

Envy-free division of multi-layered cakes

Ayumi Igarashi and Frédéric Meunier

Abstract

Dividing a multi-layered cake under non-overlapping constraints captures several scenarios, e.g., allocating multiple facilities over time where each agent can utilize at most one facility simultaneously. We establish the existence of an envy-free multi-division that is non-overlapping and contiguous within each layer when the number of agents is a prime power, solving partially an open question by Hosseini et al. (2020). Our approach follows an idea proposed by Jojić et al. (2020) for envy-free divisions, relying on a general fixed-point theorem. We further design a fully polynomial-time approximation scheme for the two-layer three-agent case, with monotone preferences. All results are actually established for divisions among groups of almost the same size. In the one-layer three-group case, our algorithm is able to deal with any predetermined sizes, still with monotone preferences. For three groups, this provides an algorithmic version of a recent theorem by Segal-Halevi and Suksompong (2021).

1 Introduction

Imagine n students taking a gymnasium course. The professor of the course wishes to divide the students into groups and assign m activities over a given period of time. Students may have different opinions about the time slot and activity to which they would like to be assigned; for instance, some may prefer to swim in the morning and play basketball in the afternoon, while the others may want to play ping pong as long as possible.

The problem of fairly distributing a resource has been often studied in the classical cake-cutting model [25], where the cake, represented by the unit interval, is to be divided among heterogeneous agents. A variety of cake-cutting techniques have been developed over the past decades; in particular, the existence and algorithmic questions concerning an *envy-free* division, where each agent receives her first choice under the given division, have turned out to involve highly non-trivial arguments [10, 26, 28, 30].

In the above example of assigning multiple activities, however, one cannot trivially reduce the problem to the one-dimensional case. Indeed, if we merely divide the m time intervals independently, this may result in a division that is not feasible, i.e., it may assign overlapping time intervals to the same agent who can perform at most one activity at a given time. In order to capture such constraints, Hosseini, Igarashi, and Searns [13] have recently introduced the multi-layered cake-cutting problem. There, n agents divide a cake formed by m different layers under the *feasibility* constraint: the pieces of different layers assigned to the same agent should be non-overlapping, i.e., these pieces should have disjoint interiors. Besides an application to assign activities, the model can capture a plethora of real-life situations. For instance, consider a situation with several tasks running all day long in parallel (e.g., a fast-food), and workers turning between these tasks and who cannot be assigned to more than one task simultaneously.

For the special case of two layers and two agents, Hosseini et al. showed that the cut-and-choose protocol achieves envy-free division that is both feasible and contiguous within each layer by a single cut located at the same position over the two layers; in other words, the division is obtained by a “long knife.” For a more general combination of positive integers m and n with $m \leq n$, it has remained an open question whether there is an envy-free multi-division that is both contiguous and feasible, even in the special case when $m = 2$ and $n = 3$. Though, when the contiguity requirement is relaxed, they showed the existence of envy-free feasible multi-divisions [12]. Note that when the number of layers strictly exceeds

the number of agents, i.e., $m > n$, there is no way to allocate the entire layered cake while satisfying feasibility.

Now, returning to our first example, recall that in that example, we wish to assign activities to *groups* of students, instead of *individuals*. In the above example with workers assigned to tasks, it also makes sense to consider a variant with groups of workers staying together all day long. In this paper, we consider a “group” generalization of the multi-layered cake-cutting problem introduced by Hosseini et al. [13]. The aim is to divide a multi-layered cake into q pieces, partition n agents into q groups of almost equal size, and assign the pieces to groups in a fair manner. Our focus is on envy-free divisions of a multi-layered cake under feasible and contiguous constraints.

We first show that when q is a prime power and $m \leq q \leq n$, an envy-free multi-division that is both feasible and contiguous exists under mild conditions on preferences (especially they are not required to be monotone). Our proof relies on a general Borsuk–Ulam-type theorem proven by Volovikov [29], whose application for this type of problems has been initiated by Jojić, Panina, and Živaljević [14]; we comment on this proof technique in the beginning of Section 3. Further, such divisions can be obtained using $q - 1$ cuts in each layer located at the same positions, i.e., the divisions are obtained by $q - 1$ long knives. (In the initial example of a gymnasium, we can ensure for free that all activities begin and end at the same time.) Note that when $q = n$, our problem reduces to the setting of Hosseini et al., which means that the present paper settles their open question partially. Unfortunately, our existence result is the best one could hope to obtain, under the general preference model with choice functions: Avvakumov and Karasev [2] and Panina and Živaljević [21] provided examples of a cake-cutting instance with choice functions for which no envy-free division exists, for every choice of q that is not a prime power.

Our existence result concerning envy-freeness answers another open question raised by Hosseini et al.: By a recursive procedure initialized with the above existence result, we prove the existence of a *proportional* multi-division that is feasible and contiguous for any q with $m \leq q \leq n$ when agents have additive valuations, and thus properly generalizes the known existence result when m is a power of 2, $m \leq n$, and $q = n$ [13]; see Appendix D for details.

Hosseini et al. also discussed the question of a procedure achieving an envy-free multi-division in the two-layered case. They considered divisions obtained by cutting the top layer with one “short knife” and dividing the rest with one “long knife.” For this particular problem, we observe that one can encode such divisions by the points of the unit square. This way, a Sperner-type argument turns out to be applicable: there is an envy-free multi-division using only one short knife and one long knife when agents have monotone preferences. By exploiting the monotonicity that arises when fixing the long knife position, we further devise a fully polynomial-time approximation scheme (FPTAS) to compute an approximate envy-free multi-division. Moreover, when there is only one layer, the algorithm can handle the case where the sizes of the groups can be set arbitrarily, which corresponds to the algorithmic version of the three-group case of the recent existence result by Segal-Halevi and Suksompong [23].

Most results of the paper are actually stated and proved for *birthday cake multi-divisions*, i.e., divisions of the cake such that whichever piece a *birthday agent* selects, there is an envy-free assignment of the remaining pieces to the remaining agents. See Appendix A for a more extensive discussion on the further related work.

2 Model

We consider the setting of Hosseini, Igarashi, and Searns [13], except that we allow slightly more general preferences and that we aim to obtain a division among groups of agents. We are given m layers, n agents, and a positive integer q with $1 \leq q \leq n$. A *cake* is the unit

interval $[0, 1]$. A *piece* of cake is a union of finitely many disjoint closed subintervals of $[0, 1]$. We refer to a subinterval of $[0, 1]$ as a *contiguous piece* of cake. An *m-layered cake* is a sequence of m cakes $[0, 1]$, each being a *layer*. A *layered piece* is a sequence $\mathcal{L} = (L_\ell)_{\ell \in [m]}$ of pieces of each layer ℓ ; a layered piece is *contiguous* if each L_ℓ is a contiguous piece of layer $\ell \in [m]$. A layered piece \mathcal{L} is *non-overlapping* if no two pieces from different layers overlap, i.e., $L_\ell \cap L_{\ell'}$ is empty or formed by finitely many points for every pair of distinct layers ℓ, ℓ' . The *length* of a layered piece is the sum of the lengths of its pieces in each layer. A *multi-division* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_q)$ is a q -tuple forming a partition of the m -layered cake into q layered pieces. (Here, “partition” is used in a slightly abusive way: while the collection covers the layered cake and the layered pieces have disjoint interiors, we allow the latter to share endpoints.) A multi-division \mathcal{A} is

- *contiguous* if \mathcal{A}_i is contiguous for each $i \in [q]$.
- *feasible* if \mathcal{A}_i is non-overlapping for each $i \in [q]$.

We focus in this work on *complete* multi-divisions where the entire layered cake must be allocated.

Each agent i has a *choice function* c_i that, given a multi-division, returns the set of *preferred* layered pieces (among which the agent is indifferent). This function returns the same set of pieces over all permutations of the entries of the multi-division: for each permutation $\rho: [q] \rightarrow [q]$, $c_i(\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_q) = c_i(\mathcal{A}_{\rho(1)}, \mathcal{A}_{\rho(2)}, \dots, \mathcal{A}_{\rho(q)})$. The choice function model is used in [17, 28] and more general than the valuation model while the latter is more standard in fair division. An agent i *weakly prefers* a layered piece \mathcal{A}_j to another layered piece $\mathcal{A}_{j'}$ within a multi-division \mathcal{A} if $\mathcal{A}_{j'} \in c_i(\mathcal{A})$ implies $\mathcal{A}_j \in c_i(\mathcal{A})$. An agent has *hungry preferences* if in any multi-division every layered piece of nonzero-length is weakly preferred to every layered piece of zero-length. An agent i has *monotone preferences* if every pair $\mathcal{A}, \mathcal{A}'$ of multi-divisions with a j such that $\mathcal{A}_j \subseteq \mathcal{A}'_j$ and $\mathcal{A}_{j'} \supseteq \mathcal{A}'_{j'}$ for all $j' \neq j$, we have $\mathcal{A}'_j \in c_i(\mathcal{A}')$ whenever $\mathcal{A}_j \in c_i(\mathcal{A})$. An agent i has *closed preferences* if the following holds: for every sequence $(\mathcal{A}^{(t)})_{t \in \mathbb{Z}_+}$ of multi-divisions converging to a multi-division $\mathcal{A}^{(\infty)}$, we have $\mathcal{A}_j^{(\infty)} \in c_i(\mathcal{A}^{(\infty)})$ whenever $\mathcal{A}_j^{(t)} \in c_i(\mathcal{A}^{(t)}) \quad \forall t \in \mathbb{Z}_+$. The convergence of layered pieces is considered according to the pseudo-metric $d(\mathcal{L}, \mathcal{L}') = \mu(\mathcal{L} \Delta \mathcal{L}')$; a sequence of multi-divisions is converging if each of its layered pieces converges. Here, μ is the Lebesgue measure and $\mathcal{L} \Delta \mathcal{L}' = ((L_\ell \setminus L'_\ell) \cup (L'_\ell \setminus L_\ell))_{\ell \in [m]}$. A multi-division $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_q)$ is *envy-free* if there exists a surjective assignment $\pi: [n] \rightarrow [q]$ such that $\mathcal{A}_{\pi(i)} \in c_i(\mathcal{A})$ for all $i \in [n]$. Note that when $q = n$, our definition coincides with the standard definition of an envy-free division that assigns each piece to a single agent. In general, our results are stated with a *birthday agent*, who is not taken into account for defining or computing multi-divisions, but is still considered for the overall envy-freeness.

We also consider a setting where each agent can specify the valuation of each layered piece. Each agent i has a *valuation function* v_i that assigns a real value $v_i(\mathcal{L})$ to any layered piece \mathcal{L} . A valuation function naturally gives rise to a choice function that among several layered pieces, returns the most valuable layered pieces. A valuation function v_i satisfies

- *monotonicity* if $v_i(\mathcal{L}) \leq v_i(\mathcal{L}')$ for any pair of layered pieces $\mathcal{L}, \mathcal{L}'$ such that $L_\ell \subseteq L'_\ell$ for every $\ell \in [m]$.
- *the Lipschitz condition* if there exists a fixed constant K such that for every pair of layered pieces $\mathcal{L}, \mathcal{L}'$, $|v_i(\mathcal{L}) - v_i(\mathcal{L}')| \leq K \times \mu(\mathcal{L} \Delta \mathcal{L}')$.

It is easy to see that monotonicity along with the Lipschitz condition implies that the hungry assumption is satisfied: the Lipschitz condition implies that all layered pieces of zero-length have the same value; monotonicity then implies that every layered piece has a value at least that value.

For an instance with agents' valuation functions, one can define concepts of approximate envy-freeness. A multi-division $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_q)$ is ε -envy-free if there exists a surjective assignment $\pi: [n] \rightarrow [q]$ such that for all $i \in [n]$, $v_i(\mathcal{A}_{\pi(i)}) + \varepsilon \geq \max_{i' \in [n]} v_i(\mathcal{A}_{i'})$. For an instance with agents' valuation functions, we assume that $v_i(\mathcal{L})$ can be accessed in constant time for any agent i and layered piece \mathcal{L} .

There are several ways to achieve feasibility and contiguity constraints. The easiest way is probably with *long knives* that cut all layers simultaneously. But we can also use *short knives* that cut only a single layer at a time. See Figure 1 for an illustration. For our result ensuring the existence of an envy-free multi-division for any number of layers, we will consider only long knives, while for the FPTAS for the two-layered cake, the multi-division will be obtained with one short and one long knife.

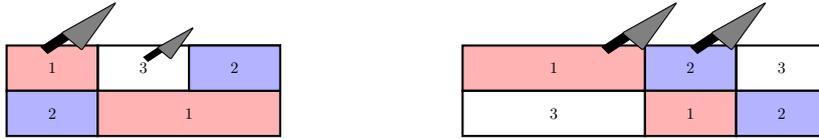


Figure 1: Multi-divisions of a two-layered cake, obtained by one long knife and one short knife and by two long knives (pictured left-to-right).

We assume basic knowledge in algebraic topology. Definitions of abstract and geometric simplicial complexes, the fact that they are somehow equivalent, and other related notions are reminded in Appendix E. The reader might consult the book by De Longueville [7] or the one by Matoušek [16], especially Chapters 1 and 6 of the latter book, for complementary material. In the sequel, we will identify geometric and abstract simplicial complexes without further mention.

3 Envy-free division using $q - 1$ long knives

Now, we formally present the first main result of this paper, stating that an envy-free multi-division using $q - 1$ long knives exists when n is a prime power.

Theorem 1. *Consider an instance of the multi-layered cake-cutting problem with m layers and n agents, $m \leq n$, with closed preferences. Let q be an integer such that $m \leq q \leq n$. If q is a prime power, then there exists a feasible and contiguous multi-division into q layered pieces so that no matter which layered piece the birthday agent chooses, the remaining agents can be assigned to the layered pieces while satisfying the following two properties:*

- *each of the remaining agents is assigned to one of her preferred layered pieces.*
- *the number of agents assigned to each layered piece, including the birthday agent, differs by at most one.*

Moreover, it can be achieved with $q - 1$ long knives.

More formally (and with the birthday agent being labeled with n), Theorem 1 ensures the existence of a feasible and contiguous multi-division \mathcal{A} into q layered pieces with the following property: for every $j^* \in [q]$, there is an assignment $\pi_{j^*}: [n] \rightarrow [q]$ with

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n - 1]$, $\mathcal{A}_{\pi_{j^*}(i)} \in c_i(\mathcal{A})$, and
- for each $j \in [q]$, $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/q \rfloor, \lceil n/q \rceil\}$.

Let us comment briefly on the special case when $m = 1$ and $q = n$. Since the agents might prefer zero-length pieces in our setting, our theorem boils down then to a recent result of Avvakumov and Karasev [2]. They showed that when n is a prime power, there always exists an envy-free division, even if we do not assume that the agents are hungry, and that this is not true anymore if n is not a prime power. (The Avvakumov–Karasev theorem was first proved for $n = 3$ by Segal-Halevi [22]—who actually initiated the study of envy-free divisions with non-necessarily hungry agents—and for prime n by Meunier and Zerbib [18].)

We note that the standard proof showing the existence of an envy-free division via Sperner’s lemma due to Su [28] may not work in the multi-layered setting even when $q = n$. In the model of standard cake-cutting, the divisions into q parallel pieces of lengths x_i ($i = 1, 2, \dots, q$) can be represented by the points of the standard simplex Δ_{q-1} , which is then triangulated with the vertices of each simplex being labeled with distinct owner agents, and colored in such a way that each “owner” agent colors the vertex with the index of her favorite bundle of the “owned” division. When agents always prefer nonzero-length pieces to zero-length ones, the coloring satisfies the boundary condition of Sperner’s lemma. This lemma guarantees then the existence of a colorful triangle, which corresponds to an approximate envy-free division.

In the same spirit of Su’s approach, one may attempt to encode feasible multi-divisions using $q - 1$ long knives, by the points of the standard simplex Δ_{q-1} and apply the usual method by using Sperner’s lemma to show the existence of an envy-free division. For instance, each q -tuple (x_1, x_2, \dots, x_q) can represent a feasible and contiguous multi-division $(\mathcal{A}_1, \dots, \mathcal{A}_q)$ where the ℓ -th layered piece of the i -th bundle \mathcal{A}_i is given by the $\eta(i, \ell)$ -piece of length $x_{\eta(i, \ell)}$ where $\eta(i, \ell) = i + \ell - 1$ (modulo m). Unfortunately, this approach may fail to work for the multi-layered cake-cutting: even when the agents have monotone preferences over the pieces, the coloring described in the previous paragraph does not satisfy in general the boundary condition of Sperner’s lemma.

We thus employ a new approach of using a general Borsuk–Ulam-type theorem, originally proven by Volovikov [29], and recently applied by Jojić, Panina, and Živaljević [14, 21] on the envy-free division of a cake. Volovikov’s theorem considers a topological space X and a sphere, both on which a group of the form $((\mathbb{Z}_p)^k, +)$ acts (with p being prime), and a map from X to the sphere commuting with the action. Under some assumptions on the connectivity of X and fixed-point freeness of the action on X , the theorem prevents the dimension of the sphere to be too small. For q being a prime power p^k , we consider a “configuration” space whose points encode at the same time feasible pieces obtained by $q - 1$ long knives and all possible assignments. The configuration space is actually the *chessboard complex* $\Delta_{2q-1, q}$, which is guaranteed to be $(q - 2)$ -connected. Roughly speaking, mapping this sufficiently connected space $X = \Delta_{2q-1, q}$ to the $(q - 1)$ -dimensional simplex recording the popularity among the q pieces, Volovikov’s theorem shows that the center of the simplex cannot be missed, and thus that an almost equal popularity of the pieces can be achieved. With this technique, the existence of an envy-free multi-division is shown for a general class of preferences that are not necessarily monotone.

