Mathematical Programming
https://doi.org/10.1007/s10107-023-01945-5

FULL LENGTH PAPER

Series B

Fair division of graphs and of tangled cakes

Ayumi Igarashi[1] · **William S. Zwicker**[2,3,4]

Received: 31 January 2021 / Accepted: 23 February 2023
© The Author(s) 2023

Abstract

Fair division has been studied in both continuous and discrete contexts. One strand of the continuous literature seeks to award each agent with a single connected piece—a subinterval. The analogue for the discrete case corresponds to the *fair division of a graph*, where allocations must be *contiguous* so that each bundle of vertices is required to induce a connected subgraph. With *envy-freeness up to one item* (*EF*1) as the fairness criterion, however, positive results for three or more agents have mostly been limited to traceable graphs. We introduce tangles as a new context for fair division. A tangle is a more complicated cake—a connected topological space constructed by gluing together several copies of the unit interval [0, 1]—and each single tangle $\mathcal{T}$ corresponds in a natural way to an infinite topological class $\mathcal{G}(\mathcal{T})$ of graphs, linking envy-free fair division of tangles to EFk fair division of graphs. In addition to the unit interval itself, we show that only five other *stringable* tangles guarantee the existence of envy-free and connected allocations for arbitrarily many agents, with the corresponding topological classes containing only traceable graphs. Any other tangle $\mathcal{T}$ has a bound r on the number of agents for which such allocations necessarily exist, and our *Negative Transfer Principle* then applies to the graphs in $\mathcal{T}$'s class; for any integer $k \geq 1$, almost all graphs in this class are non-traceable and fail to guarantee EFk contiguous allocations for $r + 1$ or more agents, even when very strict requirements are placed on the valuation functions for the agents. With bounds on the number of agents, however, we obtain positive results for some non-stringable classes. An elaboration

Ayumi Igarashi is supported by JST PRESTO under grant number JPMJPR20C1. William Zwicker is supported by the CY Initiative of Excellence (Grant "Investissements d'Avenir" ANR-16-IDEX-0008).

✉ Ayumi Igarashi
ayumi_igarashi@nii.ac.jp

William S. Zwicker
zwickerw@union.edu

1 National Institute of Informatics, Tokyo, Japan

2 Union College, Schenectady, USA

3 Murat Sertel Center for Advanced Economic Studies, Istanbul Bilgi University, Istanbul, Turkey

4 CY Cergy Paris Université, CY Advanced Studies and THEMA, Cergy-Pontoise, France

Published online: 15 April 2023

of Stromquist's moving knife procedure shows that the non-stringable lips tangle $\mathcal{L}$ guarantees envy-free allocations of connected shares for three agents. We then modify the discrete version of Stromquist's procedure in Bilò et al. (Games Econ Behav 131:197–221, 2022) to show that all graphs in the topological class $\mathcal{G}(\mathcal{L})$ (most of which are non-traceable) guarantee EF1 allocations for three agents.

Keywords Fair division · Cake cutting · Graphs · Tangles · Envy-freeness

Mathematics Subject Classification 91-08 · 91B18 · 91B32 · 55U10 · 68R10

1 Introduction

In the *continuous* setting for fair division, a single freely divisible good or "cake," often modeled by the [0, 1] interval, is to be partitioned into n pieces, with each of n agents i being allocated a different piece T_i of the partition. One concern in this setting regards the existence of allocations for the agents that are both *envy-free* ($v_i(T_i) \geq v_i(T_j)$ for all i, j, so that each agent values their assigned piece at least as highly as she values the pieces awarded to others) and *connected* (meaning each piece T_i is a single subinterval). For positive results here it suffices to assume only that our valuations v_i are both *monotone* and *continuous*.

In the following classical cake-cutting result (see Stromquist [18], Su [19], and Woodall [22]) the word "guarantees" conveys that such envy-free connected allocations exist for every possible assignment $\{v_i\}_{1 \leq i \leq n}$ of monotone continuous valuations to the agents:

Theorem 1 (Stromquist [18], Su [19], and Woodall [22]) *The closed interval* [0, 1] *guarantees the existence of connected and envy-free allocations for every finite number of agents n with monotone continuous valuations.*

The non-constructive proof uses Sperner's lemma, followed by a limit argument. We present the formal definitions of monotonicity and continuity that apply to valuation functions defined for sub-intervals of [0, 1] in Sect. 2.

