



# Almost envy-free allocations with connected bundles <sup>☆,☆☆</sup>

Vittorio Bilò <sup>a</sup>, Ioannis Caragiannis <sup>b</sup>, Michele Flammini <sup>c</sup>, Ayumi Igarashi <sup>d</sup>,  
Gianpiero Monaco <sup>e</sup>, Dominik Peters <sup>f</sup>, Cosimo Vinci <sup>c,\*</sup>, William S. Zwicker <sup>g</sup>

<sup>a</sup> University of Salento, Lecce, Italy

<sup>b</sup> Aarhus University, Aarhus, Denmark

<sup>c</sup> Gran Sasso Science Institute, L'Aquila, Italy

<sup>d</sup> National Institute of Informatics, Tokyo, Japan

<sup>e</sup> University of L'Aquila, L'Aquila, Italy

<sup>f</sup> University of Toronto, Toronto, Canada

<sup>g</sup> Union College, Schenectady, USA



## ARTICLE INFO

### Article history:

Received 7 January 2021

Available online 30 November 2021

### Keywords:

Envy-free division

Cake-cutting

Resource allocation

Algorithmic game theory

## ABSTRACT

We study the existence of allocations of indivisible goods that are envy-free up to one good (EF1), under the additional constraint that each bundle needs to be connected in an underlying item graph. If the graph is a path and the utility functions are monotonic over bundles, we show the existence of EF1 allocations for at most four agents, and the existence of EF2 allocations for any number of agents; our proofs involve discrete analogues of the Stromquist's moving-knife protocol and the Su–Simmons argument based on Sperner's lemma. For identical utilities, we provide a polynomial-time algorithm that computes an EF1 allocation for any number of agents. For the case of two agents, we characterize the class of graphs that guarantee the existence of EF1 allocations as those whose biconnected components are arranged in a path; this property can be checked in linear time.

© 2021 Elsevier Inc. All rights reserved.

## 1. Introduction

A famous literature considers the problem of *cake-cutting* (Brams and Taylor, 1996; Robertson and Webb, 1998; Procaccia, 2016). There, a divisible heterogeneous resource (a *cake*, usually formalized as the interval  $[0, 1]$ ) needs to be divided among  $n$  agents. Each agent has a valuation function over subsets of the cake, usually formalized as an atomless measure over  $[0, 1]$ . The aim is to partition the cake into  $n$  pieces, and allocate each piece to one agent, in a “fair” way. By fair, we will mean that the allocation is *envy-free*: no agent thinks that another agent's piece is more valuable than her own.

<sup>☆</sup> A preliminary version of this paper (Bilò et al., 2019) appeared in “Proceedings of the 10th Innovations in Theoretical Computer Science (ITCS)”.

<sup>☆☆</sup> D. Peters was supported by ERC grant 639945 (ACCORD). A. Igarashi is supported by KAKENHI (Grant-in-Aid for JSPS Fellows, No. 18J00997) Japan. While working on this paper, W.S. Zwicker was supported by the Oliver Smithies Visiting Fellowship at Balliol College, Oxford. M. Flammini and C. Vinci are supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Markets” (grant No. 2017R9FHSR\_002).

\* Corresponding author.

E-mail addresses: vittorio.bilo@unisalento.it (V. Bilò), iannis@cs.au.dk (I. Caragiannis), michele.flammini@gssi.it (M. Flammini), ayumi.igarashi@nii.ac.jp (A. Igarashi), gianpiero.monaco@univaq.it (G. Monaco), dominik@cs.toronto.edu (D. Peters), cosimo.vinci@gssi.it (C. Vinci), zwickerw@union.edu (W.S. Zwicker).

When there are two agents, the classic procedure of cut-and-choose can produce an envy-free division: a knife is moved from left to right, until an agent shouts to indicate that she thinks the pieces to either side are equally valuable. The other agent then picks one of the pieces, leaving the remainder for the shouter. As is easy to see, the result is an envy-free allocation. For three or more agents, finding an envy-free division has turned out to be much trickier. An early result by Dubins and Spanier (1961) used Lyapunov's Theorem and measure-theoretic techniques to show, non-constructively, that an envy-free allocation always exists. However, as Stromquist (1980) memorably writes, "their result depends on a liberal definition of a 'piece' of cake, in which the possible pieces form an entire  $\sigma$ -algebra of subsets. A player who only hopes for a modest interval of cake may be presented instead with a countable union of crumbs." In many applications of resource allocation (such as land division, or the allocation of time slots), agents have little use for a severely disconnected piece of cake.

Stromquist (1980) himself offered a solution, and gave a new non-constructive argument (using topology) which proved that there always exists an envy-free division of the cake into *intervals*. Forest Simmons later observed that the proof could be simplified by using Sperner's lemma, and this technique was subsequently presented in a paper by Su (1999). For the three-agent case, Stromquist (1980) also presented an appealing moving-knife procedure that more directly yields a connected envy-free allocation. For  $n \geq 4$  agents, no explicit procedures are known to produce a connected envy-free allocation (i.e., an allocation where the cake is cut in exactly  $n - 1$  places). However, for  $n = 4$ , several moving-knife procedures exist that only need a few cuts; for example, the Brams–Taylor–Zwicker (1997) procedure requires 11 cuts, and a protocol of Barbanel and Brams (2004) requires 5 cuts.

In many applications, the resources to be allocated are not infinitely divisible, and we face the problem of allocating *indivisible goods*. Most of the literature on indivisible goods has not assumed any kind of structure on the item space, in contrast to the rich structure of the interval  $[0, 1]$  in cake-cutting. Thus, there has been little attention on minimizing the number of "cuts" required in an allocation. However, when the items have a spatial or temporal structure, this consideration is important.

In this paper, we study the allocation of items that are arranged on a *path* or other structure, and impose the requirement that only *connected* subsets of items may be allocated to the agents. Formally, we work in the model of Bouveret et al. (2017), who assume that the items form the vertex set of a graph  $G$ , and a bundle is connected if it induces a connected subgraph of  $G$ . For example, such connectivity requirements are encountered in allocation problems in which the items correspond to indivisible pieces of an underlying region of Euclidean space (such as plots of land), and each allocated bundle (i.e., a collection of pieces) must be a connected portion of the space. We are most interested in the case when  $G$  is a path, on which connectivity requirements are natural when the items are time slots, for example. In practice, connectivity may not be a hard constraint, and agents may find bundles with few connected components acceptable, but we will not consider such relaxations.

In the work of Bouveret et al. (2017), it became apparent that techniques from cake-cutting can be usefully ported to achieve good connected allocations in the indivisible case. For example, moving-knife procedures that achieve proportionality in cake-cutting have analogues that produce allocations that satisfy the *maximin share guarantee* (Budish, 2011).<sup>1</sup>

Do envy-free procedures for cake-cutting also translate to the indivisible case? Of course, in general, it is impossible to achieve envy-freeness with indivisibilities (consider two agents and a single desirable item), but we can look for approximations. A relaxation of envy-freeness that has been very influential recently is *envy-freeness up to one good* (EF1), introduced by Budish (2011). It requires that an agent's envy towards another bundle vanishes if we remove some item from the envied bundle. In the setting without connectivity constraints and with additive valuations, the maximum Nash welfare solution satisfies EF1, as does a simple round-robin procedure (Caragiannis et al., 2019). The well-known envy-graph algorithm (Lipton et al., 2004) also guarantees EF1. However, none of these procedures respects connectivity constraints.

When items are arranged on a path, we prove that connected EF1 allocations exist when there are two, three, or four agents. As was necessary in cake-cutting, we use successively more complicated tools to establish these existence results. For two agents, there is a discrete analogue of cut-and-choose that satisfies EF1. In that procedure, a knife moves across the path, and an agent shouts when the knife reaches what we call a *lumpy tie*, that is when the bundles to either side of the knife have equal value *up to one item*. For three agents, we design an algorithm mirroring Stromquist's moving-knife procedure which guarantees EF1. For four agents, we show that Sperner's lemma can be used to prove that an EF1 allocation exists, via a technique inspired by the Simmons–Su approach, and an appropriately triangulated simplex of connected partitions of the path. For five or more agents, we were not able to establish the existence of EF1 allocations on a path, but we can show (again via Sperner's lemma) that EF2 allocations exist, strengthening a prior result of Suksompong (2019). We also show that if all agents have the same valuation function over bundles, then an egalitarian-welfare-optimal allocation, after suitably reallocating some items, is EF1.

These existence results require only that agents' valuations are monotonic (they need not be additive), and in addition, ensure that the constructed allocation satisfies the maximin share guarantee (see Appendix 9.1). Moreover, the fairness guarantee of our algorithms is slightly stronger than the standard notion of EF1: in the returned allocations, envy can be avoided by removing just an *outer* item – one whose removal leaves the envied bundle connected. Computationally speaking,

<sup>1</sup> Another paper by Suksompong (2019) works in the same model, and also found that procedures for proportionality and other concepts can be applied to the indivisible setting.

**Table 1**

Overview of our results for paths. Here,  $n$  denotes the number of agents and  $m$  the number of items. Agents' valuations are assumed to be monotone. The mark  $\checkmark$  represents that a connected allocation satisfying the corresponding fairness notion exists. When no reference is given, the result follows from other results in the table.

	# of agents	EF1	EF2
Existence	$n = 2$	$\checkmark$ (Theorem 3.2)	$\checkmark$
	$n = 3$	$\checkmark$ (Theorem 4.3)	$\checkmark$
	$n = 4$	$\checkmark$ (Theorem 6.2)	$\checkmark$
	$n \geq 5$	open problem	$\checkmark$ (Theorem 5.4)
	$n$ agents, identical valuations	$\checkmark$ (Theorem 7.1)	$\checkmark$
Complexity	$n = 2$	$O(\log m)$ (Oh et al., 2019)	$O(\log m)$
	$n = 3$	$O(m)$ (Theorem 4.3)	$O(m)$
	$n \geq 4$	$O(m^n)$	$O(m^n)$
	$n$ agents, identical valuations	$O(mn)$ (Theorem 7.1)	$O(mn)$

all our existence results are constructive in the weak sense that an EF1 allocation can be found by iterating through all  $O(m^n)$  connected allocation (this stands in contrast to cake-cutting where we cannot iterate through all possibilities). While we know of no faster algorithms to obtain an EF1 or EF2 allocation in the cases where we appeal to Sperner's lemma, our other procedures (for two or three agents, or for identical valuations) can all be implemented efficiently to produce a fair allocation in polynomial time. We summarize our results concerning paths in Table 1.

In simultaneous and independent work, Oh et al. (2019) designed protocols to find EF1 allocations in the setting without connectivity constraints, aiming for low *query complexity*. They found that adapting cake-cutting protocols to the setting of indivisible items arranged on a path are an especially potent way to achieve low query complexity. This led them to also study a discrete version of the cut-and-choose protocol which achieves connected EF1 allocations for two agents, and they found an alternative proof that an EF1 allocation on a path always exists with identical valuations. They also present a discrete analogue of the Selfridge–Conway procedure which, for three agents with additive valuations, produces an allocation of a path into bundles that have a constant number of connected components. However, they do not study connected allocations on graphs that are not paths, and they do not consider the case of (non-identical) general valuations with more than two agents.

A recurring theme in our algorithms is the specific way that the moving knives from cake-cutting are rendered in the discrete setting. While one might expect knives to be placed over the edges of the path, and 'move' from edge to edge, we find that this movement is too 'fast' to ensure EF1 (see also footnote 5 regarding EF2). Instead, our knives alternate between hovering over edges and items. When a knife hovers over an item, we imagine the knife's blade to be 'thick': the knife covers the item, and agents then pretend that the covered item does not exist. These intermediate steps are useful, since they can tell us that envy will vanish if we hide an item from a bundle.

What about graphs  $G$  other than paths? Our existential results for paths immediately generalize to traceable graphs (those that contain a Hamiltonian path), since we can run the algorithms pretending that the graph only consists of the Hamiltonian path. For the two-agent case, we completely characterize the class of graphs that guarantee the existence of EF1 allocations: Our discrete cut-and-choose protocol can be shown to work on all graphs  $G$  that admit a *bipolar numbering*, which exists if and only if the biconnected components (blocks) of  $G$  can be arranged in a path. By constructing counterexamples, we prove that no graph failing this condition (for example, a star) guarantees EF1, even for identical, additive, binary valuations. For the case of three or more agents, it is a challenging open problem to characterize the class of graphs guaranteeing EF1 (or even to find an infinite class of non-traceable graphs that guarantees EF1).

## 2. Preliminaries

For each natural number  $s \in \mathbb{N}$ , write  $[s] = \{1, 2, \dots, s\}$ . Let  $N = [n]$  be a finite set of *agents* and  $G = (V, E)$  be an undirected finite graph, where  $V = \{v_1, v_2, \dots, v_m\}$ . The vertices in  $V$  correspond to *items*. A subset  $I$  of  $V$  is *connected* if it induces a connected subgraph of  $G$ . We write  $\mathcal{C}(V)$  for the set of connected subsets of  $V$ . We call a set  $I \in \mathcal{C}(V)$  a (connected) *bundle*. Each agent  $i \in N$  has a *valuation function*  $u_i : \mathcal{C}(V) \rightarrow \mathbb{R}$  over connected bundles, which we will always assume to be *monotonic*, that is,  $X \subseteq Y$  implies  $u_i(X) \leq u_i(Y)$ . We also assume that  $u_i(\emptyset) = 0$  for each  $i \in N$ . Monotonicity implies that items are *goods*; we do not consider bads (or chores) in this paper. We say that an agent  $i \in N$  *weakly prefers* bundle  $X$  to bundle  $Y$  if  $u_i(X) \geq u_i(Y)$ .<sup>2</sup> A (connected) *allocation*  $A : N \rightarrow \mathcal{C}(V)$  assigns each agent  $i \in N$  a connected bundle  $A(i) \in \mathcal{C}(V)$  such that each item occurs in exactly one bundle, i.e.,  $\bigcup_{i \in N} A(i) = V$  and  $A(i) \cap A(j) = \emptyset$  when  $i \neq j$ . We often write  $I^i$  for the bundle  $A(i)$  assigned to agent  $i$ .

We say that the agents have *identical valuations* if, for all  $i, j \in N$  and every bundle  $I \in \mathcal{C}(V)$ , we have  $u_i(I) = u_j(I)$ . A valuation function  $u_i$  is *additive* if  $u_i(I) = \sum_{v \in I} u_i(\{v\})$  for each bundle  $I \in \mathcal{C}(V)$ . Many examples in this paper will use

<sup>2</sup> Our arguments only operate based on agents' ordinal preferences over bundles, and the (cardinal) valuation functions are only used for notational convenience. One exception, perhaps, is in Algorithm 1 where we calculate a leximin allocation, but the algorithm can be applied after choosing an arbitrary utility function consistent with the ordinal preferences.

identical additive valuations, and will take  $G$  to be a path. In this case, we use a shorthand to specify these examples; the meaning of this notation should be clear. For example, we write “2–1–3–1” to denote an instance with four items  $v_1, v_2, v_3, v_4$  arranged on a path, and where  $u_i(\{v_1\}) = 2, \dots, u_i(\{v_4\}) = 1$  for each  $i$ . For such an instance, an allocation will be written as a tuple, e.g., (2, 1–3–1) denoting an allocation allocating bundles  $\{v_1\}$  and  $\{v_2, v_3, v_4\}$ , noting that with identical valuations it does not usually matter which agent receives which bundle.

An allocation  $A$  is *envy-free* if  $u_i(A(i)) \geq u_i(A(j))$  for every pair  $i, j \in N$  of agents, that is, if every agent thinks that their bundle is at least as good as any other bundle in the allocation. It is well-known that an envy-free allocation may not exist (consider two agents and one good). The main fairness notion that we study is a version of *envy-freeness up to one good* (EF1), a relaxation of envy-freeness introduced by Budish (2011), adapted to the model with connectivity constraints. This property states that an agent  $i$  will not envy another agent  $j$  after we remove some item from  $j$ ’s bundle. Since we only allow connected bundles in our set-up, we may only remove an item from  $A(j)$  if removal of this item leaves the bundle connected.

**Definition 2.1** (EF1: *envy-freeness up to one outer good*). An allocation  $A$  satisfies EF1 if, for any pair  $i, j \in N$  of agents, either  $A(j) = \emptyset$  or there is a good  $v \in A(j)$  such that  $A(j) \setminus \{v\}$  is connected and  $u_i(A(i)) \geq u_i(A(j) \setminus \{v\})$ .

In the instance 2–1–3–1 for two agents, the allocation (2–1, 3–1) is EF1, since the left agent’s envy can be eliminated by removing the item of value 3 from the right-hand bundle. However, the allocation (2, 1–3–1) fails to be EF1 according to our definition, since eliminating either outer good of the right bundle does not prevent envy.<sup>3</sup>

**Definition 2.2.** A graph  $G$  *guarantees EF1* for  $n$  agents if, for all possible monotonic valuations for  $n$  agents, there exists some connected allocation that is EF1. A graph  $G$  *guarantees EF1* for  $n$  agents and a restricted class of valuations if, for all allowed valuations, a connected EF1 allocation exists.

For reasoning about EF1 allocations, let us introduce a few shorthands. Given an allocation  $A$  we will say that  $i \in N$  *does not envy*  $j \in N$  *up to*  $v$  if  $u_i(A(i)) \geq u_i(A(j) \setminus \{v\})$ . The *up-to-one valuation*  $u_i^- : \mathcal{C}(V) \rightarrow \mathbb{R}_{\geq 0}$  of agent  $i \in N$  is defined, for every  $I \in \mathcal{C}(V)$ , as

$$u_i^-(I) := \begin{cases} 0 & \text{if } I = \emptyset, \\ \min \{u_i(I \setminus \{v\}) : v \in I \text{ such that } I \setminus \{v\} \text{ is connected}\} & \text{if } I \neq \emptyset. \end{cases} \tag{2.1}$$

Thus, an allocation  $A$  satisfies EF1 if and only if  $u_i(A(i)) \geq u_i^-(A(j))$  for any pair  $i, j \in N$  of agents.

As we show in the appendix in Example 9.5, allocations satisfying a strengthened version of EF1 called *envy-freeness up to the least good* (EFX) (Caragiannis et al., 2019) may not exist on a path.