Tools from equivariant topology We introduce now a specific abstract simplicial complex that will play a central role in the proof of Theorem 1. The *chessboard complex* $\Delta_{2q-1, q}$ is the abstract simplicial complex whose ground set is $[2q - 1] \times [q]$ and whose simplices are the subsets $\sigma \subseteq [2q - 1] \times [q]$ such that for every two distinct pairs (r, j) and (r', j') in σ we have $r \neq r'$ and $j \neq j'$. The name comes from the following: If we interpret $[2q - 1] \times [q]$ as a $(2q - 1) \times q$ chessboard, the simplices are precisely the configurations of pairwise non-attacking rooks. See Figures 3a and 4a for an illustration of the chessboard complex $\Delta_{2q-1, q}$ when $q = 2$ and $q = 3$.

Given an additive group G of order q , we get a natural action $(\varphi_g)_{g \in G}$ of G on $\Delta_{2q-1, q}$ by identifying $[q]$ with G via a bijection $\eta: [q] \rightarrow G$: this natural action is defined by

$\varphi_g(r, j) = g \cdot (r, j) := (r, \eta^{-1}(g + \eta(j)))$. This action is *free*, namely the orbit of each point in any geometric realization of $\Delta_{2q-1, q}$ is of size q . Equivalently, the relative interiors of $\varphi_g(\sigma)$ and σ are disjoint for every simplex σ of $\Delta_{2q-1, q}$ and every element g of G distinct from the neutral element; see [16, Chapter 6].

The following lemma is an immediate consequence of Volovikov’s theorem [29], which has found many applications in topological combinatorics. For an additive group G , we denote by Δ^G the standard simplex whose vertices are the unit vectors \mathbf{e}_g (where g ranges over G) of \mathbb{R}^G . The group acts naturally on Δ^G by setting $\varphi_{g'}(\mathbf{e}_g) := \mathbf{e}_{g+g'}$ and by extending the action affinely on each face of Δ^G . When G does not satisfy the condition of the lemma, its conclusion does not necessarily hold. For $G = \mathbb{Z}_6$, counterexamples are known; see [31].

(Given two topological spaces X and Y on which the group G acts, a map $h: X \rightarrow Y$ is *G-equivariant* if $h(g \cdot x) = g \cdot h(x)$ for all $x \in X$ and all $g \in G$.)

Lemma 2. *Let $q = p^k$, where p is a prime number and k a positive integer. Denote by G the additive group $((\mathbb{Z}_p)^k, +)$. For any G -equivariant continuous map $f: \Delta_{2q-1, q} \rightarrow \Delta^G$, there exists $\mathbf{x}_0 \in \|\Delta_{2q-1, q}\|$ such that $f(\mathbf{x}_0) = \frac{1}{q} \sum_{g \in G} \mathbf{e}_g$.*

Encoding divisions via the chessboard complex The proof of Theorem 1 uses a “configuration space” encoding some possible contiguous and feasible multi-divisions with $q - 1$ long knives. This configuration space is the simplicial complex $\Delta_{2q-1, q}$, with $q = p^k$ and $G = ((\mathbb{Z}_p)^k, +)$ acting on it. The elements in G will be used to identify the layered pieces. We choose an arbitrary bijection $\eta: [q] \rightarrow G$ to ease this identification. (In case $k = 1$, it is certainly most intuitive to set $\eta(j) = j$; note that when $k \neq 1$, this definition does not make sense.) Moreover, we fix an arbitrary injective map $h: [m] \rightarrow G$ and a geometric realization of $\Delta_{2q-1, q}$. We denote by $\mathbf{v}_{r, j}$ the realization of the vertex (r, j) . We explain now how each point of $\|\Delta_{2q-1, q}\|$ encodes a multi-division with $q - 1$ long knives.

We assign to each point \mathbf{x} of $\|\Delta_{2q-1, q}\|$ a $(q-1)$ -dimensional simplex of $\Delta_{2q-1, q}$ containing it. A tie can occur, e.g., for $q = 3$, when \mathbf{x} belongs to the interior of the edge $\mathbf{v}_{3,1} - \mathbf{v}_{1,2}$; there are three triangles containing this edge in $\Delta_{5,3}$; each of them contains the vertices $\mathbf{v}_{3,1}$ and $\mathbf{v}_{1,2}$; the third vertex can be any of $\mathbf{v}_{2,3}$, $\mathbf{v}_{4,3}$, and $\mathbf{v}_{5,3}$. We make this assignment in such a way that all points with same support are assigned to the same $(q-1)$ -dimensional simplex. Moreover, we make this assignment “equivariant”: given any $g \in G$, the simplex assigned to $g \cdot \mathbf{x}$ is the image by φ_g of the simplex assigned to \mathbf{x} . This is possible because the action of G on $\Delta_{2q-1, q}$ is free. We call this assignment the procedure \mathfrak{P} .

Consider any point \mathbf{x} in $\|\Delta_{2q-1, q}\|$. Let $\mathbf{v}_{r_1, 1}, \mathbf{v}_{r_2, 2}, \dots, \mathbf{v}_{r_q, q}$ be the vertices of the simplex assigned to \mathbf{x} by \mathfrak{P} . We write then \mathbf{x} as $\sum_{j=1}^q x_{r_j} \mathbf{v}_{r_j, j}$. We set $x_k = 0$ for every $k \notin \{r_1, \dots, r_q\}$. The values of x_1, \dots, x_{2q-1} do not depend on the choice made by \mathfrak{P} : only the vertices $\mathbf{v}_{r, j}$ spanning the minimal simplex of $\Delta_{2q-1, q}$ containing \mathbf{x} get nonzero coefficients, and these coefficients are then the barycentric coordinates in this face.

Let ρ be the permutation in \mathcal{S}_q such that $r_{\rho(1)} < r_{\rho(2)} < \dots < r_{\rho(q)}$. We interpret $x_{r_{\rho(j)}}$ as the length of the j -th piece: in a way similar to the traditional encoding of the divisions (see, e.g., [28]), the j -th piece in any layer is of length $x_{r_{\rho(j)}}$. We give then the j -th piece of the ℓ -th layer the element $\eta(\rho(j)) + h(\ell)$ of G as its “bundle-name.” We get a non-overlapping layered piece by considering all pieces with a same bundle-name: if it were overlapping, it would contain two pieces named $\eta(\rho(j)) + h(\ell) = \eta(\rho(j)) + h(\ell')$ with $\ell \neq \ell'$, which is not possible because h is injective. The non-overlapping layered pieces obtained this way form the multi-division encoded by \mathbf{x} , which we denote by $\mathcal{A}(\mathbf{x}) = (\mathcal{A}_1(\mathbf{x}), \mathcal{A}_2(\mathbf{x}), \dots, \mathcal{A}_q(\mathbf{x}))$, where $\mathcal{A}_j(\mathbf{x})$ is the layered piece formed by the pieces with bundle-name $\eta(j) \in G$. Clearly, each $\mathcal{A}(\mathbf{x})$ is a feasible and contiguous multi-division that uses $q - 1$ long knives. Figures 3, 4, and 5 in Appendix B illustrate the chessboard complex $\Delta_{2q-1, q}$ for $q = 2$ and $q = 3$ and associated multi-divisions corresponding to points of $\|\Delta_{2q-1, q}\|$.

The multi-divisions $\mathcal{A}(\mathbf{x})$ enjoy some “equivariant” property.

Lemma 3. We have $\mathcal{A}_j(\mathbf{x}) = \mathcal{A}_{\eta^{-1}(g+\eta(j))}(g \cdot \mathbf{x})$ for all $\mathbf{x} \in \|\Delta_{2q-1,q}\|$, $j \in [q]$, and $g \in G$.

See Figures 4c and 5c in Appendix B for an illustration of multi-divisions $\mathcal{A}(\mathbf{x})$ and $\mathcal{A}(g \cdot \mathbf{x})$. Another important point is that $\mathcal{A}(\mathbf{x})$ depends continuously on \mathbf{x} as stated by the following lemma. The convergence of multi-divisions is defined according to the pseudo-metric $d(\cdot, \cdot)$.

Lemma 4. Let $(\mathbf{x}^{(t)})_{t \in \mathbb{Z}_+}$ be a sequence of points of $\|\Delta_{2q-1,q}\|$ converging to some limit point $\mathbf{x}^{(\infty)}$. Then $(\mathcal{A}(\mathbf{x}^{(t)}))_{t \in \mathbb{Z}_+}$ converges to $\mathcal{A}(\mathbf{x}^{(\infty)})$.

Proof of Theorem 1 Before proceeding to the proof of Theorem 1, we present the following auxiliary lemma, which will also be used in Section 4. In the applications of this lemma, the vertices of the graph H will represent on one side the (non-birthday) agents and on the other side the pieces, and its edges will represent the acceptable assignments (in terms of preferences). We denote by $\delta_H(i)$ the edges incident to a vertex i in the graph H .

Lemma 5. Let n, q be positive integers. Let a_1, a_2, \dots, a_q be nonnegative real numbers summing up to $n-1$. Consider a bipartite graph $H = ([n-1], [q]; E)$ with nonnegative weights w_e on its edges $e \in E$. Suppose $\sum_{e \in \delta_H(i)} w_e = 1$ for each $i \in [n-1]$ and $\sum_{e \in \delta_H(j)} w_e = a_j$ for each $j \in [q]$. Then for every $j^* \in [q]$ there is an assignment $\pi_{j^*}: [n] \rightarrow [q]$ such that

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H ,
- $|\pi_{j^*}^{-1}(j^*)| = \lfloor a_{j^*} \rfloor + 1$, and
- for each $j \in [q]$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor a_j \rfloor, \lceil a_j \rceil\}$.

Now we are ready to prove Theorem 1.

Proof of Theorem 1. Without loss of generality, agent n is the birthday agent. Let p be the prime number and k the integer such that $q = p^k$. Let T be a triangulation of $\Delta_{2q-1,q}$ such that $\varphi_g(\sigma) \in \mathsf{T}$ for all $\sigma \in \mathsf{T}$ and all $g \in G$. (It is invariant by the action of $G = ((\mathbb{Z}_p)^k, +)$.) Such a triangulation can be achieved by taking repeated barycentric subdivisions of $\Delta_{2q-1,q}$.

We partition the vertices of T into their G -orbits. From each orbit, we pick a vertex \mathbf{v} . We ask each non-birthday agent $i \in [n-1]$ the index j of the non-overlapping layered piece $\mathcal{A}_j(\mathbf{v})$ she prefers in $\mathcal{A}(\mathbf{v})$. (In case of a tie, she makes an arbitrary choice.) We define $f^{(i)}(\mathbf{v})$ to be $\mathbf{e}_{\eta(j)}$. We extend $f^{(i)}$ on each orbit in an equivariant way: $f^{(i)}(g \cdot \mathbf{v}) := g \cdot f^{(i)}(\mathbf{v})$. This is done unambiguously because the action of G on T is free. Lemma 3 implies that, for every vertex \mathbf{v} of the triangulation T , the integer j such that $f^{(i)}(\mathbf{v}) = \mathbf{e}_{\eta(j)}$ is the index of a layered piece preferred by agent i in the multi-division $\mathcal{A}(\mathbf{v})$.

For each non-birthday agent $i \in [n-1]$, we extend the map $f^{(i)}$ affinely on each simplex of T . Denote by $\bar{f}^{(i)}$ the affine extension of $f^{(i)}$. This way, the map $\bar{f}^{(i)}$ is a G -equivariant simplicial map from T to Δ^G . The affine extension $\bar{f}^{(i)}$ of $f^{(i)}$ is an ‘‘approximation’’ of the original preferences. We then aggregate these approximate preferences among $n-1$ non-birthday agents by setting $\bar{f} = \frac{1}{n-1} \sum_{i=1}^{n-1} \bar{f}^{(i)}$. It is a G -equivariant continuous map from $\Delta_{2q-1,q}$ to Δ^G . For each vertex \mathbf{v} of T , the point $\bar{f}(\mathbf{v})$ represents the average preference of the $n-1$ agents according to the original preferences. For each \mathbf{x} in $\|\Delta_{2q-1,q}\|$, $\bar{f}(\mathbf{x})$ represents the average preference of the $n-1$ agents according to the approximate preferences. (This averaging technique has been introduced by Gale [11] and applied by Asada et al. [1] for the birthday cake-division.)

According to Lemma 2, there exists a point \mathbf{x}_0 in $\|\Delta_{2q-1,q}\|$ such that

$$\bar{f}(\mathbf{x}_0) = \frac{1}{q} \sum_{g \in G} \mathbf{e}_g. \quad (1)$$

For every non-birthday agent $i \in [n-1]$ and every layered piece $j \in [q]$, define $w_{ij} = \bar{f}^{(i)}(\mathbf{x}_0) \cdot \mathbf{e}_{\eta(j)}$ (where the product in the right-hand term is the dot product in \mathbb{R}^G). We clearly have $\sum_{i=1}^{n-1} w_{ij} = \frac{n-1}{q}$ for all $j \in [q]$ (by equation (1)), and $\sum_{j=1}^q w_{ij} = 1$ for all $i \in [n-1]$ (because $\bar{f}^{(i)}$ has its image in Δ^G).

Now, consider the bipartite graph $H = ([n-1], [q]; E)$, with one side being the agents from 1 to $n-1$ and with the other side being the layered pieces and where the edge ij exists precisely when $w_{ij} > 0$. Applying Lemma 5 with $a_j = \frac{n-1}{q}$ for $j \in [q]$, there exists for every $j^* \in [q]$ an assignment $\pi_{j^*} : [n] \rightarrow [q]$ such that

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H ,
- $|\pi_{j^*}^{-1}(n)| = \lfloor (n-1)/q \rfloor + 1$, and
- for each $j \in [q]$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor (n-1)/q \rfloor, \lceil (n-1)/q \rceil\}$.

We claim that

$$\lfloor n/q \rfloor \leq |\pi_{j^*}^{-1}(j)| \leq \lceil n/q \rceil,$$

for all $j \in [q]$. Clearly, for each $j \in [q]$, we have $|\pi_{j^*}^{-1}(j)| \leq \lceil n/q \rceil$ since $\lceil (n-1)/q \rceil \leq \lceil n/q \rceil$ and $\lfloor (n-1)/q \rfloor + 1 = \lceil n/q \rceil$. To see the lower bound, observe that if n is not a multiple of q , we have $\lfloor (n-1)/q \rfloor = \lfloor n/q \rfloor$ and hence $|\pi_{j^*}^{-1}(j)| \geq \lfloor n/q \rfloor$ for each $j \in [q]$; if n is a multiple of q , we have $|\pi_{j^*}^{-1}(j)| \leq n/q$ for each $j \in [q]$, which together with the fact that $\sum_{j \in [q]} |\pi_{j^*}^{-1}(j)| = n$ implies that $|\pi_{j^*}^{-1}(j)| = n/q$ for each $j \in [q]$.

For every integer $N > 0$, we can choose $\mathbb{T} := \mathbb{T}_N$ so that it has a mesh size upper bounded by $1/N$ and define $\bar{f}^{(i)} := \bar{f}_N^{(i)}$. For each N , we have \mathbf{x}_0^N satisfying (1). Let H_N be the graph $([n-1], [q]; E_N)$ such that the edge ij exists precisely when $\bar{f}_N^{(i)}(\mathbf{x}_0^N) \cdot \mathbf{e}_{\eta(j)} > 0$. Compactness implies that we can select among these arbitrarily large N an infinite sequence such that (\mathbf{x}_0^N) converges to a point \mathbf{x}^* and such that H_N is always the same graph H^* . As we have seen, for every $j^* \in [q]$, there is then an assignment $\pi_{j^*} : [n] \rightarrow [q]$ with:

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H^* , and
- for each $j \in [q]$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/q \rfloor, \lceil n/q \rceil\}$.