In the *indivisible items* setting, we begin with a finite set O of items, with each individual item to be awarded in its entirety to some one agent. An allocation partitions O into n subsets A_i, with each agent i allocated a different piece A_i of the partition. This setting is fundamentally "lumpy" in a way that rules out envy-freeness as a reasonable goal for an allocation—think of two agents sharing one desirable good. However, a notion due to Budish [9] nicely walks around the lumpiness obstacle. An allocation $A = \{A_i\}_{1 \leq i \leq n}$ is *envy-free up to k goods*, aka EFk, if for each pair $i \neq j$ of agents there exists a selection of k or fewer items in agent j's share A_j such that any envy that agent i may feel for A_j would be erased by pretending those items had been removed from A_j. A variety of methods have subsequently been used to prove

that an EF1 allocation is guaranteed to exist—that is, such an allocation exists for any set O and any sequence $\{v_i\}_{1\leq i\leq n}$ of n monotone valuations for the n agents [17].[1]

Graph fair division, our third context, was introduced recently, in Bouveret et al. [6]. Here, the vertices of a connected simple graph[2] $G = (V, E)$ are viewed as indivisible items, and we insist that each agent be allocated a *contiguous* share—meaning that the subgraph induced by the subset $A_i \subseteq V$ of vertices allocated to them must be connected.[3]

Again, we assume only monotonicity for the valuations v_i that agents use to determine the worth of subsets S of vertices. More than one fairness notion has subsequently been applied to graphs, but our focus here is on EF1 and EFk, and on the following strengthening:[4]

Definition 2 An allocation A of graph vertices satisfies EFk_{outer} if for any pair i, j of agents, either $v_i(A_i) \geq v_i(A_j)$, or there is a set X of k or fewer goods in A_j such that $A_j \setminus X$ is contiguous and $v_i(A_i) \geq v_i(A_j \setminus X)$.

This was the notion used in Bilò et al. [4], where the following result can be found, though not in the exact form stated here:

Theorem 3 (Bilò et al. [4]) *The class of path graphs guarantees the existence of contiguous and $EF1_{outer}$ allocations for two, three, or four agents with monotone valuations, and guarantees the existence of contiguous and $EF2_{outer}$ allocations for five or more agents with monotone valuations.*

Igarashi [16] very recently has improved these results, guaranteeing $EF1_{outer}$ allocations of a path graph for arbitrarily many agents. Of course, adding more edges to a path graph cannot hurt the contiguity of any agent's share, so these results apply as well to the class of all *traceable* graphs, which are graphs that contain a *Hamiltonian* path, i.e., a path that visits every vertex exactly once. How about non-traceable graphs?

[♣] *"For the case of three or more agents, it is a challenging open problem to find an infinite class of non-traceable graphs that guarantee EF1"* [from [4], introduction].

Our investigations can be seen as an immediate follow-up to [4]—one aimed directly at addressing this open problem. Here, we provide two somewhat different answers (which are restated more precisely in Sect. 9):

[♣, Answer #1] There exist infinite *topological classes* of non-traceable graphs that guarantee $EF1_{outer}$ for three agents.

[1] In the case of n additive valuation functions for the n agents, one can obtain an EF1 allocation through maximizing the *Nash product*; if $A = \{A_i\}_{1\leq i\leq n}$ is an allocation of the items in the finite set O that maximizes the product $\Pi_{1\leq i\leq n} v_i(A_i)$ of agents' valuations of their shares, then A is EF1 [10].

[2] A *graph* may be a *multi-graph*, with more than one edge linking a given pair of vertices and with loops (edges from a vertex to itself). A *simple graph* excludes both of these edge types.

[3] We will use *connected* for the continuous setting, and *contiguous* for the discrete context of graphs.

[4] The positive results in this paper establish $EF1_{outer}$ while our negative results defeat the weaker EF1 (or EFk). One rationale for the *outer* version is that we are in a context that requires contiguous shares, so a non-contiguous set of vertices may have value 0, or may conceivably have no defined value.

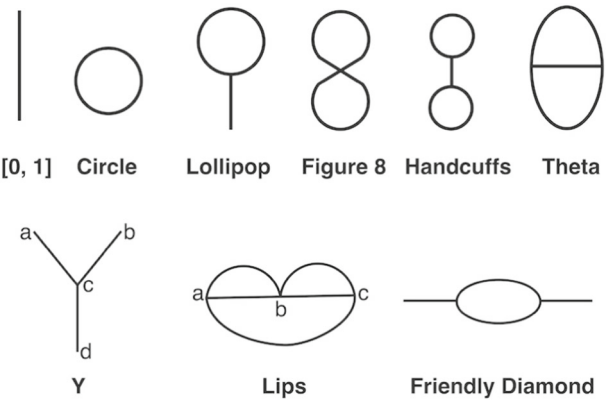

Fig. 1 Nine tangles. Top row contains the six basic stringable tangles

[♣, Answer #2] There exists no infinite *topological class* of non-traceable graphs that guarantees EFk for arbitrarily many agents (and this holds for each fixed integer $k \geq 1$).