Given an ordered sequence of the vertices  $P = (v_1, v_2, \dots, v_m)$ , and  $j, k \in [m]$  with  $j \leq k$ , we write  $P(v_j, v_k)$  for the subsequence from  $v_j$  to  $v_k$ , so  $P(v_j, v_k) = (v_j, v_{j+1}, \dots, v_{k-1}, v_k)$ . With a little abuse of notation, we often identify a subsequence  $P(v_j, v_k)$  with the bundle of the corresponding vertices. Let  $L(v_j) = P(v_1, v_{j-1})$  be the subsequence of vertices strictly left of  $v_j$  and  $R(v_j) = P(v_{j+1}, v_m)$  be the subsequence of vertices strictly right of  $v_j$ . When graph  $G$  is a path, we always implicitly assume that its vertices  $v_1, v_2, \dots, v_m$  are numbered from left to right according to the order they appear along the path, so that the set of the edges of  $G$  is  $\{\{v_j, v_{j+1}\} : 1 \leq j < m\}$ . Each connected bundle in the path clearly corresponds to a subpath or subsequence of the vertices. A *Hamiltonian path* of a graph  $G$  is a path that visits all the vertices of the graph exactly once. A graph is *traceable* if it contains a Hamiltonian path.

### 3. EF1 existence for two agents

In cake-cutting for two agents, the standard way of obtaining an envy-free allocation is the cut-and-choose protocol: Alice divides the cake into two equally-valued pieces, and Bob selects the piece he prefers; the other piece goes to Alice. The same strategy almost works in the indivisible case when items form a path; the problem is that Alice might not be able to divide the items into two exactly equal pieces. Instead, we ask Alice to divide the items into pieces that are equally valued “up to one good”. The formal version is as follows. For a sequence of vertices  $P = (v_1, v_2, \dots, v_m)$  and an agent  $i$ , we say that  $v_j$  is the *lumpy tie* over  $P$  for agent  $i$  if  $j$  is the smallest index such that

$$u_i(L(v_j) \cup \{v_j\}) \geq u_i(R(v_j)) \quad \text{and} \quad u_i(R(v_j) \cup \{v_j\}) \geq u_i(L(v_j)). \tag{3.1}$$

For example, when  $i$  has additive valuations 1–3–2–1–3–1, then the third item (of value 2) is the lumpy tie for  $i$ , since  $1 + 3 + 2 \geq 1 + 3 + 1$  and  $2 + 1 + 3 + 1 \geq 1 + 3$ . The lumpy tie always exists: taking  $j$  to be the smallest index such that

<sup>3</sup> This example shows that our definition is strictly stronger than the standard definition of EF1 without connectivity constraints. In the instance 2–1–3–1, considered without connectivity constraints, the allocation (2, 1–3–1) does satisfy EF1 since in the standard setting we are allowed to remove the middle item (with value 3) of the right bundle.




Fig. 1. Non-traceable graphs with bipolar numberings.

$u_i(L(v_j) \cup \{v_j\}) \geq u_i(R(v_j))$  (which exists as the inequality holds for  $j = m$  by monotonicity), the first part of (3.1) holds. If  $j = 1$ , the second part of (3.1) is immediate by monotonicity. If  $j > 1$ , then since  $j$  is minimal, we have  $u_i(L(v_j)) = u_i(L(v_{j-1}) \cup \{v_{j-1}\}) < u_i(R(v_{j-1})) = u_i(R(v_j) \cup \{v_j\})$  as required.

Using lumpy ties, our discrete version of the cut-and-choose protocol is specified as follows.

**Definition 3.1.** The **discrete cut-and-choose protocol for  $n = 2$  agents** on a sequence  $P = (v_1, v_2, \dots, v_m)$  proceeds as follows:

- Step 1. Alice selects her lumpy tie  $v_j$  over  $(v_1, v_2, \dots, v_m)$ .
- Step 2. Bob chooses a weakly preferred bundle among  $L(v_j)$  and  $R(v_j)$ .
- Step 3. Alice receives the bundle of all the remaining vertices, including  $v_j$ .

Intuitively, the protocol allows Alice to select an item  $v_j$  that she will receive for sure, with the advice that the two pieces to either side of  $v_j$  should have almost equal value to her. Then, Bob is allowed to choose which side of  $v_j$  he wishes to receive. In our example with valuations 1–3–2–1–3–1, Alice selects the lumpy tie of value 2, then Bob chooses the bundle 1–3–1 to the right and receives it, and Alice receives the bundle 1–3–2. The result is EF1. This is true in general, and also if valuations are not identical.

**Proposition 3.2.** When  $G$  is a path and there are  $n = 2$  agents, the discrete cut-and-choose protocol yields an EF1 allocation.

**Proof.** Clearly, the protocol returns a connected allocation. The returned allocation satisfies EF1: Bob does not envy Alice up to item  $v_j$ , since Bob receives his preferred bundle among  $L(v_j)$  and  $R(v_j)$ . Also, by (3.1), Alice does not envy Bob, since Alice either receives the bundle  $L(v_j) \cup \{v_j\}$  which she weakly prefers to Bob’s bundle  $R(v_j)$ , or she receives the bundle  $R(v_j) \cup \{v_j\}$ , which she weakly prefers to Bob’s bundle  $L(v_j)$ . □

Proposition 3.2 implies that an EF1 allocation always exists on a path. Hence, an EF1 allocation exists for every traceable graph  $G$ : simply use the discrete cut-and-choose protocol on a Hamiltonian path of  $G$ . In fact, the discrete cut-and-choose protocol works on a broader class of graphs: We only need to require that the vertices of the graph can be numbered in a way that the allocation resulting from the discrete cut-and-choose protocol is guaranteed to be connected. Since the protocol always partitions the items into an initial and a terminal segment of the sequence, such a numbering needs to satisfy the following property.

**Definition 3.3.** A *bipolar numbering* of a graph  $G$  is an ordering  $(v_1, v_2, \dots, v_m)$  of its vertices such that for all  $j \in [n]$ , the sets  $L(v_j) \cup \{v_j\}$  and  $R(v_j) \cup \{v_j\}$  are connected in  $G$ .

In a slightly different context, bipolar numberings are known as *st*-numberings and turn out to be useful in algorithms for testing planarity and for graph drawing (Lempel et al., 1967; Even and Tarjan, 1976; Tarjan, 1986). The more common (equivalent) definition is phrased to say that a numbering is bipolar if, for every  $j \in [n]$ , the vertex  $v_j$  has a neighbor that appears earlier in the sequence, and a neighbor that appears later in the sequence.

Clearly, every traceable graph has a bipolar numbering, since we can just use a Hamiltonian path. However, there are also non-traceable graphs that admit a bipolar numbering. Fig. 1 shows some examples.

**Proposition 3.4.** When there are  $n = 2$  agents, then the discrete cut-and-choose protocol run on a bipolar numbering of  $G$  yields an EF1 allocation.

**Proof.** The discrete cut-and-choose protocol returns an allocation whose bundles are either initial or terminal segments of the ordered sequence  $(v_1, v_2, \dots, v_m)$ . By definition of a bipolar numbering, such an allocation is connected, and it is EF1 by the same argument as in Proposition 3.2. □

It is clear that the discrete cut-and-choose protocol cannot be extended to graphs other than those admitting a bipolar numbering. However, it could be that a different protocol is able to produce EF1 allocations on other graphs. In the remainder of this section, we prove that this is not the case: for  $n = 2$  agents, a connected graph  $G$  guarantees the existence of an



Fig. 2. A type I trident (left) and a type II trident (right).


EF1 allocation if and only if it admits a bipolar numbering. This completely characterizes the class of graphs that guarantee EF1 existence in the two-agent case.<sup>4</sup>

For a different number of agents, the class of graphs guaranteeing an EF1 allocation will be different. In particular, the star with three leaves does not guarantee an EF1 allocation for two agents (as it does not have a bipolar numbering, see below), but one can check that this star does guarantee an EF1 allocation for three or more agents (see Example 9.6 in the appendix).

### 3.1. Characterization of graphs guaranteeing EF1 for two agents

Based on a known characterization of graphs admitting a bipolar numbering, we characterize this class in terms of forbidden substructures. We then show that these forbidden structures are also forbidden for EF1: if a graph contains such a structure, we can exhibit an additive valuation profile for which no EF1 allocation exists.

As a simple example, consider the star with three leaves, which is the smallest connected graph that does not have a bipolar numbering.



Take two agents with identical additive valuations that value each item at 1. Any connected allocation must allocate three items to one agent, and a single item to the other agent. Then the latter agent envies the former agent, even up to one good. This star is an example of a forbidden substructure called a trident, which takes one of two forms, illustrated in Fig. 2.

**Definition 3.5.** A graph  $G$  contains a trident if either

- (a) there is a vertex  $s$  whose removal from  $G$  leaves three or more connected components (a type I trident), or
- (b) there are subgraphs  $C, P_1, P_2, P_3$  of  $G$  such that (i)  $P_1, P_2, P_3$  are vertex-disjoint, (ii) each  $P_i$  contains at least two vertices, (iii)  $C$  has exactly one contact vertex  $s_i$  in common with  $P_i$ ,  $i = 1, 2, 3$ , and (iv) for  $i = 1, 2, 3$ , removal of vertex  $s_i$  from  $G$  disconnects  $P_i \setminus \{s_i\}$  from  $C \setminus \{s_i\}$  (hence from the other two  $P_j$ ) in  $G$  (a type II trident).

We will prove that a graph  $G$  fails to admit a bipolar numbering, and fails to guarantee EF1 for two agents, if and only if  $G$  contains a trident. To reason about these structures, it is useful to consider the standard concept of the *block decomposition* of a graph (see, e.g., the textbook Bondy and Murty, 2008, Sec. 5.2).

**Definition 3.6.** A decomposition of a graph  $G = (V, E)$  is a family  $\{F_1, F_2, \dots, F_t\}$  of edge-disjoint subgraphs of  $G$  such that  $\bigcup_{i=1}^t E(F_i) = E$  where  $E(F_i)$  is the set of edges of  $F_i$ . A vertex is called a *cut vertex* of a graph  $G$  if removing it increases the number of connected components of  $G$ . A graph  $G$  is *biconnected* if  $G$  is connected and does not have a cut vertex. A *block* of  $G$  is a maximal biconnected subgraph of  $G$ .

Equivalently, a block of a graph  $G$  can be defined as a maximal subgraph of  $G$  where each pair of vertices lie on a common cycle (Bondy and Murty, 2008). Given a connected graph  $G$ , we define a bipartite graph  $B(G)$  with bipartition  $(\mathcal{B}, S)$ , where  $\mathcal{B}$  is the set of blocks of  $G$  and  $S$  is the set of cut vertices of  $G$ ; a block  $B$  and a cut vertex  $v$  are adjacent in  $B(G)$  if and only if  $B$  includes  $v$ . Since every cycle of a graph is included in some block, the graph  $B(G)$  is a tree:

**Lemma 3.7** (e.g., Bondy and Murty, 2008, Prop. 5.3). *Let  $G$  be a connected graph. Then*

<sup>4</sup> Note that no non-trivial disconnected graph guarantees EF1 for two agents: If  $G$  is disconnected, take a connected component  $C$  with at least two vertices. Let both agents have additive valuations that value each item in  $C$  at 1, and value items outside of  $C$  at 0. Then, in a connected allocation, all items in  $C$  must go to a single agent, since the other agent needs to receive items from another connected component. This induces envy in the other agent that is not bounded by one good.




Fig. 3. Block decompositions of the graphs in Fig. 1.

- any two blocks of  $G$  have at most one cut vertex in common;
- the set of blocks forms a decomposition of  $G$ ; and
- the graph  $B(G)$  is a tree.

Thus, for a connected graph  $G$ , we call  $B(G)$  the *block tree* of  $G$ . It turns out that  $G$  admits a bipolar numbering if and only if  $B(G)$  is a path. For example, the graphs shown in Fig. 1 all have their blocks arranged in a path (so that  $B(G)$  is a path), as shown in Fig. 3.

**Lemma 3.8.** *A graph  $G$  admits a bipolar numbering if its block tree  $B(G)$  is a path.*

**Proof.** Lempel et al. (1967) show that  $G$  admits a bipolar numbering if there are  $s, t \in V$  such that adding an edge  $\{s, t\}$  to  $G$  makes it biconnected. If  $B(G)$  is a path, let  $B_1$  and  $B_2$  be the leaf blocks at the ends of the path  $B(G)$ . Take any  $s \in B_1$  and  $t \in B_2$ . If we add the edge  $\{s, t\}$  to  $G$ , the graph becomes biconnected. Hence,  $G$  admits a bipolar numbering.  $\square$

There is a linear-time algorithm based on depth-first search to construct a bipolar numbering for any biconnected graph (Even and Tarjan, 1976; Tarjan, 1986), and one can also calculate the block tree  $B(G)$  of a given graph in linear time (Hopcroft and Tarjan, 1973). Thus, in linear time, we can compute a bipolar numbering of a graph or report that none exists. Clearly, given a bipolar numbering, the discrete cut-and-choose protocol can also be run in linear time.

Next, we show that if  $B(G)$  is not a path, then  $G$  cannot guarantee EF1. The proof constructs explicit counter-examples, which have a very simple structure. We say that additive valuations  $u_i$  are *binary* if  $u_i(\{v\}) \in \{0, 1\}$  for every  $v \in V$ .

**Lemma 3.9.** *Let  $G$  be a connected graph.*

- If the block tree  $B(G)$  of  $G$  is not a path, then  $G$  contains a trident.
- If  $G$  contains a trident, then there exist identical, additive, binary valuations over  $G$  for two agents such that no connected allocation is EF1.

**Proof.** If  $B(G)$  is not a path, then it contains a vertex with at least three neighbors, and thus either

- there is a cut vertex  $s$  adjacent to three blocks  $B_1, B_2,$  and  $B_3$ ; or
- there is a block  $B$  adjacent to three different cut vertices  $s_1, s_2,$  and  $s_3$ .

Note that in both cases, all blocks contain at least two vertices each, as maximality guarantees that a block in a connected graph  $G$  never consists of a single vertex, unless  $G$  itself has only one vertex. Thus, in case (a),  $G$  contains a type I trident. In case (b), the cut vertices  $s_1, s_2,$  and  $s_3$  serve as the contact vertices in the earlier definition of type II tridents and are adjacent to blocks that serve as the subgraphs  $P_1, P_2,$  and  $P_3$ . This proves the first part.

To prove the second part, we construct identical additive valuations that do not admit an EF1 allocation. If  $G$  contains a type I trident, let  $s$  be the corresponding cut vertex, and choose vertices  $v_1, v_2, v_3$  from each of three different connected components that remain after  $s$  is deleted from  $G$ . The two agents have utility 1 for each of  $s, v_1, v_2,$  and  $v_3$ , and 0 for the remaining vertices. Now take any connected allocation  $(I_1, I_2)$ . One of the bundles, say  $I_1$ , includes the cut vertex  $s$ . Then  $I_2$  can contain at most one of the vertices  $v_1, v_2, v_3$ , since  $I_2$  is connected and does not contain  $s$  yet any path between distinct  $v_i$  and  $v_j$  goes through  $s$ . Hence  $u_i(I_2) \leq 1$ . Now, the bundle  $I_1$  contains  $s$  and at least two of  $v_1, v_2, v_3$ , so  $u_i(I_1) \geq 3$ . Thus, the allocation is not EF1.

Suppose  $G$  contains a type II trident consisting of subgraphs  $C, P_1, P_2, P_3$  with contact vertices  $s_1, s_2, s_3$ . Then for  $i = 1, 2, 3$  choose a vertex  $v_i \neq s_i$  from  $P_i$ . The two agents have utility 1 for each of  $s_1, s_2, s_3, v_1, v_2,$  and  $v_3$ , and 0 for the remaining vertices. Now take any connected allocation  $(I_1, I_2)$ . One of the bundles, say  $I_1$ , contains at least two contact vertices  $s_i$  and the other contains at most one contact vertex  $s_i$ . Say that  $s_1, s_2 \in I_1$ . Now,  $G \setminus \{s_1, s_2\}$  has at least three connected components, and since  $I_2$  is connected, it must be contained in one of these components. But each component contains at most two vertices with utility 1, so  $u_i(I_2) \leq 2$ . Since there are six vertices with utility 1 in total,  $u_i(I_1) \geq 4$ . Thus, the allocation is not EF1.  $\square$




Fig. 4. Stromquist's moving knife protocol.

Combining these results, we obtain the promised characterization.

**Theorem 3.10.** *The following conditions are equivalent for every connected graph  $G$ :*

1.  $G$  admits a bipolar numbering.
2.  $G$  guarantees EF1 for two agents.
3.  $G$  guarantees EF1 for two agents with identical, additive, binary valuations.
4.  $G$  does not contain a trident.
5. The block tree  $B(G)$  is a path.

**Proof.** The implication (1)  $\Rightarrow$  (2) follows from Proposition 3.4 which shows that the discrete cut-and-choose protocol yields a connected EF1 allocation when run on a bipolar numbering. The implication (2)  $\Rightarrow$  (3) is immediate. The implications (3)  $\Rightarrow$  (4) and (4)  $\Rightarrow$  (5) follow from Lemma 3.9 which proves the contrapositives. Finally, (5)  $\Rightarrow$  (1) follows from Lemma 3.8.  $\square$

The equivalence (2)  $\Leftrightarrow$  (3) is noteworthy and perhaps surprising: It is often easier to guarantee fairness when agents' valuations are identical, yet in terms of the graphs that guarantee EF1 for two agents, there is no difference between identical and non-identical valuations. Intriguingly, even for more than two agents, we do not know of a graph which guarantees EF1 for identical valuations, but fails it for non-identical valuations.

#### 4. EF1 existence for three agents: a moving-knife protocol

We will now consider the case of three agents. Stromquist (1980) designed a protocol that results in an envy-free contiguous allocation of a divisible cake. We now give a brief outline of the protocol, illustrated by Fig. 4.

A referee holds a sword over the cake. Each of the three agents holds their own knife over the portion of the cake to the right of the sword, positioning it so that this portion is divided into two pieces they judge to have the same value. Now, initially, the sword is at the left end of the cake. It starts moving at a constant speed from left to right, while the agents continuously move their knives to keep dividing the right-hand portion into equally-valued pieces. At some point (when the leftmost piece becomes valuable enough), one of the agents shouts "cut", and the cake will be cut twice: once by the sword, and once by the middle one of the three knives. Agents shout "cut" as soon as the left piece is a highest-valued piece among the three. The agent who shouts receives the left piece. The remaining agents each receive a piece containing their knife. The resulting allocation is envy-free, since the agent receiving the left piece prefers it to the other pieces, and the other agents who are not shouting receive at least half the value of the part of the cake to the right of the sword.