Consider any $j^* \in [q]$ and any N from the infinite sequence. We have $H_N = H^*$. By definition of $\bar{f}_N^{(i)}$, for each agent $i \in [n-1]$, there exists thus a vertex $\mathbf{v}^{i,j^*,N}$ of the supporting simplex of \mathbf{x}_0^N in \mathbb{T}_N such that $f_N^{(i)}(\mathbf{v}^{i,j^*,N})$ is $\mathbf{e}_{\eta(\pi_{j^*}(i))}$, meaning that $\mathcal{A}_{\pi_{j^*}(i)}(\mathbf{v}^{i,j^*,N}) \in c_i(\mathcal{A}(\mathbf{v}^{i,j^*,N}))$. We have $\mathcal{A}_{\pi_{j^*}(i)}(\mathbf{v}^{i,j^*,N}) \in c_i(\mathcal{A}(\mathbf{v}^{i,j^*,N}))$ for all $i \in [n-1]$ and arbitrarily large N . Since (\mathbf{x}_0^N) converges to \mathbf{x}^* , the sequence $(\mathbf{v}^{i,j^*,N})$ converges to the same point \mathbf{x}^* for every $i \in [n-1]$ and every $j^* \in [q]$. By the closed preferences assumption and Lemma 4, we have $\mathcal{A}_{\pi_{j^*}(i)} \in c_i(\mathcal{A}^*)$ for all $i \in [n-1]$, where $\mathcal{A}^* := \mathcal{A}(\mathbf{x}^*)$. Since we have $\pi_{j^*}(n) = j^*$ and $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/q \rfloor, \lceil n/q \rceil\}$ for all $j \in [q]$, the properties required for π_{j^*} are satisfied. \square

4 Envy-free division using one short knife and one long knife

While the main result of Section 3 is a pure existence result with a non-constructive proof based on Volovikov’s theorem, we focus in this section on the computational aspect of the problem of finding an envy-free multi-division. Actually, we show not only that an envy-free multi-division with one short and one long knife exists for a two-layered cake division among three groups of n agents with closed, monotone, and hungry preferences, but also that such a division can be efficiently computed for agents with valuations satisfying the Lipschitz condition. In fact, we prove stronger statements where both existence and computational results extend to those for birthday multi-divisions.

We start by stating the existence result.

Theorem 6. *Consider an instance of the two-layered cake-cutting problem with n agents, $n \geq 3$, with closed, monotone, and hungry preferences. Then there exists a feasible and contiguous multi-division into three layered pieces so that no matter which layered piece the birthday agent chooses, the remaining agents can be assigned to the layered pieces while satisfying the following two properties:*

- *each of the remaining agents is assigned to one of her preferred layered pieces.*
- *the number of agents assigned to each layered piece differs by at most one.*

Moreover, this multi-division requires only one short knife and one long knife.

In a more formal way (and with the birthday agent being labeled with n), Theorem 6 ensures the existence of a feasible and contiguous multi-division \mathcal{A} into three layered pieces with the following property: for every $j^* \in \{1, 2, 3\}$, there is an assignment $\pi_{j^*}: [n] \rightarrow \{1, 2, 3\}$ with

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n - 1]$, $\mathcal{A}_{\pi_{j^*}(i)} \in c_i(\mathcal{A})$, and
- for each $j \in \{1, 2, 3\}$, $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/3 \rfloor, \lceil n/3 \rceil\}$.

For three agents with valuations satisfying monotonicity and the Lipschitz condition, Deng et al. [9] designed an FPTAS to compute an approximate envy-free division. We show that in the context of two-layered cake-cutting among n agents, we can also design an FPTAS. This result, which can be seen as an algorithmic version of Theorem 6, generalizes the result of Deng et al. in three respects: First, our result holds for the birthday version; second, it holds for the group version, ensuring an envy-free division among three groups of almost equal size; third, it holds for the case of two-layered cake-cutting. Recall that monotonicity and the Lipschitz condition imply the hungry assumption under the valuation function model; see Section 2.

Theorem 7. *Consider an instance of the two-layered cake-cutting problem with n agents, $n \geq 3$, whose valuation functions satisfy monotonicity and the Lipschitz condition with constant K . Then, for any $\varepsilon > 0$, one can find in time $O(n \log^2 \frac{K}{\varepsilon})$ a feasible and contiguous multi-division into three layered pieces where no matter which layered piece the birthday agent chooses, the remaining agents can be assigned to the layered pieces while satisfying the following two properties:*

- *each of the remaining agents is assigned to one of her ε -approximate preferred layered pieces.*

- the number of agents assigned to each layered piece differs by at most one.

Moreover, this multi-division requires only one short knife and one long knife.

More formally (and with the birthday agent being labeled with n), Theorem 7 ensures that one can find in time $O(n \log^2 \frac{K}{\varepsilon})$, a feasible and contiguous multi-division \mathcal{A} into three layered pieces with the following property: for every $j^* \in \{1, 2, 3\}$, there is an assignment $\pi_{j^*}: [n] \rightarrow \{1, 2, 3\}$ with

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n - 1]$, $v_i(\mathcal{A}_{\pi_{j^*}(i)}) + \varepsilon \geq \max_{j \in \{1, 2, 3\}} v_i(\mathcal{A}_j)$, and
- for each $j \in \{1, 2, 3\}$, $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/3 \rfloor, \lceil n/3 \rceil\}$.

Note that the algorithm of Theorem 7 only computes the division and not the assignment itself. Yet, once we find a desired multi-division, we can ask each agent her approximate preferred pieces and then easily compute a desired assignment π_{j^*} in $O(n^2)$; this can be done, e.g., by a recent max-flow algorithm of Orlin [20].

In order to establish the above theorems, we encode by the points of the unit square $[0, 1]^2$ the divisions of the two-layered cake using the short knife and the long knife. The position of the long knife corresponds to the x -axis and the position of the short knife corresponds to the y -axis. Fixing bundle names appropriately, the two vertical boundaries when $x = 0$ and $x = 1$ enjoy a certain symmetry: the divisions that appear on these boundaries are the same. By exploiting this symmetry, one can apply a Sperner-type argument to show the existence of an envy-free division. A careful utilization of the monotonicity allows to make this argument algorithmic and to get the FPTAS.

With this kind of techniques, we are able to prove the following algorithmic version of the theorem of Segal-Halevi and Suksompong [23].

Theorem 8. *Consider an instance of the one-layered cake-cutting problem with n agents, $n \geq 3$, whose valuation functions satisfy monotonicity and the Lipschitz condition with constant K . Let k_1, k_2, k_3 be positive integers summing up to n . Then, for any $\varepsilon > 0$, one can find in time $O(n \log^2 \frac{K}{\varepsilon})$ a feasible and contiguous multi-division into three pieces where no matter which piece the birthday agent chooses, the remaining agents can be assigned to the pieces while satisfying the following two properties:*

- each of the remaining agents is assigned to one of her ε -approximate preferred pieces.
- the number of agents assigned to piece j is k_j for all $j \in \{1, 2, 3\}$.

More formally (and with the birthday agent being labeled with n), Theorem 8 ensures that one can find in time $O(n \log^2 \frac{K}{\varepsilon})$, a feasible and contiguous multi-division \mathcal{A} into three pieces with the following property: for every $j^* \in \{1, 2, 3\}$, there is an assignment $\pi_{j^*}: [n] \rightarrow \{1, 2, 3\}$ with

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n - 1]$, $v_i(\mathcal{A}_{\pi_{j^*}(i)}) + \varepsilon \geq \max_{j \in \{1, 2, 3\}} v_i(\mathcal{A}_j)$, and
- for each $j \in \{1, 2, 3\}$, $|\pi_{j^*}^{-1}(j)| = k_j$.

We remark that the existence of such a division follows from the proof technique by Segal-Halevi and Suksompong [23] and the classical existence result of a birthday division Woodall [30] and Asada et al. [1]. However, the algorithmic result does not directly follow because the proof in Segal-Halevi and Suksompong reduces the problem to the n -agent case (i.e., it invokes the $n - 1$ dimensional version of Sperner's lemma).

4.1 Proof sketch

We introduce the main tools and lemmas used for the proof of Theorems 6, 7, and 8.

Consider n agents who have closed, monotone, and hungry preferences. The n -th agent is assumed to be the birthday agent. We consider a two-layered cake. Its *top layer* is the layer that is weakly preferred to the other layer, within the “degenerate” multi-division with no cuts, by the majority of agents; that is, at least $\lceil \frac{n}{2} \rceil$ agents weakly prefer the top layer to the other layer. This latter layer is then referred to as the *bottom layer*.

As in Section 3, we use a “configuration space” to encode the considered multi-divisions, but in a much simpler way. Here, the configuration space is the unit square. A point (x, y) of the unit square encodes a division as in Figure 2 where x corresponds to the long knife over the two layers and y corresponds to the short knife over the top layer. We denote by $\mathcal{A}(x, y) = (\mathcal{A}_1(x, y), \mathcal{A}_2(x, y), \mathcal{A}_3(x, y))$ the multi-division in Figure 2 represented by $(x, y) \in [0, 1]^2$ where the agents first cut the cake via the short knife position y , and then cut the rest via the long knife position x . Namely, $\mathcal{A}_1(x, y)$ consists the $[0, y]$ segment of the top layer, $\mathcal{A}_2(x, y)$ consists of the $[\max\{x, y\}, 1]$ of the top layer and the $[0, x]$ segment of the bottom layer, and $\mathcal{A}_3(x, y)$ contains the remaining pieces. The multi-division $\mathcal{A}(x, y)$ is feasible and contiguous.

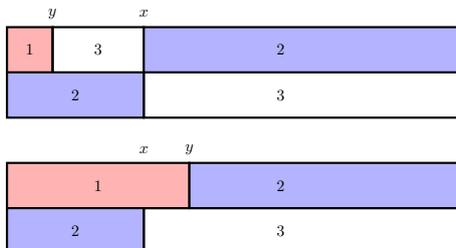


Figure 2: Three-agent multi-divisions $\mathcal{A}(x, y)$. Note that the short knife y is prioritized over the long knife x .

We divide the unit square into small squares. Let N be an integer. Partition the unit square $[0, 1]^2$ into N^2 smaller squares of side length $\frac{1}{N}$; we call each of these squares a *basic square*. We call each vertical line $L(x) = \{(x, y) \mid y \in [0, 1]\}$ with $x = \frac{k}{N}$ for $k \in \{0, 1, \dots, N\}$ a *basic line*. See Figure 6 in Appendix C.1.

We construct a continuous map \bar{f} from the unit square to the two-dimensional standard simplex, which represents the average preferences of the non-birthday agents. To this end, consider a particular triangulation \mathbb{T} of the unit square; for each basic square, one new edge is added from the top left corner vertex to the bottom right corner vertex. We denote the vertices of the triangulation by $V(\mathbb{T})$. We define $f^{(i)}: V(\mathbb{T}) \rightarrow \Delta^2$ which encodes the preferences of agent i . Namely, we ask each agent i the index j of the layered piece $\mathcal{A}_j(\mathbf{v})$ she prefers in $\mathcal{A}(\mathbf{v})$ and set $f^{(i)}(\mathbf{v})$ to be \mathbf{e}_j . In that case, we says that agent i *colors* the vertex \mathbf{v} with color j . We consider the following tie-breaking rule: a zero-length piece is never chosen and for each $\mathbf{v} = (x, y)$, in case of a tie with piece 1, each agent colors \mathbf{v} with 1; in case of a tie between pieces 2 and 3 only, each agent colors \mathbf{v} with 2 if $x \leq y$ and with 3 otherwise.

Similarly to the proof of Theorem 1, we apply Gale’s technique [11] to our problem. For each non-birthday agent $i \in [n - 1]$, we extend the map $f^{(i)}$ affinely on each simplex of \mathbb{T} . Denote by $\bar{f}^{(i)}$ the affine extension of $f^{(i)}$. We then aggregate these approximate preferences among $n - 1$ non-birthday agents as follows: define $\bar{f}: [0, 1]^2 \rightarrow \Delta^2$ by

$$\bar{f}(x, y) = \frac{1}{n - 1} \sum_{i=1}^{n-1} \bar{f}^{(i)}(x, y),$$

for each $(x, y) \in [0, 1]^2$.

The crucial step in the proof of Theorems 6 consists in establishing that $\omega = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ lies in the image of \bar{f} , and similarly for Theorem 8. We provide a brief description on how this is done, the full proof being given hereafter. Because the preferences are monotone, it is not difficult to see that the image of every basic line by \bar{f} has its endpoint $y = 0$ below the line $B = \{(z_1, z_2, z_3) \mid z_1 = \frac{1}{3}, (z_1, z_2, z_3) \in \Delta^2\}$ and its endpoint $y = 1$ above this line. Here, *above* means $z_1 \geq \frac{1}{3}$ and *below* means $z_1 \leq \frac{1}{3}$. This implies in particular that the image of every basic line by \bar{f} intersects with B (by continuity). Further, we show that the images $\bar{f}(L(0))$ and $\bar{f}(L(1))$ of the sides enjoy symmetry with respect to the line $z_2 = z_3$. The existence of two adjacent vertical lines $L(\frac{k}{N})$ and $L(\frac{k+1}{N})$ enclosing a vertical strip whose image by \bar{f} hits ω then follows. To identify these lines, both mathematically and algorithmically, we need to determine the relative position of the image of any basic line with respect to ω . This is done with the help of its intersection with L or R , where L and R are obtained by splitting B by its middle point ω , i.e., $L = \{(z_1, z_2, z_3) \mid z_2 \leq \frac{1}{3}, (z_1, z_2, z_3) \in B\}$ and $R = \{(z_1, z_2, z_3) \mid z_2 \geq \frac{1}{3}, (z_1, z_2, z_3) \in B\}$.

For Theorems 7 and 8, instead of having the general preferences as in Theorems 6, we have now valuation functions that satisfy monotonicity and the Lipschitz condition with constant K . In a similar spirit to the work of Deng et al. [9], we employ a divide-and-conquer approach to compute an approximate solution. Deng et al. use a triangulation of the standard triangle to encode possible divisions, compute an approximate solution by recursively computing smaller polygons containing such a solution, and finish when the polygon is actually a small triangle of the triangulation. We work instead on the unit square subdivided into basic squares, recursively compute thinner full-height rectangles containing an approximate solution, and end up with a very thin vertical strip in which we identify a basic square containing an approximate solution. In both approaches, the part which is kept at any iteration is identified by computing a sort of “degree” of the boundary, which is related to the intersections with L^δ and R^δ . The uniqueness which will be ensured by Lemma 13 in Section C.1 allows to completely determine these intersections via binary search.

5 Concluding remarks

Theorem 1 states the existence of an envy-free multi-division of a multi-layered cake using $q - 1$ long knives for n agents when q is a prime power. The proof relies crucially on Volovikov’s theorem, a powerful result from equivariant topology. For a number q not equal to a prime power, it is known that Volovikov’s theorem does not hold. Further, for such a q , as already noted in the introduction, there is an example of a one-layered cake-cutting instance with non-hungry choice functions for which no envy-free division among q agents exists [2, 21]. These examples show some limitations of the approach based on Volovikov’s theorem, but do not prohibit the existence of an envy-free division in the valuation function model. Indeed, Avvakumov and Karasev [3] showed that an envy-free division among any number of agents still exists when agents have identical valuations that are not necessarily hungry; it does not seem that this result can be obtained from Volovikov’s theorem. See also the last paragraph of Section 1 in [3]. In the context of multi-layered cake-cutting, the examples of [2, 21] imply that an envy-free multi-division using $q - 1$ long knives may not exist under the choice function model. However, it is still an open question whether or not a counterexample exists for some natural valuation functions.

References

- [1] Megumi Asada, Florian Frick, Vivek Pisharody, Maxwell Polevy, David Stoner, Ling Hei Tsang, and Zoe Wellner. Fair division and generalizations of Sperner- and KKM-type results. *SIAM Journal of Discrete Mathematics*, 32(1):591–610, 2018.
- [2] Sergey Avvakumov and Roman Karasev. Envy-free division using mapping degree. *Mathematika*, 67(1):36–53, Oct 2020.
- [3] Sergey Avvakumov and Roman Karasev. Equipartition of a segment. *Mathematics of Operations Research*, page Advance online publication, 2022.
- [4] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*, pages 416–427, 2016.
- [5] Anders Björner, László Lovász, Sinisa T Vrecica, and Rade T Živaljević. Chessboard complexes and matching complexes. *Journal of the London Mathematical Society*, 49(1):25–39, 1994.
- [6] John Cloutier, Kathryn L Nyman, and Francis E Su. Two-player envy-free multi-cake division. *Mathematical Social Sciences*, 59(1):26–37, 2010.
- [7] Mark De Longueville. *A course in topological combinatorics*. Springer Science & Business Media, 2012.
- [8] Argyrios Deligkas, Aris Filos-Ratsikas, and Alexandros Hollender. Two’s company, three’s a crowd: Consensus-halving for a constant number of agents. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021.
- [9] Xiaotie Deng, Qi Qi, and Amin Saberi. Algorithmic solutions for envy-free cake cutting. *Operations Research*, 60(6):1461–1476, 2012.
- [10] Lester E Dubins and Edwin H Spanier. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1):1–17, 1961.
- [11] David Gale. Equilibrium in a discrete exchange economy with money. *International Journal of Game Theory*, 13(1):61–64, 1984.
- [12] Hadi Hosseini, Ayumi Igarashi, and Andrew Searns. Fair division of time: Multi-layered cake cutting. *CoRR*, abs/2004.13397, 2020. URL <http://arxiv.org/abs/2004.13397>.
- [13] Hadi Hosseini, Ayumi Igarashi, and Andrew Searns. Fair division of time: Multi-layered cake cutting. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 182–188, 2020.
- [14] Duško Jojić, Gaiane Panina, and Rade Živaljević. Splitting necklaces, with constraints. *SIAM Journal on Discrete Mathematics*, 35(2):1268–1286, 2021.
- [15] Nicolas Lebert, Frédéric Meunier, and Quentin Carbonneaux. Envy-free two-player m -cake and three-player two-cake divisions. *Operations Research Letters*, 41(6):607–610, 2013.
- [16] Jiří Matoušek. *Using the Borsuk–Ulam theorem: lectures on topological methods in combinatorics and geometry*. Springer Science & Business Media, 2003.
- [17] Frédéric Meunier and Francis E Su. Multilabeled versions of Sperner’s and Fan’s lemmas and applications. *SIAM Journal on Applied Algebra and Geometry*, 3(3):391–411, 2019.