Here, the *topological class* of a graph G consists of all graphs that can be obtained from G by adding zero or more vertices of degree 2 along G's edges; note that the class of path graphs is itself a topological class. But why focus on classes rather than on individual graphs? Dominik Peters (private communication) had identified several small, non-traceable graphs (not discussed in [4]) that seemed to also guarantee EF1 for three agents, but these appeared to be sporadic, fitting no discernible pattern.[5]

Our methods here rely on introducing a fourth setting for fair division, involving continuous spaces built from several copies of the [0, 1] interval. This fourth context—which we call *tangles*—seems to be a new one for fair division, although a very similar model was proposed at about the same time in Bei and Suksompong [3]. Examples include a division of a road network, an irrigation ditch, and so on, where a resource to be divided can be identified with a tangle.

Our original motivation for studying tangles stemmed from what they might tell us about fair division of graphs, which proves to be quite a bit, at least when the notion of fairness under consideration is some version of envy-freeness. In much the same way that the [0, 1] interval corresponds to the class of all *path graphs* (obtained by adding vertices at 0, at 1, and at finitely many points in the interval's interior) each individual tangle $\mathcal{T}$ corresponds to an infinite topological class $\mathcal{G}(\mathcal{T})$ of graphs, almost all of which are *simple graphs* (see footnote 2). The *Negative Transfer Principle* proved here then converts any negative result for the envy-free fair division of $\mathcal{T}$ into a corresponding result for EFk fair division of almost all graphs in $\mathcal{G}(\mathcal{T})$. But the fair division of tangles turns out to be pretty interesting on its own, and so this paper can equally well be classified as an introduction to a new context for fair division.

Most of our results flow from tangles to graphs, and this determines the organization of the rest of the paper. In Sect. 2, we introduce fair division of tangles. In Sect. 3, we

[5] Any graph at all guarantees EF1 for n agents, when n is at least as great as the number of vertices, but sporadic examples may exist for other reasons as well.

use *cutsets* to prove the Tangle Boundedness Theorem (Theorem 11): while each of the six *basic stringable* tangles satisfies a version of Theorem 1 (because each is essentially similar to the unit interval), every other tangle comes with an upper bound on the number of agents for which connected and envy-free allocations are guaranteed. A part of that proof is deferred to Sect. 4, where a case argument shows that every realizable degree sequence for a tangle leads either to one of the six stringable tangles, or to a key inequality used in the proof of Theorem 11.

For the case of two agents, we provide a complete characterization: envy-free and connected allocations are guaranteed if and only if the tangle has a *bipolar ordering*. It follows, for example, that the friendly diamond and lips tangles both guarantee connected, envy-free allocations for two agents. The result and its proof are quite parallel to the characterization via *bipolar numbering* for the class of graphs guaranteeing EF1 among two agents in Bilò et al. [4], and thus represent a transfer from the discrete to the continuous realm. Details are in Sect. 5.

Our more interesting positive results for tangles are covered in Sect. 6, where we show that the lips tangle $\mathcal{L}$ (see Fig. 1) guarantees connected and envy-free allocations for three agents, via an argument that elaborates Stromquist's well-known moving knife method [18] over $[0, 1]$; the additional complexity for $\mathcal{L}$ arises from the need to take explicit account of the ways in which $\mathcal{L}$'s arcs are connected. A discretized version of Stromquist's argument from [4] applies to path graphs, where it guarantees contiguous and EF1 allocations of path graphs for three agents. In Sect. 7, we combine elements of that argument with the version of Stromquist for the lips tangle $\mathcal{L}$ (from Sect. 6) to show that every graph in the topological class $\mathcal{G}(\mathcal{L})$ of lips (in which almost all graphs are non-traceable) similarly guarantees EF1 allocations for three agents.

In Sect. 8, we define a natural directed partial ordering over the infinite *topological class* $\mathcal{G}(\mathcal{T})$ of graphs associated with a tangle. This ordering generates a *final segment filter* over $\mathcal{G}(\mathcal{T})$, accompanied by a notion of "almost all." Almost all members of $\mathcal{G}(\mathcal{T})$ are then simple graphs. Moreover, if $\mathcal{T}$ is not one of the six basic stringable tangles, then

- almost all these graphs are non-traceable, and
- there exists an upper bound n' such that for any positive integer k and any $n > n'$ almost all graphs in $\mathcal{G}(\mathcal{T})$ fail to guarantee contiguous EFk allocations for n agents.