Let  $G$  be a path,  $P = (v_1, v_2, \dots, v_m)$ . There are several difficulties in translating Stromquist's continuous procedure to the discrete setting for  $G$ . First, agents need to divide the piece to the right of the sword in half, and this might not be possible exactly given indivisibilities; but this can be handled using our concept of lumpy ties from Section 3. Next, when the sword moves one item to the right, the lumpy ties of the agents may need to jump several items to the right, for example, because the new member of the leftmost bundle is very valuable. To ensure EF1, we will need to smoothen these jumps, so that the middle piece grows one item at a time. Also, it will be helpful to have the sword move in half-steps: it alternates between being placed between items (so it cuts the edge between the items), and being placed over an item, in which case the sword covers the item and agents ignore that item. Finally, while the sword covers an item, we will only terminate if at least two agents shout to indicate that they prefer the leftmost piece; this will ensure that there is an agent who is flexible about which of the bundles they are assigned. The algorithm moves in steps, and alternates between moving the sword, and updating the lumpy ties.

In our formal description of the algorithm, we do not use swords and knives. Instead, we maintain three bundles  $L$ ,  $M$ , and  $R$  that can be seen as resulting from a certain configuration this cutting implements. We also need a few definitions. For a subsequence of vertices  $P(v_s, v_r) = (v_s, v_{s+1}, \dots, v_r)$  and an agent  $i$ , recall that  $v_j$  ( $s \leq j \leq r$ ) is the lumpy tie over  $P(v_s, v_r)$  for  $i$  if  $j$  is the smallest index such that

$$u_i(L(v_j) \cup \{v_j\}) \geq u_i(R(v_j)) \quad \text{and} \quad u_i(R(v_j) \cup \{v_j\}) \geq u_i(L(v_j)). \tag{4.1}$$



Here, the definitions of  $L(v_j)$  and  $R(v_j)$  apply to the subsequence  $P(v_s, v_r)$ . The lumpy tie always exists by the discussion after equation (3.1). Each of the three agents has a lumpy tie over  $P(v_s, v_r)$ ; a key concept for us is the *median lumpy tie* which is the median of the lumpy ties of the three agents, where the median is taken with respect to the ordering of  $P(v_s, v_r)$ . We say that  $i \in N$  is a *left agent* (respectively, a *middle agent* or a *right agent*) over  $P(v_s, v_r)$  if the lumpy tie for  $i$  appears strictly before (respectively, is equal to, or appears strictly after) the median lumpy tie. Note that by definition of the median, there is at most one left agent, at most one right agent, and at least one middle agent. Suppose that the median lumpy tie over the subsequence  $P(v_s, v_r)$  is  $v_j$ , and let  $i$  be an agent. Then using the definitions of lumpy tie and left/right agents, we find that

$$\begin{aligned} u_i(L(v_j)) &\geq u_i(R(v_j) \cup \{v_j\}) && \text{if } i \text{ is a left agent, and} \\ u_i(R(v_j)) &\geq u_i(L(v_j) \cup \{v_j\}) && \text{if } i \text{ is a right agent.} \end{aligned} \tag{4.2}$$

Given the median lumpy tie  $v_j$  over  $P(v_s, v_r)$ , and a two-agent set  $S = \{i, k\} \subseteq N$ , we define  $\text{Lumpy}(S, v_j, P(v_s, v_r))$  to be the allocation of the items in  $P(v_s, v_r)$  to  $S$  such that

- if  $i$  is a left agent and  $k$  is a right agent, then  $i$  receives  $L(v_j)$  and  $k$  receives  $R(v_j) \cup \{v_j\}$ ;
- if  $i$  is a middle agent, then agent  $k$  receives  $k$ 's preferred bundle among  $L(v_j)$  and  $R(v_j)$ , and agent  $i$  receives the other bundle along with  $v_j$ .

Using (4.1) and (4.2), we see that  $\text{Lumpy}(S, v_j, P(v_s, v_r))$  is an EF1 allocation:

**Lemma 4.1** (*Median Lumpy Ties Lemma*). *Let  $S = \{i, k\} \subseteq N$  and let  $v_j$  be the median lumpy tie over  $P(v_s, v_r)$ . Then  $\text{Lumpy}(S, v_j, P(v_s, v_r))$  is an EF1 allocation of the items in  $P(v_s, v_r)$  to  $S$ . Further, each agent in  $S$  weakly prefers their bundle to  $L(v_j)$  and  $R(v_j)$ .*

The algorithm is specified in Definition 4.2. It alternately moves a left pointer  $\ell$  (in Steps 2 and 3) and a right pointer  $r$  (in Step 4). It also maintains bundles  $L$ ,  $M$ , and  $R$  during the execution of the algorithm.

**Definition 4.2.** The **discrete moving-knife protocol for  $n = 3$  agents** on a sequence  $P = (v_1, v_2, \dots, v_m)$  proceeds as follows. We say that an agent  $i \in N$  is a *shouter* if  $u_i(L) \geq u_i(M)$  and  $u_i(L) \geq u_i(R)$ .

- *Step 1.* Initialize  $\ell = 0$  and set  $r$  so that  $v_r$  is the median lumpy tie over the subsequence  $P(v_2, v_m)$ . Initialize  $L = \emptyset$ ,  $M = \{v_2, v_3, \dots, v_{r-1}\}$ , and  $R = \{v_{r+1}, v_{r+2}, \dots, v_m\}$ .
- *Step 2.* Add an additional item to  $L$ , i.e., set  $\ell = \ell + 1$  and  $L = \{v_1, v_2, \dots, v_\ell\}$ . If no agent shouts, go to Step 3. If some agent  $s_{\text{left}}$  shouts,  $s_{\text{left}}$  receives the left bundle  $L$ . Allocate the remaining items according to  $\text{Lumpy}(N \setminus \{s_{\text{left}}\}, v_r, P(v_{\ell+1}, v_m))$ .
- *Step 3.* Delete the leftmost point of the middle bundle, i.e., set  $M = \{v_{\ell+2}, v_{\ell+3}, \dots, v_{r-1}\}$ . If the number of shouters is smaller than two, go to Step 4. If at least two agents shout, we show (next page) that there is a shouter  $s$  who is a middle agent over  $P(v_{\ell+1}, v_m)$ . Then, allocate  $L$  to a shouter  $s_{\text{left}}$  distinct from  $s$ . Let the agent  $c$  distinct from  $s$  and  $s_{\text{left}}$  choose his preferred bundle among  $\{v_{\ell+1}\} \cup M$  and  $\{v_r\} \cup R$ . Agent  $s$  receives the other bundle.
- *Step 4.* If  $v_r$  is the median lumpy tie over  $P(v_{\ell+2}, v_m)$ , directly move to the following cases (a)–(d). If  $v_r$  is not the median lumpy tie over  $P(v_{\ell+2}, v_m)$ , set  $r = r + 1$ ,  $M = \{v_{\ell+2}, v_{\ell+3}, \dots, v_{r-1}\}$ , and  $R = \{v_{r+1}, v_{r+2}, \dots, v_m\}$ ; then, go to cases (a)–(d).
  - If at least two agents shout, find a shouter  $s$  who did not shout at the previous step. If there is a shouter  $s_{\text{left}}$  who shouted at the previous step,  $s_{\text{left}}$  receives  $L$ ; else, give  $L$  to an arbitrary shouter  $s_{\text{left}}$  distinct from  $s$ . The agent  $c$  distinct from  $s$  and  $s_{\text{left}}$  choose his preferred bundle among  $\{v_{\ell+1}\} \cup M$  and  $\{v_r\} \cup R$ , breaking ties in favor of the former option. Agent  $s$  receives the other bundle.
  - If  $v_r$  is the median lumpy tie over  $P(v_{\ell+2}, v_m)$  and only one agent  $s_{\text{left}}$  shouts, give  $L \cup \{v_{\ell+1}\}$  to  $s_{\text{left}}$  and allocate the rest according to  $\text{Lumpy}(N \setminus \{s_{\text{left}}\}, v_r, P(v_{\ell+2}, v_m))$ .
  - If  $v_r$  is the median lumpy tie over  $P(v_{\ell+2}, v_m)$  but no agent shouts, go to Step 2.
  - Otherwise  $v_r$  is not the median lumpy tie over  $P(v_{\ell+2}, v_m)$ : Repeat Step 4.

**Theorem 4.3.** *The moving-knife protocol finds an EF1 allocation for three agents and runs in  $O(m)$  time, when  $G$  is a path.*

**Proof.** The algorithm is well-defined – there is one place where this is not immediate: If two agents shout in Step 3, the algorithm description claims that there is a shouter who is a middle agent over the subsequence  $P(v_{\ell+1}, v_m)$ . Suppose for the moment that there is a shouter  $i$  who is a *right agent*. Due to (4.2), we have  $u_i(R) \geq u_i(\{v_{\ell+1}\} \cup M \cup \{v_r\})$ . Since  $i$  is a shouter, we have  $u_i(L) \geq u_i(R)$ , so  $u_i(L) \geq u_i(\{v_{\ell+1}\} \cup M \cup \{v_r\})$ . But  $i$  did not shout in the previous Step 2 (when no-one shouted), so either  $u_i(R) > u_i(L)$  or  $u_i(\{v_{\ell+1}\} \cup M) > u_i(L)$ , and either case is a contradiction. Hence neither of the at least two shouters of Step 3 is a right agent, so at least one shouter is a middle agent, since there is at most one left agent.

The algorithm terminates and returns an allocation, since the bundle  $L$  grows throughout the algorithm until eventually, at least two agents will think that  $L$  is a best bundle and thus will shout and thereby terminate the algorithm. We will now consider every possible way that the algorithm could have terminated, and show that the resulting allocation is EF1.

**Step 2.**

- Agent  $s_{\text{left}}$  receives  $L$  and does not envy the other agents (up to good  $v_r$ ) since  $s_{\text{left}}$  is a shouter.
- An agent  $i$  who is not a shouter does not envy  $s_{\text{left}}$  because  $i$  prefers either  $M$  or  $R$  to  $L$ , and hence by Lemma 4.1 receives a bundle preferred to  $L$ .  
Agent  $i$  also does not envy the other agent  $j \neq s_{\text{left}}$  up to one good by Lemma 4.1.
- An agent  $i \neq s_{\text{left}}$  who is a shouter does not envy  $s_{\text{left}}$  up to one good: If this is the first time Step 2 was performed, then  $L = \{v_1\}$ , so  $i$  does not envy  $s_{\text{left}}$  up to  $v_1$ . Otherwise, the last step was an iteration of Step 4(c), where by definition of Step 4(c) no-one shouted. Since  $i$  did not shout during Step 4(c), and Step 2 did not change the bundles  $M$  and  $R$ , then  $i$  strictly prefers either  $M$  or  $R$  to the left bundle  $L \setminus \{v_\ell\}$  of Step 4(c). By Lemma 4.1, agent  $i$  gets a bundle at least as good as  $M$  or  $R$ . Thus,  $i$  does not envy  $s_{\text{left}}$  up to  $v_\ell$ .  
Also by Lemma 4.1, agent  $i$  does not envy the other agent  $j \neq s_{\text{left}}$  up to one good.

**Step 3.**

- Agent  $s_{\text{left}}$  receives  $L$  and, because  $s_{\text{left}}$  shouted, does not envy the bundle  $\{v_{\ell+1}\} \cup M$  up to good  $v_{\ell+1}$ , and does not envy the bundle  $\{v_r\} \cup R$  up to good  $v_r$ .
- Agent  $c$  gets his preferred bundle among  $\{v_{\ell+1}\} \cup M$  and  $\{v_r\} \cup R$ , and so does not envy agent  $s$  who receives the other bundle. Further, agent  $c$  does not envy agent  $s_{\text{left}}$  since  $c$  did not shout at the last Step 2 (where no-one shouted), which, since bundle  $L$  did not change in Step 3, means that  $c$  prefers either  $\{v_{\ell+1}\} \cup M$  or  $R$  to  $L$ , and hence also prefers his chosen bundle to  $L$ .
- Agent  $s$  is a middle agent, so the lumpy tie of  $s$  over  $P(v_{\ell+1}, v_m)$  is  $v_r$ , and hence by (4.1),

$$u_s(\{v_r\} \cup R) \geq u_s(\{v_{\ell+1}\} \cup M). \tag{4.3}$$

Now, agent  $s$  did not shout at the preceding Step 2 (when no-one shouted). However,  $s$  does shout after deleting  $v_{\ell+1}$  from  $M$ . Since  $L$  and  $R$  have not changed, the reason  $s$  did not shout at Step 2 was that  $L$  is worse than the middle bundle during Step 2, so

$$u_s(\{v_{\ell+1}\} \cup M) > u_s(L). \tag{4.4}$$

Combining (4.3) and (4.4), we also have

$$u_s(\{v_r\} \cup R) > u_s(L).$$

Since  $s$  receives either  $\{v_{\ell+1}\} \cup M$  or  $\{v_r\} \cup R$ , agent  $s$  does not envy agent  $s_{\text{left}}$  receiving  $L$ . Finally, from (4.3), agent  $s$  weakly prefers  $\{v_r\} \cup R$  to  $\{v_{\ell+1}\} \cup M$ . Thus, if  $c$  picks  $\{v_{\ell+1}\} \cup M$ , then  $s$  does not envy  $c$ . On the other hand, if  $c$  picks the bundle  $\{v_r\} \cup R$ , then  $s$  does not envy  $c$  up to good  $v_r$ : we have  $u_s(L) \geq u_s(R)$  since  $s$  shouts, and so by (4.4), also

$$u_s(\{v_{\ell+1}\} \cup M) > u_s(R).$$

**Step 4(a).** We first prove that if  $i$  is a shouter who did not shout in the previous step, then

$$u_i(\{v_r\} \cup R) > u_i(L) \geq u_i(M). \tag{4.5}$$

In the previous step (which was either Step 3 or Step 4), the middle bundle was  $M \setminus \{v_{r-1}\}$  and the right bundle was  $\{v_r\} \cup R$ . (While Step 4 allows for the possibility that the middle and right bundles are not changed in Step 4, this is not the case if we enter Step 4(a): if the bundles are unchanged and two agents shout, these agents already shouted in Step 3, contradicting that we did not terminate then.) Since  $i$  did not shout with the middle and right bundles of the previous step, we have

$$u_i(M \setminus \{v_{r-1}\}) > u_i(L) \quad \text{or} \quad u_i(\{v_r\} \cup R) > u_i(L).$$

Since  $i$  is a shouter,  $u_i(L) \geq u_i(M)$ , so that the first case is impossible by monotonicity. Hence  $u_i(\{v_r\} \cup R) > u_i(L)$ , showing (4.5), when combined with  $u_i(L) \geq u_i(M)$ .

- Agent  $s_{\text{left}}$  receives  $L$  and does not envy other agents up to one good like in Step 3.

- Agent  $c$  gets his preferred bundle among  $\{v_{\ell+1}\} \cup M$  and  $\{v_r\} \cup R$ , and so does not envy agent  $s$  who receives the other bundle. Agent  $c$  also does not envy  $s_{\text{left}}$ : If  $c$  is not a shouter, then  $c$  does not envy  $s_{\text{left}}$  because  $c$  prefers either  $M$  or  $R$  to  $L$ , and hence prefers his picked piece to  $L$ . If  $c$  is a shouter, then all three agents are shouters, and by choice of  $c$ , this means that  $c$  was not a shouter at the previous step, when there was at most one shouter. By (4.5),  $u_c(\{v_r\} \cup R) > u_c(L)$ , and hence

$$\max\{u_c(\{v_{\ell+1}\} \cup M), u_c(\{v_r\} \cup R)\} \geq u_c(L),$$

so that  $c$  does not envy  $s_{\text{left}}$ .

- Agent  $s$  does not envy others up to one good:
  - Suppose agent  $c$  strictly prefers  $\{v_r\} \cup R$  to  $\{v_{\ell+1}\} \cup M$ . Then agent  $c$ 's lumpy tie over  $P(v_{\ell+1}, v_m)$  appears at or after  $v_r$  by definition of the lumpy tie. As we argued before, the bundles  $M$  and  $R$  were changed in the execution of Step 4, and  $r$  was increased by 1. Thus,  $v_r$  appears strictly after the median lumpy tie over  $P(v_{\ell+1}, v_m)$ . Thus,  $c$  is the right agent over  $P(v_{\ell+1}, v_m)$ . Hence  $s$  is either a left or middle agent over  $P(v_{\ell+1}, v_m)$  since there is at most one right agent. Using (4.1) or (4.2), this implies

$$u_s(\{v_{\ell+1}\} \cup M) \geq u_s(\{v_r\} \cup R), \tag{4.6}$$

so that  $s$  does not envy  $c$ .

By definition of  $s$ , agent  $s$  did not shout in the previous step. By (4.5),  $u_s(\{v_r\} \cup R) \geq u_s(L)$ , so together with (4.6), we have  $u_s(\{v_{\ell+1}\} \cup M) \geq u_s(L)$ , so  $s$  does not envy  $s_{\text{left}}$ .

- Suppose  $c$  weakly prefers  $\{v_{\ell+1}\} \cup M$  to  $\{v_r\} \cup R$ . Then  $s$  receives the bundle  $\{v_r\} \cup R$  (since  $c$  breaks ties in favor of  $\{v_{\ell+1}\} \cup M$ ). By choice of  $s$ , agent  $s$  did not shout at the last step. So by (4.5), we have  $u_s(\{v_r\} \cup R) > u_s(L)$  so that  $s$  does not envy  $s_{\text{left}}$ , and also by (4.5), we have  $u_s(\{v_r\} \cup R) > u_s(M)$  so that  $s$  does not envy  $c$  up to item  $v_{\ell+1}$ .

