avoiding Labeling with Color Lists: *Given a hypergraph $H = (V, E)$ with color lists $L(v) \subseteq [k], v \in V$, find a labeling $\ell(v) \in L(v)$ for each $v \in V$ that minimizes the number of hyperedges containing at least j distinct colors.*

In particular, for $j = k$ we try to avoid hyperedges containing all k colors; we call this problem Hypergraph Rainbow-Avoiding Labeling with Color Lists. Our result implies that a natural LP for this problem cannot distinguish between instances that can be labeled so that each hyperedge contains at most 2 colors, and instances where some hyperedge must contain all k colors. We prove that it is in fact NP-hard to decide whether there is a feasible labeling such that every hyperedge contains at most 2 colors, for $k = 3$.

Organization. In Section 2, we state the simplex coloring lemma conjectured by [5]. In Section 3, we prove the simplex coloring lemma. In Section 4, we present a labeling of the hypergraph of [5] with at most 4 colors on each hyperedge. In Section 5, we discuss the applications of these results to Hypergraph Labeling with Color Lists and present our improved $(k-1)$-approximation. In Section 6, we describe the Hypergraph j-Colors-Avoiding Labeling with Color Lists problem and our hardness result for it. Finally, we discuss an application to fair division in Section 7.

2 Preliminaries

A note on vector notation: We denote vectors in boldface, such as $\mathbf{v} \in \mathbb{R}^k$. The coordinates of $\mathbf{v}$ are written in italics, such as $\mathbf{v} = (v_1, \ldots, v_k)$. By $\mathbf{e}_i$, we denote the canonical basis vectors $(0, \ldots, 1, \ldots, 0)$.

2.1 The Simplex-Lattice Hypergraph and subdivisions of the simplex.

Let $q \geq 1$ be an integer and consider the $(k-1)$-dimensional simplex defined by

$$\Delta_{k,q} = \left\{ \mathbf{x} = (x_1, x_2, \ldots, x_k) \in \mathbb{R}^k : \mathbf{x} \geq 0, \sum_{i=1}^{k} x_i = q \right\}.$$

The Simplex-Lattice Hypergraph. We consider a vertex set of all the points in $\Delta_{k,q}$ with integer coordinates:

$$V_{k,q} = \left\{ \mathbf{a} = (a_1, a_2, \ldots, a_k) \in \mathbb{Z}^k : \mathbf{a} \geq 0, \sum_{i=1}^{k} a_i = q \right\}.$$

The *Simplex-Lattice Hypergraph* is a k-uniform hypergraph $H_{k,q} = (V_{k,q}, E_{k,q})$ whose hyperedges (which we also call *cells* due to their geometric interpretation) are indexed by $\mathbf{b} \in \mathbb{Z}_+^k$ such that $\sum_{i=1}^{k} b_i = q - 1$: we have

$$E_{k,q} = \left\{ e(\mathbf{b}) : \mathbf{b} \in \mathbb{Z}^k, \mathbf{b} \geq 0, \sum_{i=1}^{k} b_i = q - 1 \right\}$$

where $e(\mathbf{b}) = \{\mathbf{b} + \mathbf{e}_1, \mathbf{b} + \mathbf{e}_2, \ldots, \mathbf{b} + \mathbf{e}_k\} = \{(b_1 + 1, b_2, \ldots, b_k), (b_1, b_2 + 1, \ldots, b_k), \ldots, (b_1, b_2, \ldots, b_k + 1)\}$. We sometimes omit the indices k, q when there is no danger of confusion. For each vertex $\mathbf{a} \in V_{k,q}$, we have a list of admissible colors $L(\mathbf{a})$, which is

$$L(\mathbf{a}) = \{i \in [k] : a_i > 0\}.$$

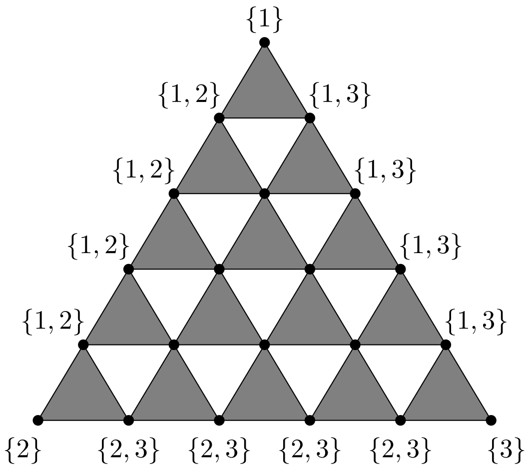

Figure 1: The Simplex Lattice Hypergraph for $k = 3, q = 5$, with hyperedges shaded in gray. The gray triangles together with the white triangles form a simplicial subdivision. The lists of admissible colors are given on the boundary; for internal vertices the lists are all $\{1, 2, 3\}$.

The reader may notice that this is a setup reminiscent of Sperner's Lemma [11]. (See Figure 1.) However, Sperner's Lemma concerns *subdivisions* of the simplex; $H_{k,q}$ is not a subdivision of the simplex since its hyperedges viewed as geometric cells do not cover the full volume of $\Delta_{k,q}$.

Simplicial subdivisions. A simplicial subdivision of $\Delta_{k,q}$ is a collection of simplices ("cells") Σ such that

- The union of the cells in Σ is the simplex $\Delta_{k,q}$.
- For any two cells $\sigma_1, \sigma_2 \in \Sigma$, their intersection is either empty or a full face of a certain dimension shared by σ_1, σ_2.

We describe a concrete subdivision of $\Delta_{k,q}$ in Section 6.

2.2 Colorings of simplicial subdivisions. First, let us recall the statement of Sperner's Lemma [11]. We call a labeling $\ell : V \to [k]$ Sperner-admissible if $\ell(\mathbf{a}) \in L(\mathbf{a})$ for each $\mathbf{a} \in V$; i.e. , if $\ell(\mathbf{a}) = j$ then $a_j > 0$.

Lemma 2.1. (Sperner's Lemma) *For every Sperner-admissible labeling of the vertices of a simplicial subdivision of $\Delta_{k,q}$, there is a cell whose vertices receive all k colors.*

We remark that this does not say anything about the Simplex-Lattice Hypergraph: Even if the subdivision uses the point set $V_{k,q}$, the rainbow cell given by Sperner's Lemma might not be a member of $E_{k,q}$ since $E_{k,q}$ consists only of scaled copies of $\Delta_{k,q}$ without rotation; it is not a full subdivision of the simplex. (See Figure 2.)

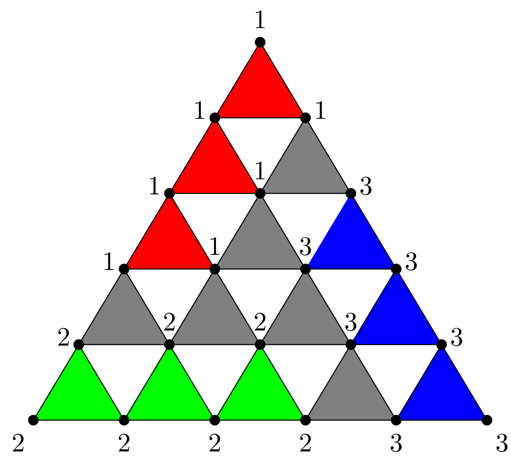

Figure 2: A Sperner-admissible labeling for $k = 3$ and $q = 5$. The set E of hyperedges consists of the shaded triangles. The gray triangles are non-monochromatic hyperedges. At least one triangle (not necessarily in E) must be 3-colored (rainbow).

Instead of rainbow cells, the statement proposed (and proved for $k = 3$) in [5] involves non-monochromatic cells.

Proposition 2.1. (Simplex-Lattice Coloring Lemma) *For any Sperner-admissible labeling $\ell : V_{k,q} \to [k]$, there are at least $\binom{q+k-3}{k-2}$ hyperedges $e \in E_{k,q}$ that are non-monochromatic under ℓ.*

The first-choice labeling. In particular, the proposition is that a Sperner-admissible labeling minimizing the number of non-monochromatic cells is a "first-choice one" which labels each vertex $\mathbf{a}$ by the smallest coordinate i such that $a_i > 0$. Under this labeling, all the hyperedges $e(\mathbf{b})$ such that $b_1 > 0$ are labeled monochromatically by 1. The only hyperedges that receive more than 1 color are those where $b_1 = 0$, and the number of such hyperedges is exactly $\binom{q+k-3}{k-2}$ (see [5]).

3 A proof of the Simplex-Lattice Coloring Lemma

Here we give a proof of Proposition 2.1.

Proof. Consider the set of hyperedges $E_{k,q}$: observe that it can be written naturally as

$$E_{k,q} = \{e(\mathbf{b}) : \mathbf{b} \in V_{k,q-1}\}.$$

I.e., the hyperedges can be identified one-to-one with the vertices in $V_{k,q-1}$. Recall that $e(\mathbf{b}) = \{\mathbf{b}+\mathbf{e}_1, \mathbf{b}+\mathbf{e}_2, \ldots, \mathbf{b}+\mathbf{e}_k\}$. Two hyperedges $e(\mathbf{b}), e(\mathbf{b}')$ share a vertex if and only if $\mathbf{b}'+\mathbf{e}_j = \mathbf{b}+\mathbf{e}_i$ for some pair $i, j \in [k]$; or in other words if $\mathbf{b}, \mathbf{b}'$ are nearest neighbors in $V_{k,q-1}$ (differ by ± 1 in exactly two coordinates).