We refer to this last result as the *Negative Transfer Principle*. Finally, in Sect. 9 we tie the results from earlier sections more precisely to our two answers (stated earlier) to question ♣.

1.1 Related work on graph fair division

The problem of fairly dividing a graph has been recently proposed in Bouveret et al. [6]. Under this model, the vertices of a graph correspond to the indivisible items and the edges correspond to the adjacency relation among the items. Agents are interested in receiving a contiguous bundle of items. There are several natural applications that can be represented by this model. For example, the problem of dividing cities connected by a road network among several parties, as when an island is partitioned and each side wishes to drive among its allocated cities while staying within its own territory.

Alternately, consider offices allocated to several departments of an organization, where the edges represent corridors in the building, and the goal is for each department to receive contiguous offices. Buildings typically have multiple offices along each corridor, modeled as multiple degree 2 vertices inserted along a graph edge. We would like to emphasize that in the graph fair division problem, the vertices are allocated among agents and the edges have value 0, while in our problem of dividing a tangle, the tangle edges are allocated among agents and singular points, corresponding to vertices of a graph, have value 0.

When the items are indivisible, classical fairness notions of the cake-cutting problem, such as envy-freeness and proportionality, cannot be guaranteed in the graph division setting: consider two agents dividing one desirable item. A line of work has therefore explored two notions of approximate fairness. One focuses on an approximate notion of envy-freeness, namely, EF1 [2, 4, 11, 15]. Another focuses on an approximate notion of proportionality, called the *maximin fair share criterion* (MMS), wherein the maximin fair share is defined for the set of all connected divisions [2, 6, 14, 20]. Our work focuses on EF and establishes a link between an EF1 allocation of a graph and an EF allocation of a tangle.

Bei et al. [2] showed that the complete bipartite graph $K_{a,b}$ guarantees EF1 for n agents whenever $a, b \geq n$. As $K_{a,b}$ is non-traceable whenever $|a - b| \geq 2$, classes of bipartite graphs provide an alternative to [♣, Answer #1]. It is open whether or not these results extend to their topological classes – we do not know whether adding degree 2 subdivision vertices along edges of $K_{a,b}$ always preserves its EF1 guarantee – but the question in [4] was not restricted to topological classes. Note that when $a - b \geq 2$, $K_{a,b}$ fails to guarantee EF1 for $b + 1$ agents; this follows because some agent gets none of the vertices on the b side, limiting that agent's share to a single vertex on the a side. Some other agent must get at least two vertices on the a side along with one on the b side to connect those two, for a total of three or more vertices.

A similar model to ours has been investigated by Bei and Suksompong [3], who proposed a model of *graphical cakes*; the divisible resource corresponds to edges of an undirected graph, and each agent must be allocated to a connected piece of the resource. Unlike tangles, however, graphical cakes allow a singular point to be allocated to several agents. For this model, in the case of two agents with additive valuations, Bei and Suksompong characterized the class of graphical cakes guaranteeing a proportional allocation (which, for two agents, is equivalent to an envy-free allocation when valuations are additive) as those admitting an "oriented labeling." Intuitively, this notion is an analogue of a bipolar numbering applied to the ordering of edges.

Tangles do not allow singular points to be shared in this way – an essential restriction if we wish to draw links with graph fair division. This difference with Bei and Suksompong [3] is crucial in terms of the existence of EF allocations, and as a result there is little overlap between their results and ours. In particular, under the assumption of Bei and Suksompong, envy-freeness can be trivially achieved when the underlying tangle admits an *Eulerian walk* which visits every tangle edge exactly once. For example, consider a circle attached to three other circles at three different points. In our setting, EF cannot be guaranteed because the center circle is a generalized gap $\geq$ 2 cutset of size 1 (see Sect. 3) whereas in Bei and Suksompong's setting, EF can be

guaranteed because the tangle admits an Eulerian walk. Another example is the lips tangle. Under our model, we need to establish a non-trivial extension of Stromquist's moving knife algorithm to show that the lips tangle guarantees EF for three agents; further, we observe in Sect. 3 that the lips tangle does not guarantee EF for four or more agents. On the other hand, in Bei and Suksompong's setting, the lips tangle guarantees EF for any number of agents, again because the lips tangle admits an Eulerian walk.