**Step 4(b).**

- Agent  $s_{\text{left}}$  gets  $L \cup \{v_{\ell+1}\}$  and does not envy the other agents (up to good  $v_r$ ) as  $s_{\text{left}}$  shouts.
- Any agent  $i \neq s_{\text{left}}$  is not a shouter, and thus prefers either  $M$  or  $R$  to  $L$ . Hence by Lemma 4.1 receives a bundle preferred to  $L$ , and so does not envy  $s_{\text{left}}$  up to item  $v_{\ell+1}$ .  
Agent  $i$  also does not envy the other agent  $j \neq s_{\text{left}}$  up to one good by Lemma 4.1.

Thus, the allocation returned by any of the steps satisfies EF1. Our algorithm can be implemented in  $O(m)$  time: Each of steps 2, 3, and 4 will be executed at most  $m$  times (since  $\ell$  and  $r$  can only be incremented  $m$  times). The execution of each step takes constant time: In each step, we need to check which agents shout, and this can be done in a constant number of queries to agents' valuations; also, in Step 4 we need to calculate the lumpy ties of the agents, but this can be done in amortized constant time, since during the execution of the algorithm, the position of each agent's lumpy tie can only move to the right. Finally, when enough agents shout, we can clearly compute and return the final allocation in  $O(m)$  time.  $\square$

**5. EF2 existence for any number of agents**

For two or three agents, we have seen algorithms that are guaranteed to find an EF1 allocation on a path (and on traceable graphs). Both algorithms were adaptations of procedures that identify envy-free divisions in the cake-cutting problem. For the case of four or more agents, we face a problem: there are no known procedures that find connected envy-free division in cake-cutting if the number of agents is larger than three. However, in the divisible setting, a non-constructive existence result is known: Su (1999) proved, using Sperner's lemma, that for any number of agents, a connected envy-free division of a cake always exists. One might try to use this result as a black box to obtain a fair allocation for the indivisible problem on a path: Translate an indivisible instance with additive valuations into a divisible cake (where each item corresponds to a region of the cake), obtain an envy-free division of the cake, and round it to get an allocation of the items. Suksompong (2019) followed this approach and showed that the result is an allocation where any agent  $i$ 's envy  $u_i(A(j)) - u_i(A(i))$  is at most  $2u_{\max}$ , where  $u_{\max}$  is the maximum valuation for a single item.

In this section, rather than using Su's (1999) result as a black box, we directly apply Sperner's lemma to the indivisible problem. This allows us to obtain a stronger fairness guarantee: We show that on paths (and on traceable graphs), there always exists an EF2 allocation.<sup>5</sup> An allocation is EF2 if any agent's envy can be avoided by removing up to two items from the envied bundle. Again, we only allow removal of items if this operation leaves a connected bundle.

**Definition 5.1** (EF2: envy-freeness up to two outer goods). An allocation  $A$  satisfies EF2 if, for any pair  $i, j \in N$  of agents, either  $|A(j)| \leq 1$ , or there are two goods  $u, v \in A(j)$  such that  $A(j) \setminus \{u, v\}$  is connected and  $u_i(A(i)) \geq u_i(A(j) \setminus \{u, v\})$ .

<sup>5</sup> To see that EF2 is a stronger property than bounding envy up to  $2u_{\max}$ , consider a path of four items and two agents with additive valuations 1–10–2–2. The allocation (1, 10–2–2) is not EF2, but the first agent has an envy of  $13 < 20 = 2u_{\max}$ .




Fig. 5. Connected partitions as a subdivided simplex.

Let us first give a high-level illustration with three agents of how Sperner’s lemma can be used to find low-envy allocations.

Given a path  $P = (a, b, c, d)$ , the family of connected partitions of  $P$  can naturally be arranged as the vertices of a subdivided simplex, as in Fig. 5.

For each of these partitions, each agent  $i$  labels the corresponding vertex by the index of a bundle from that partition that  $i$  most-prefers. For example, the top vertex will be labeled as “index 1” by all agents, since they all most-prefer the leftmost bundle in  $(abcd, \emptyset, \emptyset)$ . Now, Sperner’s lemma will imply that at least one of the simplices (say the shaded one) is “fully-labeled”, which means that the first agent most-prefers the leftmost bundle at one vertex, the second agent most-prefers the middle bundle at another vertex, and the third agent most-prefers the rightmost bundle at the last vertex. Notice that the partitions at the corner points of the shaded simplex are all “similar” to each other (they can be obtained from each other by moving only one item). Hence, we can “round” the corner-partitions into a common allocation  $A^*$ , say by picking one of the corner partitions arbitrarily and then allocating bundles to agents according to the labels. The resulting allocation has the property that any agents’ envy can be eliminated by moving at most one good.<sup>6</sup>

The argument sketched above does not yield an EF1 nor even an EF2 allocation. Intuitively, the problem is that the connected partitions at the corners of the fully-labeled simplex are “too far apart”, so that no matter how we round the corner partitions into a common allocation  $A^*$ , some agents’ bundles will have changed too much, and so we cannot prevent envy even up to one or two goods. In the following, we present a solution to this problem, by considering a finer subdivision: we introduce  $n - 1$  knives which move in half-steps (rather than full steps), and which might ‘cover’ an item so that it appears in none of the bundles. The result is that the partial partitions in the corners of the fully-labeled simplex are closer together, and can be successfully rounded into an EF2 allocation  $A^*$ .

In our approach, we use a specific triangulation (Kuhn’s triangulation, Kuhn, 1960). This triangulation has the needed property that the partitions at the corners of sub-simplices are close together, and adjacent partitions can be obtained from each other in a natural way. While this type of triangulation has also been used in cake-cutting, e.g., by Deng et al. (2012), there it was only used to speed up algorithms (compared to the barycentric subdivision used by Su (1999)), not to obtain better fairness properties.

### 5.1. Sperner’s lemma

We start by formally introducing Sperner’s lemma (cf. Flegg, 1974). Let  $\text{conv}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k)$  denote the convex hull of  $k$  vectors  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ . An  $n$ -simplex is an  $n$ -dimensional polytope which is the convex hull of its  $n + 1$  main vertices. A  $k$ -face of the  $n$ -simplex is the  $k$ -simplex formed by the span of any subset of  $k + 1$  main vertices. A triangulation  $T$  of a simplex  $S$  is a collection of sub- $n$ -simplices whose union is  $S$  with the property that the intersection of any two of them is either the empty set, or a face common to both. Each of the sub-simplices  $S^* \in T$  is called an elementary simplex of the triangulation  $T$ . We denote by  $V(T)$  the set of vertices of the triangulation  $T$ , i.e., the union of vertices of the elementary simplices of  $T$ .

<sup>6</sup> One can generalize this argument to show that on paths, there exists an allocation  $A$  satisfying a weak form of EF1: for any  $i, j \in [n]$ , we have  $u_i(I_i \cup \{g_i\}) \geq u_i(I_j \setminus \{g_j\})$  for some items  $g_i, g_j$  such that  $I_i \cup \{g_i\}$  and  $I_j \setminus \{g_j\}$  are connected. For additive valuations, this implies that envy is bounded by  $u_i(g_i) + u_i(g_j) \leq 2u_{\max}$ , which is the result of Suksompong (2019).

Let  $T$  be some fixed triangulation of an  $(n - 1)$ -simplex  $S = \text{conv}(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n)$ . A *labeling function* is a function  $L : V(T) \rightarrow [n]$  that assigns a number in  $[n]$  (called a *color*) to each vertex of the triangulation  $T$ . A labeling function  $L$  is called *proper* if

- For each main vertex  $\mathbf{v}_i$  of the simplex,  $L$  assigns color  $i$  to  $\mathbf{v}_i$ :  $L(\mathbf{v}_i) = i$ ; and
- $L(\mathbf{v}) \neq i$  for any vertex  $\mathbf{v} \in V(T)$  belonging to the  $(n - 2)$ -face of  $S$  not containing  $\mathbf{v}_i$ .

Sperner’s lemma states that if  $L$  is a proper labeling function, then there exists an elementary simplex of  $T$  whose vertices have all different labels.

We will consider a generalized version of Sperner’s lemma, proved, for example, by Bapat (1989). In this version, there are  $n$  labeling functions  $L_1, \dots, L_n$ , and we are looking for an elementary simplex that is *fully-labeled* for some way of assigning labeling functions to vertices, where we must use each labeling function exactly once. The formal definition is as follows.

**Definition 5.2** (*Fully-labeled simplex*). Let  $T$  be a triangulation of an  $(n - 1)$ -simplex, and let  $L_1, \dots, L_n$  be labeling functions. An elementary simplex  $S^*$  of  $T$  is *fully-labeled* if we can write  $S^* = \text{conv}(\mathbf{v}_1^*, \mathbf{v}_2^*, \dots, \mathbf{v}_n^*)$  such that there exists a permutation  $\phi : [n] \rightarrow [n]$  with

$$L_i(\mathbf{v}_i^*) = \phi(i) \quad \text{for each } i \in [n].$$

The generalized version of Sperner’s lemma that we consider, taken from Bapat (1989), guarantees the existence of a fully-labeled simplex.

**Lemma 5.3** (*Generalized Sperner’s Lemma*). Let  $T$  be a triangulation of an  $(n - 1)$ -simplex  $S$ , and let  $L_1, \dots, L_n$  be proper labeling functions. Then there is a fully-labeled simplex  $S^*$  of  $T$ .

### 5.2. Existence of EF2 allocations

Suppose that our graph  $G$  is a path  $P = (1, 2, \dots, m)$ , where the items are named by integers. We assume that  $m \geq n$ , so that there are at least as many items as agents (when  $m < n$  it is easy to find EF1 allocations). Our aim is to cut the path  $P$  into  $n$  intervals (bundles)  $I_1^*, I_2^*, \dots, I_n^*$ . Throughout the argument, we use superscripts to denote indices of bundles; index 1 refers to the leftmost bundle and index  $n$  refers to the rightmost bundle.

*Construction of the triangulation.* Consider the  $(n - 1)$ -simplex<sup>7</sup>

$$S_m = \{ \mathbf{x} \in \mathbb{R}^{n-1} : \frac{1}{2} \leq x^1 \leq x^2 \leq \dots \leq x^{n-1} \leq m + \frac{1}{2} \}. \tag{5.1}$$

We construct a triangulation  $T_{\text{half}}$  of  $S_m$  whose vertices  $V(T_{\text{half}})$  are the points  $\mathbf{x} \in S_m$  such that each  $x^j$  is either integral or half-integral, namely,

$$V(T_{\text{half}}) = \{ \mathbf{x} \in S_m : x^j \in \{ \frac{1}{2}, 1, \frac{3}{2}, 2, \frac{5}{2}, \dots, m, m + \frac{1}{2} \} \text{ for all } j \in [n] \}.$$

For reasons that will become clear shortly, we call a vector  $\mathbf{x} \in V(T_{\text{half}})$  a *knife position*.

Using Kuhn’s triangulation (Kuhn, 1960; Scarf, 1982; Deng et al., 2012), we construct  $T_{\text{half}}$  so we can write each elementary simplex  $S' \in T_{\text{half}}$  as  $S' = \text{conv}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$  and there is a permutation  $\pi : [n] \rightarrow [n]$  with

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \frac{1}{2} \mathbf{e}^{\pi(i)} \quad \text{for each } i \in [n - 1], \tag{5.2}$$

where  $\mathbf{e}^j = (0, \dots, 1, \dots, 0)$  is the  $j$ -th unit vector.

We give an interpretation of (5.2) shortly. Each vertex  $\mathbf{x} = (x^1, x^2, \dots, x^{n-1}) \in V(T_{\text{half}})$  of the triangulation  $T_{\text{half}}$  corresponds to a partial partition  $A(\mathbf{x}) = (I^1(\mathbf{x}), I^2(\mathbf{x}), \dots, I^n(\mathbf{x}))$  of  $P$  where  $I^j(\mathbf{x}) := \{ y \in \{1, 2, \dots, m\} : x^{j-1} < y < x^j \}$ , writing  $x^0 = \frac{1}{2}$  and  $x^n = m + \frac{1}{2}$  for convenience. Intuitively,  $\mathbf{x}$  specifies the location of  $n - 1$  knives that cut  $P$  into  $n$  pieces. If  $x^j$  is integral, that is  $x^j \in \{1, \dots, m\}$ , then the  $j$ -th knife ‘covers’ the item  $x^j$ , which is then part of neither  $I^j(\mathbf{x})$  nor  $I^{j+1}(\mathbf{x})$ . This is why  $A(\mathbf{x})$  is a *partial* partition. Since there are only  $n - 1$  knives but  $m \geq n$  items, not all items are covered, so at least one bundle is non-empty.

Property (5.2) means that, if we visit the knife positions  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$  at the corners of an elementary simplex in the listed order, then at each step exactly one of the knives moves by half a step, and each knife moves only at one of the steps.

<sup>7</sup> The simplex  $S_m$  is affinely equivalent to the standard  $(n - 1)$ -simplex  $\Delta_{n-1} = \{(l_1, \dots, l_n) \geq 0 : \sum l_i = 1\}$  via  $x_i = m \cdot (l_1 + l_2 + \dots + l_i) + \frac{1}{2}$ . In these coordinates,  $l_i$  is the length of the  $i$ -th piece (times  $1/m$ ).

*Construction of the labeling functions.* We now construct, for each agent  $i \in [n]$ , a labeling function  $L_i : V(T_{\text{half}}) \rightarrow [n]$ . The function  $L_i$  takes as input a vertex  $\mathbf{x}$  of the triangulation  $T_{\text{half}}$  (interpreted as the partial partition  $A(\mathbf{x})$ ), and returns a color in  $[n]$ . The color will specify the index of a bundle in  $A(\mathbf{x})$  that agent  $i$  likes the most. Formally,

$$L_i(\mathbf{x}) \in \{j \in [n] : u_i(I^j(\mathbf{x})) \geq u_i(I^k(\mathbf{x})) \text{ for all } k \in [n]\}.$$

If there are several most-preferred bundles in  $A(\mathbf{x})$ , ties can be broken arbitrarily. However, we insist that the index  $L_i(\mathbf{x})$  always corresponds to a non-empty bundle (this can be ensured since  $A(\mathbf{x})$  always contains a non-empty bundle, and  $u_i$  is monotonic).

The labeling functions  $L_i$  are proper. For each  $j \in [m]$ , the main vertex  $\mathbf{v}_j$  of the simplex  $S_m$  has the form  $\mathbf{v}_j = (\frac{1}{2}, \dots, \frac{1}{2}, m + \frac{1}{2}, \dots, m + \frac{1}{2})$ , where the first  $j - 1$  entries are  $\frac{1}{2}$  and the rest are  $m + \frac{1}{2}$ . In the partition  $A(\mathbf{v}_j)$ , the bundle  $I^j(\mathbf{v}_j)$  contains all the items, so is most-preferred (since  $u_i$  is monotonic and by our tie-breaking), and so  $L_i(\mathbf{v}_j) = j$ . Further, any vertex  $\mathbf{x}$  belonging to the  $(n - 2)$ -face of  $S_m$  not containing  $\mathbf{v}_j$  satisfies  $x^{j-1} = x^j$ , and thus in partition  $A(\mathbf{x})$ , bundle  $I^j(\mathbf{x})$  is empty, hence is *not* selected, and so  $L_i(\mathbf{x}) \neq j$ .

By the generalized version of Sperner's lemma (Lemma 5.3), there exists an elementary simplex  $S^* = \text{conv}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$  of the triangulation  $T_{\text{half}}$  which is fully-labeled, so that, for some permutation  $\phi : [n] \rightarrow [n]$ , we have  $L_i(\mathbf{x}_i) = \phi(i)$  for all  $i \in [n]$ .

*Translation into partial partitions.* The fully-labeled elementary simplex  $S^*$  corresponds to a sequence  $(A_1, A_2, \dots, A_n)$  of partial partitions of  $P$ , which we call the *Sperner sequence*, where  $A_i = (I_i^1, \dots, I_i^n) := A(\mathbf{x}_i)$  for each  $i \in [n]$ . An example of a Sperner sequence is shown in Fig. 6. From the labeling, for each agent  $i \in [n]$ , since  $L_i(\mathbf{x}_i) = \phi(i)$ , the bundle with index  $\phi(i)$  in the partition  $A_i$  is a best bundle for  $i$ :

$$u_i(I_i^{\phi(i)}) \geq u_i(I_i^j) \quad \text{for each } j \in [n]. \tag{5.3}$$

Now, for each  $j \in [n]$ , we define the *basic bundle*  $B^j := I_1^j \cap \dots \cap I_n^j$  to be the bundle of items that appear in the  $j$ -th bundle of every partition in the Sperner sequence. The set of basic bundles is a partial partition. Let us analyze the items between basic bundles.

From (5.2), each of the  $n - 1$  knives moves exactly once, by half a step, while passing through the Sperner sequence  $(A_1, A_2, \dots, A_n)$ . Thus, the numbers  $x_1^j, \dots, x_n^j$  take on two different values, one of which is integral and the other half-integral. We write  $y^j$  for the integral value (so  $y^j = x_i^j$  for some  $i \in [n]$ ), and call  $y^j$  a *boundary item*. The  $j$ -th knife covers the item  $y^j$  in some, but not all, of the partial partitions in the Sperner sequence. Now, there are two cases:

- (a)  $x_1^j = \dots = x_i^j = y^j - \frac{1}{2}$  and  $x_{i+1}^j = \dots = x_n^j = y^j$  for some  $i \in [n]$ , so that  $y^j$  never occurs in the  $j$ -th bundle in the Sperner sequence but sometimes occurs in the  $(j + 1)$ -th bundle, or
- (b)  $x_1^j = \dots = x_i^j = y^j$  and  $x_{i+1}^j = \dots = x_n^j = y^j + \frac{1}{2}$  for some  $i \in [n]$ , so that  $y^j$  sometimes occurs in the  $j$ -th bundle in the Sperner sequence but never occurs in the  $(j + 1)$ -th bundle.

Since  $y^j$  is sometimes covered by a knife, it is not part of any basic bundle. Note that

$$B^j \subseteq I_i^j \subseteq \{y^{j-1}\} \cup B^j \cup \{y^j\} \quad \text{for every } i, j \in [n]. \tag{5.4}$$

*Rounding into a complete partition.* We now construct a complete partition of the path  $P$  into the bundles  $(I_*^1, I_*^2, \dots, I_*^n)$  which are defined as follows:

$$I_*^j := I_1^j \cup \dots \cup I_n^j \quad \text{for each } j \in [n].$$

Thus, the bundle  $I_*^j$  contains the basic bundle  $B^j$ , plus all of the boundary items  $y^{j-1}$  or  $y^j$  that occur in the  $j$ -th bundle at some point of the Sperner sequence. Precisely, for each boundary item  $y^j$ ,  $j \in [n - 1]$ , the item  $y^j$  is placed in bundle  $I_*^{j+1}$  in case (a) above, and it is placed in bundle  $I_*^j$  in case (b). Thus, every item is allocated to exactly one bundle.