- [18] Frédéric Meunier and Shira Zerbib. Envy-free cake division without assuming the players prefer nonempty pieces. *Israel Journal of Mathematics*, 234(2):907–925, 2019.
- [19] Kathryn Nyman, Francis E Su, and Shira Zerbib. Fair division with multiple pieces. *Discrete Applied Mathematics*, 283:115–122, 2020.
- [20] James B. Orlin. Max flows in $O(nm)$ time, or better. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, page 765–774, 2013.
- [21] Gaiane Panina and Rade T Živaljević. Envy-free division via configuration spaces. *CoRR*, abs/2102.06886, 2021. URL <http://arxiv.org/abs/2102.06886>.
- [22] Erel Segal-Halevi. Fairly dividing a cake after some parts were burnt in the oven. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1276–1284, 2018.
- [23] Erel Segal-Halevi. Fair multi-cake cutting. *Discrete Applied Mathematics*, 291:15–35, 2021.
- [24] Erel Segal-Halevi and Warut Suksompong. How to cut a cake fairly: A generalization to groups. *The American Mathematical Monthly*, 128(1):79–83, 2021.
- [25] Hugo Steinhaus. Sur la division pragmatique. *Econometrica*, 17:315–319, 1949.
- [26] Walter Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.
- [27] Walter Stromquist. Envy-free cake divisions cannot be found by finite protocols. *The Electronic Journal of Combinatorics*, 15(R11), 2008.
- [28] Francis E Su. Rental harmony: Sperner’s lemma in fair division. *The American Mathematical Monthly*, 106(10):930–942, 1999.
- [29] Alexey Y Volovikov. On a topological generalization of the Tverberg theorem. *Mathematical Notes*, 59(3):324–326, 1996.
- [30] Douglas R Woodall. Dividing a cake fairly. *Journal of Mathematical Analysis and Applications*, 78(1):233–247, 1980.
- [31] Rade T Živaljević. User’s guide to equivariant methods in combinatorics. ii. *Publications de l’Institut Mathématique. Nouvelle Série*, 64:107–132, 1998.

Ayumi Igarashi
 The University of Tokyo
 Tokyo, Japan
 Email: igarashi@mist.i.u-tokyo.ac.jp

Frédéric Meunier
 École Nationale des Ponts et Chaussées
 Paris, France
 Email: frederic.meunier@enpc.fr

A Related work

Early literature on cake-cutting has established the existence of an envy-free contiguous division under *closed* (if the i -th piece is preferred in a convergent sequence of divisions, it is preferred in the limit) and *hungry* preferences (pieces of nonzero-length are preferred over pieces of zero-length). While classical works [26, 30] applied non-constructive topological proofs, Su [28] provided a more combinatorial argument by explicitly using Sperner’s lemma. Recently, the problem of dividing a partially unappetizing cake has attracted a great deal of attention. Here, some agents may find that a part of the cake is unappetizing and prefer nothing, while others may find it tasty (agents may have *non-hungry* preferences). Even in such cases, an envy-free division only using $n - 1$ cuts has been shown to exist for a particular number n of agents under the assumption of closed preferences [2, 14, 18, 21, 22]. The most general result obtained so far is the one by Avvakumov and Karasev [2], who showed the existence of an envy-free division for the case when n is a prime power. Jojić, Panina, and Živaljević [14] gave an alternative proof of the result of Avvakumov and Karasev, by using Volovikov’s theorem [29]. Panina and Živaljević [21] pushed further this technique to refine this kind of results allowing non-hungry preferences.

An important tool in the proof of Avvakumov and Karasev, as well as in the proof of Jojić, Panina, and Živaljević, is *Gale’s averaging trick*, introduced by Gale [11], in the context of an exchange economy. As far as we know, the first paper applying Gale’s averaging trick in the context of envy-free cake-cutting is a paper by Asada et al. [1]. Roughly speaking, it consists in considering an aggregated preference function to which a topological result is applied, and then in recovering information for each agent by applying some flow argument. In the present work, we use Volovikov’s theorem and the Gale’s averaging trick.

In his classical work on cake-cutting, Woodall [30] also proved that an envy-free division can be obtained without knowing one agent’s preference. For example, the cut-and-choose protocol does not need the chooser’s preference to obtain an envy-free division. More generally, for any number n of agents, there is a division of the cake into n contiguous pieces such that whichever piece a birthday agent selects, there is an envy-free assignment of the remaining pieces to the remaining agents. Asada et al. [1] gave a simple combinatorial proof that shows the existence of such division.

The vast majority of the literature on cake-cutting is concerned with an allocation among single agents. A notable exception is the work of Segal-Halevi and Suksompong [24], who introduced the standard cake-cutting problem among groups of agents. They established the existence of an envy-free division among groups of varying sizes, showing that it is possible to partition a cake into q contiguous pieces as well as agents into q groups of any desired sizes, and assign the pieces to the groups so that no agent prefers a piece assigned to another group to the piece assigned to her own group.

In general, there is no finite protocol that computes an exact envy-free division even for three agents [27], though such protocol exists when relaxing the contiguity requirement [4]. Nevertheless, for three agents with monotone valuations, Deng et al. [9] proved that an ε -approximate envy-free division can be computed in logarithmic time of $\frac{1}{\varepsilon}$, while obtaining PPAD-hardness of the same problem for choice functions whose choice is given explicitly by polynomial time algorithms. Our method for establishing the FPTAS exploits monotonicity in agents’ preferences and employs a divide-and-conquer approach similar to that of Deng et al. A difference is that while Deng et al. used a triangulation of the two-dimensional standard simplex, we subdivide the unit square into small squares and compute by binary-search thinner and thinner full-height rectangles containing an approximate solution. Our algorithmic result holds for a more general version, considering simultaneously a birthday agent, groups, and two layers.

In a different context, Deligkas et al. [8] also observed that partial monotonicity is useful to design a similar binary-search algorithm to find a consensus-halving among two agents.

There, one agent has a monotone preference and another agent has a continuous preference. We would like to emphasize, however, that our result for a two-layered cake does not assume continuity of the birthday agent while they assume continuity.

Several papers also studied the fair division problem in which agents divide multiple cakes [6, 15, 19, 23]. This model requires each agent to receive at least one nonempty piece of each cake [6, 15, 19] or receive pieces on as few cakes as possible [23]. On the other hand, our setting requires the allocated pieces to be non-overlapping. Thus, the existence/non-existence of envy-free divisions in one setting does not imply those for another.

B Omitted figures and proofs from Section 3

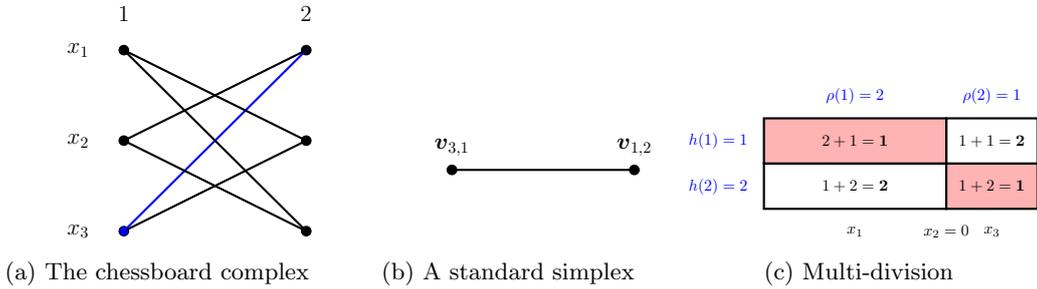


Figure 3: Illustration of the chessboard complex $\Delta_{2q-1,q}$ when $q = 2$. The blue edge of Figure 3a corresponds to a standard simplex, depicted on Figure 3b, which represents the set of multi-divisions of the form described in Figure 3c when $m = 2$ and $h(\ell) = \ell$ for the ℓ -th layer. In Figure 3c, we have $G = (\mathbb{Z}_2, +)$ and $r_{\rho(1)} = r_2 = 1 < r_{\rho(2)} = r_1 = 3$.

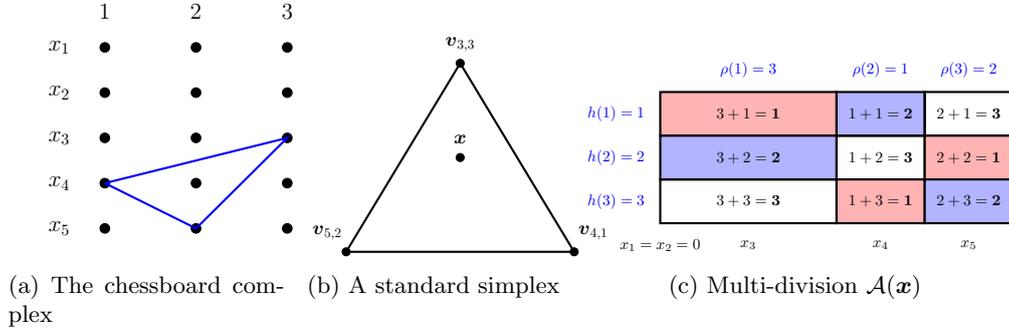


Figure 4: Illustration of the division encoding and of the statement of Lemma 3 in the case of $G = (\mathbb{Z}_3, +)$. Figure 4a is an illustration of the chessboard complex $\Delta_{2q-1,q}$. The blue edges of Figure 4a correspond to a standard simplex, depicted on Figure 4b, which represents multi-divisions of the form described in Figure 4c. In Figure 4c, $r_{\rho(1)} = r_3 = 3 < r_{\rho(2)} = r_1 = 4 < r_{\rho(3)} = r_2 = 5$. The number in the j -th piece of the ℓ -th layer corresponds to $\eta(\rho(j)) + h(\ell)$, where η is the identity map (because when $k = 1$, we can identify $(\mathbb{Z}_p)^k$ and $[p]$).

Lemma 2. *Let $q = p^k$, where p is a prime number and k a positive integer. Denote by G the additive group $((\mathbb{Z}_p)^k, +)$. For any G -equivariant continuous map $f: \Delta_{2q-1,q} \rightarrow \Delta^G$, there exists $x_0 \in \|\Delta_{2q-1,q}\|$ such that $f(x_0) = \frac{1}{q} \sum_{g \in G} e_g$.*

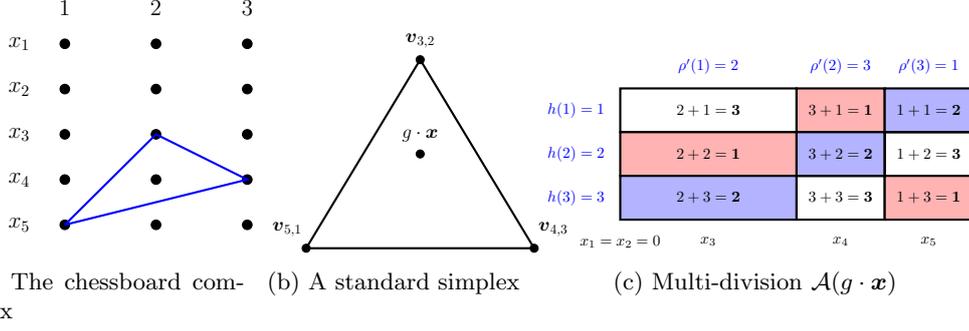


Figure 5: Illustration of the division encoding and of the statement of Lemma 3 in the case of $G = (\mathbb{Z}_3, +)$. The blue edges of Figure 5a correspond to the image of the standard simplex of Figure 4a when $g = 2$. This image is depicted on Figure 5b and corresponds to multi-divisions of the form described in Figure 5c. In Figure 5c, $r'_{\rho'(1)} = r'_2 = 3 < r'_{\rho'(2)} = r'_3 = 4 < r'_{\rho'(3)} = r'_1 = 5$ (with the notation of the proof). The number in the j -th piece of the ℓ -th layer corresponds to $\eta(\rho'(j)) + h(\ell)$, where η is the identity map (because when $k = 1$, we can identify $(\mathbb{Z}_p)^k$ and $[p]$).

Proof. Volovikov's theorem states the following; see [16, Section 6.2, Notes]. *Let G be the additive group $((\mathbb{Z}_p)^k, +)$ and let X and Y be two topological spaces on which G acts in a fixed-point free way. If X is d -connected and Y is a d -dimensional sphere, then there is no G -equivariant continuous map $X \rightarrow Y$.* An action is *fixed-point free* if each orbit has at least two elements. The simplicial complex $\Delta_{2q-1, q}$ is $(q-2)$ -connected (see [5]) and the action is fixed-point free because it is free. The simplicial complex $\partial\Delta^G$ is a $(q-2)$ -dimensional sphere and the action is fixed-point free as we explain now. It is enough to show that for every simplex σ there is an element $g \in G$ such that $g \cdot \sigma$ is distinct from σ because by considering the support of any point \mathbf{x} , this would show that the image of \mathbf{x} by g would be distinct from \mathbf{x} . Consider any simplex σ , pick an arbitrary vertex $g' \in G$ of σ and an arbitrary vertex $g'' \in G$ of $\partial\Delta^G$ not in σ (which exists because the full simplex is not present). Set $g := g'' - g'$. The image of σ by g contains g'' and is thus a simplex distinct from σ .

Suppose for a contradiction that the image of f misses $\frac{1}{q} \sum_{g \in G} \mathbf{e}_g$, which is actually the barycenter of Δ^G . Define then $h(\mathbf{x})$ as the intersection of $\partial\Delta^G$ with the half-line originating at the barycenter of Δ^G and going through $f(\mathbf{x})$. This map contradicts Volovikov's theorem. \square

Lemma 3. *We have $\mathcal{A}_j(\mathbf{x}) = \mathcal{A}_{\eta^{-1}(g+\eta(j))}(g \cdot \mathbf{x})$ for all $\mathbf{x} \in \|\Delta_{2q-1, q}\|$, $j \in [q]$, and $g \in G$.*

Proof. Let $\mathbf{x} \in \|\Delta_{2q-1, q}\|$ and $g \in G$. Write \mathbf{x} as $\sum_{j=1}^q x_{r_j} \mathbf{v}_{r_j, j}$. By the definition of the action of G on $\Delta_{2q-1, q}$, we have $g \cdot \mathbf{v}_{r, j} = \mathbf{v}_{r, \eta^{-1}(g+\eta(j))}$ and thus we have

$$g \cdot \mathbf{x} = \sum_{j=1}^q x_{r'_j} \mathbf{v}_{r'_j, j}, \quad \text{with} \quad r'_j = r_{\eta^{-1}(-g+\eta(j))}. \quad (2)$$

(Note that the sets of q distinct indices r_j and r'_j are the same up to a permutation.) The equivariance of \mathfrak{P} allows the vertices of the simplex chosen by \mathfrak{P} for $g \cdot \mathbf{x}$ to be the $\mathbf{v}_{r'_j, j}$. Denoting by ρ' the permutation such that $r'_{\rho'(1)} < \dots < r'_{\rho'(q)}$, the j -th piece of layer ℓ gets thus $\eta(\rho'(j)) + h(\ell)$ as bundle-name in $\mathcal{A}(g \cdot \mathbf{x})$. Since the sets of indices r_j and r'_j are the same, we have $r_{\rho(j)} = r'_{\rho'(j)}$ for all j , which implies with Equation (2) that $\rho(j) = \eta^{-1}(-g + \eta(\rho'(j)))$. Hence, the j -th piece of layer ℓ gets $g + \eta(\rho(j)) + h(\ell)$ as

bundle-name in $\mathcal{A}(g \cdot \mathbf{x})$. Moreover, in both multi-divisions $\mathcal{A}(\mathbf{x})$ and $\mathcal{A}(g \cdot \mathbf{x})$, the j -th piece of layer ℓ is of the same length $x_{r_{\rho(j)}} = x_{r'_{\rho'(j)}}$. The layered piece with bundle-name $\eta(j)$ for $j \in [q]$ in $\mathcal{A}(\mathbf{x})$ gets thus bundle-name $g + \eta(j)$ in $\mathcal{A}(g \cdot \mathbf{x})$. \square

Lemma 4. *Let $(\mathbf{x}^{(t)})_{t \in \mathbb{Z}_+}$ be a sequence of points of $\|\Delta_{2q-1,q}\|$ converging to some limit point $\mathbf{x}^{(\infty)}$. Then $(\mathcal{A}(\mathbf{x}^{(t)}))_{t \in \mathbb{Z}_+}$ converges to $\mathcal{A}(\mathbf{x}^{(\infty)})$.*

Proof. Consider a $(q-1)$ -dimensional simplex σ of $\Delta_{2q-1,q}$ whose interior contains infinitely many $\mathbf{x}^{(t)}$. (There can actually be several simplices of this kind—this is what makes the proof not completely obvious—but they all contain $\mathbf{x}^{(\infty)}$.) Denote by $\mathbf{v}_{r_1,1}, \mathbf{v}_{r_2,2}, \dots, \mathbf{v}_{r_q,q}$ its vertices. Let ρ be the permutation such that $r_{\rho(1)} < r_{\rho(2)} < \dots < r_{\rho(q)}$. We will define soon a multi-division $\tilde{\mathcal{A}}^\sigma$ such that $d(\tilde{\mathcal{A}}^\sigma, \mathcal{A}(\mathbf{x}^{(\infty)})) = 0$, and that $\lim d(\mathcal{A}(\mathbf{x}^{(t)}), \tilde{\mathcal{A}}^\sigma) = 0$, where the limit is taken over the t for which $\mathbf{x}^{(t)} \in \sigma$. Since this will hold for every simplex σ containing infinitely many $\mathbf{x}^{(t)}$ in its interior, we will get the desired conclusion.