Consider a labeling $\ell : V_{k,q} \to [k]$. For each $i \in [k]$, let C_i denote the set of points in $V_{k,q-1}$ representing the monochromatic hyperedges in color i,

$$C_i = \{\mathbf{b} \in V_{k,q-1} : \forall \mathbf{v} \in e(\mathbf{b}); \ell(\mathbf{v}) = i\}.$$

Define an injective mapping $\phi_i : C_i \to V_{k,q-2}$ as follows:

$$\phi_i(\mathbf{b}) = \mathbf{b} - \mathbf{e}_i.$$

The image is indeed in $V_{k,q-2}$: if $\mathbf{b} \in C_i$, we have $b_i > 0$, or else $e(\mathbf{b})$ would contain a vertex $\mathbf{a}$ such that $a_i = 0$ and hence $e(\mathbf{b})$ could not be monochromatic in color i. Therefore, $\mathbf{b} - \mathbf{e}_i \in \mathbb{Z}_+^k$ and $(\mathbf{b} - \mathbf{e}_i) \cdot \mathbf{1} = q-2$ which means $\mathbf{b} - \mathbf{e}_i \in V_{k,q-2}$.

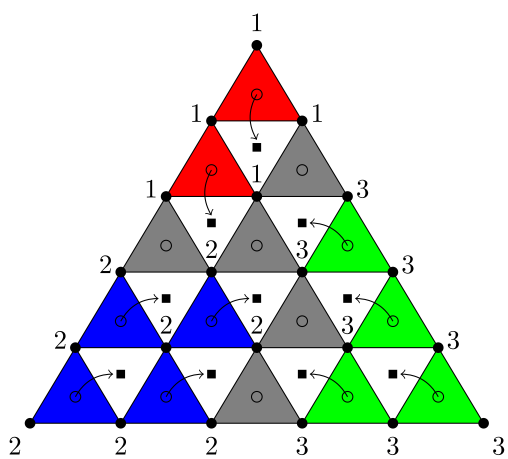

Figure 3: The mappings $\phi_i : C_i \to V_{k,q-2}$. The hyperedges are represented by the empty circles; C_i is the subset of them monochromatic in color i. The black squares represent $V_{k,q-2}$; note that each point in $V_{k,q-2}$ is the image of at most one monochromatic hyperedge.

Further, we claim that $\phi_i[C_i] \cap \phi_j[C_j] = \emptyset$ for every $i \neq j$. If not, there would be $\mathbf{b} \in C_i$ and $\mathbf{b}' \in C_j$ such that $\mathbf{b} - \mathbf{e}_i = \mathbf{b}' - \mathbf{e}_j$. Then, the point $\mathbf{a} = \mathbf{b} + \mathbf{e}_j = \mathbf{b}' + \mathbf{e}_i$ would be an element of both the hyperedge $e(\mathbf{b})$ and the hyperedge $e(\mathbf{b}')$. This contradicts the assumption that $e(\mathbf{b})$ is monochromatic in color i and $e(\mathbf{b}')$ is monochromatic in color j. So the sets $\phi_i[C_i]$ are pairwise disjoint subsets of $V_{k,q-2}$. By the definition of ϕ_i, we clearly have $|\phi_i[C_i]| = |C_i|$. We conclude that the total number of monochromatic hyperedges is

$$\sum_{i=1}^{k} |C_i| = \sum_{i=1}^{k} |\phi_i[C_i]| \leq |V_{k,q-2}|.$$

The total number of hyperedges is $|E_{k,q}| = |V_{k,q-1}|$. Considering that $|V_{k,q}| = \binom{q+k-1}{k-1}$ (the number of partitions of q into a sum of k nonnegative integers), we obtain that the number of non-monochromatic hyperedges is

$$|E_{k,q}| - \sum_{i=1}^{k} |C_i| \geq |V_{k,q-1}| - |V_{k,q-2}|$$
$$= \binom{q+k-2}{k-1} - \binom{q+k-3}{k-1} = \binom{q+k-3}{k-2}.$$

4 A labeling of $H_{k,q}$ with at most 4 colors on each hyperedge

We recall that Sperner's lemma states that any Sperner-admissible labeling of a subdivision of the simplex must contain a simplex with all k colors. The hypergraph $H_{k,q}$ defined in Section 2.1 is not a subdivision since it covers only a subset of the large simplex. It is easy to see that the conclusion of Sperner's lemma does not hold for $H_{k,q}$ — for example for $k = 3$, we can label a 2-dimensional triangulation so that exactly one triangle has 3 different colors, and this triangle is not in $E_{3,q}$. (See Figure 2.) Hence, each triangle in $E_{3,q}$ has at most 2 colors. By an extension of this argument, we can label $H_{k,q}$ so that each hyperedge in $E_{k,q}$ contains at most $k-1$ colors. The question we ask in this section is, what is the minimum ℓ^* such that there is a Sperner-admissible labeling with at most ℓ^* colors on each hyperedge in $E_{k,q}$? We prove the following result.

Lemma 4.1. *For any $k \geq 4$ and $q \geq k^2$, there is a Sperner-admissible labeling of $H_{k,q} = (V_{k,q}, E_{k,q})$ such that every hyperedge in $E_{k,q}$ contains at most 4 colors.*

Proof. We define a labeling $\ell : V_{k,q} \to [k]$ as follows:

- Given $\mathbf{a} \in V_{k,q}$, let $\pi : [k] \to [k]$ be a permutation such that $a_{\pi(1)} \geq a_{\pi(2)} \geq \ldots \geq a_{\pi(k)}$ (and if $a_{\pi(i)} = a_{\pi(i+1)}$, we order π so that $\pi(i) < \pi(i+1)$).
- Define $t(\mathbf{a})$ to be the maximum $t \in [k]$ such that $\forall 1 \leq j \leq t$, $a_{\pi(j)} \geq k - j + 1$. We define the "Top coordinates" of $\mathbf{a}$ to be $Top(\mathbf{a}) = (\pi(1), \ldots, \pi(t(\mathbf{a})))$ (an ordered set).

- We define the color of $\mathbf{a}$ to be $\ell(\mathbf{a}) = \pi(t(\mathbf{a}))$, the index of the "last Top coordinate".

First, we verify that this is a well-defined Sperner-admissible labeling. Since $\sum_{i=1}^{k} a_i = q \geq k^2$, we have $a_{\pi(1)} = \max a_i \geq k$ and hence $1 \leq t(\mathbf{a}) \leq k$. For each $\mathbf{a} \in V_{k,q}$, we have: $a_{\ell(\mathbf{a})} = a_{\pi(t(\mathbf{a}))} \geq k - t(\mathbf{a}) + 1 > 0$, since $t(\mathbf{a}) \leq k$. Therefore, ℓ is Sperner-admissible.

Now, consider a hyperedge $e(\mathbf{b}) = (\mathbf{b} + \mathbf{e}_1, \mathbf{b} + \mathbf{e}_2, \ldots, \mathbf{b} + \mathbf{e}_k)$ where $\mathbf{b} \geq 0, \sum_{i=1}^{k} b_i = q - 1$. We claim that $\ell(\mathbf{b} + \mathbf{e}_i)$ attains at most 4 different values for $i = 1, \ldots, k$. Without loss of generality, assume that $b_1 \geq b_2 \geq \ldots \geq b_k$. Define ℓ^* to be the color that would be assigned to $\mathbf{b}$ (if $\mathbf{b}$ were a vertex in $V_{k,q}$; in fact it is not but we can still apply our definition): ℓ^* is the maximum value in $[k]$ such that for all $1 \leq j \leq \ell^*$, $b_j \geq k - j + 1$. Hence, we have $Top(\mathbf{b}) = \{1, 2, \ldots, \ell^*\}$.

Let $i \in [k]$, $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$, and let π be the permutation such that $a_{\pi(1)} \geq \ldots \geq a_{\pi(k)}$ as above. (Recall that for $\mathbf{b}$, we assumed that the respective permutation is the identity.) We consider the following cases:

- If $1 \leq i < \ell^*$, then we claim that $\ell(\mathbf{a})$ for $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$ is equal to $\ell(\mathbf{b} + \mathbf{e}_i) = \ell^*$. In the rule for selecting $t(\mathbf{a})$, one of the first $\ell^* - 1$ coordinates has been incremented compared to $\mathbf{b}$, which possibly pushes i forward in the ordering of the Top coordinates. However, the other coordinates remain unchanged, the condition $a_{\pi(j)} \geq k - j + 1$ is still satisfied for $1 \leq j \leq \ell^*$, and $Top(\mathbf{a}) = Top(\mathbf{b})$. In particular ℓ^* is still the last coordinate included in $Top(\mathbf{a})$ and hence $\ell(\mathbf{a}) = \ell^*$.

- If $i = \ell^*$, then $\ell(\mathbf{a}) = \ell(\mathbf{b} + \mathbf{e}_{\ell^*})$ is still one of the coordinates in $Top(\mathbf{b})$, possibly different from ℓ^* (due to a change in order, although we still have $Top(\mathbf{a}) = Top(\mathbf{b})$) — let us call this color ℓ_2^*.