*An EF2 allocation.* We first show that the partition  $(I_*^1, I_*^2, \dots, I_*^n)$  is such that agents' expectations about the value of the bundles  $I_*^j$  are approximately correct (up to two items):

$$u_i(I_*^j) \geq u_i(I_i^j) \geq u_i(B^j) \quad \text{for every agent } i \in [n] \text{ and every } j \in [n]. \tag{5.5}$$

This follows by monotonicity of  $u_i$ , since  $I_*^j = I_1^j \cup \dots \cup I_n^j \supseteq I_i^j \supseteq B^j$  by (5.4).

Now, based on the partition, we define an allocation  $A_*$  by  $A_*(i) = I_*^{\phi(i)}$  for each agent  $i \in [n]$ . Then  $A_*$  satisfies EF2: For any pair  $i, j \in [n]$  of agents, we have




Fig. 6. Example of the Sperner sequence  $A_1, \dots, A_4$  for  $n = 4$ , as well as the derived partition  $A_*$ . Vertical lines indicate the positions  $x_i^1, x_i^2, x_i^3$  of the knives,  $i = 1, \dots, 4$ . Shaded in gray, for  $i = 1, \dots, 4$ , is the bundle  $I_i^{\phi(i)}$  selected by agent  $i$  as their favorite bundle in  $A_i$ .

$$\begin{aligned}
 u_i(A_*(i)) &= u_i(I_*^{\phi(i)}) \geq u_i(I_i^{\phi(i)}) && \text{by (5.5)} \\
 &\geq u_i(I_i^{\phi(j)}) && \text{by (5.3)} \\
 &\geq u_i(B^{\phi(j)}) && \text{by (5.5)} \\
 &= u_i(A_*(j) \setminus \{y^{j-1}, y^j\}). && \text{by (5.4)}
 \end{aligned}$$

Hence, we have proved the main result of this section:

**Theorem 5.4.** *On a path, for any number of agents with monotone valuation functions, a connected EF2 allocation exists.*

### 6. EF1 existence for four agents

We have seen that Sperner’s lemma can be used to show EF2 existence for any number of agents. Why does our proof in the previous section only establish EF2, and not EF1? The reason is that agents’ expectations about the contents of a bundle might differ by up to two goods from what the bundle will actually contain. In the notation of the previous section, an agent  $i$  may be presented with a partial partition  $I_i$  where the  $j$ -th bundle  $I_i^j$  is the basic bundle, i.e.,  $I_i^j = B^j$ . The agent then selects their favorite bundle from  $I_i$ , implicitly assuming that the  $j$ -th bundle in the rounded partition  $I_*$  will also equal  $B^j$ , i.e., that  $I_*^j = B^j$ . However, it may happen that in fact  $I_*^j = \{y^{j-1}\} \cup B^j \cup \{y^j\}$ , and then  $i$  envies the agent who receives bundle  $j$  by a margin of two goods.

For four agents, we can adapt our argument to achieve EF1. To do this, we both change the way we round the Sperner sequence into an allocation, and define new labeling functions that better anticipate how a partial partition will be rounded into the final allocation. In this way, agents’ expectations about bundles can only be wrong up to one good. In crude terms, agents will expect that each of the two interior bundles will be assigned at least one of the boundary items, and the rounding method ensures that this will indeed happen.

Let  $n = 4$ . Formally, to define the labeling function, for each agent  $i \in [n]$  we construct a virtual valuation function  $\hat{u}_i(\mathbf{x}, j)$  which assigns a value to each bundle  $j \in [n]$  of a partial allocation as specified by a vertex  $\mathbf{x} \in V(T_{\text{half}})$ . The way these virtual valuations are defined differs based on the index  $j$ ; in particular, end bundles ( $j = 1, 4$ ) are treated differently from interior bundles ( $j = 2, 3$ ). The virtual valuations are defined as follows, for each  $\mathbf{x} \in V(T_{\text{half}})$  and each  $i \in [n]$ , where the middle row (6.2) applies to  $j = 2$  and  $j = 3$ :

$$\hat{u}_i(\mathbf{x}, 1) = \begin{cases} u_i(\{1, \dots, x^1 - 1\}) & \text{if } x^1 \in \mathbb{Z}, \\ u_i(\{1, \dots, x^1 - \frac{3}{2}\}) & \text{if } x^1 \notin \mathbb{Z}. \end{cases} \tag{6.1}$$

$$\hat{u}_i(\mathbf{x}, j) = \begin{cases} u_i^-(\{x^{j-1}, \dots, x^j\}) & \text{if } x^{j-1} \in \mathbb{Z} \text{ and } x^j \in \mathbb{Z}, \\ u_i(I^j(\mathbf{x})) & \text{otherwise.} \end{cases} \tag{6.2}$$

$$\hat{u}_i(\mathbf{x}, 4) = \begin{cases} u_i(\{x^3 + 1, \dots, m\}) & \text{if } x^3 \in \mathbb{Z}, \\ u_i(\{x^3 + \frac{3}{2}, \dots, m\}) & \text{if } x^3 \notin \mathbb{Z}. \end{cases} \tag{6.3}$$

Thus, for an interior bundle  $j = 2, 3$ , if both the items  $x^{j-1}$  and  $x^j$  to either side of the bundle are covered by a knife, an agent expects that one of these items (the less-valuable one) will be put into bundle  $I_*^j$  of the final rounded allocation (recall the definition of  $u_i^-$  in equation (2.1)). For exterior bundles,  $j = 1$  (resp.  $j = 4$ ), if the item  $x^1$  (resp.  $x^3$ ) is not covered by a knife, the agent does not expect the interior item (next to the knife) to belong to the final bundle  $I_*^j$ , even though it belongs to the observed bundle  $I_i^j$ . Otherwise, the virtual allocations are equal to  $u_i(I^j(\mathbf{x}))$ , so the agent expects that  $I_*^j = I_i^j$ . Later, we show that these expectations are correct up to one item.

Using these virtual valuations, we define labeling functions  $\hat{L}_i : V(T_{\text{half}}) \rightarrow [n]$  so that

$$\hat{L}_i(\mathbf{x}) \in \{j \in [n] : \hat{u}_i(\mathbf{x}, j) \geq \hat{u}_i(\mathbf{x}, k) \text{ for all } k \in [n]\}.$$

One can check that these valuation functions are still proper.

Again, by Sperner’s lemma, there exists an elementary simplex  $S^* = \text{conv}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$  of the triangulation  $T_{\text{half}}$  which is fully-labeled according to our new labeling function: there is a permutation  $\phi : [n] \rightarrow [n]$ , with  $\hat{L}_i(\mathbf{x}_i) = \phi(i)$  for all  $i \in [n]$ . Again, this elementary simplex induces a Sperner sequence  $(A_1, \dots, A_n)$  of partial partitions.

To shorten a case distinction, we assume that  $y^2 \in I_1^2 \cup I_2^2 \cup I_3^2 \cup I_4^2$ , i.e., that the boundary item  $y^2$  appears in the second but not in the third bundle in the Sperner sequence. This assumption is without loss of generality, since by the left-right symmetry of the definition of virtual valuations, if necessary we can reverse the path  $P$  and consider the same elementary simplex with vertices ordered in reverse  $(\mathbf{x}_4, \mathbf{x}_3, \mathbf{x}_2, \mathbf{x}_1)$ ; it will still be fully-labeled.

With this assumption made throughout the rest of the argument, we now round the Sperner sequence into a complete partition  $(I_*^1, I_*^2, I_*^3, I_*^4)$  of  $P$  defined as follows:

$$I_*^1 := I_1^1 \cup \dots \cup I_4^1, \quad I_*^2 := I_1^2 \cup \dots \cup I_4^2, \quad I_*^3 := B^3 \cup \{y^3\}, \quad I_*^4 := B^4.$$

Depending on the placement of the boundary item  $y^1$ , we will either have  $I_*^1 = B^1$  or  $I_*^1 = B^1 \cup \{y^1\}$ ; and either  $I_*^2 = \{y^1\} \cup B^2 \cup \{y^2\}$  or  $I_*^2 = B^2 \cup \{y^2\}$ . With these choices, each interior bundle ( $j = 2, 3$ ) receives at least one of the boundary items adjacent to it.

The main part of showing that the partition  $(I_*^1, I_*^2, I_*^3, I_*^4)$  can be made into an EF1 allocation is an analogue of (5.5), which shows that agents’ expectations about their bundle are approximately correct. The following analogous proposition is proved by case analysis.

**Proposition 6.1.** *For each  $i \in [n]$  and each  $j \in [n]$ , we have  $u_i(I_*^j) \geq \hat{u}_i(\mathbf{x}_i, j) \geq u_i^-(I_*^j)$ .*

**Proof.** We consider each bundle  $j = 1, 2, 3, 4$  separately.

- Suppose  $j = 1$ .
  - Suppose  $x_i^1 \in \mathbb{Z}$ . Then  $y^1 = x_i^1$  and  $B^1 = \{1, \dots, x_i^1 - 1\}$ . Thus  $u_i(I_*^1) \geq \hat{u}_i(\mathbf{x}_i, 1) = u_i(B^1) \geq u_i^-(I_*^1)$ , since  $I_*^1$  is either  $B^1$  or  $B^1 \cup \{y^1\}$ .
  - Suppose  $x_i^1 \notin \mathbb{Z}$ . Then  $\hat{u}_i(\mathbf{x}_i, 1) = u_i(\{1, \dots, x_i^1 - \frac{3}{2}\})$ . Now, either
    - \*  $y^1 = x_i^1 - \frac{1}{2}$  so that  $y^1 \in I_1^1$ , and so  $I_*^1 = B^1 \cup \{y^1\} = \{1, \dots, x_i^1 - \frac{1}{2}\}$ , or
    - \*  $y^1 = x_i^1 + \frac{1}{2}$  so that  $y^1 \notin I_*^1$ , and so  $I_*^1 = B^1 = \{1, \dots, x_i^1 - \frac{1}{2}\}$ .
 In either case,  $I_*^1 = \{1, \dots, x_i^1 - \frac{3}{2}, x_i^1 - \frac{1}{2}\}$ , so  $u_i(I_*^1) \geq \hat{u}_i(\mathbf{x}_i, 1) \geq u_i^-(I_*^1)$ .
- Suppose  $j = 2$ , and suppose that  $I_*^2 = B^2 \cup \{y^2\}$ 
  - Suppose  $x_i^1 \in \mathbb{Z}$  and  $x_i^2 \in \mathbb{Z}$ . So  $x_i^1 = y^1$  and  $x_i^2 = y^2$ . Then  $\hat{u}_i(\mathbf{x}_i, 2) = u_i^-(\{y^1, \dots, y^2\})$ . Thus  $u_i(I_*^2) \geq \hat{u}_i(\mathbf{x}_i, 2) \geq u_i^-(I_*^2)$  since  $I_*^2 = B^2 \cup \{y^2\}$ .
  - Otherwise  $\hat{u}_i(\mathbf{x}_i, 2) = u_i(I_i^2(\mathbf{x}))$ . Since  $y^1 \notin I_1^2(\mathbf{x})$  (because  $y^1 \in I_*^1$ ), we have that  $I_i^2(\mathbf{x})$  is either  $B^2$  or  $B^2 \cup \{y^2\}$ . So  $u_i(I_*^2) \geq \hat{u}_i(\mathbf{x}_i, 2) = u_i(I_i^2(\mathbf{x})) \geq u_i^-(I_*^2)$  since  $I_*^2 = B^2 \cup \{y^2\}$ .
- Suppose  $j = 2$ , and suppose that  $I_*^2 = \{y^1\} \cup B^2 \cup \{y^2\}$ .
  - Suppose  $x_i^1 \in \mathbb{Z}$  and  $x_i^2 \in \mathbb{Z}$ . So  $x_i^1 = y^1$  and  $x_i^2 = y^2$ . Then  $\hat{u}_i(\mathbf{x}_i, 2) = u_i^-(\{y^1, \dots, y^2\})$ . Thus  $u_i(I_*^2) \geq \hat{u}_i(\mathbf{x}_i, 2) = u_i^-(I_*^2)$  since  $I_*^2 = \{y^1\} \cup B^2 \cup \{y^2\}$ .
  - Otherwise  $\hat{u}_i(\mathbf{x}_i, 2) = u_i(I_i^2(\mathbf{x}))$ . First note that  $I_i^2(\mathbf{x}) \neq B^2$ : this is because both  $y^1$  and  $y^2$  appear in the second bundle of the Sperner sequence (by the case and the wlog assumption), so that  $x_i^1 \leq y^1$  and  $y^2 \leq x_i^2$ . Since at least one of  $x_i^1$  or  $x_i^2$  is not integral, at least one of  $y^1$  or  $y^2$  must be in  $I_i^2(\mathbf{x})$ . Hence  $I_i^2(\mathbf{x})$  is either  $\{y^1\} \cup B^2 \cup \{y^2\}$  or  $\{y^1\} \cup B^2$  or  $B^2 \cup \{y^2\}$ . In each case,  $u_i(I_*^2) \geq \hat{u}_i(\mathbf{x}_i, 2) = u_i(I_i^2(\mathbf{x})) \geq u_i^-(I_*^2)$  since  $I_*^2 = \{y^1\} \cup B^2 \cup \{y^2\}$ .
- Suppose  $j = 3$ .
  - Suppose  $x_i^2 \in \mathbb{Z}$  and  $x_i^3 \in \mathbb{Z}$ . So  $x_i^2 = y^2$  and  $x_i^3 = y^3$ . Then  $\hat{u}_i(\mathbf{x}_i, 3) = u_i^-(\{y^2, \dots, y^3\})$ . Thus  $u_i(I_*^3) \geq \hat{u}_i(\mathbf{x}_i, 3) \geq u_i^-(I_*^3)$  since  $I_*^3 = B^3 \cup \{y^3\}$ .






Fig. 7. If  $i$  envies  $j$  even up to one good, Algorithm 1 takes an item out of bundle  $I^j$  and moves it in  $i$ 's direction.

- Otherwise, since  $y^2$  does not appear in  $I_1^3(\mathbf{x})$  (by our wlog assumption), we have that  $I_i^3(\mathbf{x})$  is either  $B^3$  or  $B^3 \cup \{y^3\}$ .  
 Now  $u_i(I_*^3) \geq \hat{u}_i(\mathbf{x}_i, 3) = u_i(I_i^3(\mathbf{x})) \geq u_i^-(I_*^3)$  since  $I_*^3 = B^3 \cup \{y^3\}$ .
- Suppose  $j = 4$ .
  - Suppose  $x_i^3 \in \mathbb{Z}$ . Then  $y^3 = x_i^3$  and  $B^4 = \{x_i^3 + 1, \dots, m\}$ . Thus  $u_i(I_*^4) \geq \hat{u}_i(\mathbf{x}_i, 4) = u_i(B^4) \geq u_i^-(I_*^4)$ , since  $I_*^4 = B^4$ .
  - Suppose  $x_i^3 \notin \mathbb{Z}$ . Then  $\hat{u}_i(\mathbf{x}_i, 4) = u_i(\{x_i^3 + \frac{3}{2}, \dots, m\})$ . Now, either
    - \*  $y^3 = x_i^3 + \frac{1}{2}$  so  $I_*^4 = B^4 = \{x_i^3 + \frac{3}{2}, \dots, m\}$ , or
    - \*  $y^3 = x_i^3 - \frac{1}{2}$  so  $I_*^4 = B^4 = \{x_i^3 + \frac{1}{2}, \dots, m\}$ .
 In either case,  $u_i(I_*^4) \geq \hat{u}_i(\mathbf{x}_i, 4) = u_i(\{x_i^3 + \frac{3}{2}, \dots, m\}) \geq u_i^-(I_*^4)$ .  $\square$

Now again, based on the partition, we can define an allocation  $A_*$  by  $A_*(i) = I_*^{\phi(i)}$  for each agent  $i \in [n]$ . Thus, each agent  $i$  receives the bundle in the complete partition corresponding to  $i$ 's most-preferred index  $\phi(i)$ . We prove that  $A_*$  satisfies EF1: For any pair  $i, j \in [n]$  of agents, we have

$$\begin{aligned}
 u_i(A_*(i)) &= u_i(I_*^{\phi(i)}) \geq \hat{u}_i(\mathbf{x}_i, \phi(i)) && \text{by Proposition 6.1} \\
 &\geq \hat{u}_i(\mathbf{x}_i, \phi(j)) && \text{since } \hat{L}_i(\mathbf{x}_i) = \phi(i) \\
 &\geq u_i^-(I_*^{\phi(j)}) = u_i^-(A_*(j)). && \text{by Proposition 6.1}
 \end{aligned}$$

Hence, we have proved the main result of this section:

**Theorem 6.2.** *On a path, for four agents with monotone valuation functions, a connected EF1 allocation exists.*

For five or more agents, we were not able to construct labeling functions and a rounding scheme which ensure that agents' expectations are correct up to one item. In the four-agent case, each interior bundle is adjacent to an exterior bundle (which helps in the construction), but for five agents, there is a middle bundle whose neighboring bundles are also interior.

### 7. EF1 existence for identical valuations

A special case of the fair division problem is the case of *identical valuations*, where all agents have the same valuation for the goods: for all agents  $i, j \in N$  and every bundle  $I \in \mathcal{C}(V)$ , we have  $u_i(I) = u_j(I)$ . We then write  $u(I)$  for the common valuation of bundle  $I$ . The case of identical valuations often allows for more positive results and an easier analysis. Indeed, we can prove that, for identical valuations and any number of agents, an EF1 allocation connected on a path is guaranteed to exist and can be found in polynomial time.

Now, one might guess that in the restricted case of identical valuations, egalitarian allocations are EF1. However, the leximin-optimal connected allocation may fail EF1: Consider a path with five items and additive valuations 1–3–1–1–1 shared by three agents. The unique leximin allocation is (1, 3, 1–1–1), which induces envy even up to one good. The same allocation also uniquely maximizes Nash welfare, so the Nash optimum also does not guarantee EF1. In contrast, when requiring bundles to satisfy matroid constraints (rather than connectivity constraints), the Nash optimum is EF1 with identical valuations (Biswas and Barman, 2018).