The multi-division $\tilde{\mathcal{A}}^\sigma$ is defined as if σ were the $(q-1)$ -dimensional simplex chosen by \mathfrak{P} for $\mathbf{x}^{(\infty)}$: the quantity $x_{r_{\rho(j)}}^{(\infty)}$ is the length of the j -th piece in any layer (note that ρ depends on σ); the j -th piece of the ℓ -th layer gets $\eta(\rho(j)) + h(\ell)$ as a bundle-name; $\tilde{\mathcal{A}}_j^\sigma$ is obtained by taking all pieces with bundle-name $\eta(j)$. Denote by $L_\ell^j(\mathbf{x}^{(t)})$ the j -th piece in the ℓ -th layer in multi-division $\mathcal{A}(\mathbf{x}^{(t)})$ and by \tilde{L}_ℓ^j the j -th piece in the ℓ -th layer in multi-division $\tilde{\mathcal{A}}^\sigma$. We have

$$d(L_\ell^j(\mathbf{x}^{(t)}), \tilde{L}_\ell^j) \leq \left| \sum_{j'=1}^{j-1} (x_{r_{\rho(j')}}^{(t)} - x_{r_{\rho(j')}}^{(\infty)}) \right| + \left| \sum_{j'=1}^j (x_{r_{\rho(j')}}^{(t)} - x_{r_{\rho(j')}}^{(\infty)}) \right|$$

(by comparing the positions of the left- and right-hand endpoints of $L_\ell^j(\mathbf{x}^{(t)})$ and \tilde{L}_ℓ^j), whose right-hand side converges to 0 when t goes to $+\infty$ by definition of the $x_i^{(\infty)}$'s. Since $L_\ell^j(\mathbf{x}^{(t)})$ and \tilde{L}_ℓ^j get the same bundle-name, we have $\lim d(\mathcal{A}(\mathbf{x}^{(t)}), \tilde{\mathcal{A}}^\sigma) = 0$, where the limit is taken over the t for which $\mathbf{x}^{(t)} \in \sigma$.

The point $\mathbf{x}^{(\infty)}$ belongs to the interior of a unique face τ of σ . Let $\mathbf{v}_{r'_1,1}, \mathbf{v}_{r'_2,2}, \dots, \mathbf{v}_{r'_q,q}$ be the vertices of the $(n-1)$ -dimensional simplex chosen by \mathfrak{P} for $\mathbf{x}^{(\infty)}$. Denote by J the set of indices j for which $\mathbf{v}_{r'_j,j}$ is a vertex of τ . (Note that $|J| = \dim \tau + 1$.) Since τ is unique and its vertices uniquely determined, we have $r_j = r'_j$ for all $j \in J$. Let ρ' be the permutation so that $r'_{\rho'(1)} < r'_{\rho'(2)} < \dots < r'_{\rho'(q)}$. Among them, we have the elements r'_j with $j \in J$, which are the same as the elements r_j with $j \in J$. Denoting by $j_1 < \dots < j_{|J|}$ the elements from $\rho^{-1}(J)$ and by $j'_1 < \dots < j'_{|J|}$ the elements from $\rho'^{-1}(J)$, we have thus $r_{\rho(j_a)} = r'_{\rho'(j'_a)}$ for every $a \in [|J|]$. This implies, since $r_j = r'_j$ for all $j \in J$, that $\rho(j_a) = \rho'(j'_a)$ for all $a \in [|J|]$. Therefore, the a -th piece of nonzero length (of any layer) in $\tilde{\mathcal{A}}^\sigma$ has length $x_{r_{\rho(j_a)}}^{(\infty)}$, which equals $x_{r'_{\rho'(j'_a)}}^{(\infty)}$, length of the a -th piece of nonzero length (of any layer) in $\tilde{\mathcal{A}}(\mathbf{x}^{(\infty)})$. Its bundle-name in $\tilde{\mathcal{A}}^\sigma$ is $\eta(\rho(j_a)) + h(\ell)$, which is equal to $\eta(\rho'(j'_a)) + h(\ell)$, its bundle-name in $\tilde{\mathcal{A}}(\mathbf{x}^{(\infty)})$. Hence, the pieces of nonzero length in any layer are the same in $\mathcal{A}(\mathbf{x}^{(\infty)})$ and $\tilde{\mathcal{A}}^\sigma$ and get the same bundle-names. The pseudo-metric $d(\cdot, \cdot)$ between two pieces of zero-length is 0. Therefore, $d(\tilde{\mathcal{A}}^\sigma, \mathcal{A}(\mathbf{x}^{(\infty)})) = 0$. \square

Lemma 5. *Let n, q be positive integers. Let a_1, a_2, \dots, a_q be nonnegative real numbers summing up to $n-1$. Consider a bipartite graph $H = ([n-1], [q]; E)$ with nonnegative weights w_e on its edges $e \in E$. Suppose $\sum_{e \in \delta_H(i)} w_e = 1$ for each $i \in [n-1]$ and $\sum_{e \in \delta_H(j)} w_e = a_j$ for each $j \in [q]$. Then for every $j^* \in [q]$ there is an assignment $\pi_{j^*}: [n] \rightarrow [q]$ such that*

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H ,
- $|\pi_{j^*}^{-1}(j^*)| = \lfloor a_{j^*} \rfloor + 1$, and
- for each $j \in [q]$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor a_j \rfloor, \lceil a_j \rceil\}$.

Proof. Take any $j^* \in [q]$. Consider the following polytope:

$$P = \left\{ \mathbf{x} \in \mathbb{R}_+^E \left| \sum_{e \in \delta_H(i)} x_e = 1 \text{ for all } i \in [n-1] \text{ and } \lfloor a_j \rfloor \leq \sum_{e \in \delta_H(j)} x_e \leq \lceil a_j \rceil \text{ for all } j \in [q] \right. \right\}.$$

The polytope P is nonempty because it contains $\mathbf{w} = (w_e)_{e \in E}$. Note that since the incidence matrix defining P is totally unimodular, the vertices of P are integral. Now, we claim that there exists an integral vertex $\bar{\mathbf{w}} \in P$ with $\sum_{e \in \delta_H(j^*)} \bar{w}_e = \lfloor a_{j^*} \rfloor$. To see this, if a_{j^*} is an integer, the existence of such an integral vertex is obvious. If a_{j^*} is not an integer, then just note that all integral vertices of P cannot make the previous sum equal to $\lceil a_{j^*} \rceil$ since there exists a point in P not satisfying this equality, namely \mathbf{w} . The coefficient \bar{w}_e belongs to $\{0, 1\}$ for every $e \in E$ because $\sum_{e \in \delta_H(i)} \bar{w}_e = 1$ for each $i \in [n-1]$ and because $\bar{\mathbf{w}}$ is integral. Thus every $i \in [n-1]$ has a unique neighbor $j \in [q]$ with $\bar{w}_{ij} = 1$. Defining $\pi_{j^*}(i)$ as this j for $i \in [n-1]$ and $\pi_{j^*}(n)$ as j^* , we get a desired assignment. \square

C Omitted figures and proofs from Section 4

C.1 Preliminaries

We first establish a useful property of \bar{f} . In standard one-layered cake cutting into three pieces with monotone preferences, agents' preferences are monotone when moving one knife from left to right while fixing the other. In two-layered cake cutting, we have partial monotonicity, that is, agents' preferences exhibit this property when fixing the long knife. A map $g: [0, 1]^2 \rightarrow \Delta^2$ is *vertically monotone* with respect to $x \in [0, 1]$ if the first coordinate of g is nondecreasing when fixing x , i.e., $g_1(x, y') \geq g_1(x, y)$ for every pair $y, y' \in [0, 1]$ with $y' \geq y$.

Lemma 9. *The map \bar{f} is vertically monotone with respect to every $x = \frac{k}{N}$ for $k \in \{0, 1, \dots, N\}$.*

Proof. Fix $x = \frac{k}{N}$ for $k \in \{0, 1, \dots, N\}$. Consider any agent $i \in [n-1]$. On the basic line $\mathbb{L}(x)$, if a vertex is colored with 1, all the vertices with bigger y are colored with 1 as well. This property is due to the tie-breaking rule and the monotonicity of preferences: as y increases from 0 to 1, the piece $\mathcal{A}_1(x, y)$ gets bigger, while the second piece $\mathcal{A}_2(x, y)$ and the third piece $\mathcal{A}_3(x, y)$ do not grow. Thus, since each $\bar{f}^{(i)}$ is an affine extension of $f^{(i)}$, the function $\bar{f}^{(i)}$ is vertically monotone with respect to x . Further, since all the maps $\bar{f}^{(i)}$ are vertically monotone with respect to x , their average \bar{f} also satisfies the same property. \square

Note that in general, this property may not hold for other $x \in [0, 1]$.

The following lemma states that the images of $\mathbb{L}(0)$ and $\mathbb{L}(1)$ are symmetric with respect to the line $\{(z_1, z_2, z_3) \in \Delta^2 \mid z_2 = z_3\}$.

Lemma 10. *For every $y \in [0, 1]$, the following holds:*

$$\bar{f}_1(0, y) = \bar{f}_1(1, y), \quad \bar{f}_2(0, y) = \bar{f}_3(1, y), \quad \bar{f}_3(0, y) = \bar{f}_2(1, y).$$

Proof. Since the partitions induced by $(0, y)$ and $(1, y)$ are symmetric, we have:

- if one of $(0, y)$ and $(1, y)$ is colored with 1 by agent i (i.e., $f^{(i)}(0, y) = e_1$ or $f^{(i)}(1, y) = e_1$), the other is also colored with 1 by i .
- if one of $(0, y)$ and $(1, y)$ is colored with 2 by agent i (i.e., $f^{(i)}(0, y) = e_2$ or $f^{(i)}(1, y) = e_2$), the other is colored with 3 by i .

(When $0 \leq y < 1$, the above symmetry holds because of the tie-breaking rule; in case of ties between 2 and 3, 2 is chosen at $(0, y)$ since $y \geq 0$ while 3 is chosen at $(1, y)$ since $y < 1$. When $y = 1$, it holds because $(0, 1)$ receives color 1 and $(1, 1)$ receives color 1.) \square

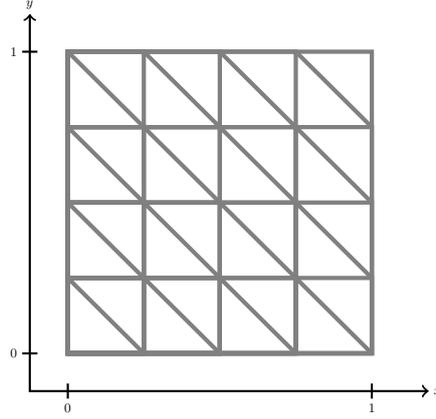


Figure 6: Triangulation of the unit square.

To avoid unnecessary case distinctions, we consider a slightly perturbed version of B , which we denote by B^δ and which is defined by $B^\delta = \{ (z_1, z_2, z_3) \mid z_1 = \frac{1}{3} - \delta, (z_1, z_2, z_3) \in \Delta^2 \}$. We extend the previous notions attached to B . A point (z_1, z_2, z_3) is *above* (resp. *below*) B^δ if $z_1 \geq \frac{1}{3} - \delta$ (resp. $z_1 \leq \frac{1}{3} - \delta$). The point ω^δ is the middle point of B^δ (with coordinates $(\frac{1}{3} - \delta, \frac{1}{3} + \frac{\delta}{2}, \frac{1}{3} + \frac{\delta}{2})$ thus), and L^δ and R^δ are defined by splitting B^δ by ω^δ , i.e., $L^\delta = \{ (z_1, z_2, z_3) \mid z_2 \leq \frac{1}{3} + \frac{\delta}{2}, (z_1, z_2, z_3) \in B^\delta \}$ and $R^\delta = \{ (z_1, z_2, z_3) \mid z_2 \geq \frac{1}{3} + \frac{\delta}{2}, (z_1, z_2, z_3) \in B^\delta \}$.

Now, set $\delta := \frac{1}{6n^2}$. The next two lemmas states that with this small δ , while we get rid of some “degeneracy” (images of vertices located on B), we can safely replace ω by ω^δ in the task of finding a triangle of \mathbb{T} containing ω in its image by \bar{f} .

Lemma 11. *No vertex of \mathbb{T} is mapped by \bar{f} on B^δ . Moreover, if a vertex is mapped by \bar{f} above (resp. below) B^δ , it is mapped above (resp. below) B .*

Proof. Consider a vertex \mathbf{v} of \mathbb{T} . By definition of \bar{f} , there is $\ell_1 \in \{0, 1, \dots, n-1\}$ such that $\bar{f}_1(\mathbf{v}) = \frac{\ell_1}{n-1}$. The quantities $\bar{f}_1(\mathbf{v}) - (\frac{1}{3} - \delta)$ and $6\ell_1 - (2(n-1) - \frac{n-1}{n^2})$ are equal up to a multiplication by $6(n-1)$ and are thus either both positive, or both zero, or both negative. Since $n \geq 2$, the quantity $\frac{n-1}{n^2}$ is not an integer, which implies that $\bar{f}_1(\mathbf{v}) - (\frac{1}{3} - \delta)$ is not 0. This already proves the first part of the statement.

To prove the second part, assume first that $\bar{f}(\mathbf{v})$ is above B^δ , namely that $\bar{f}_1(\mathbf{v}) \geq \frac{1}{3} - \delta$. The same computation as above shows then that $6\ell_1 \geq 2(n-1) - \frac{n-1}{n^2}$. Since $n-1 < n^2$, we have $6\ell_1 \geq 2(n-1)$. This latter inequality translates into $\bar{f}_1(\mathbf{v}) \geq \frac{1}{3}$. The case $\bar{f}(\mathbf{v})$ below B^δ is immediate: the inequality $\bar{f}_1(\mathbf{v}) \leq \frac{1}{3} - \delta$ implies $\bar{f}_1(\mathbf{v}) \leq \frac{1}{3}$. \square

Lemma 12. *For every triangle τ of \mathbb{T} , if ω^δ belongs to $\bar{f}(\tau)$, then so does ω .*

Proof. Consider two vertices \mathbf{v}_1 and \mathbf{v}_2 of \mathbb{T} . The sign of the determinant

$$\varphi(z_1, z_2, z_3) = \begin{vmatrix} z_1 & z_2 & z_3 \\ \bar{f}_1(\mathbf{v}_1) & \bar{f}_2(\mathbf{v}_1) & \bar{f}_3(\mathbf{v}_1) \\ \bar{f}_1(\mathbf{v}_2) & \bar{f}_2(\mathbf{v}_2) & \bar{f}_3(\mathbf{v}_2) \end{vmatrix}$$

records the relative position of a point $\mathbf{z} = (z_1, z_2, z_3) \in \Delta^2$ with respect to the line going through $\bar{f}(\mathbf{v}_1)$ and $\bar{f}(\mathbf{v}_2)$. Since $\bar{f}_j(\mathbf{v})$ is of the form $\frac{\ell_j}{n-1}$ with $\ell_j \in \{0, 1, \dots, n-1\}$ for every $j \in \{1, 2, 3\}$ and every vertex \mathbf{v} of \mathbb{T} , there are three integers a, b, c smaller than n^2 in absolute value such that

$$\varphi(\omega) = \frac{1}{3(n-1)^2}(a+b+c) \quad \text{and} \quad \varphi(\omega^\delta) = \frac{1}{3(n-1)^2} \left((a+b+c) - \frac{1}{4n^2}(2a-b-c) \right).$$

Since each of a, b , and c is smaller than n^2 in absolute value, we have $|2a-b-c| < 4n^2$. In case $\varphi(\omega^\delta) = 0$, then this prevents $a+b+c$ to be different from 0 (it is an integer number), and we have $\varphi(\omega) = 0$ as well. In case $\varphi(\omega^\delta) > 0$, then this prevents $a+b+c \leq -1$ to hold, and we have $\varphi(\omega) \geq 0$. In case $\varphi(\omega^\delta) < 0$, then this prevents $a+b+c \geq 1$ to hold, and we have $\varphi(\omega) \leq 0$.

Now, take a triangle τ of \mathbb{T} such that $\omega^\delta \in \bar{f}(\tau)$. No vertex of τ is mapped on ω^δ by \bar{f} : otherwise, this point would be mapped on B^δ , which is not possible by Lemma 11. So, there are two possibilities: either τ has an edge whose image by \bar{f} is a nondegenerate segment containing ω^δ in its relative interior; or the image of τ by \bar{f} is a nondegenerate triangle containing ω^δ in its relative interior.