Elkind et al. [13] investigated the graphical cake model under the assumptions that singular points can be shared (as in the work of Bei and Suksompong [3]) or that singular points cannot be shared (as in our work). They observed that this difference is critical with respect to the notion of maximin share fairness. Specifically, a maximin share allocation may not exist in the former model even when the graph is a star. On the other hands, such an allocation does exist when the graph is a forest in the latter model. See Proposition 3.1 and Theorem 3.2 of [13].

Of course, a graphical cake may have an envy-free division for some particular assignment of valuation functions, but not for a different one. Deligkas et al. [12] recently studied (in both models) the complexity of deciding the existence of an envy-free division for a given assignment of valuation functions, each of which values each edge uniformly. Such valuations are somewhat more general than the edgewise constant valuations (defined in Sect. 2 of this paper) because a single tangle edge might correspond to several graph edges linked by vertices of degree 2.

2 A Fourth Context for Fair Division: Tangles

We begin by introducing the formal definition of a tangle.

Definition 4 A *tangle* $\mathcal{T}$ is any connected topological space that can be obtained as a quotient space in the following way: take finitely many disjoint copies of the closed interval [0, 1] and apply an equivalence relation that never identifies two points unless each was a 0 or a 1.[6]

Informally, we are connecting copies of [0, 1] by gluing at some endpoints. For example, the lollipop tangle of Fig. 1 can be obtained from two copies of [0, 1] by gluing together any three of the four endpoints.[7] When two tangles are homeomorphic as topological spaces, we consider them to be identical.[8] Other tangles (see Fig. 1) include the unit interval, the circle, letter Y, number 8, handcuffs, letter Θ, the lips, and friendly diamond. Informally, one might say that tangles are graphs made continuous.

Definition 5 Let $\mathcal{T}$ be a tangle. For any point $x \in \mathcal{T}$ a small neighborhood of x looks like k copies of the half open interval [0, 1) with all the 0s identified—that is, k

[6] Tangles are essentially the same as the topologies of finite, one-dimensional CW complexes. Equivalently, to get a tangle take any "drawing" of a connected graph in $\Re^3$ that uses portions of circular arcs as edges, has finitely many vertices, and avoids any edge intersections (except at vertices shared by the edges concerned), treat it as a subset of $\Re^3$ and endow it with the subspace topology induced by that of $\Re^3$.

[7] See, e.g., the Wikipedia article on quotient spaces, also known as identification spaces.

[8] In particular, this means that different instructions for building tangles from multiple copies of [0, 1] can yield identical tangles; the stick of the lollipop could be built from more than one copy, as could the loop. Unlike a CW complex, a tangle does not remember how it was built.

open rays spreading out from x. We will refer to k as the *degree* of x in $\mathcal{T}$, denoted $deg(x)$; points of degree 1 or degree at least 3 are *singular*, and those of degree 2 are *nonsingular*. An *edge* of $\mathcal{T}$ is a connected component in the space obtained when all the singular points of $\mathcal{T}$ are removed. A *subtangle* $\mathcal{S}$ of $\mathcal{T}$ is a connected subspace containing one or more of $\mathcal{T}$'s singular points, together with some of $\mathcal{T}$'s tangle edges joining these points; such an $\mathcal{S}$ is itself a tangle.

Note that any point x that belongs to one of the copies of $[0, 1]$ (from Definition 4) and is not an endpoint of that copy must wind up as a nonsingular point, so each tangle contains only a finite number of singular points. It is easy to see that each edge of $\mathcal{T}$ is homeomorphic to the open interval $(0, 1)$ (with the exception of the circle tangle, which has no singular points and one edge), and that the number of edges corresponds to the minimum number of copies of $[0, 1]$ needed to build $\mathcal{T}$, in the sense of Definition 4.

Any tangle $\mathcal{T}$ can be converted into a connected graph, which—except for the case of the circle tangle—has no degree 2 vertices:

Definition 6 The graph $G_{\mathcal{T}}$ induced by any tangle $\mathcal{T}$ other than the circle tangle $\mathcal{O}$ is obtained by replacing each singular point of $\mathcal{T}$ with a vertex and each tangle edge of $\mathcal{T}$ between two singular points with a graph edge between the corresponding vertices; $G_{\mathcal{O}}$ is defined to be the graph with a single vertex x and one edge from x to x (a loop).