Maximizing an egalitarian objective seemed promising because it ensures that no-one is too badly off, and therefore has not much reason to envy others. The problem is that some bundles might be too desirable. To fix this, we could try to reallocate items so that no bundle is too valuable. This is exactly the strategy of our algorithm: It starts with a leximin allocation, and then moves items from high-value bundles to lower-value bundles, until the result is EF1. In more detail, the algorithm identifies one agent  $i$  who is worst-off in the leximin allocation, and then adjusts the allocation so that  $i$  does not envy any other bundle up to one good. The algorithm does this by going through all bundles in the allocation, outside-in, and if  $i$  envies a bundle  $I^j$  even up to one good, it moves one item from  $I^j$  inwards (in  $i$ 's direction), see Fig. 7. As we will show, a key invariant preserved by the algorithm is that the value of  $I^i$  never increases, and  $i$  remains worst-off. Thus, since  $i$  does not envy others up to one good, the allocation at the end is EF1.

Formally, a *leximin allocation* is an allocation which maximizes the lowest utility of an agent; subject to that it maximizes the second-lowest utility, and so on. In particular, if the highest achievable minimum utility is  $u_L$ , then the leximin allocation is such that every agent has utility at least  $u_L$ , and the number of agents with utility exactly  $u_L$  is minimum.

**Theorem 7.1.** *For identical valuations on a path, Algorithm 1 finds an EF1 allocation.*

**Algorithm 1** Find a connected EF1 allocation of a path  $P$  with identical and monotonic valuations.

**Input:** a path  $P = (v_1, v_2, \dots, v_m)$ , and  $n$  agents with identical and monotonic valuations  $u$   
**Output:** an EF1 connected allocation of  $P$   
 1: Let  $A = (I^1, \dots, I^n)$  be a leximin allocation of  $P$   
 2: Fix an agent  $i$  with minimum utility in  $A$ , i.e.,  $u(I^i) \leq u(I^j)$  for all  $j \in [n]$   
 3: **for**  $j = 1, \dots, i - 1$  **do**  
 4:   **if**  $i$  envies  $I^j$  even up to one good, i.e.,  $u(I^i) < u^-(I^j)$  **then**  
 5:     repeatedly delete the rightmost item of  $I^j$  and add it to  $I^{j+1}$  until  $u(I^i) \geq u^-(I^j)$   
 6: **for**  $j = n, \dots, i + 1$  **do**  
 7:   **if**  $i$  envies  $I^j$  even up to one good, i.e.,  $u(I^i) < u^-(I^j)$  **then**  
 8:     repeatedly delete the leftmost item of  $I^j$  and add it to  $I^{j-1}$  until  $u(I^i) \geq u^-(I^j)$   
 9: **return**  $A$

**Proof.** For an allocation  $A = (I^1, \dots, I^n)$ , write  $u_L(A) := \min_{j \in [n]} u(I^j)$  for the minimum utility obtained in  $A$ , and write  $L(A) := \{j \in [n] : u(I^j) = u_L(A)\}$  for the set of agents (losers) who obtain this utility. For the leximin allocation  $A_{\text{leximin}}$  obtained at the start of the algorithm, write  $u_L^* := u_L(A_{\text{leximin}})$  and  $L^* := L(A_{\text{leximin}})$ . Note that by leximin-optimality, for every allocation  $A$  we must have  $u_L(A) \leq u_L^*$ , and if  $u_L(A) = u_L^*$  then  $|L(A)| \geq |L^*|$ . Let  $i \in L^*$  be the agent fixed at the start of the algorithm.

*Claim 1.* Throughout the algorithm,  $u_L(A) = u_L^*$  and  $L(A) = L^*$ .

The claim is true before we start the for-loops. Suppose the claim holds up until some iteration of the first for-loop, and we now move an item from  $I^j$  to  $I^{j+1}$ , obtaining the new bundles  $I_{\text{new}}^j$  and  $I_{\text{new}}^{j+1}$  in the new allocation  $A_{\text{new}}$ . Then  $u(I_{\text{new}}^j) \geq u^-(I^j) > u(I^i) = u_L^*$ , where the strict inequality holds by the if- and until-clauses. Since no agent other than  $j$  has become worse-off in  $A_{\text{new}}$ , it follows that  $u_L(A_{\text{new}}) \geq u_L(A) = u_L^*$ . As noted, by optimality of  $u_L^*$ , we have  $u_L(A_{\text{new}}) \leq u_L^*$ . Hence  $u_L(A_{\text{new}}) = u_L^*$ . Thus, by optimality of  $L^*$ , we have  $|L(A_{\text{new}})| \geq |L^*|$ . Because agent  $j$  has not become a loser (since  $u(I_{\text{new}}^j) > u_L^*$  as shown before) and no other agent has become a loser, we have  $L(A_{\text{new}}) \subseteq L(A) = L^*$ . Thus  $L(A_{\text{new}}) = L^*$ , as required. The second for-loop is handled similarly.

*Claim 2.* After both for-loops terminate, agent  $i$  does not envy any agent up to one good.

For any  $j \neq i$ , agent  $i$  does not envy  $j$  up to one good immediately after the relevant loop has handled  $j$ , and at no later stage of the algorithm does  $I^j$  change.

It follows that the allocation  $A$  returned by the algorithm is EF1: By Claim 1, we have  $i \in L(A)$ , so that  $u(I^i) \geq u(I^j)$  for all  $j \in [n]$ . By Claim 2, agent  $i$  does not envy any other agent up to one good, so that  $u(I^i) \geq u^-(I^k)$  for all  $k \in [n]$ . Hence, for all  $j, k \in [n]$ , we have  $u(I^j) \geq u^-(I^k)$ , that is, no agent envies another agent up to one good.  $\square$

Algorithm 1 can be implemented to run in polynomial time, because with identical valuations, one can use dynamic programming to find a leximin allocation in time  $O(m^2n^2)$ , and the remainder of Algorithm 1 takes time  $O(mn)$ , as each item is moved at most  $n$  times. A slight speed-up can be achieved by observing that the proof of Theorem 7.1 only needed that the initial allocation optimizes the egalitarian welfare  $u_L$  and minimizes the cardinality of the set  $L$  of losers. Such an allocation can be found by dynamic programming in time  $O(m^2n)$ , and, after some refinements on the implementation of the dynamic programming approach, the running time can be lowered to  $O(mn)$  (see Algorithm 2 in the appendix).

The reallocation stage of our algorithm bears some similarity to Suksompong’s (2019, Thm. 2) proof that a  $u_{\text{max}}$ -equitable allocation exists. Oh et al. (2019, Lem. C.2) proved independently, using an inductive argument, that EF1 allocations on a path exist for identical valuations, and can be found in polynomial time. More recently, Misra et al. (2021) presented another algorithm for this task in the context of aiming for connected allocations satisfying *equitability up to one good* (EQ1).

**8. Concluding remarks**

We have studied the existence of EF1 allocations under connectivity constraints imposed by an undirected graph. We have shown that for two, three, or four agents, an EF1 allocation exists if the graph is traceable and if the agents have monotone valuations. For any number of agents, we also proved that traceable graphs guarantee the existence of an EF2 allocation. The latter two results are proved using Sperner’s lemma, which has been used many times in economics and game theory to show the existence of equilibria and fair allocations (Sarf, 1982; Su, 1999). Unusually, in our application we were able to use Sperner’s lemma in a setting with indivisibilities. We leave as an open question whether EF1 allocations on a path exist for five and more agents.

Our procedures for identical valuations as well as for the cases of two or three agents can be efficiently implemented, so that we can find EF1 allocations in polynomial time. For our results based on Sperner’s lemma, it is not clear how to compute EF1 and EF2 allocations efficiently. On the other hand, as is the case with the computation of other structures whose existence follows from Sperner’s lemma (such as Nash equilibrium or more broadly PPAD problems), our proof based on Sperner’s lemma allows for a “path-following” algorithm through the subdivided simplex. This type of algorithm has been observed to be practically efficient in other contexts (Sarf, 1967), and may also be efficient in our allocation setting on practical instances. Formally, we do not know of a computational complexity result for the problem of finding EF1 or EF2 allocations on a path. For divisible cake-cutting, it is PPAD-complete to find an  $\epsilon$ -approximate envy-free allocation (Deng et al., 2012), implying that it is unlikely that there is an algorithm that runs in time polynomial in  $n$  and  $\log \frac{1}{\epsilon}$ . However, the

PPAD-hardness proof of Deng et al. (2012) uses non-monotone valuations, and thus does not easily extend to our setting, where we assume monotone valuations. For general graphs, Deligkas et al. (2021, Thm. 2) recently showed that deciding whether an EF1 allocation exists on a given instance is NP-hard. Their result applies when the underlying graph is a star, and holds even if the  $n$  agents have binary additive valuations.

Regarding strategic aspects, existing results from the literature imply that there are no EF1 allocation rules which are strategyproof.<sup>8</sup> Amanatidis et al. (2017a) characterized all strategyproof allocation mechanisms when there are  $n = 2$  agents with additive valuations over indivisible items (with no connectivity constraints). They then proved that no mechanism in their class guarantees EF1 (Amanatidis et al., 2017a, Sec. 4.2) for  $m \geq 5$  items. It follows that there is also no strategyproof EF1 mechanism that respects connectivity constraints. We can also obtain such a result by reduction from divisible cake-cutting. Fix some  $\varepsilon > 0$ , and suppose we had a mechanism for allocating a path of  $M$  items among  $n$  agents while being strategyproof and EF1. Then we can use this mechanism as a mechanism for cake-cutting: Given continuous agent valuations over the interval  $[0, 1]$ , approximate these by additive valuations over the path of  $M$  items and run the mechanism on this instance. For sufficiently large  $M$ , the resulting mechanism for cake-cutting will be  $\varepsilon$ -strategyproof (in the sense that a misreport can increase utility by at most  $\varepsilon$ ) and  $\varepsilon$ -envy-free (in the sense that envy is bounded by  $\varepsilon$ ). However, the literature on cake-cutting contains impossibilities about strategyproofness and envy-freeness when requiring connected pieces (Bei et al., 2017, Theorem 1, Bei et al., 2018, Theorem 3), and the proofs also establish impossibility for the  $\varepsilon$ -versions of these properties for small enough  $\varepsilon$ . Hence, for large enough  $M$ , no strategyproof EF1 mechanism for the indivisible setting can exist. By using the result of Bei et al. (2018), we can obtain an impossibility for  $n = 2$  and even for binary additive valuations where each agent approves an interval of items beginning with the left-most item. Finally, Peters (2019, Chapter 12) gives a simple direct proof that there are no strategyproof EF1 mechanisms, even for  $n = 2$  agents and  $m = 5$  items on a line.

We gave a forbidden minor type characterization of all graphs that guarantee the existence of EF1 allocations for two agents. It is natural to also consider the case of more than two agents. However, there are several difficulties in extending the characterization result beyond two agents. First, one cannot generalize the notion of a bipolar ordering in a meaningful way; indeed, a *tripolar ordering*, requiring each initial, middle, and last segment to be connected in a given graph, reduces to the notion of a Hamiltonian path because every consecutive pair of such an ordering must be connected. Second, our two-agent characterization heavily depends on having a simple envy-free protocol: the cut-and-choose procedure. Unfortunately, for three agents, the known protocol becomes much more complex (see Section 4), and for four agents, there is no known explicit protocol that constructs an envy-free division. Nevertheless, Igarashi and Zwicker (2021) recently proposed a forbidden minor type conjecture for the continuous variant of our problem. It would be interesting to explore the discrete version of their conjecture.

In the setting without connectivity constraints, it is possible to achieve efficiency and fairness simultaneously: the maximum Nash welfare solution yields an allocation that is both EF1 and Pareto-optimal (Caragiannis et al., 2019). In our model, this is unfortunately impossible, since on a path there are instances where there is no connected allocation which is EF1 and Pareto-optimal, and it is NP-hard to decide whether such an allocation exists (Igarashi and Peters, 2019).

In this paper, we have only considered *goods*, with monotonic valuations. The setting where some or all items are undesirable (so-called *chores*) is also of interest (Aziz et al., 2019; Bogomolnaia et al., 2016; Meunier and Zerbib, 2019; Segal-Halevi, 2018; Bouveret et al., 2019; Hühne and van Stee, 2021). On a path, a connected allocation satisfying proportionality up to one good (PROP1) always exists (Aziz et al., 2019), but the existence of EF1 or EF2 allocations in this domain is open. For cake-cutting, when agents consider some parts of the cake undesirable, Sperner’s lemma does not directly produce a connected envy-free allocation (Segal-Halevi, 2018), but other methods can prove the existence of such allocations in most cases (Segal-Halevi, 2018; Meunier and Zerbib, 2019).

## 9. Appendix

### 9.1. Maximin share allocations

The *maximin share guarantee* of an agent  $i \in N$  is

$$\text{MMS}_i := \max_{(P^1, P^2, \dots, P^n) \in \Pi_n} \min_{j \in [n]} u_i(P^j),$$

where  $\Pi_n$  denotes the space of all partitions of  $V$  into  $n$  connected bundles. An allocation  $A$  is a *maximin share (MMS) allocation* if  $u_i(A(i)) \geq \text{MMS}_i$  for each agent  $i \in N$ . (Note that the maximum is taken only over connected partitions, so the MMS value could be lower than the standard definition from the model without connectivity constraints.) Bouveret et al. (2017) showed that an MMS allocation exists if the underlying graph  $G$  is a tree. On the other hand, MMS allocations need not exist on a cycle (Bouveret et al., 2017; Lonc and Truszczynski, 2020) or on a complete graph (Kurokawa et al., 2018). The computational complexity of determining the existence of MMS allocations under several graph constraints has also been investigated (Greco and Scarcello, 2020).

<sup>8</sup> In this paragraph, we follow the exposition of Peters (2019, Chapter 12).

Since we have seen that EF1 or EF2 allocations are guaranteed to exist on a path, it is natural to ask whether we can additionally require MMS: on a path, does there always exist an allocation that satisfies EF1 and MMS?

First, let us note that not every EF1 allocation is also MMS. For 3–1–1–1–3 and three agents, the MMS value is 3 via the partition (3, 1–1–1, 3), but the EF1 allocation (3–1, 1, 1–3) gives the middle agent a utility of only 1. In fact, one can show that this example is worst possible, for subadditive valuations. Valuations  $u_i$  are *subadditive* if, for any bundles  $I, I'$ , we have  $u_i(I \cup I') \leq u_i(I) + u_i(I')$ . An allocation satisfies  $\alpha$ -MMS for some  $\alpha > 0$  if  $u_i(A(i)) \geq \alpha \cdot \text{MMS}_i$  for each agent  $i \in N$ . As MMS allocations need not exist in general,  $\alpha$ -MMS allocations have been widely investigated (Amanatidis et al., 2017b, 2018; Kurokawa et al., 2018).

**Proposition 9.1.** *For subadditive valuations, an EF1-allocation on a path guarantees 1/3-MMS.*

**Proof.** Let  $A$  be an EF1 allocation, write  $I^j = A(j)$  for all  $j \in [n]$ , and fix some agent  $i$ . For each  $j \in [n] \setminus \{i\}$ , let  $g_j \in I^j$  be an item such that  $u_i(I^j) \geq u_i(I^j \setminus \{g_j\})$ . We show that  $u_i(I^i) \geq \frac{1}{3} \text{MMS}_i$ .

Let  $P = (P^1, \dots, P^n)$  be a partition of the items into  $n$  bundles such that  $u_i(P^j) \geq \text{MMS}_i$  for each  $j \in [n]$ . Since there are  $n$  bundles in  $P$  but only  $n - 1$  items  $g_j$ , there must be some bundle  $P^k$  such that  $g_j \notin P^k$  for all  $j \in [n] \setminus \{i\}$ ; we show that  $u_i(P^k) \leq 3 \cdot u_i(I^i)$ .

Suppose for a contradiction that there are three distinct agents  $j_1, j_2, j_3 \in [n] \setminus \{i\}$  such that  $P^k \cap I^{j_r} \neq \emptyset$  for  $r = 1, 2, 3$ . Since  $P^k$  and the  $I^{j_r}$ 's are all intervals of a path, the middle interval must be completely contained in  $P^k$ , that is,  $I^{j_r} \subseteq P^k$  for some  $r$ . Hence  $g_{j_r} \in P^k$ , contradicting the choice of  $P^k$ . So  $P^k$  intersects at most two bundles from  $A$  other than  $I^i$ . Thus, for some  $j_1, j_2 \in [n] \setminus \{i\}$ , we have  $P^k \subseteq I^{j_1} \cup I^i \cup I^{j_2} \setminus \{g_{j_1}, g_{j_2}\}$ , and thus by subadditivity,

$$u_i(P^k) \leq u_i(I^{j_1} \setminus \{g_{j_1}\}) + u_i(I^i) + u_i(I^{j_2} \setminus \{g_{j_2}\}) \leq 3 \cdot u_i(I^i).$$

Hence, we have  $u_i(I^i) \geq \frac{1}{3} u_i(P^k) \geq \frac{1}{3} \text{MMS}_i$ , as required.  $\square$

Interestingly, using a similar proof, one can show that the two agents receiving the outer bundles of the path both get at least half of their MMS value. This is also tight; consider 1–1–2–2–1–1 for four agents, and the EF1 allocation (1,1–2,2–1,1).

If we do not restrict valuations to be subadditive, then EF1 does not guarantee  $\alpha$ -MMS for any  $\alpha > 0$ : Consider a path  $P = (v_1, v_2, v_3)$  of three items, and two agents with identical valuations  $u$  defined so that  $u(I) = 1$  if  $I \supseteq \{v_1, v_2\}$  or  $I \supseteq \{v_3\}$ , and  $u(I) = 0$  otherwise. Then the MMS value is 1 via the partition  $(v_1 - v_2, v_3)$ , but the allocation  $(v_1, v_2 - v_3)$  is EF1 and gives the left agent utility 0.

For graphs that are not paths, Proposition 9.1 does not hold. For a complete graph (i.e., in the absence of connectivity constraints), EF1 only implies  $1/n$ -MMS (Caragiannis et al., 2019; Amanatidis et al., 2018).