Suppose that the first possibility occurs: the point ω^δ lies in the relative interior of $[\bar{f}(\mathbf{v}_1), \bar{f}(\mathbf{v}_2)]$ for some vertices \mathbf{v}_1 and \mathbf{v}_2 (with distinct images by \bar{f}). It means $\varphi(\omega^\delta) = 0$. We have noticed that $\varphi(\omega) = 0$ holds then, which means that ω lies on the line supported by the segment $[\bar{f}(\mathbf{v}_1), \bar{f}(\mathbf{v}_2)]$. Lemma 11 then shows that the relative positions of ω^δ and ω with respect to $\bar{f}(\mathbf{v}_1)$ and $\bar{f}(\mathbf{v}_2)$ are the same, which means that ω is also in $[\bar{f}(\mathbf{v}_1), \bar{f}(\mathbf{v}_2)]$ (but not necessarily in its relative interior).

Suppose then that the second possibility occurs: the point ω^δ lies in the relative interior of $\text{conv}(\bar{f}(\mathbf{v}_1), \bar{f}(\mathbf{v}_2), \bar{f}(\mathbf{v}_3))$ for some vertices $\mathbf{v}_1, \mathbf{v}_2$, and \mathbf{v}_3 (with pairwise distinct images by \bar{f}). For each pair of vertices $\mathbf{v}_i, \mathbf{v}_j$ with $i \neq j$, the determinant $\varphi(\omega^\delta)$ is either positive or negative (but not 0); we have noticed that in the first case $\varphi(\omega) \geq 0$ holds and that in the second case $\varphi(\omega) \leq 0$ holds. This means that, for every edge e of τ , the points ω and ω^δ have the same relative position with respect to the line containing $\bar{f}(e)$. This implies that ω lies in the image of τ by \bar{f} (but not necessarily in its relative interior). \square

The next lemma will be used to ensure the existence of adjacent basic lines forming a strip whose image by \bar{f} hits ω^δ . Uniqueness will be crucial for the algorithmic determination of the relative position of the image of a basic line with respect to ω^δ . Note that this relative position is completely determined by whether $\bar{f}(x, y)$ lies on L^δ or on R^δ .

Lemma 13. *For every $x = \frac{k}{N}$ with $k \in \{0, 1, \dots, N\}$, there exists $y \in [0, 1]$ such that $\bar{f}(x, y)$ lies on B^δ , and this y is unique.*

Proof. We first establish that $\bar{f}(x, 0)$ lies “below” B^δ and that $\bar{f}(x, 1)$ lies “above” B^δ . To see that $\bar{f}(x, 0)$ lies “below” B^δ , i.e., that $\bar{f}_1(x, 0) \leq \frac{1}{3} - \delta$, note that the piece $\mathcal{A}_1(x, 0)$ cannot be strictly preferred on the boundary of $y = 0$ since it is of zero-length. Thus, by the tie-breaking rule, we have $f^{(i)}(x, 0) = \mathbf{e}_2$ or $f^{(i)}(x, 0) = \mathbf{e}_3$ for each agent $i \in [n-1]$, meaning that $\bar{f}_1(x, 0) = \frac{1}{n-1} \sum_{i=1}^{n-1} f_1^{(i)}(x, 0) = 0$.

We show now that $\bar{f}(x, 1)$ lies “above” B^δ , i.e., that $\bar{f}_1(x, 1) \geq \frac{1}{3} - \delta$. Let X be the set of agents whose preferred piece at vertex $(0, 1)$ is 1, i.e., $X = \{i \in [n-1] \mid f^{(i)}(0, 1) = \mathbf{e}_1\}$.

Observe that $|X| \geq \lceil \frac{n}{2} \rceil$ by the tie-breaking rule and by the assumption that the majority of agents weakly prefer the top layer to the bottom layer. By monotonicity of preferences, every agent in X answers $\mathcal{A}_1(x, 1)$ as a preferred piece. Hence,

$$\bar{f}_1(x, 1) = \frac{1}{n-1} \sum_{i \in X} f_1^{(i)}(x, 1) = \frac{|X|}{n-1} \geq \frac{\lceil \frac{n}{2} \rceil}{n-1} > \frac{1}{3} - \delta.$$

By Lemmas 9 and 11, there is a $k' \in \{0, 1, \dots, N-1\}$ such that $\bar{f}_1(x, \frac{k'}{N}) < \frac{1}{3} - \delta$ for every $k \in \{0, 1, \dots, k'\}$ and such that $\bar{f}_1(x, \frac{k}{N}) > \frac{1}{3} - \delta$ for every $k \in \{k'+1, k'+2, \dots, N\}$. There exists thus a single y as in the statement, determined by the intersection of the segment $[\bar{f}(x, \frac{k'}{N}), \bar{f}(x, \frac{k'+1}{N})]$ with the line B^δ . \square

Next, we show the relation between the line B^δ and the image of the boundary of each basic square. It is a well-known fact from topology that a sufficiently regular closed curve, as is the image of the boundary of a basic square by \bar{f} , intersects with any generic straight line an even number of times. The following lemma shows that we can be more precise in our case. We provide a direct proof. An alternate proof based on the mentioned fact from topology would also be possible and short.

Lemma 14. *For a basic square S , either the image of the boundary by \bar{f} does not intersect with B^δ , or there are exactly two points of S , located on two distinct edges, whose image by \bar{f} lies on B^δ .*

Proof. Let $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ be the top left, top right, bottom right, bottom left corner vertices of S , respectively. (We follow a clockwise order.) Recall that \bar{f} is affine, so the image of S 's boundary by \bar{f} is given by the union of $[\bar{f}(\mathbf{a}), \bar{f}(\mathbf{b})]$, $[\bar{f}(\mathbf{b}), \bar{f}(\mathbf{c})]$, $[\bar{f}(\mathbf{c}), \bar{f}(\mathbf{d})]$, and $[\bar{f}(\mathbf{d}), \bar{f}(\mathbf{a})]$. Any of these segments intersects with B^δ if and only if one endpoint is above B^δ and the other is below B^δ . None of these endpoints are located on B^δ by the choice of δ (Lemma 11). There is thus an even number of edges of S whose image by \bar{f} intersects with B^δ (just because such an edge corresponds to going from one side to the other of B^δ).

Suppose for a contradiction this number is equal to four. It means that each segment is intersecting with B^δ , which implies that $\bar{f}(\mathbf{a})$ and $\bar{f}(\mathbf{c})$ are on one side of B^δ , and $\bar{f}(\mathbf{b})$ and $\bar{f}(\mathbf{d})$ are on the other side. Thus, if $\bar{f}(\mathbf{a})$ and $\bar{f}(\mathbf{c})$ are above B^δ , then $\bar{f}(\mathbf{b})$ and $\bar{f}(\mathbf{d})$ are below B^δ ; and if $\bar{f}(\mathbf{a})$ and $\bar{f}(\mathbf{c})$ are below B^δ , then $\bar{f}(\mathbf{b})$ and $\bar{f}(\mathbf{d})$ are above B^δ . In any case, this is a contradiction because of \bar{f} being vertically monotone (see Lemma 9). \square

The following lemma ensures that if the image of the boundary of a basic square intersects with both L^δ and R^δ , it contains a point whose image by \bar{f} is ω^δ .

Lemma 15. *Consider a basic square whose boundary has its image by \bar{f} intersecting with L^δ and with R^δ . Then the image of the basic square by \bar{f} contains ω^δ .*

Proof. Consider a basic square S . Let $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ be the top left, top right, bottom right, bottom left corner vertices of S , respectively. By Lemma 11, none of the points $\bar{f}(\mathbf{a}), \bar{f}(\mathbf{b}), \bar{f}(\mathbf{c}), \bar{f}(\mathbf{d})$ are located on B^δ .

Suppose first that $\bar{f}(\mathbf{a})$ and $\bar{f}(\mathbf{c})$ both lie above B^δ . Since \bar{f} is vertically monotone (see Lemma 9), $\bar{f}(\mathbf{b})$ also lies above B^δ and only the segments with $\bar{f}(\mathbf{d})$ as an endpoint can intersect with B^δ . Thus $\bar{f}(\mathbf{d})$ lies below B^δ . The two points in ∂S whose image by \bar{f} belongs to B^δ (see Lemma 14) are located on $[\mathbf{c}, \mathbf{d}]$ and $[\mathbf{d}, \mathbf{a}]$. By assumption, one image is located on L^δ and the other on R^δ . Since \bar{f} is piecewise affine, there is a point \mathbf{x} in the triangle of T with vertices $\mathbf{a}, \mathbf{c}, \mathbf{d}$ such that $\bar{f}(\mathbf{x}) = \omega^\delta$. When $\bar{f}(\mathbf{a})$ and $\bar{f}(\mathbf{c})$ both lie below B^δ , we get the result similarly.

Finally, suppose that $\bar{f}(\mathbf{a})$ and $\bar{f}(\mathbf{c})$ do not lie on the same side of B^δ . It implies that the segment $[\bar{f}(\mathbf{a}), \bar{f}(\mathbf{c})]$ intersects with B^δ . If it contains the point ω^δ , we are done. We can

thus assume that the intersection of $[\bar{f}(\mathbf{a}), \bar{f}(\mathbf{c})]$ and B^δ is distinct from this point. Consider the case when the intersection is on R^δ . Let τ be the triangle of \mathbb{T} containing the edge of S whose image intersects with L^δ . This triangle has also $[\mathbf{a}, \mathbf{c}]$ as an edge. Since \bar{f} is piecewise affine, there is a point \mathbf{x} in τ such that $\bar{f}(\mathbf{x}) = \omega^\delta$. The case when the intersection is on L^δ is dealt with similarly. \square

See Figure 7 for examples of the image of a basic square by \bar{f} .

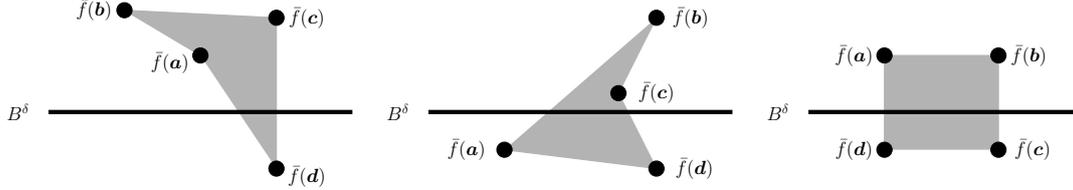


Figure 7: Illustration of the image of a basic square by \bar{f} .

C.2 Existence of an envy-free multi-division

In this section, we prove Theorem 6. The main lemma is the following one. Figure 8 illustrates its proof.

Lemma 16. *There is a point of $[0, 1]^2$ whose image by \bar{f} is ω^δ .*

Proof. By Lemmas 10 and 13, $\bar{f}(L(0)) \cup \bar{f}(L(1))$ intersects with L^δ and with R^δ . Further, by Lemma 13, for each $x = \frac{k}{N}$ with $k \in \{1, 2, \dots, N\}$, the line $\bar{f}(L(x))$ intersects with B^δ at a unique point. Thus, there are two adjacent basic lines $L(x_1)$ and $L(x_2)$ with $x_2 = x_1 + \frac{1}{N}$ or $x_2 = x_1 - \frac{1}{N}$ such that $\bar{f}(L(x_1))$ intersects with L^δ and $\bar{f}(L(x_2))$ intersects with R^δ .

Let e_1 be the vertical edge of $L(x_1)$ such that $\bar{f}(e_1)$ intersects with L^δ and e_2 the vertical edge of $L(x_2)$ such that $\bar{f}(e_2)$ intersects with R^δ . Let S_1 be the basic square that lies between $L(x_1)$ and $L(x_2)$ and contains e_1 and S_2 the basic square that lies between $L(x_1)$ and $L(x_2)$ and contains e_2 . By Lemma 14, for $i = 1, 2$, $|\{\mathbf{x} \in \partial S_i \mid \bar{f}(\mathbf{x}) \in B^\delta\}| = 2$. If the image of the boundary of S_1 by \bar{f} intersects with L^δ and R^δ , Lemma 15 implies that the image of S_1 by \bar{f} contains ω^δ , which proves the claim. Similarly, if the image of the boundary of S_2 by \bar{f} intersects with L^δ and R^δ , we obtain the desired claim.

Thus, assume that the image of S_1 by \bar{f} intersects with L^δ only, and the image of S_2 by \bar{f} intersects with R^δ only. By Lemma 14, this means that there are two edges in S_1 whose image by \bar{f} intersects with L^δ . Likewise, there are two edges in S_2 whose image by \bar{f} intersects with R^δ . Consider the vertical rectangle C of width $\frac{1}{N}$ that consists of the sequence of basic squares lying between S_1 and S_2 . We claim that C admits two horizontal edges of a same basic square whose image by \bar{f} intersects with both L^δ and R^δ and hence contains ω^δ by Lemma 15.

To prove this, recall that by Lemma 13, the image of $L(x_1)$ by \bar{f} intersects with B^δ at a single point. Thus, there is no vertical edge in $L(x_1)$ except for e_1 whose image by \bar{f} intersects with B^δ . Similarly, there is no vertical edge in $L(x_2)$ except for e_2 whose image by \bar{f} intersects with B^δ . Further, the image of the top horizontal edge of C does not intersect with B^δ since \bar{f} is affine on the edge and the corner vertices of the edge have their images lying above B^δ . For the same reasons, the image of the bottom horizontal edge of C does not intersect with B^δ .

By Lemma 14 and by the fact that the image of each S_i ($i = 1, 2$) by \bar{f} intersects with B^δ , all squares in C have two edges whose image intersect with B^δ . Hence, the image of every horizontal edge in C except for the top and bottom ones by \bar{f} intersects with L^δ or R^δ .

This means that when moving from S_1 to S_2 , one necessarily meets a basic square whose horizontal edges have their image by \bar{f} intersecting both L^δ and R^δ . Thus, by Lemma 15, we have found a basic square S^* whose image by \bar{f} contains ω^δ . \square

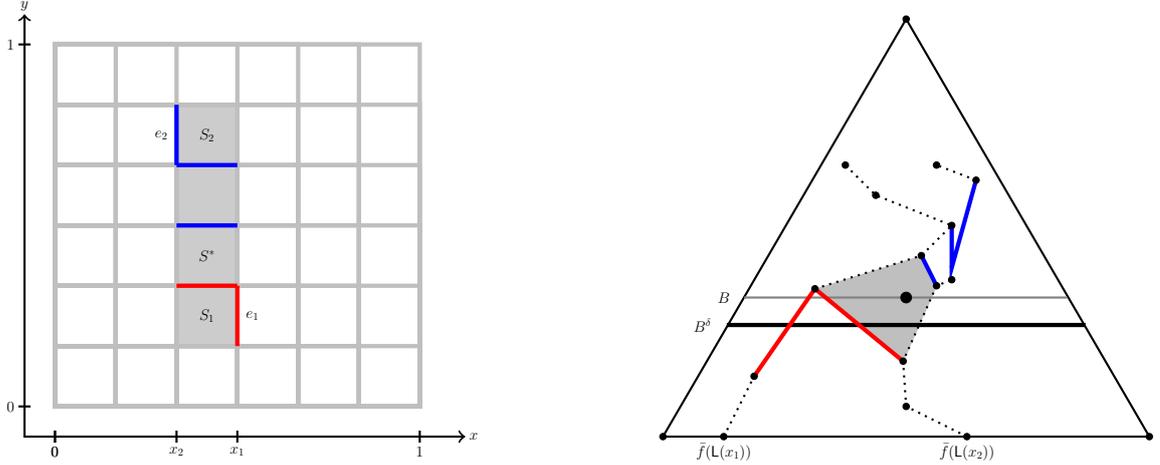


Figure 8: Illustration of the proof of Lemma 16. The left figure depicts the unit square where the gray region corresponds to the rectangle C ; the blue edges correspond to the edges whose image by f intersects with R^δ ; and the red edges correspond to the edges whose image by \bar{f} intersects with L^δ . The right figure depicts the image by \bar{f} of basic lines $L(x_1)$ and $L(x_2)$ as well as the colored edges of the left figure on the standard simplex Δ^2 .

Proof of Theorem 6. The proof is almost verbatim the end of the proof of Theorem 1.

By Lemma 16, there is a point of $[0, 1]^2$ whose image by \bar{f} is ω^δ . According to Lemma 12, this implies the existence of a point $(x^*, y^*) \in [0, 1]^2$ such that $\bar{f}(x^*, y^*) = \omega$. For every non-birthday agent $i \in [n - 1]$ and every layered piece $j \in \{1, 2, 3\}$, we define $w_{ij} = \bar{f}_j^{(i)}(x^*, y^*)$. By definition of \bar{f} and the fact that $\bar{f}_j(x^*, y^*) = \frac{1}{3}$ for all $j = 1, 2, 3$, we have

$$\sum_{i=1}^{n-1} w_{ij} = \sum_{i=1}^{n-1} \bar{f}_j^{(i)}(x^*, y^*) = (n-1) \bar{f}_j(x^*, y^*) = \frac{n-1}{3}$$

for all $j = 1, 2, 3$. Further,

$$\sum_{j=1}^3 w_{ij} = \sum_{j=1}^3 \bar{f}_j^{(i)}(x^*, y^*) = 1$$

for all $i \in [n - 1]$ because $\bar{f}^{(i)}$ has its image in Δ^2 . Consider the bipartite graph $H = ([n - 1], \{1, 2, 3\}; E)$ where the edge ij exists precisely when $w_{ij} > 0$. Applying Lemma 5 to H with $a_j = \frac{n-1}{3}$, for every $j^* \in \{1, 2, 3\}$, there exists an assignment $\pi_{j^*}: [n] \rightarrow \{1, 2, 3\}$ such that

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n - 1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H ,
- $|\pi_{j^*}^{-1}(n)| = \lfloor (n-1)/3 \rfloor + 1$, and
- for each $j \in \{1, 2, 3\}$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor (n-1)/3 \rfloor, \lceil (n-1)/3 \rceil\}$.