- If $\ell^* < i \leq k$, then it is possible that in $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$, we obtain additional Top coordinates ($Top(\mathbf{a}) \supset Top(\mathbf{b})$). First of all, it could be $a_i = b_i + 1$ itself which is now included among the Top coordinates, and possibly additional coordinates that previously satisfied the condition $b_j \geq k - j + 1$ but were not selected due to the condition being false for b_{ℓ^*+1}. If this does not happen and we have $Top(\mathbf{a}) = Top(\mathbf{b})$, the color of $\mathbf{a}$ is still $\ell(\mathbf{a}) = \ell^*$ (because the ordering of the Top coordinates remains the same).

  Assume now that $\mathbf{a}$ has additional Top coordinates beyond $Top(\mathbf{b})$. Observe the following: by the definition of ℓ^*, there is no coordinate $j > \ell^*$ such that $b_j = b_{\ell^*} - 1$; otherwise j would have been still chosen in $Top(\mathbf{b})$. The only way $Top(\mathbf{b} + \mathbf{e}_i)$ can grow beyond $Top(\mathbf{b})$ is that $b_i = b_{\ell^*} - 2$ and $a_i = b_i + 1 = b_{\ell^*} - 1$. In this case, we have $i = \pi(\ell^* + 1)$, because a_i is the maximum coordinate among $\{a_j : j > \ell^*\}$, and still smaller than a_{ℓ^*}. Therefore, since we have $b_{\ell^*} \geq k - \ell^* + 1$ (by inclusion of ℓ^* in $Top(\mathbf{b})$), we also have $a_i \geq k - \ell^* \geq k - \pi^{-1}(i) + 1$; hence, i will be included in $Top(\mathbf{a})$. Now, $Top(\mathbf{a})$ may grow further; it will include at least all the coordinates j such that $b_j = b_{\ell^*} - 2$, by the same argument. However, note that the construction of $Top(\mathbf{a})$ will proceed in the same way for every $\mathbf{b} + \mathbf{e}_i$ such that $b_i = b_{\ell^*} - 2$. This is because all the coordinates equal to $b_{\ell^*} - 2$ will be certainly included in $Top(\mathbf{a})$, and coordinates smaller than $b_{\ell^*} - 2$ remain the same in each of these cases (equal to the coordinates of $\mathbf{b}$). Therefore, the set $Top(\mathbf{a})$ will be the same in all these cases; let us call this set Top_+.

  The color assigned to $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$ is the index of the last coordinate included in Top_+. Since Top_+ is the same whenever $b_i = b_{\ell^*} - 2$, this will be the coordinate j^* minimizing b_j (and maximizing j to break ties) among all $j \in Top_+$, unless $i = j^*$ in which case the last included coordinate might be another one. This gives potentially two additional colors, let us call them ℓ_3^*, ℓ_4^*, that are assigned to $\mathbf{a} = \mathbf{b} + \mathbf{e}_i$ for all $i > \ell^*$ where $b_i = b_{\ell^*} - 2$. For other choices of $i > \ell^*$, we have $Top(\mathbf{b} + \mathbf{e}_i) = Top(\mathbf{b})$ and the color assigned to $\mathbf{b} + \mathbf{e}_i$ is $\ell(\mathbf{b} + \mathbf{e}_i) = \ell^*$.

To summarize, all the colors that appear in the labeling of $e(\mathbf{b})$ are included in $\{\ell^*, \ell_2^*, \ell_3^*, \ell_4^*\}$.

5 Applications to Hypergraph Labeling

In this section, we discuss several labeling problems on hypergraphs and the consequences that our results have for them. The basic problem that we study in this paper is the following.

Hypergraph Labeling with Color Lists. *Given a hypergraph $H = (V, E)$ with weights $w(e), e \in E$ and color lists $L(v) \subseteq [k], v \in V$, find a labeling $\ell(v) \in L(v)$ for each $v \in V$ that minimizes the total weight of non-monochromatic hyperedges.*

A related problem was proposed by Chekuri and Ene [2] as a generalization of the Uniform Metric Labeling problem of Kleinberg and Tardos [7].

Hypergraph Labeling. *Given a hypergraph $H = (V, E)$ with edge weights $w(e) \geq 0$ and vertex assignment costs $c(v, i) \geq 0$, find a labeling $\ell : V \to [k]$ minimizing the total assignment cost plus the total weight of* *hyperedges that receive more than 1 label:*

$$\sum_{v \in V} c(v, \ell(v)) + \sum_{e \in E : |\ell[e]| > 1} w(e).$$

Clearly, this problem subsumes Hypergraph Labeling with Color Lists, by considering $w(e) = 1$ and assignment costs in $\{0, \infty\}$ (or a huge finite value instead of ∞). In fact, Hypergraph Labeling with Color Lists is equivalent to Hypergraph Labeling by the following reduction (attributed to Julia Chuzhoy by [1]): Given an instance of Hypergraph Labeling, for each $v \in V$ and each $i \in [k]$, create a new vertex (v, i) with the color list $L((v, i)) = [k] \setminus \{i\}$. Place an edge of weight $c(v, i)$ between (v, i) and v. Vertex v gets the color list $L(v) = [k]$. Then it is easy to see that if vertex v is labeled ℓ, the optimal labeling of the vertices $(v, i), i \in [k]$ is such that we will pay exactly $c(v, \ell)$ for cutting the edge between v and (v, ℓ). Hence, Hypergraph Labeling with Color Lists and Hypergraph Labeling are approximation-equivalent.

A simplex coloring lemma was conjectured in [5] with the aim of proving hardness of approximation for Hypergraph Labeling. We prove this lemma in this paper (Proposition 2.1). In effect, Proposition 2.1 shows an integrality gap arbitrarily close to $k-1$ for a certain LP relaxation of the Hypergraph Labeling problem. Using the general reduction of [6], we obtain the following (more details about are given in Section 5.2 below).

Corollary 5.1. *Assuming the Unique Games Conjecture, it is NP-hard to achieve a $(k-1-\epsilon)$-approximation for the* Hypergraph Labeling *problem on k-uniform hypergraphs with label set $[k]$, for any fixed $\epsilon > 0$.*

Chekuri and Ene [2] gave a Δ-approximation for Hypergraph Labeling with hyperedge size bounded by Δ. The algorithm that achieves this is the Kleinberg-Tardos algorithm for Metric Labeling [7]. By a more careful analysis, we show that the same algorithm also achieves a $(k-1)$-approximation whenever the label set is $[k]$ (see Section 5.3).

Theorem 5.1. *There is a $(k-1)$-approximation for the* Hypergraph Labeling *problem with label set $[k]$.*

Hence, this approximation algorithm is optimal in terms k and almost tight (up to an additive 1) in terms of Δ. We remark that we do not expect a $(\Delta - 1)$-approximation for Hypergraph Labeling: The special case of $\Delta = 2$ is the Uniform Metric Labeling problem: This problem admits a 2-approximation [7] and it is Unique-Games-hard to achieve a $(2 - \epsilon)$-approximation [8].

In the following, we give some more details behind these results.

5.1 LP relaxations of Hypergraph Labeling.

Chekuri and Ene [2] gave a linear-programming relaxation called LE-Rel for partitioning problems involving submodular functions, based on the Lovász extension of a submodular function. In particular, for Hypergraph Labeling the resulting LP reads as follows (see [5] for more discussion).

LE-Rel for Hypergraph Labeling

$$\min \sum_{v \in V} \sum_{i=1}^{k} c(v, i) x_{v,i} + \sum_{e \in E} w(e) \left(1 - \sum_{i=1}^{k} \min_{v \in e} x_{v,i}\right)$$

$$\sum_{i=1}^{k} x_{v,i} = 1 \quad \forall v \in V$$

$$x_{v,i} \geq 0 \quad \forall v \in V, i \in [k]$$

Formally, this is not in the form of a linear program but it is easy to see that the expression $\min_{v \in e} x_{v,i}$ can be replaced by a new variable $z_{e,i}$ with constraints $z_{e,i} \leq x_{v,i} \forall v \in e$. We prefer to keep the form above for compactness.

This LP is equivalent to the “Local Distribution LP” for Min-CSP problems considered in [6]. In the Local Distribution LP, we have $x_{v,i}$ variables as above, and also $y_{e,\alpha}$ variables for each hyperedge $e \in E$ and each possible assignment $\alpha \in [k]^e$. The hyperedge variables $y_{e,\alpha}$ can be interpreted as a distribution over labelings of the respective hyperedge e. The hyperedge variables must be consistent with the vertex variables in the sense that all assignments such that $\alpha_v = i$ should add up to $\sum_{\alpha \in [k]^e : \alpha_v = i} y_{e,\alpha} = x_{v,i}$. The Local Distribution LP reads as follows.

Local Distribution LP

$$\min \sum_{v \in V} \sum_{i=1}^{k} c(v, i) x_{v,i} + \sum_{e \in E, \alpha \in [k]^e} w(e) y_{e,\alpha} \Phi_e(\alpha)$$

$$\sum_{\alpha \in [k]^e, \alpha_v = i} y_{e,\alpha} = x_{v,i} \quad \forall v \in e \in E, i \in [k]$$

$$\sum_{i=1}^{k} x_{v,i} = 1 \quad \forall v \in V$$

$$x_{v,i}, y_{e,\alpha} \geq 0 \quad \forall v \in V, i \in [k], e \in E, \alpha \in [k]^e$$

In particular, the Local Distribution LP for Hypergraph Labeling is obtained by using the cost function $\Phi_e(\alpha) = 0$ if $\alpha = (i, i, \ldots, i)$ for some $i \in [k]$, and $\Phi_e(\alpha) = 1$ otherwise (which is the hypergraph cut function for a single hyperedge). We refer the reader to [5] for a proof that these two LPs are equivalent for Hypergraph Labeling, in the sense that given a feasible assignment of the variables $x_{v,i}$, the optimal assignment of the variables $y_{e,\alpha}$ in the Local Distribution LP is one that achieves exactly the objective value of LE-Rel .