While we have seen that EF1 on a path does not immediately imply MMS, it does imply MMS in many cases. The following lemma will be useful to show that the allocations produced by our arguments in the main text all satisfy the MMS guarantee.

**Lemma 9.2.** *Suppose there are  $n \geq 2$  agents, and the items are arranged on a path. Take any  $n - 1$  items  $y^1 < \dots < y^{n-1}$ , and define the bundles  $B^1, \dots, B^n$  as follows:*

$$B^1 = L(y^1), B^2 = P(y^1 + 1, y^2 - 1), \dots, B^{n-1} = P(y^{n-2} + 1, y^{n-1} - 1), B^n = R(y^{n-1}).$$

*Then for any agent  $i$ , there is some  $r \in [n]$  such that  $u_i(B^r) \geq \text{MMS}_i$ .*

**Proof.** Let  $P = (P^1, \dots, P^n)$  be a connected partition of the items (ordered left-to-right) so that  $u_i(P^j) \geq \text{MMS}_i$  for all  $j \in [n]$ . Since there are  $n$  bundles in  $P$  but only  $n - 1$  items  $y^1, \dots, y^{n-1}$ , there exists a bundle  $P^k$  in  $P$  that does not contain any  $y^j$ . Writing  $Y = \{y^1, \dots, y^{n-1}\}$ , we see that there is some  $r \in [n]$  such that

$$(P^1 \cup \dots \cup P^{k-1}) \cap Y = \{y^1, \dots, y^{r-1}\} \quad \text{and} \quad (P^{k+1} \cup \dots \cup P^n) \cap Y = \{y^r, \dots, y^{n-1}\}.$$

Thus, we have  $P^k \subseteq P(y^{r-1} + 1, y^r - 1) = B^r$  so that  $u_i(B^r) \geq u_i(P^k) \geq \text{MMS}_i$ .  $\square$

**Theorem 9.3.** *For a path, the EF1 allocations constructed by any of our methods guarantee MMS.*

**Proof.** *Discrete cut-and-choose protocol for two agents.* Suppose Alice's lumpy tie is  $v_j$ . Then, using the definition of lumpy tie, a connected partition witnessing Alice's MMS value is either  $P_1 = (L(v_j), R(v_j) \cup \{v_j\})$  or  $P_2 = (L(v_j) \cup \{v_j\}, R(v_j))$ . At the end of the procedure, Alice receives either  $L(v_j) \cup \{v_j\}$  or  $R(v_j) \cup \{v_j\}$ . For either of these options, there is a bundle in  $P_1$  and a bundle in  $P_2$  which are weakly worse. So Alice receives a bundle that satisfies her MMS value. For Bob, he receives his preferred bundle among  $L(v_j)$  or  $R(v_j)$ . These two bundles are of the shape described in Lemma 9.2 with  $y^1 = v_j$ , so Bob's choice satisfies his MMS value.

*Moving-knife protocol for three agents.* The allocation returned by the algorithm of Theorem 4.3 guarantees MMS. To see this, first suppose the algorithm terminates in Step 3 or Step 4(a). At that step, the bundles  $L$ ,  $M$ , and  $R$  are of the shape described in Lemma 9.2 with  $y^1 = v_{\ell+1}$  and  $y^2 = v_r$ . The agent  $s_{\text{left}}$  who receives  $L$  thinks that  $L$  is best among  $L, M, R$  (since he shouted), so by Lemma 9.2 agent  $s_{\text{left}}$  receives his MMS value. The proof of Theorem 4.3 shows that no other agent envies  $s_{\text{left}}$  (even without removing an item), so that every other shouter receives value at least  $u_i(L)$ , meaning that player  $s$  receives at least her MMS value (since she shouted, she finds  $L$  weakly better than  $M$  and  $R$ ). Finally, agent  $c$  does not envy any other agent (even without removing an item), and so automatically receives at least his MMS value. Next, suppose the algorithm terminates in Step 2. As the proof of Theorem 4.3 shows, the agent  $s_{\text{left}}$  does not envy either of the other players, and hence receives at least his MMS value. Apply Lemma 9.2 with  $y^1 = v_\ell$  and  $y^2 = v_r$ . It follows that for every agent, one of  $L \setminus \{v_\ell\}$ ,  $M$ , or  $R$  provides at least the MMS value. But we know from the proof that any agent prefers either  $M$  or  $R$  to  $L \setminus \{v_\ell\}$ , and we know that all agents  $i \neq s_{\text{left}}$  receive a bundle weakly preferred to both  $M$  and  $R$ , so  $i$  receives his MMS value. Finally, suppose the algorithm terminates in Step 4(b). Similarly to the argument for Step 2, apply Lemma 9.2 with  $y^1 = v_{\ell+1}$  and  $y^2 = v_r$ . Then, each player thinks that one of  $L, M, R$  provides the MMS value. By the proof of Theorem 4.3, agent  $s_{\text{left}}$  receives a bundle weakly preferred to each of  $L, M, R$ , and agents  $i \neq s_{\text{left}}$  prefer either  $M$  or  $R$  to  $L$ , and receive a bundle that is weakly preferred to both  $M$  and  $R$ , so they also receive their MMS value.

*Identical valuations.* Algorithm 1 gives each agent a utility of at least  $u_L^*$ . By their definitions, the MMS-value is the same as the optimal egalitarian welfare under identical valuations.

*EF2 via Sperner’s lemma.* For each agent  $i$ , by Lemma 9.2, there exists a basic bundle whose value is at least  $\text{MMS}_i$ . We showed that the allocation  $A_*$  is such that agent  $i$  weakly prefers the bundle  $i$  receives in  $A_*$  to any basic bundle. Hence,  $A_*$  is an MMS allocation.

*EF1 for four agents via Sperner’s lemma.* For each vertex  $\mathbf{x}_i$  of the full-labeled simplex  $S^*$ , invoke Lemma 9.2 with  $y^1 = \lfloor x_i^1 \rfloor$ ,  $y^2 = \lfloor x_i^2 \rfloor$ ,  $y^3 = \lceil x_i^3 \rceil$ . By case-analysis one can check that  $\hat{u}_i(\mathbf{x}_i, j) \geq u_i(B^j)$  for each  $j = 1, 2, 3, 4$ , where the  $B^j$ ’s are defined like in Lemma 9.2. By Proposition 6.1, we have that  $u_i(A_*(i)) \geq \hat{u}_i(\mathbf{x}_i, \phi(i)) = \max_{j \in [n]} \hat{u}_i(\mathbf{x}_i, j) = \max_{j \in [n]} u_i(B^j) \geq \text{MMS}_i$ .  $\square$

### 9.2. Example of an instance with no EFX allocation

We define EFX as follows.

**Definition 9.4 (EFX).** An allocation  $A$  satisfies EFX (Envy-freeness up to any outer good) if the envy is bounded up to the least valuable outer good, i.e., for any pair  $i, j \in N$  of agents, and for every good  $u \in A(j)$  such that  $A(j) \setminus \{u\}$  is connected, we have  $u_i(A(i)) \geq u_i(A(j) \setminus \{u\})$ .

**Example 9.5.** Consider the instance 2–3–1–3 for three agents. This instance admits no connected EFX allocation: It is clear that no allocation in which some bundle is empty satisfies EFX. In (2, 3, 1–3), the left agent envies the right agent even after removing the outer good of value 1; in (2, 3–1, 3), the left agent envies the middle agent even after removing the outer good of value 1; and in (2–3, 1, 3), the middle agent envies the left agent even after removing the outer good of value 3. One can also consider the instance 1–1–3–3 for two agents.  $\square$

### 9.3. Example of a non-traceable graph that guarantees EF1

**Example 9.6.** Consider a star with three leaves. We will divide the graph among three agents. Consider the allocation where each agent chooses the most favorite leaf-vertex among the unallocated vertices in order, with the last agent in that order being assigned to the central vertex of the star. The resulting allocation satisfies EF1, since the envy towards agents allocated to a single item can be bounded up to one good, and the first and second agent do not envy the third agent if one removes the central vertex from his bundle.  $\square$

### 9.4. Efficient computation of SMMS allocations for identical valuations

In this section, we discuss how to compute the initial allocation that is needed for Algorithm 1 to obtain an EF1 allocation under identical valuations (Theorem 7.1).

Given a path  $P = (v_1, v_2, \dots, v_m)$ ,  $n$  agents with identical monotonic valuations  $u$ , and an allocation  $A = (I^1, \dots, I^n)$ , write  $u_L(A) := \min_{j \in N} u(I^j)$  to denote the egalitarian welfare of  $A$ , i.e., the minimum valuation in  $A$  among all agents, and  $L(A) := \{j \in [n] : u(I^j) = u_L(A)\}$  for the set of agents who obtain this minimum utility. We refer to the agents in  $L(A)$  as the losers.

**Definition 9.7 (SMMS allocation).** An allocation of a path for agents with identical monotonic valuations satisfies strong maximin share (SMMS) if it minimizes the number of losers among all allocations maximizing the egalitarian welfare.

In the following, we provide an efficient dynamic programming algorithm for computing an SMMS allocation. We first present a simple algorithm that runs in time  $O(m^2n)$ , and then we refine it to improve the running time to  $O(mn)$ .

*A suboptimal algorithm.* In order to use a dynamic programming approach, we start by deriving a recurrence relation characterizing the SMMS allocations.

Given two partial allocations  $A, A'$ , we write  $A \succeq A'$  if either  $u_L(A) > u_L(A')$  holds, or both  $u_L(A) = u_L(A')$  and  $|L(A)| \leq |L(A')|$  hold; furthermore, we write  $A \sim A'$  if both  $A \succeq A'$  and  $A \preceq A'$  hold. We observe that an allocation  $A$  of path  $P$  for  $n$  agents is SMMS iff it is “optimal” according to the ordering relation  $\succeq$ , i.e., iff  $A \succeq A'$  for any allocation  $A'$  (of path  $P$  for  $n$  agents).

For any  $h \in [m + 1]$  and  $j \in [m]$ , let  $u[h, j] := u(P(v_h, v_j))$  be the utility assigned by  $u$  to the path segment from  $v_h$  to  $v_j$ . If  $h > j$  then we use the convention that  $P(v_h, v_j) = \emptyset$  and  $u[h, j] = 0$ . Given  $i \in [n]$  and  $j \in [m] \cup \{0\}$ , let  $A[i, j]$  be an SMMS allocation of the subpath  $P(v_1, v_j)$  for  $i$  agents. Also, let  $Egal[i, j]$  and  $L[i, j]$  denote the egalitarian welfare and the number of losers of the SMMS allocation  $A[i, j]$ . Note that our ultimate aim is to find  $A[n, m]$ .

For any  $i \in [n] \setminus \{1\}$ ,  $j \in [m] \cup \{0\}$  and  $h \in [j + 1]$ , let  $A[i, h, j]$  be an allocation of the subpath  $P(v_1, v_j)$  for  $i$  agents that is optimal according to  $\succeq$  after constraining the  $i$ -th bundle to be equal to the subpath  $P(v_h, v_j)$ ; furthermore, let  $Egal[i, h, j]$  and  $L[i, h, j]$  denote the egalitarian welfare and the number of losers of allocation  $A[i, h, j]$ . Observe that we allow  $h$  to reach the value  $j + 1$  in order to model the case in which the  $i$ -th agent gets an empty bundle. By definition, it holds that

$$Egal[i, h, j] = \min\{Egal[i - 1, h - 1], u[h, j]\}, \tag{9.1}$$

$$L[i, h, j] = \begin{cases} L[i - 1, h - 1] & \text{if } Egal[i - 1, h - 1] < u[h, j], \\ L[i - 1, h - 1] + 1 & \text{if } Egal[i - 1, h - 1] = u[h, j], \\ 1 & \text{if } Egal[i - 1, h - 1] > u[h, j], \end{cases} \tag{9.2}$$

where  $Egal[1, h - 1] = u[1, h - 1]$  and  $L[1, h - 1] = 1$ .

One can easily observe that, for any fixed  $i \in [n]$  and  $j \in [m] \cup \{0\}$ , an optimal allocation  $A[i, j]$  of subpath  $P(v_1, v_j)$  for  $i$  agents can be computed according to the following recurrence relation:

$$A[i, j] := \begin{cases} \emptyset & \text{if } j = 0, \\ P(v_1, v_j) & \text{if } i = 1, \\ \text{the best allocation in } \{A[i, h, j] : h \in [j + 1]\} & \text{otherwise,} \end{cases} \tag{9.3}$$

where the quality of each allocation  $A[i, h, j]$  depends on  $Egal[i, h, j]$  and  $L[i, h, j]$  only.

The recurrence relation (9.3) can be used to design a dynamic programming algorithm that computes the SMMS allocation  $A[n, m]$  in time  $O(m^2n)$ . To do this, we will iteratively compute an integer  $k[i, j]$  such that  $A[i, j] \sim A[i, k[i, j], j]$  for all  $i \in [n]$  and for all  $j \in [m]$ .

To do this, in each round  $(i, j)$ , we identify the allocations  $A[i, h, j]$  for each  $h \in [j + 1]$ , and compute the corresponding values  $Egal[i, h, j]$  and  $L[i, h, j]$  using (9.1) and (9.2). Using the computed values  $Egal[i, h, j]$  and  $L[i, h, j]$  we can then use (9.3) to find the index  $k[i, j]$  such that  $A[i, k[i, j], j] = A[i, j]$  is optimal, and we store the resulting values  $Egal[i, j]$  and  $L[i, j]$  (which will be used in the subsequent rounds). Then we proceed to the next round. Finally, at the end of the last round  $(n, m)$ , we can recursively reconstruct the optimal allocation  $A[n, m]$  by using the indices of type  $k[i, j]$  previously stored.

By (9.3), the resulting algorithm returns an SMMS allocation, and its time complexity is  $O(m^2n)$ , given by the number of rounds (which is the number of pairs  $(i, j)$  which is  $O(nm)$ ) multiplied by the complexity of each round (checking each value of  $h$  which is in  $O(m)$ ).

*An improved algorithm.* By exploiting the monotonicity properties of the valuation functions, we can improve the above algorithm and lower its running time. In particular, we will see how to execute each round  $(i, j)$  in constant amortized time, thus lowering the overall time complexity to  $O(mn)$ .

The problem with the existing algorithm is we need to check all possible values of  $k[i, j]$ . Instead we will introduce three quantities  $k_1[i, j]$ ,  $k_2[i, j]$ , and  $k_3[i, j]$ , each of which can be computed quickly, and prove that one of the three values provides a suitable value of  $k[i, j]$ . Specifically, for any  $i \in [n]$  and  $j \in [m]$ , let<sup>9</sup>

$$\begin{aligned} k_1[i, j] &:= \max\{\sup\{h \geq 1 : Egal[i - 1, h - 1] < u[h, j]\}, 1\}, \\ k_2[i, j] &:= \max\{h \geq 1 : Egal[i - 1, h - 1] \leq u[h, j]\}, \\ k_3[i, j] &:= k_2[i, j] + 1. \end{aligned}$$

By the monotonicity of the utility function, we have that, for any fixed value of  $h$ ,  $u[h, j]$  is non-decreasing in  $j$ . This implies that for  $t = 1, 2, 3$  and fixed  $i \in [n]$ , the value  $k_t[i, j]$  is non-decreasing in  $j \in [m]$ . Hence the integers of type  $k_t[i, j]$  can be recursively written as

<sup>9</sup> In the definition of  $k_1[i, j]$ , if the set  $D := \{h \geq 1 : Egal[i - 1, h - 1] < u[h, j]\}$  is empty, taking the maximum between  $\sup D$  and 1 guarantees that  $k_1[i, j]$  is equal to 1 in this extreme case.

$$\begin{aligned}
 k_1[i, j] &= \max\{\sup\{h \geq k_1[i, j - 1] : \text{Egal}[i - 1, h - 1] < u[h, j]\}, k_1[i, j - 1]\}, \\
 k_2[i, j] &= \max\{h \geq k_2[i, j - 1] : \text{Egal}[i - 1, h - 1] \leq u[h, j]\}, \\
 k_3[i, j] &= k_2[i, j] + 1,
 \end{aligned}
 \tag{9.4}$$

where  $k_t[i, 0] := 1$  for  $t = 1, 2, 3$ .

In Lemma 9.9 we will show that  $A[i, j]$  can be set equal to the best allocation among the three allocations of type  $A[i, k_t[i, j], j]$  (with  $t = 1, 2, 3$ ). We first outline some preliminary properties in Lemma 9.8.

**Lemma 9.8.** *For any  $i \in [n]$  and  $j \in [m]$ , we have*

- (i)  $A[i, j] \geq A[i, j - 1]$
- (ii)  $\text{Egal}[i, h, j]$  is non-decreasing in  $h \leq k_1[i, j]$ , it is constant in  $k_1[i, j] < h \leq k_2[i, j]$ , and it is non-increasing in  $h > k_2[i, j]$ .

**Proof.** We first show (i). Given  $i \in [n]$  and  $j \in [m]$ , let  $A'[i, j]$  be the allocation obtained from  $A[i, j - 1]$  by adding item  $j$  to the last bundle of  $A[i, j - 1]$ . By the optimality of  $A[i, j]$ , we have that  $A[i, j] \geq A'[i, j] \geq A[i, j - 1]$ , and this shows (i).

Now, we show (ii). We have the following properties: (a)  $\text{Egal}[i - 1, h - 1]$  is non-decreasing in  $h$  (by (i)), and (b)  $u[h, j]$  is non-increasing in  $h$  (by the monotonicity of the valuation function). Thus, we get the following additional properties, that immediately imply (ii):

- $\text{Egal}[i - 1, h - 1] < u[h, j]$  for any  $h \leq k_1[i, j]$  (by definition of  $k_1[i, j]$  and because of (a) and (b)), and then  $\text{Egal}[i, h, j] = \text{Egal}[i - 1, h - 1]$  is non-decreasing in  $h \leq k_1[i, j]$  (by (a));
- $u[h, j] < \text{Egal}[i - 1, h - 1]$  for any  $h > k_2[i, j]$  (by definition of  $k_2[i, j]$  and because of (a) and (b)), and then  $\text{Egal}[i, h, j] = u[h, j]$  is non-increasing in  $h > k_2[i, j]$  (by (b));
- $\text{Egal}[i, h, j] = \text{Egal}[i - 1, h - 1] = u[h, j]$  in  $k_1[i, j] < h \leq k_2[i, j]$  (by definition of both  $k_1[i, j]$  and  $k_2[i, j]$ ), and then  $\text{Egal}[i, h, j]$  is necessarily constant in  $k_1[i, j] < h \leq k_2[i, j]$  (by (a) and (b)).  $\square$

**Lemma 9.9** (Characterization of SMMS Allocations). *Given  $i \in [n]$  and  $j \in [m]$ , at least one index  $k[i, j] \in \{k_1[i, j], k_2[i, j], k_3[i, j]\}$  guarantees that  $A[i, k[i, j], j]$  is an SMMS allocation of subpath  $P(v_1, v_j)$  for  $i$  agents (i.e.,  $A[i, j] \sim A[i, k[i, j], j]$ ).*

**Proof.** Let  $i \in [n]$ ,  $j \in [m]$ , and let  $k \in [j + 1]$  be an index such that  $A[i, k, j]$  is optimal (i.e., SMMS). We consider three cases, depending on the value of  $k$ .