Similar to the proof of Theorem 1, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/3 \rfloor, \lceil n/3 \rceil\}$ for all $j \in \{1, 2, 3\}$. The rest of the proof is the same as the last two paragraphs in the proof of Theorem 1. (Especially, we make N go to infinity.) \square

We remark that it is unclear whether the statement of Theorem 6 remains true if we consider more general choice functions. The problem appears from the fact that otherwise the top side $\{(x, 1) \mid x \in [0, 1]\}$ of the unit square could also receive colors 2 or 3.

C.3 FPTAS for a two-layered cake

In this section, we prove Theorem 7.

Take any $\varepsilon > 0$. We define N to be $\lceil \frac{6K}{\varepsilon} \rceil$ and consider the triangulation \mathbb{T} and the map \bar{f} for this N (their definitions depend on N ; see Section 4.1).

We first show that, by the Lipschitz condition of the valuation functions, any vertex of a basic square whose image by \bar{f} contains ω translates into a desired approximate multi-division.

Lemma 17. *Every vertex \mathbf{v} of a basic square whose image by \bar{f} contains ω encodes a multi-division $\mathcal{A}(\mathbf{v})$ such that for every $j^* \in \{1, 2, 3\}$, there is an assignment $\pi_{j^*} : [n] \rightarrow \{1, 2, 3\}$ with:*

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, $v_i(\mathcal{A}_{\pi_{j^*}(i)}(\mathbf{v})) + \varepsilon \geq \max_{j \in \{1, 2, 3\}} v_i(\mathcal{A}_j(\mathbf{v}))$, and
- for each $j \in \{1, 2, 3\}$, $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/3 \rfloor, \lceil n/3 \rceil\}$.

Proof. Consider a basic square containing a point (x^*, y^*) such that $\bar{f}(x^*, y^*) = \omega$, and let \mathbf{v} be any of its vertices. We define $w_{ij} = \bar{f}_j^{(i)}(x^*, y^*)$ for every non-birthday agent $i \in [n-1]$ and every layered piece $j \in \{1, 2, 3\}$. Similar to the proof of Theorem 6, we apply Lemma 5 to the bipartite graph $H = ([n-1], \{1, 2, 3\}; E)$, where the edge ij exists precisely when $w_{ij} > 0$ and $a_j = \frac{n-1}{3}$ for $j \in \{1, 2, 3\}$. Then, for every $j^* \in \{1, 2, 3\}$, we get an assignment $\pi_{j^*} : [n] \rightarrow \{1, 2, 3\}$ such that

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H , and
- for each $j \in \{1, 2, 3\}$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/3 \rfloor, \lceil n/3 \rceil\}$.

Consider any $j^* \in \{1, 2, 3\}$ and any $i \in [n-1]$. Since $\pi_{j^*}(i)$ is a neighbor of i in H , there exists a vertex \mathbf{v}^{i, j^*} of the supporting simplex of (x^*, y^*) in \mathbb{T} such that $f^{(i)}(\mathbf{v}^{i, j^*}) = \mathbf{e}_{\pi_{j^*}(i)}$, meaning that $v_i(\mathcal{A}_{\pi_{j^*}(i)}(\mathbf{v}^{i, j^*})) \geq \max_{j \in \{1, 2, 3\}} v_i(\mathcal{A}_j(\mathbf{v}^{i, j^*}))$. By the Lipschitz condition, we have for every j :

$$|v_i(\mathcal{A}_j(\mathbf{v})) - v_i(\mathcal{A}_j(\mathbf{v}^{i, j^*}))| \leq K \times \mu(\mathcal{A}_j(\mathbf{v}) \Delta \mathcal{A}_j(\mathbf{v}^{i, j^*})) \leq K \times \frac{3}{N} \leq \frac{\varepsilon}{2}.$$

(Here, the ‘3’ in the penultimate term comes from the following: consider two vertices of a same basic square; the layered pieces $j = 1$ in the multi-divisions they encode have a symmetric difference of length 0 or $\frac{1}{N}$; the layered pieces $j = 2$ in the multi-divisions they encode have a symmetric difference of length 0, $\frac{1}{N}$, or $\frac{2}{N}$; same thing for the layered pieces $j = 3$.) We have thus:

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, $v_i(\mathcal{A}_{\pi_{j^*}(i)}(\mathbf{v})) + \varepsilon \geq \max_{j \in \{1, 2, 3\}} v_i(\mathcal{A}_j(\mathbf{v}))$, and

- for each $j \in \{1, 2, 3\}$, we have $|\pi_{j^*}^{-1}(j)| \in \{\lfloor n/3 \rfloor, \lceil n/3 \rceil\}$. \square

The problem thus boils down to finding a basic square whose image by \bar{f} includes ω in logarithmic time of N . In order to design the FPTAS, we exploit the uniqueness of Lemma 13 (which relies on Lemma 9, which eventually relies on the preferences being monotone). This enables us to compute the intersection of the image of a basic line with B^δ , in logarithmic time of N .

For an interval X such that $\min X = \frac{i_1}{N}$ and $\max X = \frac{i_2}{N}$ for $i_1, i_2 \in \{0, 1, \dots, N\}$, we write $\text{med } X = \lceil \frac{i_1 + i_2}{2N} \rceil$.

Consider Algorithm 1. It consists of two stages. First, by binary search (Lines 2 – 6), it identifies the vertical rectangle $X \times [0, 1]$ of horizontal length $\frac{1}{N}$ where the image of one vertical side by \bar{f} intersects with L^δ and the image of the other side by \bar{f} intersects with R^δ . After the first stage of the algorithm, it identifies the vertical edges e_1 and e_2 where these intersections occur; at this stage, we have identified basic squares S_1 and S_2 as in Figure 8. The algorithm then by binary search (Lines 15 – 19) updates a vertical rectangle $X \times Y$. Even if the intersection test is done only with at most six edges at Line 16, Lemma 19 ensures that the boundary of that rectangle keeps its image intersecting with both L^δ and R^δ . The rectangle therefore ends as a basic square of size $\frac{1}{N} \times \frac{1}{N}$ with the same property. By Lemma 15, such a square contains a point whose image is ω^δ , and thus a point whose image is ω by Lemma 12; this implies by Lemma 17 that any vertex of the square corresponds to a desired approximate division. The proof of Theorem 7 consists mainly in verifying that Algorithm 1 computes correctly such a basic square and in analyzing its running time.

According to Lemma 13, the intersection of the image of any basic line with B^δ is unique. It can be computed by binary search.

Lemma 18. *Given any $x = \frac{k}{N}$ for $k \in \{0, 1, \dots, N\}$, one can compute in time $O(n \log N)$ the vertical edge of $\mathbb{L}(x)$ whose image by \bar{f} intersects with L^δ or R^δ .*

Proof. The point $\bar{f}(x, 0)$ lies below B^δ and the point $\bar{f}(x, 1)$ lies above B^δ . Since \bar{f} is continuous, and $\bar{f}(v)$ can be computed in $O(n)$ time, binary search allows to compute the vertical edge of $\mathbb{L}(x)$ whose image intersects with B^δ in $O(n \log N)$ time. \square

The next lemma ensures that an invariant is kept over the second **while** loop. The proof of this lemma is actually very close to that of the end of Lemma 16.

Lemma 19. *At each iteration of the second **while** loop of Lines 15 – 19, the intersections of $\bar{f}(\partial(X \times Y))$ with L^δ and with R^δ are both nonempty.*

Proof. Denote by F the set of edges incident with the corner vertices of $X \times Y$. Notice that except at the end of the last iteration, F contains exactly six edges. We prove by induction that both $\bar{f}(F) \cap L^\delta$ and $\bar{f}(F) \cap R^\delta$ are nonempty. At the first iteration, this is true because F contains all the vertical edges of S_1 and S_2 . We turn now to the subsequent iterations.

To deal with these iterations, we need to establish that all edges inside $X \times Y$, just before the first iteration, have their image by \bar{f} intersecting with B^δ . Before the first iteration of the while loop, the images of the top and bottom edges in F do not intersect with B^δ since \bar{f} is affine on each edge, the corner vertices of the top edge have their images lying above B^δ , and the corner vertices of the bottom edge have their images lying below B^δ . The image of each basic line intersects exactly once with B^δ (Lemma 13) and exactly two edges of each basic square in $X \times Y$ have their image by \bar{f} intersecting with B^δ (Lemma 14). Hence, all the horizontal edges inside $X \times Y$ have their images intersecting with B^δ . This property is obviously kept along all iterations.

Now, consider an arbitrary iteration. Denote by F^{top} the three top edges in F , and by F^{bottom} the three bottom edges in F . By induction, $\bar{f}(F^{\text{top}}) \cup \bar{f}(F^{\text{bottom}})$ intersects with both L^δ and R^δ . Let e_0 be the horizontal edge with the second coordinate equal to $\text{med } Y$.

Algorithm 1: Computing an approximate envy-free multi-division

```
1 set  $X := [0, 1]$ ;
2 while  $\max X - \min X > \frac{1}{N}$  do
3   if  $\bar{f}(L(\min X)) \cup \bar{f}(L(\text{med } X))$  intersects with  $L^\delta$  and with  $R^\delta$  then // This
4     can be done using binary search; see Lemma 18
5      $\lfloor$ set  $X := [\min X, \text{med } X]$ ;
6   else
7      $\lfloor$ set  $X := [\text{med } X, \max X]$ ;
7 compute a basic square  $S_1$  in  $X \times [0, 1]$  which admits a vertical edge  $e_1$  whose
  image by  $\bar{f}$  intersects with  $L^\delta$ ;
  // This can be done using binary search; see Lemma 18
8 compute a basic square  $S_2$  in  $X \times [0, 1]$  which admits a vertical edge  $e_2$  whose
  image by  $\bar{f}$  intersects with  $R^\delta$ ;
  // This can be done using binary search; see Lemma 18
9 if  $S_1 = S_2$  then
10    $\lfloor$ return  $S_1$ ;
11 else
12   set  $x_1 := \min X$  and  $x_2 = \max X$ ;
13   set  $y_1 := \min\{y \mid (x, y) \in S_1 \cup S_2\}$  and  $y_2 := \max\{y \mid (x, y) \in S_1 \cup S_2\}$ ;
14   set  $Y := [y_1, y_2]$ ;
15   while  $\max Y - \min Y > \frac{1}{N}$  do
16     if  $\bigcup_e \bar{f}(e)$  intersects with  $L^\delta$  and with  $R^\delta$ , where  $e$  ranges over the edges
17       incident to the corner vertices of  $X \times [\min Y, \text{med } Y]$  then
18          $\lfloor$ set  $Y := [\min Y, \text{med } Y]$ ;
19       else
20          $\lfloor$ set  $Y := [\text{med } Y, \max Y]$ ;
20   return  $X \times Y$ ;
```

Before any iteration (even the last one), e_0 is inside $X \times Y$ and thus $\bar{f}(e_0)$ has a nonempty intersection with B^δ . If it misses one of L^δ and R^δ , one of $\bar{f}(F^{\text{top}})$ and $\bar{f}(F^{\text{bottom}})$ (at least) does not miss it. So, in any case, $\bar{f}(F^{\text{top}} \cup \{e_0\})$ or $\bar{f}(F^{\text{bottom}} \cup \{e_0\})$ intersects with both L^δ and R^δ . \square

We are ready to prove Theorem 7.

Proof of Theorem 7. We start by proving the correctness of the algorithm. In the first **while** loop of Lines 2 – 6, Lemmas 10 and 13 together ensure that, for the initial interval $X = [0, 1]$, the intersections of $\bar{f}(L(\min X)) \cup \bar{f}(L(\max X))$ with L^δ and with R^δ are both nonempty. By Lemma 13 applied on $\bar{f}(L(\text{med } X))$, this property is kept for each iteration. This means that one can find S_1 and S_2 satisfying the conditions in Lines 7 and 8. If $S_1 = S_2$, Algorithm 1 clearly computes a desired square by Lemma 15. Otherwise, Lemmas 15 and 19 applied at the end of the algorithm ensure that the rectangle $X \times Y$ is a basic square whose image by \bar{f} contains ω^δ and therefore ω by Lemma 12. By Lemma 17, any vertex of such a square induces a desired approximate division.

It remains to analyze the running time of Algorithm 1. The **while** loop in Lines 2 – 6 makes $O(\log N)$ iterations and Line 3 can be implemented in time $O(n \log N)$ by Lemma 18. Thus, Lines 2 – 6 can be implemented in $O(n \log^2 N)$. Again, by Lemma 18, Lines 7 and 8 can be implemented in time $O(n \log N)$. Line 16 can be implemented in $O(n)$ time, because this can be checked by computing $\bar{f}(v)$ of each vertex v that belongs to an edge incident to the corner vertices of $X \times [\min Y, \text{med } Y]$. Further, the **while** loop in Lines 15 – 19 makes $O(\log N)$ iterations. Thus, the algorithm computes a basic square $X \times Y$ of size $\frac{1}{N} \times \frac{1}{N}$ in $O(n \log^2 N)$ time. \square

C.4 FPTAS for a one-layered cake

In this section, we provide a proof sketch for Theorem 8. Again, we assume that the n -th agent is the birthday agent. We keep the same tools as in Section 4.1: the unit square, the triangulation \mathbb{T} , and the function \bar{f} . We also take the same N as in Section C.3. With monotonicity of preferences, the one-layered case is equivalent to the two-layered problem where the bottom layer is empty, which means that all agents weakly prefer the top layer to the bottom. This will enable us to establish that any point $(a_1, a_2, a_3) \in \Delta^2$ lies in the image of \bar{f} , i.e., \bar{f} is surjective. This surjectivity will allow eventually, via standard arguments, to show that any choice for k_1, k_2, k_3 is achievable.

Proof sketch of Theorem 8. Let $\omega = \frac{1}{n-1}(k_1 - \frac{1}{3}, k_2 - \frac{1}{3}, k_3 - \frac{1}{3}) \in \Delta^2$. We define $B = \{(z_1, z_2, z_3) \mid z_1 = \omega_1, (z_1, z_2, z_3) \in \Delta^2\}$. Again, L and R are obtained by splitting B by ω . Like Lemmas 11 and 12, we choose a sufficiently small $\delta > 0$ to obtain a perturbed version B^δ of B so that no vertex in \mathbb{T} has its image by \bar{f} lying on B^δ and so that if the middle point ω^δ of B^δ belongs to $\bar{f}(\tau)$ for some triangle τ in \mathbb{T} , then ω also belongs to $\bar{f}(\tau)$. Define L^δ and R^δ , analogously as in Section C.2, by splitting B^δ by ω^δ .

We have $\bar{f}_1(x, 1) = 1$ for every $x = \frac{k}{N}$ with $k \in \{0, 1, \dots, N\}$ since every agent in $[n-1]$ answers $\mathcal{A}_1(x, 1)$ as a preferred piece. We have $\bar{f}_1(x, 0) = 0$ for every $x = \frac{k}{N}$ with $k \in \{0, 1, \dots, N\}$; indeed, such a piece is never chosen by the fact that $\mathcal{A}_1(x, 0)$ is of zero-length and by the tie-breaking rule of Section 4.1. (When the valuation functions are monotone and Lipschitz, the hungry assumption is satisfied; see Section 2.) Similarly, we have $\bar{f}_2(0, y) = 0$ (resp. $\bar{f}_3(1, y) = 0$) for every $y = \frac{k}{N}$ with $k \in \{0, 1, \dots, N\}$. (From this, we could already conclude via standard techniques from topology that \bar{f} is surjective, but we provide an algorithmic proof in the sequel.)

As we have just seen, $\bar{f}(x, 1)$ and $\bar{f}(x, 0)$ lie above and below B^δ for each $x = \frac{k}{N}$ with $k \in \{0, 1, \dots, N\}$, respectively. This implies, together with the vertical monotonicity of

\bar{f} (Lemma 9), that $\bar{f}(L(x))$ intersects with B^δ at a single point for every $x = \frac{k}{N}$ with $k \in \{0, 1, \dots, N\}$. Thus, we obtain a statement corresponding to Lemma 13.

Also the images $\bar{f}(L(0))$ and $\bar{f}(L(1))$ of the sides are symmetric with respect to $z_2 = z_3$ (Lemma 10) because $\bar{f}(L(0))$ (resp. $\bar{f}(L(1))$) coincides with the boundaries of Δ^2 where the third (resp. the second) coordinate is 0. It is easy to see that Lemmas 14 and 15 remain true for the new B^δ . Algorithm 1 still applies, Lemmas 18 and 19 remain true, and the proof of Theorem 7 is still valid. Thus, we are able to find a basic square whose image by \bar{f} contains ω in $O(n \log^2 N)$.