5.2 Hardness of approximation of Hypergraph Labeling. In this section, we explain the connection between our combinatorial results and hardness of approximation for hypergraph labeling problems. This connection is based on a hardness reduction for Min-CSP recently developed in [6], which in turn builds on a hardness reduction for Multiway Cut problems discovered in [8]. A Min-CSP instance consists of a hypergraph $H = (V, E)$ with weights w_e and predicate cost functions $\Phi_e : [k]^e \to [0, 1]$ for $e \in E$. The goal is to find an assignment $\ell : V \to [k]$ that minimizes $\sum_{(v_{i_1},\dots,v_{i_j})=e\in E} w_e \Phi_e(\ell(v_{i_1}), \dots, \ell(v_{i_j}))$. Observe that the hypergraph labeling problems considered in this paper are exactly in this form. (Except for the list-coloring constraint $\ell(v) \in L(v)$, which can be simulated by a unary predicate $\Phi_{\{v\}}(\ell) = 0$ iff $\ell \in L(v)$ and $\Phi_{\{v\}}(\ell) = 1$ otherwise, with a prohibitively large weight $w_{\{v\}}$.)

The hardness reduction of [6] takes any integrality gap instance for a problem of a certain type and turns it into a hardness of approximation result for the same problem. We summarize this reduction in the following theorem. We denote by NAE_k the Not-All-Equal predicate on k variables, $NAE_k(x_1, \dots, x_k) = 0$ if $x_1 = x_2 = \dots = x_k$ and 1 otherwise. We call NAE_2 the Not-Equal predicate.

THEOREM 5.2. *[6] Suppose $\mathcal{I}$ is a Min-CSP instance including the Not-Equal predicate $NAE_2(x_1, x_2)$. Assume the Unique Games Conjecture. If the optimum value of $\mathcal{I}$ is s and the optimum value of its Local Distribution LP is c where $s > c \geq 0$, then for any $\epsilon > 0$, it is NP-hard to distinguish between instances whose optimal value is at least $s - \epsilon$ and those whose optimal value is at most $c + \epsilon$.*

Note that the reduction works only for Min-CSP problems involving the Not-Equal predicate; however, this is exactly the predicate appearing in problems where we want to avoid "cutting" edges or hyperedges, such as the problems considered in this paper. Let us outline now how our hardness results follow from this reduction.

The coloring conjecture of [5] was proposed with the aim of designing an integrality gap instance for Hypergraph Labeling. The instance is exactly the hypergraph with color lists which is the subject of Proposition 2.1. We interpret the color lists as unary predicates with cost 0 for an admissible color and ∞ for a forbidden color. Note that the predicate on each hyperedge is the Not-All-Equal (NAE_k) predicate — equal to 1 unless all k variables are equal, in which case it is 0. The intended fractional solution is simply $\mathbf{x}_v = \frac{1}{q}\mathbf{v}$. The LE-Rel relaxation pays a cost of $1/q$ for each hyperedge, for a total cost of $c = \frac{1}{q}\binom{q+k-2}{k-1}$. According to Proposition 2.1, the cost of the optimum solution is $s = \binom{q+k-3}{k-2}$. The ratio of these two quantities tends to $k - 1$ as $q \to \infty$. Therefore, Theorem 5.2 implies that it is Unique-Games-hard to achieve a $(k - 1 - \epsilon)$-approximation for any constant $\epsilon > 0$, for Hypergraph Labeling on hypergraphs with edges of size at most k and label set $[k]$.

Note that Theorem 5.2 requires the $NAE_2(x, y)$ predicate to be part of the instance. We can include this predicate, if we allow hyperedges of size 2 to be part of the hard instance. Alternatively, if we want to obtain a k-uniform hypergraph, we can simulate NAE_2 by NAE_k predicates as follows: we add $k - 2$ dummy vertices $d_1, \dots, d_{k-2}$ for each $NAE_2(x, y)$ constraint, and we replace $NAE_2(x, y)$ by $NAE_k(x, y, d_1, \dots, d_{k-2})$. It is easy to see that this does not change the value of the optimum. This proves Corollary 5.1.

In addition, we observe the following. The hard instances arising from the hardness reduction of [6] can be viewed again as instances of Hypergraph Labeling with Color Lists, with hyperedges of two types. Hyperedges of the first type (corresponding to the "edge test") have the structure of hyperedges of $H_{k,q}$. According to Lemma 4.1, the hypergraph $H_{k,q}$ Hypergraph Labeling with Color Lists can be labeled in such a way that every hyperedge uses at most 4 colors. If we label the vertices of $\mathcal{I}$ in accordance with the labeling of $H_{k,q}$, each hyperedge of the first type will again use at most 4 colors. Hyperedges of the second type (corresponding to the "vertex test") contain only 2 vertices (or are converted into hyperedges of size k using the construction above). Clearly, these hyperedges use at most 2 colors. Therefore, the hardness result of Corollary 5.1 holds even for instances of Hypergraph Labeling with Color Lists that can be labeled in such a way that every hyperedges uses at most 4 colors.

5.3 $(k-1)$-approximation for Hypergraph Labeling. Chekuri and Ene [2] presented a Δ-approximation algorithm for Hypergraph Labeling, where Δ is the maximum size of a hyperedge. This algorithm solves the LE-Rel relaxation and then rounds the fractional solution to an integral one, using the randomized rounding technique of Kleinberg and Tardos [7].

We show here that this rounding technique also achieves a $(k-1)$-approximation for Hypergraph Labeling where $[k]$ is the label set and $k \geq 3$, thus proving Theorem 5.1. (For $k = 2$, the problem can be solved exactly as a special case of submodular minimization.)

Algorithm 1 Kleinberg-Tardos-Rounding($\mathbf{x}_v : v \in V$)

```
S_1, ..., S_k ← ∅
while ⋃_{i=1}^k S_i ≠ V do
  pick i ∈ [k] uniformly at random
  pick λ ∈ [0,1] uniformly at random
  for all v ∈ V \ ⋃_{i=1}^k S_i do
    if x_{v,i} > λ then
      S_i ← S_i ∪ {v}
    end if
  end for
end while
return ℓ : V → [k] where ℓ(v) = i whenever v ∈ S_i
```

We analyze the Kleinberg-Tardos rounding procedure in a sequence of claims (building upon the analysis of [2]).

LEMMA 5.1. *The expected assignment cost of vertex $v \in V$ is exactly $\sum_{i=1}^{k} c(v,i)x_{v,i}$.*

Proof. It is known that the probability that vertex v is labeled $\ell(v) = i$ by Kleinberg-Tardos rounding is exactly $x_{v,i}$ [7]. Hence the expected assignment cost is $\sum_{i=1}^{k} c(v,i)x_{v,i}$.

DEFINITION 5.1. *For a hyperedge $e \in E$, we say that a*

- *$Capture_e(i)$ event happens if i is the index chosen by the algorithm and $\lambda < \min_{v\in e} x_{v,i}$.*
- *$Cut_e(i)$ event happens if i is the index chosen by the algorithm and $\min_{v\in e} x_{v,i} \le \lambda < \max_{v\in e} x_{v,i}$.*
- *$Touch_e(i)$ event happens if i is the index chosen by the algorithm and $\lambda < \max_{v\in e} x_{v,i}$; $Touch_e(i) = Capture_e(i) \vee Cut_e(i)$.*

LEMMA 5.2. *A hyperedge $e \in E$ ends up monochromatic, unless the first event that happens for e is $Cut_e(i)$ for some $i \in [k]$, and the second event that happens is* not *$Capture_e(i)$.*

Proof. If the first event that happens for e is $Capture_e(i)$ for some $i \in [k]$, then e becomes immediately monochromatic in color i. If the first event is $Cut_e(i)$, then some vertices of e are labeled i but not all. There must happen at least one more event for e before the algorithm terminates. If this event is $Capture_e(i)$ then e becomes monochromatic in color i. Therefore the only way to become non-monochromatic is that the second event is not $Capture_e(i)$.

We note that considering the first event for each hyperedge would be sufficient to derive a k-approximation for Hypergraph Labeling. However, to get the (optimal) $(k-1)$-approximation we have to be more careful and that is the reason for considering the second event.

LEMMA 5.3. *The probability that the first event for $e \in E$ is $Cut_e(i)$ is*

$$Pr[Cut_e(i) \mid \bigvee_{j=1}^{k} Touch_e(j)] = \frac{\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i}}{\sum_{j=1}^{k} \max_{v\in e} x_{v,j}}.$$

Proof. The successive rounds are independent. Therefore, we can analyze the first event for $e \in E$ by conditioning on the event $\bigvee_{j=1}^{k} Touch_e(j)$. This event happens iff the chosen threshold is below $\max_{v\in e} x_{v,j}$, where $j \in [k]$ is random; hence

$$\Pr[\bigvee_{j=1}^{k} Touch_e(j)] = \frac{1}{k}\sum_{j=1}^{k} \max_{v\in e} x_{v,j}.$$

The event $Cut_e(i)$ happens iff the chosen coordinate is i and $\min_{v\in e} x_{v,i} \le \lambda < \max_{v\in e} x_{v,i}$. This happens with probability

$$\Pr[Cut_e(i)] = \frac{1}{k}(\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i}).$$

Since this is a sub-event of $\bigvee_{j=1}^{k} Touch_e(j)$, we obtain the lemma.