- (a) Suppose  $k \leq k_1[i, j]$ . By exploiting the monotonicity properties of Lemma 9.8 and the definition of  $k_1[i, j]$ , we will show that  $A[i, k_1[i, j], j] \geq A[i, k, j]$ . By Lemma 9.8, the set  $H_1$  of integers  $h \leq k_1[i, j]$  such that  $\text{Egal}[i, h, j] = \text{Egal}[i, j]$  is an integer interval having  $k_1[i, j]$  as its maximum, thus both  $k$  and  $k_1[i, j]$  belong to  $H_1$ . Furthermore, for any  $h \in H_1$ , the egalitarian welfare of each allocation  $A[i, h, j]$  is equal to that of allocation  $A[i - 1, h - 1]$ , and the set of losers in  $A[i, h, j]$  is the same as in  $A[i - 1, h - 1]$  (indeed, agent  $i$  is not a loser since  $\text{Egal}[i - 1, h - 1] < u[h, j]$ ). Thus, since  $A[i - 1, k_1[i, j] - 1] \geq A[i - 1, h - 1]$  (by Lemma 9.8(i)), we necessarily have that  $A[i, k_1[i, j], j] \geq A[i, h, j]$  for any  $h \in H_1$ , and this shows the optimality of allocation  $A[i, k_1[i, j], j]$ .
- (b) Suppose  $k_1[i, j] < k \leq k_2[i, j]$ . Then we can analogously show that  $A[i, k_2[i, j], j] \geq A[i, k, j]$ . The set  $H_2$  of values  $h$  with  $k_1[i, j] < h \leq k_2[i, j]$  and  $\text{Egal}[i, h, j] = \text{Egal}[i, j]$  is an integer interval having  $k_2[i, j]$  as its maximum, thus both  $k$  and  $k_2[i, j]$  belong to  $H_2$ . For any  $h \in H_2$ , the valuation  $u[h, j]$  for the  $i$ -th bundle of  $A[i, h, j]$  is equal to the egalitarian welfare  $\text{Egal}[i, h - 1]$  of the allocation restricted to the first  $i - 1$  bundles (by definition of  $k_1[i, j]$  and  $k_2[i, j]$ , and by Lemma 9.8(ii)). Thus, the number of losers  $L[i, h, j]$  in allocation  $A[i, h, j]$  is equal to  $L[i - 1, h - 1]$  (i.e., the quantity of losers among the first  $i - 1$  agents) plus 1 (i.e., agent  $i$ ). As  $A[i - 1, k_2[i, j] - 1] \geq A[i - 1, h - 1]$  (by Lemma 9.8(i)) and  $\text{Egal}[i - 1, k_2[i, j] - 1] = \text{Egal}[i - 1, h - 1]$ , we necessarily have that  $L[i, k_2[i, j], j] \leq L[i, h, j]$  for any  $h \in H_2$ . We conclude that the best allocation of type  $A[i, h, j]$  with  $h \in H_2$  is achieved by  $h := k_2[i, j]$ , and this shows the optimality of  $A[i, k_2[i, j], j]$ .
- (c) Suppose  $k_3[i, j] \leq k$ . Then we can also show that  $A[i, k_3[i, j], j] \geq A[i, k, j]$ . The set  $H_3$  of values  $h \geq k_3[i, j]$  with  $\text{Egal}[i, h, j] = \text{Egal}[i, j]$  is an interval whose minimum is  $k_3[i, j]$ , thus both  $k$  and  $k_3[i, j]$  belong to  $H_3$ . As  $u[k_3[i, j], j] < \text{Egal}[i - 1, k_3[i, j] - 1]$  (by definition of  $k_3[i, j]$ ), we have that agent  $i$  is a loser in  $A[i, k_3[i, j], j]$ , and it is the only one. Thus, allocation  $A[i, k_3[i, j], j]$  must be necessarily optimal.  $\square$

By exploiting (9.3) and the characterization of optimal allocations given in Lemma 9.9, we can derive an efficient dynamic programming algorithm (Algorithm 2) that computes an SMMS allocation.

**Theorem 9.10.** *Algorithm 2 computes an SMMS allocation in  $O(mn)$  time.*

**Proof.** We first show that the output of Algorithm 2 is an SMMS allocation. In lines 1–5 of the algorithm we initialize, for  $i = 1$  or  $j = 0$ , the maximum egalitarian welfare  $\text{Egal}[i, j]$  and the number of losers  $L[i, j]$  of the optimal allocation

**Algorithm 2** Find an SMMS allocation of a path  $P$  with identical and monotonic valuations.

**Input:** a path  $P(v_1, v_m)$ ,  $n$  agents with identical and monotonic valuations  $u$   $\{u[h, j]$  denotes valuation  $u(P(v_h, v_j))$ , with the convention that  $P(v_h, v_j) = \emptyset$  and  $u[h, j] = 0$  if  $h > j$ .\}

**Output:** an SMMS allocation of  $P$

```

1: for  $j = 0, \dots, m$  do
2:    $Egal[1, j] \leftarrow u[1, j]$ ,  $L[1, j] \leftarrow 1$ 
3: for  $i = 1, \dots, n$  do
4:    $Egal[i, 0] \leftarrow 0$ ,  $L[i, 0] \leftarrow 0$ 
5:    $k_t[i, 0] \leftarrow 1$  for  $t = 1, 2, 3$ 
   ( $Egal[i, j]$  and  $L[i, j]$  will be computed in such a way to be respectively the egalitarian welfare and the number of losers of an optimal allocation of subpath  $P(v_1, v_j)$  for  $i$  agents; each  $k_t[i, j]$  will be defined as in (9.4), and will be used to compute an index  $k[i, j]$  such that  $A[i, k[i, j], j]$  is an optimal allocation of subpath  $P(v_1, v_j)$  for  $i$  agents.)
6: for  $i = 2, \dots, n$  do
7:   for  $j = 1, \dots, m$  do
8:      $k_1[i, j] \leftarrow \max\{\sup\{h \geq k_1[i, j-1] : Egal[i-1, h-1] < u[h, j]\}, k_1[i, j-1]\}$ 
9:      $k_2[i, j] \leftarrow \max\{h \geq k_2[i, j-1] : Egal[i-1, h-1] \leq u[h, j]\}$  {By the monotonicity of the utility function, to compute  $k_1[i, j]$ , with  $t = 1$  (resp.  $t = 2$ ), it is sufficient to analyze in increasing order all the integers  $h > k_t[i, j-1]$  until a value  $h$  with  $Egal[i-1, h-1] \geq u[h, j]$  (resp.  $Egal[i-1, h-1] > u[h, j]$ ) is reached; then,  $k_t[i, j]$  can be set equal to  $h-1$ . We conclude that  $k_t[i, j]$  can be found in  $k_t[i, j] - k_t[i, j-1]$  iterations (for any fixed  $i \in [n]$ ,  $j \in [m]$ , and  $t = 1, 2$ .)}
10:     $k_3[i, j] \leftarrow k_2[i, j] + 1$ 
11:     $M_t \leftarrow \min\{Egal[i-1, k_t[i, j]-1], u[k_t[i, j], j]\}$  for  $t = 1, 2, 3$ 
12:     $L_t \leftarrow \begin{cases} L[i-1, h-1] & \text{if } Egal[i-1, h-1] < u[h, j], \\ L[i-1, h-1] + 1 & \text{if } Egal[i-1, h-1] = u[h, j], \\ 1 & \text{if } Egal[i-1, h-1] > u[h, j]. \end{cases}$ 
    (We observe that  $M_t$  and  $L_t$  are respectively the egalitarian welfare and the number of losers of allocation  $A[i, k_t[i, j], j]$ .)
13:     $t^* \leftarrow$  the index  $t$  maximizing  $M_t$ , and then, minimizing  $L_t$ 
14:     $Egal[i, j] \leftarrow M_{t^*}$ ,  $L[i, j] \leftarrow L_{t^*}$ ,  $k[i, j] \leftarrow k_{t^*}[i, j]$  ( $t^*$  is computed in such a way that  $A[i, k_{t^*}[i, j], j]$  is the optimal allocation among those of type  $A[i, k_t[i, j], j]$ , and by Lemma 9.9,  $A[i, k[i, j], j]$  is an optimal allocation for subpath  $P(v_1, v_j)$  and  $i$  agents.)
15:  $I^n \leftarrow P(v_{k[n, m]}, v_m)$ 
16:  $j \leftarrow k[n, m] - 1$ 
17: for  $i = n-1, \dots, 2, 1$  do
18:    $I^i \leftarrow P(v_{k[i, j]}, v_j)$ 
19:    $j \leftarrow k[i, j] - 1$ 
20: return  $A = (I^1, \dots, I^n)$ 

```

$A[i, j]$  (of path  $P(v_1, v_j)$  for  $i$  agents). By using the characterization provided in Lemma 9.9, in lines 6–14 we iteratively compute an index  $k[i, j]$  such that  $A[i, k[i, j], j]$  is an SMMS allocation of subpath  $P(v_1, v_j)$  for  $i$  agents, and we compute the corresponding maximum egalitarian welfare  $Egal[i, j]$  and number of losers  $L[i, j]$ . Finally, in lines 15–20 we recursively reconstruct the optimal allocation  $A[n, m]$  that is returned as output.

Now, we show that the time complexity of Algorithm 2 is  $O(mn)$ . Observe that the body of the nested for-loops in lines 6–14 can be performed in time  $T(i, j) = c \sum_{t=1}^2 (k_t[i, j] - k_t[i, j-1])$ , where  $c$  is a constant that does not depend on  $i$  and  $j$ . Indeed, this running time depends on the computation in lines 8–9 of each index  $k_t[i, j]$  for  $t = 1, 2$ , and to compute it we can simply analyze all the indices from  $k_t[i, j-1] + 1$  to  $k_t[i, j] + 1$  only. We conclude that the time complexity  $T$  of the nested for-loops in lines 6–14 satisfies

$$\begin{aligned}
 T &= \sum_{i=1}^n \sum_{j=1}^m T(i, j) = \sum_{i=1}^n \sum_{j=1}^m c \sum_{t=1}^2 (k_t[i, j] - k_t[i, j-1]) = \sum_{t=1}^2 c \sum_{i=1}^n \sum_{j=1}^m (k_t[i, j] - k_t[i, j-1]) \\
 &= \sum_{t=1}^2 c \sum_{i=1}^n (k_t[i, m] - k_t[i, 0]) \leq 2c \sum_{i=1}^n m = 2cnm \in O(mn).
 \end{aligned}$$

Since the time complexity of the other parts of the algorithm is clearly  $O(mn)$ , it follows that Algorithm 2 terminates in  $O(mn)$  time.  $\square$

**References**

Amanatidis, G., Birmpas, G., Christodoulou, G., Markakis, E., 2017a. Truthful allocation mechanisms without payments: characterization and implications on fairness. In: Proceedings of the 2017 ACM Conference on Economics and Computation (EC), pp. 545–562.

Amanatidis, G., Markakis, E., Nikzad, A., Saberi, A., 2017b. Approximation algorithms for computing maximin share allocations. ACM Trans. Algorithms 13 (4), 52:1–52:28.

Amanatidis, G., Birmpas, G., Markakis, E., 2018. Comparing approximate relaxations of envy-freeness. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), pp. 42–48.

Aziz, H., Caragiannis, I., Igarashi, A., Walsh, T., 2019. Fair allocation of indivisible goods and chores. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 53–59.

Bapat, R.B., 1989. A constructive proof of a permutation-based generalization of Sperner’s lemma. Math. Program. 44 (1), 113–120.

Barbanel, J.B., Brams, S.J., 2004. Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. Math. Soc. Sci. 48 (3), 251–269.



- Bei, X., Chen, N., Huzhang, G., Tao, B., Wu, J., 2017. Cake cutting: envy and truth. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI). AAAI Press, pp. 3625–3631.
- Bei, X., Huzhang, G., Suksompong, W., 2018. Truthful fair division without free disposal. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI). AAAI Press, pp. 63–69.
- Bilò, V., Caragiannis, I., Flammini, M., Igarashi, A., Monaco, G., Peters, D., Vinci, C., Zwicker, W.S., 2019. Almost envy-free allocations with connected bundles. In: Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS), pp. 14:1–14:21.
- Biswas, A., Barman, S., 2018. Fair division under cardinality constraints. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), pp. 91–97.
- Bogomolnaia, A., Moulin, H., Sandmirsky, F., Yanovskaya, E., 2016. Dividing goods and bads under additive utilities. arXiv:1610.03745.
- Bondy, J.A., Murty, U.S.R., 2008. Graph Theory, 1st edition. Springer.
- Bouveret, S., Cechlárová, K., Elkind, E., Igarashi, A., Peters, D., 2017. Fair division of a graph. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), pp. 135–141.
- Bouveret, S., Cechlárová, K., Lesca, J., 2019. Chore division on a graph. *Auton. Agents Multi-Agent Syst.* 33 (5), 540–563.
- Brams, S.J., Taylor, A.D., 1996. Fair Division: From Cake-Cutting to Dispute Resolution. Cambridge University Press.
- Brams, S.J., Taylor, A.D., Zwicker, W.S., 1997. A moving-knife solution to the four-person envy-free cake-division problem. *Proc. Am. Math. Soc.* 125 (2), 547–554.
- Budish, E., 2011. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* 119 (6), 1061–1103.
- Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J., 2019. The unreasonable fairness of maximum Nash welfare. *ACM Trans. Econ. Comput.* 7 (3), 12:1–12:32.
- Deligkas, A., Eiben, E., Ganian, R., Hamm, T., Ordyniak, S., 2021. The parameterized complexity of connected fair division. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), pp. 139–145.
- Deng, X., Qi, Q., Saberi, A., 2012. Algorithmic solutions for envy-free cake cutting. *Oper. Res.* 60 (6), 1461–1476.
- Dubins, L.E., Spanier, E.H., 1961. How to cut a cake fairly. *Am. Math. Mon.* 68 (1), 1–17.
- Even, S., Tarjan, R.E., 1976. Computing an st-numbering. *Theor. Comput. Sci.* 2 (3), 339–344.
- Flegg, H., Graham, 1974. From Geometry to Topology. Crane, Russak & Co.
- Greco, G., Scarcello, F., 2020. The complexity of computing maximin share allocations on graphs. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI), pp. 2006–2013.
- Höhne, F., van Stee, R., 2021. Allocating contiguous blocks of indivisible chores fairly. *Inf. Comput.*, 104739.
- Hopcroft, J., Tarjan, R., 1973. Algorithm 447: efficient algorithms for graph manipulation. *Commun. ACM* 16 (6), 372–378.
- Igarashi, A., Peters, D., 2019. Pareto-optimal allocation of indivisible goods with connectivity constraints. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI), pp. 2045–2052.
- Igarashi, A., Zwicker, W.S., 2021. Fair division of graphs and of tangled cakes. arXiv:2102.08560.
- Kuhn, H.W., 1960. Some combinatorial lemmas in topology. *IBM J. Res. Dev.* 4 (5), 518–524.
- Kurokawa, D., Procaccia, A.D., Wang, J., 2018. Fair enough: guaranteeing approximate maximin shares. *J. ACM* 65 (2), 8.
- Lempel, A., Even, S., Cederbaum, I., 1967. An algorithm for planarity testing of graphs. *Theory Graphs: Int. Symp.*, 215–232.
- Lipton, R.J., Markakis, E., Mossel, E., Saberi, A., 2004. On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce (EC), pp. 125–131.
- Lonc, Z., Truszczynski, M., 2020. Maximin share allocations on cycles. *J. Artif. Intell. Res.* 69, 613–655.
- Meunier, F., Zerbib, S., 2019. Envy-free cake division without assuming the players prefer nonempty pieces. *Isr. J. Math.* 234 (2), 907–925.
- Misra, N., Sonar, C., Vaidyanathan, P.R., Vaish, R., 2021. Equitable division of a path. arXiv:2101.09794.
- Oh, H., Procaccia, A.D., Suksompong, W., 2019. Fairly allocating many goods with few queries. In: Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI), pp. 2141–2148.
- Peters, D., 2019. Fair Division of the Commons. PhD thesis, DPhil thesis. University of Oxford.
- Procaccia, A.D., 2016. Cake cutting algorithms. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (Eds.), *Handbook of Computational Social Choice*. Cambridge University Press, chapter 13.
- Robertson, J., Webb, W., 1998. *Cake-Cutting Algorithms: Be Fair If You Can*. AK Peters/CRC Press.
- Scarf, H., 1967. The approximation of fixed points of a continuous mapping. *SIAM J. Appl. Math.* 15 (5), 1328–1343.
- Scarf, H.E., 1982. The computation of equilibrium prices: an exposition. In: *Handbook of Mathematical Economics*, vol. 2, pp. 1007–1061.
- Segal-Halevi, E., 2018. Fairly dividing a cake after some parts were burnt in the oven. In: Proceedings of the 17th Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 1276–1284.
- Stromquist, W., 1980. How to cut a cake fairly. *Am. Math. Mon.* 87 (8), 640–644.
- Su, F.E., 1999. Rental harmony: Sperner's lemma in fair division. *Am. Math. Mon.* 106 (10), 930–942.
- Suksompong, W., 2019. Fairly allocating contiguous blocks of indivisible items. *Discrete Appl. Math.* 260, 227–236.
- Tarjan, R.E., 1986. Two streamlined depth-first search algorithms. *Fundam. Inform.* 9, 85–94.