In much the same way that that the simplest tangle of all—the $[0, 1]$ interval—corresponds to the class of all path graphs, every other tangle is similarly associated with an infinite *topological class* of graphs, generated by inserting additional vertices along the edges of $G_{\mathcal{T}}$:

Definition 7 If to any graph G we insert zero or more new vertices of degree 2 along edges, the result is a *subdivision* of G, and the added vertices are *subdivision vertices*. The *topological class* of a graph G is the infinite collection $\mathcal{G}(G)$ of all graphs isomorphic to some subdivisions of G; the *topological class* $\mathcal{G}(\mathcal{T})$ of a tangle $\mathcal{T}$ is just $\mathcal{G}(G_{\mathcal{T}})$.

Proposition 8 *The topological classes of tangles partition the collection of all connected finite graphs having at least one edge.*

Proof Given an arbitrary connected graph $H = (V, E)$ having at least one edge, assume for the moment that H has at least one vertex of degree unequal to 2. For any degree 2 vertex v of H we can delete v from H by removing v from V, removing the two edges $\{v, x\}$ and $\{v, y\}$ of H incident to v from E, and adding an additional $\{x, y\}$ edge to E.[9] Note that $x = y$ is possible, and in this case we are adding an additional loop from x to itself. In "step 1" we will delete all degree 2 vertices in this way, letting $S(H)$ denote the graph so obtained. In step 2, interpret $S(H)$ as a set of instructions for constructing a tangle: replace each edge of $S(H)$ by a copy of $[0, 1]$ and glue these copies together according to the manner in which edges are incident to each other in $S(H)$, yielding a tangle we will call $\mathcal{T}_H$. Note that forming $G_{\mathcal{T}_H}$ "undoes" step 2, with

[9] "Additional" because there may already exist some $\{x, y\}$ edges in E.

the graph edges that had been converted into tangle edges being turned back into graph edges, so that $S(H) \cong G_{\mathcal{T}_H}$. By next inserting degree 2 subdivision vertices along appropriate edges we can undo step 1, obtaining a graph H' with $H' \cong H$. Thus H is in the topological class of the tangle $\mathcal{T}_H$. If all vertices of H have degree 2 then H must be a subdivision of the graph having exactly one vertex x (of degree 2) and one edge (from x to x); in this case, the above argument should be modified by deleting all but one of H's degree 2 vertices when forming $S(H)$, rather than deleting all of them.

To see that topological classes are pairwise disjoint, assume H and H' are connected graphs with at least one edge, with $H \cong H'$, $H \in \mathcal{G}(\mathcal{T})$, and $H' \in \mathcal{G}(\mathcal{T}')$. Then H is isomorphic to a subdivision of $G_{\mathcal{T}}$ and H' is isomorphic to a subdivision of $G_{\mathcal{T}'}$. The isomorphism from H to H' induces an isomorphism from $G_{\mathcal{T}}$ to $G_{\mathcal{T}'}$, which in turn induces a homeomorphism from $\mathcal{T}$ to $\mathcal{T}'$, whence $\mathcal{G}(\mathcal{T}) = \mathcal{G}(\mathcal{T}')$. □

Our set-up for fair division of tangles parallels the one corresponding to [0, 1] presented in the previous section; we assume that each agent i has a valuation v_i defined for all connected pieces of $\mathcal{T}$ that is both *monotone* and *continuous* (see Definition 9). These assumptions are what we need for the proof of Theorems 11 (a) and 32.[10] They are weaker than assuming that the valuations are countably additive, nonatomic measures—in particular, nothing resembling additivity is assumed.

A (real-valued) valuation function v_i is *monotone* if $v_i(X) \leq v_i(Y)$ whenever $X \subseteq Y$. A *segment* of a tangle $\mathcal{T}$ is a subinterval (open, closed, or half-open) of a single one of the copies of [0, 1] used to form $\mathcal{T}$. We denote by $\mathcal{FI}[\mathcal{T}]$ the collection of all unions of finitely many such segments. For $X \in \mathcal{FI}[\mathcal{T}]$ a *minimal representation*—one using the fewest number of segments possible—will guarantee that any two segments overlap in at most two points. Note that any connected subset of $\mathcal{T}$ is a union of at most $2k$ such segments,[11] where k is the number of tangle edges, and is determined (up to boundary points) by the locations of the endpoints of these segments. The following definition of continuity may seem technical, but it captures the intuitive idea that a small enough change in a connected subset of a tangle (due to a small change in endpoint locations) should only make a small change in the value of that subset.