LEMMA 5.4. *Conditioned on the first event for $e \in E$ being $Cut_e(i)$, the probability that the second event is $Capture_e(i)$ is*

$$Pr[Capture_e(i) \mid \bigvee_{j=1}^{k} Touch_e(j)] = \frac{\min_{v\in e} x_{v,i}}{\sum_{j=1}^{k} \max_{v\in e} x_{v,j}}.$$

Proof. If the first event was $Cut_e(i)$ then another event must still happen because not all vertices of e have been labeled. The next event is independent of what happened before and again we can condition on $\bigvee_{j=1}^{k} Touch_e(j)$ as above. The event $Capture_e(i)$ happens iff the chosen coordinate is i and $\lambda < \min_{v\in e} x_{v,i}$. This happens with probability

$$\Pr[Capture_e(i)] = \frac{1}{k}\min_{v\in e} x_{v,i}.$$

Since this is a sub-event of $\bigvee_{j=1}^{k} Touch_e(j)$ and $\Pr[\bigvee_{j=1}^{k} Touch_e(j)] = \frac{1}{k}\sum_{j=1}^{k} \max_{v\in e} x_{v,j}$, we obtain the lemma.

LEMMA 5.5. *The probability that a hyperedge e ends up non-monochromatic is at most*

$$\sum_{i=1}^{k} \frac{\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i}}{\sum_{j=1}^{k} \max_{v\in e} x_{v,j}} \left(1 - \frac{\min_{v\in e} x_{v,i}}{\sum_{j=1}^{k} \max_{v\in e} x_{v,j}}\right).$$

Proof. We combine the previous three lemmas: The probability of the first event being $Cut_e(i)$ is $\frac{\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i}}{\sum_{j=1}^{k} \max_{v\in e} x_{v,j}}$, and conditioned on that, the probability of the second event *not* being $Capture_e(i)$ is $1 - \frac{\min_{v\in e} x_{v,i}}{\sum_{j=1}^{k} \max_{v\in e} x_{v,j}}$. In other cases, the hyperedge ends up monochromatic.

LEMMA 5.6. *Let* $k \geq 3$ *and* $cost(e) = 1 - \sum_{j=1}^{k} \min_{v\in e} x_{v,j}$. *Then the probability that e ends up non-monochromatic is at most* $(k-1)cost(e)$.

Proof. Let us analyze the maximum possible value that the bound provided by Lemma 5.5 could achieve. Observe that $\sum_{i=1}^{k} \max_{v\in e} x_{v,i} \geq 1$ because for any fixed vertex, $\sum_{i=1}^{k} x_{v,i} = 1$. Let $\kappa \geq 0$ be such that $\sum_{i=1}^{k} \max_{v\in e} x_{v,i} = 1 + \kappa \cdot cost(e)$. By Lemma 5.5 (after discarding some factors smaller than 1), the probability of e becoming non-monochromatic is at most

$$\begin{aligned} & \sum_{i=1}^{k} (\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i}) \\ = \ & (\sum_{i=1}^{k} \max_{v\in e} x_{v,i} - 1) + (1 - \sum_{i=1}^{k} \min_{v\in e} x_{v,i}) \\ = \ & \kappa \cdot cost(e) + cost(e) = (\kappa+1)cost(e). \end{aligned}$$

Observe that $\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i} \leq cost(e)$ for every $i \in [k]$, because $cost(e) = \sum_{j=1}^{k} (x_{w,j} - \min_{v\in e} x_{v,j})$ for the vertex w achieving $\max_{v\in e} x_{v,i}$. Therefore, $\kappa \leq k-1$. On the other hand, if $\kappa \leq k-2$, then we are done because the bound above is at most $(k-1)cost(e)$. So we can assume $\kappa \in (k-2, k-1]$.

Assume without loss of generality that the coordinates are ordered so that

$$\max_{v\in e} x_{v,1} - \min_{v\in e} x_{v,1} \geq \ldots \geq \max_{v\in e} x_{v,k} - \min_{v\in e} x_{v,k}.$$

Recall that $\sum_{i=1}^{k} \min_{v\in e} x_{v,i} = 1 - cost(e)$. Considering how this sum could be distributed among the k terms, the worst case for the bound of Lemma 5.5 is that $\min_{v\in e} x_{v,i} = 0$ for all $i < k$, and $\min_{v\in e} x_{v,k} = 1 - cost(e)$. Similarly, we have $\sum_{i=1}^{k} (\max_{v\in e} x_{v,i} - \min_{v\in e} x_{v,i}) = (\kappa+1)cost(e)$, and the worst case is that the first $k-1$ terms are as large as possible, which is $cost(e)$, while the last term is $(\kappa+1)cost(e) - (k-1)cost(e) = (\kappa - k + 2)cost(e)$. In this case, the bound of Lemma 5.5 would become

$$\begin{aligned} & \frac{cost(e)}{1+\kappa cost(e)} \left((k-1) + (\kappa - k + 2) \left(1 - \frac{1-cost(e)}{1+\kappa cost(e)} \right) \right) \\ & = \frac{cost(e)}{1+\kappa cost(e)} \left((k-1) + (\kappa - k + 2) \frac{(\kappa+1)cost(e)}{1+\kappa cost(e)} \right). \end{aligned}$$

We assumed that $k \geq 3$ and $\kappa \in (k-2, k-1]$. Using this, $\kappa - k + 2 \in (0, 1]$ and $(\kappa+1)cost(e) \leq 2\kappa cost(e) \leq \kappa(k-1)cost(e)$. So the above bound is at most $(k-1)cost(e)$.

This completes the proof of Theorem 5.1.

6 Rainbow-avoiding Hypergraph Labeling

Further, we consider the following extension of Hypergraph Labeling with Color Lists.

Hypergraph j-Colors-Avoiding Labeling with Color Lists. *Given a hypergraph* $H = (V, E)$ *with color lists* $L(v) \subseteq [k], v \in V$, *and a parameter* $j \geq 2$, *find a labeling* $\ell(v) \in L(v)$ *minimizing the number of hyperedges that receive at least* j *distinct labels.*

In particular, Hypergraph Labeling with Color Lists is the case of $j = 2$. If $j = k$ then the goal is to minimize the number of *rainbow hyperedges*, containing all k colors. We call this problem Hypergraph Rainbow-Avoiding Labeling with Color Lists.

6.1 LP relaxation of Hypergraph j-Colors-Avoiding Labeling with Color Lists. It is natural to use the Local Distribution LP with the appropriate predicate $\Phi_e^{(j)}(\alpha) = 1$ if $\alpha \in [k]^e$ uses at least j distinct labels, and 0 otherwise. However, we show that this LP performs very poorly and subsequently we prove a hardness result for this problem. Sperner's Lemma plays a role in this result.

LEMMA 6.1. *For any* $k \geq j = 3$, *there are instances of* Hypergraph j-Colors-Avoiding Labeling with Color Lists *such that the value of the Local Distribution LP is* 0 *but there is a rainbow hyperedge (containing all k colors) for any feasible labeling.*

In particular, as prove in Lemma 6.3, there are instances of Hypergraph j-Colors-Avoiding Labeling with Color Lists corresponding to a subdivision of a simplex, with a *fractional labeling* using locally at most 2 colors for each hyperedge. This fractional labeling corresponds to a fractional solution of zero cost for the Local Distribution LP for Hypergraph j-Colors-Avoiding Labeling with Color Lists. On the other hand, for any Sperner-admissible labeling of a simplicial subdivision, there must be a hyperedge using *all k* colors (by Sperner's Lemma). This will prove Lemma 6.1. Hence, the Local Distribution LP is unable to distinguish between instances that can be labeled so that each hyperedge contains at most 2 colors, and instances that must contain a rainbow hyperedge.

Fractional labelings of a subdivision of the simplex. To be more precise, we clarify what we mean by "fractional labelings" of a hypergraph H in the sense of the Local Distribution LP. A fractional labeling consists of a probability distribution for each vertex $v \in V$ over colors, $(x_{v1}, \ldots, x_{vk})$ where $x_{vi} \geq 0$ and $\sum_{i=1}^{k} x_{vi} = 1$. In addition, each hyperedge $e \in E$ receives a distribution over labelings $(y_{e\alpha} : \alpha \in [k]^e)$, $y_{e\alpha} \geq 0$, $\sum_{\alpha \in [k]^e} y_{e\alpha} = 1$. This distribution must be consistent with the vertex-based distributions in the sense that the probability of a vertex receiving a certain color should be the same under both distributions: $\sum_{\alpha:\alpha_v = j} y_{e\alpha} = x_{vj}$.

In this section, we are concerned with *subdivisions* of the simplex, i.e. hypergraphs that form a simplicial complex covering the full volume of the simplex. The hypergraph $H_{k,q} = (V_{k,q}, E_{k,q})$ introduced in Section 2 does not induce a subdivision since the respective cells do not cover the entire simplex. In the following, we describe a concrete subdivision of the simplex $\Delta_{k,q}$ whose vertex set is $V_{k,q}$ (following a construction of [9], see also [4]).