Let (x^*, y^*) be a point in this square mapped by \bar{f} on ω . To see that such a square induces a desired multi-division, define $w_{ij} = \bar{f}_j^{(i)}(x^*, y^*)$ for every non-birthday agent $i \in [n-1]$ and a layered piece $j \in \{1, 2, 3\}$. By definition of \bar{f} and the fact that $\bar{f}_j(x^*, y^*) = \omega_j = \frac{1}{n-1}(k_j - \frac{1}{3})$ for all $j = 1, 2, 3$, we have

$$\sum_{i=1}^{n-1} w_{ij} = \sum_{i=1}^{n-1} \bar{f}_j^{(i)}(x^*, y^*) = (n-1)\bar{f}_j(x^*, y^*) = k_j - \frac{1}{3}$$

for all $j = 1, 2, 3$. Further, $\sum_{j=1}^3 w_{ij} = \sum_{j=1}^3 \bar{f}_j^{(i)}(x^*, y^*) = 1$, for all $i \in [n-1]$ because $\bar{f}^{(i)}$ has its image in Δ^2 .

Now, consider the bipartite graph $H = ([n-1], \{1, 2, 3\}; E)$ where the edge ij exists precisely when $\bar{f}_j^{(i)}(x^*, y^*) > 0$. Applying Lemma 5 to H with $a_j = k_j - \frac{1}{3}$ for $j \in \{1, 2, 3\}$, for every $j^* \in \{1, 2, 3\}$, there exists an assignment $\pi_{j^*}: [n] \rightarrow \{1, 2, 3\}$ such that

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, the vertex $\pi_{j^*}(i)$ is a neighbor of i in H ,
- $|\pi_{j^*}^{-1}(n)| = k_{j^*}$, and
- for each $j \in \{1, 2, 3\}$, we have $|\pi_{j^*}^{-1}(j)| \in \{k_j - 1, k_j\}$.

The last two conditions imply that $|\pi_{j^*}^{-1}(j)| = k_j$ for each $j \in \{1, 2, 3\}$ since $\sum_{j \in \{1, 2, 3\}} |\pi_{j^*}^{-1}(j)| = n = \sum_{j \in \{1, 2, 3\}} k_j$. Then, similar to the proof of Lemma 17, for any vertex $\mathbf{v} \in V(\mathbb{T})$ that belongs to the basic square containing (x^*, y^*) , we have a multi-division $\mathcal{A}(\mathbf{v})$ such that for every $j^* \in \{1, 2, 3\}$, it admits a desired assignment π_{j^*} where

- $\pi_{j^*}(n) = j^*$,
- for each $i \in [n-1]$, $v_i(\mathcal{A}_{\pi_{j^*}(i)}(\mathbf{v})) + \varepsilon \geq \max_{j \in \{1, 2, 3\}} v_i(\mathcal{A}_j(\mathbf{v}))$, and
- for each $j \in \{1, 2, 3\}$, $|\pi_{j^*}^{-1}(j)| = k_j$. □

D Proportional multi-division

In this section, we discuss proportionality. Consider an instance of the multi-layered cake-cutting problem with m layers and n agents, $m \leq n$. Let q be an integer such that $m \leq q \leq n$. A multi-division into q layered pieces is *proportional* if there exists a surjective assignment π of the agents to the pieces such that each agent i is assigned a layered piece of value at least her *proportional fair share*, that is, $\frac{1}{q}$ of i 's valuation for the entire layered cake.

A valuation function v_i is *additive* if for any pair of layered pieces \mathcal{L} and \mathcal{L}' whose interiors are disjoint, we have $v_i(\mathcal{L} \cup \mathcal{L}') = v_i(\mathcal{L}) + v_i(\mathcal{L}')$, where the union $\mathcal{L} \cup \mathcal{L}'$ of two layered pieces $\mathcal{L} = (L_\ell)_{\ell \in [m]}$ and $\mathcal{L}' = (L'_\ell)_{\ell \in [m]}$ is defined to be $(L_\ell \cup L'_\ell)_{\ell \in [m]}$. Under additive valuations, it is easy to see that envy-freeness implies proportionality.

Proposition 20. *Consider an instance of the multi-layered cake-cutting problem with m layers and n agents, $m \leq n$, with additive valuations. Any envy-free multi-division into q layered pieces with the corresponding assignment is proportional.*

Proof. Take any envy-free multi-division $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_q)$ and the corresponding assignment $\pi: [n] \rightarrow [q]$. Consider any agent i . By envy-freeness, $v_i(\mathcal{A}_{\pi(i)}) \geq v_i(\mathcal{A}_j)$ for all $j \in [q]$. Summing these inequalities, we have $q \cdot v_i(\mathcal{A}_{\pi(i)}) \geq \sum_{j=1}^q v_i(\mathcal{A}_j)$, which by additivity, implies that the valuation of the piece assigned to i is greater or equal to the proportional fair share for i . \square

Proposition 20 and Theorem 1 imply that there exists a proportional multi-division into q layered pieces that is feasible and contiguous and satisfies the property that the number of agents assigned to each layered piece differs by at most one when q is a prime power, $m \leq q \leq n$, and agents have additive valuations. It turns out that such a division exists even when q is not a prime power.

Theorem 21. *Consider an instance of the multi-layered cake-cutting problem with m layers and n agents, $m \leq n$, with additive valuations. Let q be an integer such that $m \leq q \leq n$. Then there exist a feasible and contiguous multi-division into q layered pieces and an assignment of the agents to the layered pieces so that*

- *each agent is assigned to a layered piece of value at least her proportional fair share.*
- *the number of agents assigned to each layered piece differs by at most one.*

Proof. Given an instance of the multi-layered cake-cutting problem with m layers and n agents, $m \leq n$, recursively repeat the following steps until q reduces to a prime number:

Step 1. Choose any prime power q' in the prime decomposition of q . Construct a new cake C' of q' layers. Each layer ℓ' of the new cake is obtained by merging $c_{\ell'}$ consecutive layers of the original instance where $c_{\ell'} \in \{\lfloor m/q' \rfloor, \lceil m/q' \rceil\}$ for each $\ell' \in [q']$ and $\sum_{\ell'=1}^{q'} c_{\ell'} = m$. Define a new additive valuation of each agent $i \in [n]$ for the new instance. Namely, for a layered piece $\mathcal{L}' = (L'_{\ell'})_{\ell' \in [q']}$ of the new instance, i assigns to each $L'_{\ell'}$ the sum of values of the pieces that correspond to $L'_{\ell'}$ in $(\sum_{k=1}^{\ell'-1} c_k + 1)$ -th, \dots , $(\sum_{k=1}^{\ell'} c_k)$ -th layers of the original instance; and i assigns the sum of such values to \mathcal{L}' .

Step 2. Apply Theorem 1 to this cake C' with the whole set of agents, the number of layers being q' , and the number of groups being q' , and obtain an envy-free multi-division $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2, \dots, \mathcal{A}'_{q'})$ among q' groups of agents and an assignment $\pi': [n] \rightarrow [q']$ satisfying the conditions in Theorem 1. See Figure 9 for an illustration of the merge and division \mathcal{A}' when $m = 5$, $n = 13$, and $q = 6$.

Step 3. For each $j' \in [q']$, create a new multi-layered cake $C_{j'}$ by concatenating the allocated layered pieces in $\mathcal{A}'_{j'} = (L'_{\ell'})_{\ell' \in [q']}$ into a multi-layered cake of $\lceil m/q' \rceil$ layers: Each ℓ -th layer of the multi-layered cake $C_{j'}$ is obtained by concatenating the ℓ -th layers of the original pieces corresponding to each $L'_{\ell'}$ for $\ell' \in [q]$; in case $c_{\ell'} < \ell$, complete $L'_{\ell'}$ with empty layers with valuation 0 for all agents so that we can still consider that all layers of $C_{j'}$ are $[0, 1]$ intervals. See Figure 9 for an illustration of the concatenation. For each $j' \in [q']$, apply induction (i.e., Theorem 21) on $C_{j'}$ with the agent set $\pi'^{-1}(j')$ and the number of groups equal to q/q' . By collecting the multi-divisions and assignments for all $j' \in [q']$, obtain the final multi-division \mathcal{A} and the final assignment $\pi: [n] \rightarrow [q]$ for the original instance.

We show that \mathcal{A} and π form a desired solution by an induction on the size of the prime decomposition of q . If q is a prime power, then by Proposition 20, this is the case. Suppose that $q = dq'$ where q' is a prime power in the prime decomposition of q . Clearly, \mathcal{A} is contiguous and feasible by the induction hypothesis and by the fact that a new multi-layered cake $C_{j'}$ in Step 3 is obtained by concatenating contiguous pieces of m distinct layers of the original instance. Further, since \mathcal{A}' obtained in Step 2 is envy-free, every agent i has value at least $\frac{1}{q'}\alpha_i$ for the assigned layered piece $\mathcal{A}'_{\pi'(i)}$ in C' , where α_i denotes i 's value for the entire original layered cake (Proposition 20); also, by induction hypothesis, she obtains a piece of value at least $\frac{1}{d}$ of i 's valuation for $\mathcal{A}'_{\pi'(i)}$. Thus, the piece $\mathcal{A}_{\pi(i)}$ assigned to i has a value at least her proportional fair share. Finally, we show that the number of agents assigned to each layered piece \mathcal{A}_j is $\lfloor n/q \rfloor$ or $\lceil n/q \rceil$. By Theorem 1, $|\pi'^{-1}(j')| \in \{\lfloor n/q' \rfloor, \lceil n/q' \rceil\}$ for each $j' \in [q']$. By the induction hypothesis, the number of agents $|\pi^{-1}(j)|$ assigned to a layered piece \mathcal{A}_j is of the form $\lfloor |\pi'^{-1}(j')|/d \rfloor$ or $\lceil |\pi'^{-1}(j')|/d \rceil$ for some $j' \in [q']$. Thus, we have

$$\left\lfloor \frac{1}{d} \left\lfloor \frac{n}{q'} \right\rfloor \right\rfloor \leq |\pi^{-1}(j)| \leq \left\lceil \frac{1}{d} \left\lceil \frac{n}{q'} \right\rceil \right\rceil.$$

Now, since $q = dq'$ and since

$$\left\lfloor \frac{1}{c} \left\lfloor \frac{a}{b} \right\rfloor \right\rfloor = \left\lfloor \frac{a}{bc} \right\rfloor \quad \text{and} \quad \left\lceil \frac{1}{c} \left\lceil \frac{a}{b} \right\rceil \right\rceil = \left\lceil \frac{a}{bc} \right\rceil$$

hold for all integers a, b, c , we have $|\pi^{-1}(j)| \in \{\lfloor n/q \rfloor, \lceil n/q \rceil\}$ for each $j \in [q]$. \square

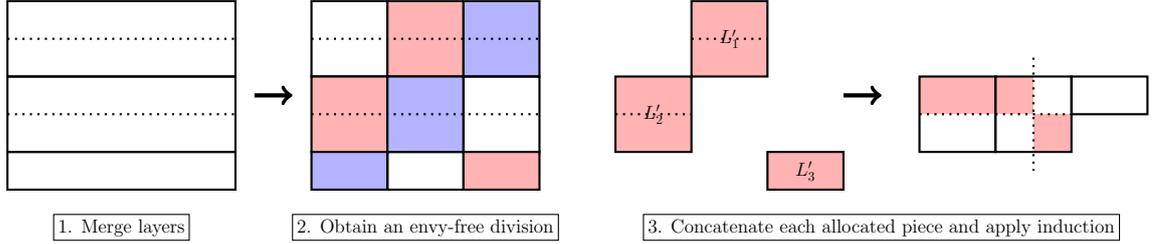


Figure 9: Illustration of the proof of Theorem 21 when $m = 5$, $n = 13$, and $q = 6 (= 2 \cdot 3)$. Here, we take $q' = 3$.

E Some basic notions from combinatorial topology

E.1 Abstract and geometric simplicial complexes

Given a finite set V , a collection K of subsets of V is an *abstract simplicial complex* if whenever $\sigma \in \mathsf{K}$ and $\tau \subseteq \sigma$, then $\tau \in \mathsf{K}$. The elements of K are the *simplices* of the abstract simplicial complex. The elements of V occurring in at least one simplex of K are the *vertices* of K and their set is denoted by $V(\mathsf{K})$.

A finite collection Γ of geometric simplices (convex hulls of affinely independent points) is a *geometric simplicial complex* if the following two conditions are satisfied:

- the faces of every σ in Γ are all in Γ .
- for every pair σ, σ' of simplices in Γ , the intersection of σ and σ' is a face of both σ and σ' .

The empty set is considered as a face of any simplex. The *underlying space* of a geometric simplicial complex Γ is the set of all points contained in at least one simplex of Γ . It is denoted by $\|\Gamma\|$ and we have thus $\|\Gamma\| = \bigcup_{\sigma \in \Gamma} \sigma$.

There is a strong relation between abstract and geometric simplicial complexes. The vertex sets of the simplices of a geometric simplicial complex form an abstract simplicial complex. The former is then a *geometric realization* of the latter. On the other hand, any abstract simplicial complex admits a geometric realization by taking generic representatives of its vertices in a sufficiently large Euclidean space. The equivalence between abstract and geometric simplicial complexes makes that one often switches from one point of view to the other without further mention, and the same notation K may be used for an abstract simplicial complex and its geometric realization.

E.2 Simplicial maps and triangulations

Given two abstract simplicial complexes K_1 and K_2 , a map $\lambda: V(K_1) \rightarrow V(K_2)$ is a *simplicial map* of K_1 into K_2 if $\lambda(\sigma) \in K_2$ for every $\sigma \in K_1$. A bijective simplicial map λ of K_1 into K_2 is an *isomorphism* if its inverse map is also a simplicial map. It is an *automorphism* if moreover $K_1 = K_2$.

Any such simplicial map λ induces a natural continuous map $\|\Gamma_1\| \rightarrow \|\Gamma_2\|$, where Γ_1 and Γ_2 are geometric realizations of respectively K_1 and K_2 , as we explain now. The *affine extension* of λ , denoted by $\bar{\lambda}$, is the map $\|\Gamma_1\| \rightarrow \|\Gamma_2\|$ obtained by extending the original λ affinely to the relative interiors of the simplices of Γ_1 : given a simplex σ in Γ_1 with vertices $\mathbf{v}_1, \dots, \mathbf{v}_k$, a point $\mathbf{x} = \sum_{i=1}^k \alpha_i \mathbf{v}_i$ in the relative interior of σ with $\alpha_1, \dots, \alpha_k \geq 0$, and $\sum_{i=1}^k \alpha_i = 1$, we define $\bar{\lambda}(\mathbf{x})$ as the point $\sum_{i=1}^k \alpha_i \lambda(\mathbf{v}_i)$. It can be proved that for every such simplicial map λ , the map $\bar{\lambda}$ is continuous. If λ is an isomorphism, then $\bar{\lambda}$ is a homeomorphism.

A *triangulation* of a topological space X is a simplicial complex (abstract or geometric) whose underlying space is homeomorphic to X . By a natural extension, a triangulation of a simplicial complex is defined as a triangulation of its underlying space.

E.3 Group acting on simplicial complexes and topological spaces

Let (G, \cdot) be a finite group, whose neutral element is denoted by e .

An *action* of G on an abstract simplicial complex K is a collection $(\varphi_g)_{g \in G}$ of simplicial automorphisms φ_g of K such that

- $\varphi_e = \text{id}_K$ (where id_K is the simplicial identity map of K), and
- $\varphi_g \circ \varphi_h = \varphi_{g \cdot h}$ for all $g, h \in G$.

Similarly, an *action* of G on a topological space X is a collection $(\psi_g)_{g \in G}$ of homeomorphisms $\psi_g: X \rightarrow X$ such that

- $\psi_e = \text{id}_X$, and
- $\psi_g \circ \psi_h = \psi_{g \cdot h}$ for all $g, h \in G$.

For a point $x \in X$, the set $\{\psi_g(x) : g \in G\}$ is the *orbit* of the point x under the G -action.

Note that if $(\varphi_g)_{g \in G}$ is an action of G on an abstract simplicial complex K with geometric realization Γ , then $(\bar{\varphi}_g)_{g \in G}$ is an action of G on $\|\Gamma\|$. This latter action is considered without further mention when the action has only been introduced for an abstract simplicial complex and when the underlying space of a geometric realization is then considered.

Sometimes, when there is no ambiguity, $\varphi_g(v)$ and $\varphi_g(x)$ are simply denoted by $g \cdot v$ and $g \cdot x$.

E.4 Connectivity

A topological space X is d -connected if every continuous map $\bar{f}: S^k \rightarrow X$ with $k \in \{-1, 0, \dots, d\}$ can be extended to a continuous map $\bar{f}: B^{k+1} \rightarrow X$. Here, S^k denotes the k -dimensional sphere and B^k the k -dimensional ball. In this context, S^{-1} is interpreted as \emptyset and B^0 as a single point. Therefore, (-1) -connected means nonempty.