Definition 9 For $X \in \mathcal{FI}[\mathcal{T}]$, $\ell(X)$ denotes X's total *length*—the sum of the lengths of the segments in a minimal representation of X. The pseudometric d on $\mathcal{FI}[\mathcal{T}]$ is defined by $d(X, Y) = \ell(X \Delta Y)$.[12]. A valuation function v (defined for each connected subset of $\mathcal{T}$) is *continuous* if for each sequence $\{X_n\}$ of connected subsets of $\mathcal{T}$ and each single connected subset Y, if $d(X_n, Y) \to 0$ as $n \to \infty$, then $v(X_n) \to v(Y)$ as $n \to \infty$.

Continuity implies that the value $v_i(X)$ is independent of which of X's boundary points are elements of X.

[10] It may seem surprising that we do not need to require $v_i(X) \geq 0$ or $v_i(\emptyset) = 0$. However, any monotone continuous valuation function v_i satisfies $v_i(X) = v_i^\star(X) + K$ (for all connected pieces X) where K is a (possibly negative) constant and $v_i^\star$ satisfies both of these additional requirements; note that v_i and $v_i^\star$ induce identical preference relations over connected subsets.

[11] This is because the intersection of a connected subset S of a tangle $\mathcal{T}$ with any single tangle edge consists of one or two segments (or is empty).

[12] Here the *symmetric difference* $X \Delta Y$ of sets X and Y is given by $X \Delta Y = (X \backslash Y) \cup (Y \backslash X)$

All of our negative results for tangles hold even for *identical edgewise constant valuations*, a subclass of identical additive valuations.[13] For a tangle $\mathcal{T}$ built from k copies of $[0, 1]$ corresponding to k tangle edges $e_1, e_2, \ldots, e_k$, a valuation function v_i is *edgewise constant* if each edge is valued uniformly, so that there exist constants $\{c_j\}_{j=1}^k$ for which each connected set X satisfies $v_i(X) = \sum_{j=1}^k c_j \ell(X \cap e_j)$.

Our main concern here regards the existence of connected and *envy-free* allocations for n agents, in which $\mathcal{T}$ is partitioned into n pieces, each agent i is awarded a different piece T_i of the partition, each piece is topologically connected (as a subset of the topological space $\mathcal{T}$), and each agent values their assigned piece at least as highly as the pieces awarded to others: $v_i(T_i) \geq v_i(T_j)$ for all i, j.

The simplest tangle of all—the unit interval—has served as the mathematical model of a *cake* in the study of *cake-cutting*, the informal term used within one major thread in the literature on fair division. What happens, then, when some other tangle stands in for $[0, 1]$ in Theorem 1; more precisely, for which tangles, and which values of n, are connected and envy-free allocations for n agents guaranteed? We begin by identifying a special sub-class of tangles.

3 Tangle boundedness theorem

Intuitively, a *stringable* tangle is made by arranging a single length of string on a tabletop so that it does not touch itself, except that either one or both of the two ends may butt up against each other, or against some other part of the string (or do both). A moment's thought reveals that this notion is not a perfect analogue of a Hamiltonian path, nor of an Eulerian path, in the continuous context.

Definition 10 A tangle $\mathcal{T}$ is *stringable* if there exists a continuous surjection $F : [0, 1] \to \mathcal{T}$ from the closed unit interval to $\mathcal{T}$ that is injective when restricted to the open interval $(0, 1)$. The *basic stringable tangles* are the six shown in the top row of Fig. 1: the closed unit interval itself, the circle, the lollipop, number 8, handcuffs, and letter Θ.

Each of these can be formed by arranging a string in the manner described above. In what follows, we demonstrate that no other stringable tangles exist. Until then, the term "basic" is used to designate these six in particular (and has no other special meaning). Each stringable tangle can substitute for $[0, 1]$ in Theorem 1, but only because stringable tangles are almost the same as $[0, 1]$. The extension is unremarkable, and turns out to add nothing to our understanding of graph division. More striking, and consequential for graphs, is that stringable tangles are the *only* ones that satisfy Theorem 1:

Theorem 11 (Tangle Boundedness Theorem) *Let $\mathcal{T}$ be a tangle.*

(a) If $\mathcal{T}$ is stringable, then connected and envy-free allocations are guaranteed to exist for any finite number of agents with monotone continuous valuations.

[13] "Identical" here refers to the fact that the same valuation is shared by all agents.

(b) *If $\mathcal{T}$ is not among the six basic stringable tangles, then there exists an upper bound n_0 on the number of agents for which envy-free allocations are guaranteed, even if agents have identical, edgewise constant valuations with at most two nonzero densities.*[14]

(c) *The only stringable tangles are the six basic ones.*

Part (c) is an immediate consequence of parts (a) and (b), while part (b) follows immediately from Lemmas 13, 15, and 16 below. (Note, however, that the proof of Lemma 16 is deferred to Sect. 4, where the relevant material on degree sequences is discussed). Here, we will prove part (a).