A regular simplicial subdivision of $\Delta_{k,q}$. Consider first a simplex parameterized as follows:

$$R_{k,q} = \{\mathbf{y} \in \mathbb{R}_+^{k-1} : 0 \leq y_1 \leq y_2 \leq \ldots \leq y_{k-1} \leq q\}.$$

We will describe a subdivision of $R_{k,q}$ with vertex set $W_{k,q} = \{\mathbf{v} \in \mathbb{Z}_+^{k-1} : 0 \leq v_1 \leq \ldots \leq v_{k-1} \leq q\}$. Once we complete this construction, we map the simplicial subdivision back to $\Delta_{k,q}$ by the following mapping: $\phi(\mathbf{y}) = \mathbf{x} \in \mathbb{R}^k$ where $x_i = y_i - y_{i-1}$ for $i \in [k]$, under the convention that $y_0 = 0$ and $y_k = q$. It is easy to verify that ϕ maps $R_{k,q}$ bijectively to $\Delta_{k,q}$, and $\phi[W_{k,q}] = V_{k,q}$.

A cell of the subdivision of $R_{k,q}$ is indexed by a vertex $\mathbf{w} \in W_{k,q-1}$ and a permutation $\pi : [k-1] \to [k-1]$. The permutation π should be *consistent* with $\mathbf{w}$ in the sense that whenever $w_i = w_{i+1}$, we have $\pi(i) < \pi(i+1)$. For any such pair $(\mathbf{w}, \pi)$, the respective cell is defined as

$$\begin{aligned}\sigma(\mathbf{w}, \pi) \quad = \quad & \{\mathbf{y} \in \mathbb{R}_+^{k-1} : 0 \leq (\mathbf{y} - \mathbf{w})_{\pi(1)} \leq (\mathbf{y} - \mathbf{w})_{\pi(2)} \\ & \leq \ldots \leq (\mathbf{y} - \mathbf{w})_{\pi(k-1)} \leq 1\}.\end{aligned}$$

Let us verify that this is a simplicial subdivision: First, for any π consistent with $\mathbf{w}$, we get that $\mathbf{y} \in \sigma(\mathbf{w}, \pi)$ must have coordinates ordered increasingly: either $w_i < w_{i+1}$ which means that $y_i \leq w_i + 1 \leq w_{i+1} \leq y_{i+1}$, or $w_i = w_{i+1}$ which means $y_i \leq y_{i+1}$ by the consistent ordering property. Also, it is easy to verify that $\mathbf{y} \in [0, q]^{k-1}$, since $\mathbf{w} \in \{0, 1, \ldots, q-1\}^{k-1}$. Hence $\sigma(\mathbf{w}, \pi) \subset R_{k,q}$.

For any point $\mathbf{y} \in R_{k,q}$, we have $0 \leq y_1 \leq y_2 \leq \ldots \leq y_{k-1} \leq q$. This point is contained in a cell $\sigma(\mathbf{w}, \pi)$ given by $w_i = \lfloor y_i \rfloor$ and π describing the ordering of the coordinates of $\mathbf{y} - \mathbf{w}$. Note that if $w_i = w_{i+1}$, the ordering is consistent with that of the coordinates of y_i; hence, this is a valid cell $\sigma(\mathbf{w}, \pi)$. The coordinates of $\mathbf{y} - \mathbf{w}$ are between $[0, 1]$ and ordered according to π, so we have indeed $\mathbf{y} \in \sigma(\mathbf{w}, \pi)$. On the other hand, the cell containing $\mathbf{y}$ is uniquely defined, except when some coordinates of $\mathbf{y} - \mathbf{w}$ are equal (which allows several consistent permutations) or when $\mathbf{y}$ has integer coordinates (which allows the choice of $w_i = y_i - 1$). The subsets satisfying these conditions form faces of various dimensions that are shared between cells.

Note that the cells $\sigma(\mathbf{w}, \pi)$ are congruent (obtained by translation and rotation of the same shape). However, when we map them from $R_{k,q}$ to $\Delta_{k,q}$ using the linear map ϕ, various shapes arise; the cells $\phi[\sigma(\mathbf{w}, \pi)]$ are no longer congruent. In fact we are not aware of any subdivision of $\Delta_{k,q}$ using the vertex set $V_{k,q}$ and congruent cells.

What is important in the following is not the particular form of our subdivision, but the properties spelled out in the following lemma.

LEMMA 6.2. *For any $k \geq 2, q \geq 1$, there is a subdivision of the simplex $\Delta_{k,q}$ with a set of simplicial cells $\Sigma_{k,q}$ such that*

- *the vertex set of the subdivision is $V_{k,q}$,*
- *the number of cells in $\Sigma_{k,q}$ is q^{k-1},*
- *each cell $\sigma \in \Sigma_{k,q}$ has the same volume, $\mu(\sigma) = q^{-(k-1)} \mu(\Delta_{k,q})$,*
- *$\forall \mathbf{x}, \mathbf{x}' \in \sigma \in \Sigma_{k,q}; \|\mathbf{x} - \mathbf{x}'\|_\infty \leq 2$.*

Proof. The properties above are satisfied by the subdivision described above. The vertex set is clearly $\phi[W_{k,q}] = V_{k,q}$. Each cell $\sigma(\mathbf{w}, \pi)$ has the same shape and volume: the volume is $1/(k-1)!$, since $(k-1)!$ such cells fill up a unit cube $[0, 1]^{k-1}$. Since ϕ is a linear map, the cells $\phi[\sigma(\mathbf{w}, \pi)]$ still have equal volume (although not the same shape). The cells $\sigma(\mathbf{w}, \pi)$ fill up the simplex $R_{k,q}$, whose volume is $q^{k-1}/(k-1)!$. Therefore, the number of cells is q^{k-1}. Finally, each cell $\sigma(\mathbf{w}, \pi)$ is contained in a translation of a unit cube $[0, 1]^{k-1}$. The map ϕ transforms coordinates by $x_i = (\phi(\mathbf{y}))_i = y_i - y_{i-1}$. Therefore, two points in the same cell $\phi[\sigma(\mathbf{w}, \pi)]$ can differ by at most 2 in each coordinate.

Constructing a fractional labeling. We recall that Sperner's lemma states that for any Sperner-admissible labeling of $V_{k,q}$, there must be a *rainbow cell* in $\Sigma_{k,q}$: one whose vertices are labeled with all k colors. The question we ask here is — is this still true for fractional labelings, in the sense that for any fractional labeling there must be a simplex $e \in \Sigma_{k,q}$ which contains a rainbow labeling with at least some nonzero weight? The answer is negative in a strong sense: For sufficiently large q, there are fractional labelings such that every cell uses a combination of labelings using at most 2 colors each.

LEMMA 6.3. *For every $k \geq 2$ and $q \geq 2k^3$, there is a subdivision $\Sigma_{k,q}$ of the simplex $\Delta_{k,q}$ using the vertex set $V_{k,q}$, and a fractional labeling $(x_{vj}, y_{e\alpha})$ of $(V_{k,q}, \Sigma_{k,q})$ such that $y_{e\alpha} = 0$ whenever α uses more than 2 colors.*

Proof. Consider a subdivision $(V_{k,q}, \Sigma_{k,q})$ as given by Lemma 6.2. We define the fractional labeling of each vertex $\mathbf{v} \in V_{k,q}$ according to its coordinates: $x_{\mathbf{v},j} = \frac{1}{q} v_j$. We have $x_{\mathbf{v},j} \geq 0$ and $\sum_{j=1}^{k} x_{\mathbf{v},j} = 1$ as desired.

Consider a cell $e \in \Sigma_{k,q}$ with vertices $\mathbf{v}_1, \ldots, \mathbf{v}_k$. Let us write these vertices as $\mathbf{v}_i = \mathbf{x} + \mathbf{y}_i$ where $x_j = \min_{1 \leq i \leq k} v_{ij}$ and $\mathbf{y}_i \geq 0$. Since the coordinates of each $\mathbf{v}_i$ sum up to q, the coordinates of each $\mathbf{y}_i$ have the same sum as well; let us denote it $\tilde{q}$. By Lemma 6.2, we have $\mathbf{y}_i \in [0,2]^k$. Hence, $\tilde{q} \leq 2k$. We define $\tilde{q}$-tuples of labels in $[k]$ as follows: for $1 \leq i \leq k$, $\beta_i \in [k]^{\tilde{q}}$ is chosen so that

$$\mathbf{y}_i = \Big(|\{j \in [\tilde{q}] : \beta_{ij} = 1\}|, \ldots, |\{j \in [\tilde{q}] : \beta_{ij} = k\}| \Big).$$

This is possible since each $\mathbf{y}_i$ is a vector with nonnegative integer coordinates summing up to $\tilde{q}$. Similarly, we define $\gamma \in [k]^{q-\tilde{q}}$ such that

$$\mathbf{x} = \Big(|\{j \in [q-\tilde{q}] : \gamma_j = 1\}|, \ldots, |\{j \in [q-\tilde{q}] : \gamma_j = k\}| \Big).$$

Since $q \geq 2k^3$, there is a coordinate in $\mathbf{x}$ of value $x_r \geq 2k^2 > (k-1)\tilde{q}$. This means there is label r which appears at least $(k-1)\tilde{q}$ times in γ. Let us extend γ to a q-tuple $\tilde{\gamma} \in ([k] \cup \{*\})^q$ by adding $\tilde{q}$ coordinates labeled '$*$'. Since we have $(k-1)\tilde{q}$ appearances of label r and $\tilde{q}$ appearances of label $*$ in $\tilde{\gamma}$, we can interleave the appearances of r and $*$ as follows: We define q-tuples $\tilde{\gamma}^1, \ldots, \tilde{\gamma}^k \in ([k] \cup \{*\})^q$ that are permutations of $\tilde{\gamma}$ such that in each position $j \in [q]$, the labels $\tilde{\gamma}^1_j, \ldots, \tilde{\gamma}^k_j$ are either all equal, or they are equal to r except for one which is equal to $*$, as shown here (for $\tilde{q} = 2$):

$$\begin{aligned}
\tilde{\gamma}^1 &= (\ldots\ldots\ldots\ * \, r\, r\, r\, r\, r\, r\, r\, * \, r\, r\, r\, r\, r\, r\, r) \\
\tilde{\gamma}^2 &= (\ldots\ldots\ldots\ r\, * \, r\, r\, r\, r\, r\, r\, r\, * \, r\, r\, r\, r\, r\, r) \\
\tilde{\gamma}^3 &= (\ldots\ldots\ldots\ r\, r\, * \, r\, r\, r\, r\, r\, r\, r\, * \, r\, r\, r\, r\, r) \\
\tilde{\gamma}^4 &= (\ldots\ldots\ldots\ r\, r\, r\, * \, r\, r\, r\, r\, r\, r\, r\, * \, r\, r\, r\, r) \\
\tilde{\gamma}^5 &= (\ldots\ldots\ldots\ r\, r\, r\, r\, * \, r\, r\, r\, r\, r\, r\, r\, * \, r\, r\, r) \\
\tilde{\gamma}^6 &= (\ldots\ldots\ldots\ r\, r\, r\, r\, r\, * \, r\, r\, r\, r\, r\, r\, r\, * \, r\, r) \\
\tilde{\gamma}^7 &= (\ldots\ldots\ldots\ r\, r\, r\, r\, r\, r\, * \, r\, r\, r\, r\, r\, r\, r\, * \, r) \\
\tilde{\gamma}^8 &= (\ldots\ldots\ldots\ r\, r\, r\, r\, r\, r\, r\, * \, r\, r\, r\, r\, r\, r\, r\, *)
\end{aligned}$$

Finally, we replace the $*$ symbols by β_{ij} as follows: each $\tilde{\gamma}^i$ has $\tilde{q}$ appearances of $*$, and we replace these by the $\tilde{q}$ coordinates of β_i, in an arbitrary order. We call the resulting q-tuples $\alpha^1, \ldots, \alpha^k$. Note that each k-tuple $(\alpha^1_j, \ldots, \alpha^k_j)$ uses at most 2 labels. Also, the number of appearances of each label ℓ in α^i is the sum of its appearances in β_i and γ, and hence equals the ℓ-coordinate of $\mathbf{x} + \mathbf{y}_i = \mathbf{v}_i$.

Now we define the fractional labeling of a cell e: for each $1 \leq j \leq q$, we set $y_{e,(\alpha^1_j,\ldots,\alpha^k_j)} = \frac{1}{q}$. For all other $\alpha \in [k]^k$, $y_{e\alpha} = 0$. By construction, we have

$$\begin{aligned}
\sum_{j \in [q] : \alpha^i_j = \ell} y_{e,(\alpha^1_j,\ldots,\alpha^k_j)} &= \frac{1}{q} |\{j \in [q] : \alpha^i_j = \ell\}| \\
&= \frac{1}{q} (\mathbf{x} + \mathbf{y}_i)_\ell = \frac{1}{q} (\mathbf{v}_i)_\ell = x_{\mathbf{v},\ell}
\end{aligned}$$

so the distributions are consistent. This proves Lemma 6.3.

6.2 Hardness of Hypergraph Rainbow-Avoiding Labeling with Color Lists. Since the Local Distribution LP cannot distinguish between instances where each hyperedge can be labeled with at most 2 colors, and instances where a rainbow hyperedge must exist, one can ask whether it is indeed computationally hard to distinguish between these two cases. Considering the general reduction for Min-CSP problems given by [6], one can formulate a hardness result assuming the Unique Games Conjecture. However, this hardness result is not entirely satisfactory, not only because of the UGC assumption, but also because of the lack of perfect completeness inherent in the UGC: We would obtain that it is UG-hard to distinguish between instances that must contain "many" rainbow hyperedges and instances that have a labeling with only "few" rainbow hyperedges. Instead, we prove by a direct reduction that it is NP-hard to distinguish between instances that must contain a rainbow hyperedge and instances that do not contain any, for $k = 3$. We conjecture that for any $k \geq 3$, it is NP-hard to distinguish between instances that must contain a rainbow hyperedge, and instances where each hyperedge can be labeled with at most 2 colors.

THEOREM 6.1. *It is NP-hard to decide for an instance of* Hypergraph Rainbow-Avoiding Labeling with Color Lists *(for $k = 3$) whether there is a labeling without any rainbow hyperedges.*

Proof. We recall the Monotone NAE-3SAT problem: Given a 3-CNF formula without negations, decide whether there is an assignment such that every clause has both a true variable and a false variable. This problem is NP-complete by [10].

Given an instance $\mathcal{I}$ of Monotone NAE-3SAT, we produce an instance $\mathcal{I}'$ of Hypergraph Rainbow-Avoiding

Labeling with Color Lists as follows. For each variable x_i, we generate a gadget based on Sperner's triangle: We have 6 vertices $V_i = \{v_i^{(1)}, v_i^{(2)}, v_i^{(3)}, v_i^{(1,2)}, v_i^{(1,3)}, v_i^{(2,3)}\}$. The color list for each vertex is given naturally by the superscript. In addition, we generate 3 hyperedges: $e_i^{(1)} = \{v_i^{(1)}, v_i^{(1,2)}, v_i^{(1,3)}\}$, $e_i^{(2)} = \{v_i^{(2)}, v_i^{(1,2)}, v_i^{(2,3)}\}$, $e_i^{(3)} = \{v_i^{(3)}, v_i^{(1,3)}, v_i^{(2,3)}\}$. Note that a fourth hyperedge, $e_i^* = \{v_i^{(1,2)}, v_i^{(2,3)}, v_i^{(1,3)}\}$, would complete a triangulation of the triangle $\{v_i^{(1)}, v_i^{(2)}, v_i^{(3)}\}$; however, e_i^* is *not* part of the instance we generate. (See Figure 4.) Instead, for each clause $x_i \vee x_j \vee x_k$, we generate a hyperedge $e_{ijk}^* = \{v_i^{(1,2)}, v_j^{(2,3)}, v_k^{(1,3)}\}$. We claim that $\mathcal{I}$ is satisfiable if and only if $\mathcal{I}'$ has an optimum of 0 as an instance of Hypergraph Rainbow-Avoiding Labeling with Color Lists (i.e., there is a labeling with no rainbow hyperedges).

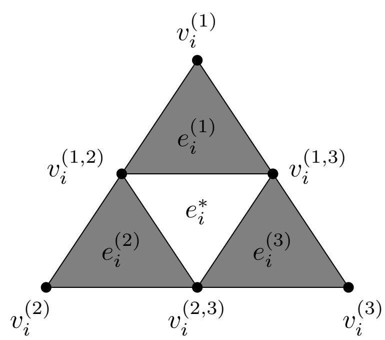

Figure 4: The gadget for variable x_i.

First, observe that by Sperner's Lemma, for any feasible labeling of V_i, one of the hyperedges $e_i^{(1)}, e_i^{(2)}, e_i^{(3)}, e_i^*$ must contain all 3 colors. Thus, the only way that a labeling of $\mathcal{I}'$ can be rainbow-free is that e_i^* contains all 3 colors. There are only two feasible rainbow labelings of $e_i^* = \{v_i^{(1,2)}, v_i^{(2,3)}, v_i^{(1,3)}\}$. In this order, the two labelings are $(1,2,3)$ and $(2,3,1)$. We associate these two labelings with the variable x_i being True or False, respectively: an assignment of the variables x_i defines a labeling, and a labeling where no hyperedge $e_i^{(j)}$ is rainbow defines an assignment to the variables x_i. We claim that this assignment is satisfying if and only if the labeling is rainbow-avoiding.

In one direction, if there is an all-true clause, then the respective hyperedge e_{ijk}^* is labeled $(1,2,3)$. If there is an all-false clause, then the respective hyperedge is labeled $(2,3,1)$. Therefore, a rainbow-avoiding labeling implies a satisfying assignment. Conversely, no other rainbow labeling of e_{ijk}^* is possible, given the color lists in $\mathcal{I}'$: therefore, a satisfying assignment to $\mathcal{I}$ implies a rainbow-avoiding labeling of $\mathcal{I}'$.

7 Application to fair division

One of the classical applications of Sperner's Lemma is *fair division* [12]: Suppose we have some divisible resource that should be allocated to k players in a way that satisfies each of them. Formally, a division of the resource is represented by a vector $(x_1, x_2, \dots, x_k) \in \Delta_k = \{\mathbf{x} \geq 0 : \sum_{i=1}^k x_i = 1\}$. Each player has a preference function $p_i : \Delta_k \to [k]$ that describes which one is her preferred piece under division $(x_1, x_2, \dots, x_k)$. We assume that preference functions satisfy the following condition.

Definition 7.1. *A preference function $p : \Delta_k \to [k]$ is greedy, if $x_\ell > 0$ whenever $p(x_1, \dots, x_k) = \ell$.*

In other words, a player always prefers a non-empty piece. Apart from that, the preferences might depend on all k pieces in any way. Sperner's Lemma implies that under this assumption, there is a division such that each player prefers a different piece; more precisely, there is a point $\mathbf{x} \in \Delta_k$ and a permutation $\pi : [k] \to [k]$ such that for each $i \in [k]$, $\mathbf{x}$ is in the closure of the set $C_{i,\pi(i)}$ where player i prefers $\pi(i)$: $C_{i,\pi(i)} = \{\mathbf{y} \in \Delta_k : p_i(\mathbf{y}) = \pi(i)\}$ (see [12]).

Our goal here is to provide an interpretation of our Lemma 4.1 in the context of fair division. Suppose that the resource is not quite continuously divisible, but instead comes in q discrete pieces, for some large integer q. If we try to find an allocation $(a_1, a_2, \dots, a_k)$ where the a_i's are integers adding up to q, it can obviously fail: The preference functions p_i could be all the same and then each player prefers the same piece. Instead, we can try to find a division $(a_1, a_2, \dots, a_k)$ such that $a_i \in \mathbb{Z}_+$ and $\sum_{i=1}^k a_i = q - 1$, we leave 1 unit of the resource unallocated and we consider a player *nearly satisfied* if the remaining unit is divisible so that the player prefers her own piece.

Definition 7.2. *A player i is nearly satisfied with the j-th piece in an integral division $(a_1, \dots, a_k)$, $\sum_{i=1}^k a_i = q - 1$, if there is a division $(x_1, \dots, x_k)$ where $\lfloor x_i \rfloor = a_i, \sum_{i=1}^k x_i = q$ and $p_i(x_1, \dots, x_k) = j$.*

Our result shows that even this is impossible, in a strong sense: For any k, there is a preference function such that for any such division, at least $k-4$ players are not even nearly satisfied. A possible interpretation of this is that one should be careful when applying the fair division theorem in a discrete setting.

Corollary 7.1. *Let $\Delta_{k,q} = \{\mathbf{x} \in \mathbb{R}_+^k : \sum_{i=1}^k x_i = q\}$. For any $k \geq 4$ and $q \geq k^2$, there exists a greedy preference function $p : \Delta_{k,q} \to [k]$ such that for any integral division $(a_1, \dots, a_k)$, $\sum_{i=1}^k a_i = q - 1$, at most 4 players are nearly satisfied.*

Proof. Consider $V_{k,q} = \{\mathbf{x} \in \mathbb{Z}_+^k : \sum_{i=1}^k x_i = q\}$ and the labeling $\ell : V_{k,q} \to [k]$ provided by Lemma 4.1. We extend this to a preference function $p : \Delta_{k,q} \to [k]$ by considering the Voronoi tiling of $\Delta_{k,q}$ determined by $V_{k,q}$. I.e., we label each point $\mathbf{x} \in \Delta_{k,q}$ by the color of its nearest neighbor in $V_{k,q}$: We let $p(\mathbf{x}) = \ell(\text{argmin}_{\mathbf{v}\in V_{k,q}} \|\mathbf{x} - \mathbf{v}\|_1)$. (If ties arise, then let's say we take the color of the lexicographically minimum nearest vertex; this does not affect the statement of the corollary.) Observe that if $x_i = 0$, then the nearest vertex $\mathbf{v} \in V_{k,q}$ cannot have $v_i \geq 1$; in that case we could decrease v_i by 1, increase another coordinate such that $v_j < x_j$ by 1 and obtain a closer vertex in $V_{k,q}$. Therefore, the color of $\mathbf{v}$ cannot be i. In other words, p is a greedy preference function.

Suppose $\mathbf{x} = (x_1, \ldots, x_k)$ such that $\lfloor x_i \rfloor = a_i$ and $\sum_{i=1}^k a_i = q - 1$. Let $\mathbf{v} \in V_{k,q}$ be the vertex closest to $\mathbf{x}$. Clearly, we must have $v_i \in \{a_i, a_i + 1\}$; if $v_i \geq a_i + 2$, there must be another coordinate such that $v_j < x_j$. Then we can define $v'_i = v_i - 1$, $v'_j = v_j + 1$, and we obtain a new vertex $\mathbf{v}'$ closer to $\mathbf{x}$. Similarly we can deal with the case where $v_i \leq a_i - 1$. Therefore, $v_i \in \{a_i, a_i + 1\}$ and since $\sum_{i=1}^k a_i = q - 1$, we get that exactly one of the coordinates of $\mathbf{v}$ is equal to $v_i = a_i + 1$ and the other coordinates are $v_j = a_j$. In other words, $p(\mathbf{x}) = \ell(\mathbf{v})$ where $\mathbf{v}$ is a vertex of the hyperedge $e(\mathbf{a})$.

Consider now the setting where all k players have the same preference function p as defined above. For any integral division $(a_1, \ldots, a_k)$, each player is nearly satisfied only with a piece ℓ such that $p(\mathbf{x}) = \ell$ for some point $\mathbf{x}$ such that $\lfloor x_i \rfloor = a_i$. By the discussion above, $p(\mathbf{x}) = \ell(\mathbf{v})$ for some vertex $\mathbf{v} \in e(\mathbf{a})$; however, this hyperedge contains only 4 colors. Therefore, only 4 players can be nearly satisfied with their allocated piece.

We remark that the specific preference function arising from the proof of Lemma 4.1 is not entirely unnatural — it can be viewed as "picking the smallest of the large pieces", perhaps with the purpose of obtaining a large piece while still appearing modest. We showed that under this preference function and the division approach outlined above, almost all the players are necessarily going to be disappointed.

8 Conclusions and open questions

We have proved several results about colorings of a discretization of the simplex. Our first result (Proposition 2.1) can be viewed as being at the opposite end of the spectrum from Sperner's Lemma: Instead of the existence of a rainbow cell, we prove a lower bound on the number of non-monochromatic cells. Due to the motivating Hypergraph Labeling problem, we consider a special hypergraph embedded in the simplex rather than a full subdivision. A natural question is whether an analogous statement holds for simplicial subdivisions. More generally, we might "interpolate" between Sperner's Lemma and our result, and ask: How many cells must contain at least j colors? It is clear that these questions depend on the structure of the subdivision, and some assumption of regularity would be needed to obtain a general result. More specifically, we can ask these questions about the concrete subdivision defined in Section 6.

- For a Sperner-admissible labeling of a "regular simplicial subdivision" (e.g., the one defined in Section 6), what is the minimum possible number of non-monochromatic cells? What is the minimum possible number of cells containing at least j colors?

We conjecture that for constant $j \leq k$ and $q \to \infty$, the number of cells containing at least j colors is $\Omega(q^{k-j})$. (For $j = k-1$, this can be shown by an argument similar to the proof of Sperner's Lemma.) We remark that while obtaining a bound of $\Omega(q^{k-2})$ would be relatively easy in the case of Proposition 2.1, it is crucial for our application that we get the tight multiplicative constant as well.

Another question is, what is the minimum number of colors per hyperedge for labelings of the Simplex-Lattice Hypergraph $H_{k,q}$ (defined in Section 2). We have proved that 4 colors suffice but it is possible that 2 colors are enough.

- Is there a Sperner-admissible labeling of the hypergraph $H_{k,q}$, for sufficiently large q, such that each hyperedge uses at most 2 colors?

This would have a consequence for fair division as in Section 7. We remark that such a labeling can be designed for $k = 4$ and a sufficiently large q (we omit the proof).

Acknowledgement. JV would like to thank Moses Charikar and Anupam Gupta for questions regarding the hardness of the Hypergraph Rainbow-Avoiding Labeling with Color Lists problem.

References

[1] Chandra Chekuri, Sanjeev Khanna, Joseph Naor and Leonid Zosin. A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM J. Discrete Math*, 608–625, 2005.

[2] Chandra Chekuri and Alina Ene. Submodular cost allocation problem and applications. In *Proc. of ICALP*, 354–366, 2011.

[3] Chandra Chekuri and Alina Ene. Approximation algorithms for submodular multiway partition. In *Proc. of IEEE FOCS*, 807–816, 2011.

[4] Herbert Edelsbrunner and Daniel R. Grayson. Edgewise subdivision of a simplex. *Discrete & Computational Geometry*, 24(4):707–719, 2000.

[5] Alina Ene and Jan Vondrák. Hardness of submodular cost allocation: Lattice matching and a simplex coloring conjecture. In *Proc. of APPROX*, 144–159, 2014.

[6] Alina Ene, Jan Vondrák, and Yi Wu. Local distribution and the symmetry gap: Approximability of multiway partitioning problems. In *Proc. of ACM-SIAM SODA*, 306–325, 2013.

[7] Jon Kleinberg and Eva Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. *Journal of the ACM*, 49(5):616–639, 2002.

[8] Rajsekar Manokaran, Joseph Naor, Prasad Raghavendra, and Roy Schwartz. SDP gaps and UGC hardness for multiway cut, 0-extension, and metric labeling. In *Proc. of ACM STOC*, 11–20, 2008.

[9] Douglas W. Moore. Simplicial mesh generation with applications, 1992. Ph.D. thesis, Cornell University.

[10] Thomas J. Schaeffer. The complexity of satisfiability problems. *Proc. of ACM STOC*, 216–226, 1978.

[11] Emanuel Sperner. Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes. *Math. Sem. Univ. Hamburg*, 6:265–272, 1928.

[12] Francis Edward Su. Rental harmony: Sperner's lemma in fair division. *Amer. Math. Monthly*, 106:930–942, 1999.