Proof of Theorem 11 **(*a*)** For part (a), let $F\colon [0,1] \to \mathcal{T}$ witness $\mathcal{T}$'s stringability, and $\{v_i\}_{1\le i\le n}$ be agents' valuations over connected subsets of $\mathcal{T}$. These induce corresponding valuations $\{v_i^\star\}_{1\le i\le n}$ on $[0,1]$ defined by $v_i^\star(X) = v_i(F[X])$ for each subinterval X. Applying Theorem 1 delivers an allocation $\{T_i^\star\}_{1\le i\le n}$ of connected subsets of $[0,1]$ to the agents, that is envy-free with respect to the $v_i^\star$.

The corresponding allocation $\{F[T_i^\star]\}_{1\le i\le n}$ of connected subsets of $\mathcal{T}$ may not quite form a partition of $\mathcal{T}$, because there may exist a point $x \in \mathcal{T}$ that is the image of two distinct points in $[0,1]$, which belong to two different sets $T_i^\star$. Then at least one of these two points must be an endpoint. Remove one such endpoint from its set $T_i^\star$ and remove its image from $F[T_i^\star]$. If there is a second such point x, do this again. The $T_i^\star$ pieces are now disjoint, and they remain connected.

Recall that the value $v_i(X)$ is independent of which of X's boundary points are elements of X. Thus, we have that $v_i^\star(T_i^\star) = v_i(F[T_i^\star]) = v_i(T_i)$ so that our allocation $\{T_i\}_{1\le i\le n}$ inherits its envy-freeness from that of $\{T_i^\star\}_{1\le i\le n}$. □

For the proof of part (b) of Theorem 11, much rests on the notion of *gap* ≥ 2 *cutset* introduced below:

Definition 12 Let $\mathcal{T}$ be a tangle. If $X \subseteq \mathcal{T}$ has t elements and $\mathcal{T} \setminus X$ is disconnected with $t+k$ connected components then we will say that the *gap* $\gamma(X)$ of X is equal to k and that X is a *gap k cutset of cardinality t*; if $\gamma(X) \ge 2$, then X is a *gap* ≥ 2 *cutset of cardinality t*.

For the point labeled c in the $\mathcal{Y}$ tangle, $\{c\}$ is a gap ≥ 2 cutset of cardinality 1; for the lips tangle, the points labeled a, b, and c comprise a gap ≥ 2 cutset of cardinality 3. See Fig. 2.

Lemma 13 *Gap ≥ 2 Lemma Let $\mathcal{T}$ be a tangle. Suppose $\mathcal{T}$ has a gap ≥ 2 cutset of cardinality t. Then, for each $n \ge t+1$, $\mathcal{T}$ does not guarantee connected envy-free allocations for n agents, even when they have identical edgewise constant valuations with at most two nonzero densities.*

Proof Let $\mathcal{T}$ be a tangle, $X \subseteq \mathcal{T}$ have t elements, $\mathcal{T} \setminus X$ have $t+k$ connected components $D_1, \ldots, D_{t+k-1}, D_{t+k}$ with $k \ge 2$, and $n \ge t+1$ be the number of

[14] Theorem 11 implies a negative result for the alternative fairness notion of proportionality, since proportionality and envy-freeness are equivalent when agents have identical additive measures.

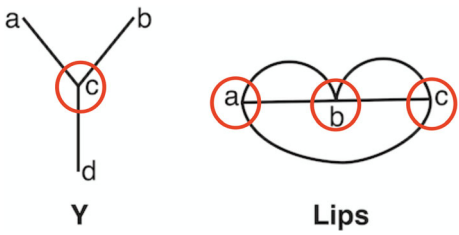

Fig. 2 Examples of a gap ≥ 2 cutset

agents. We will assume that X contains no points of degree 2—this entails no loss of generality, because adding a degree 2 point to a cutset X can never increase X's gap. It follows that the interior of each individual tangle edge is entirely contained in a single component D_j. We now construct an identical edgewise constant valuation with at most two nonzero densities for which no connected EF allocation exists. For each $i = 1, 2, \ldots, t + k - 1$ spread $\frac{t+1}{t+k-1} \leq 1$ units of value uniformly over one of D_i's tangle edges (with 0 value spread over the other edges). Spread $n - t \geq 1$ units of value uniformly over one of D_{t+k}'s tangle edges (with 0 value spread over the other edges). Note that the identical edgewise constant valuation so generated employs at most two non-zero different densities, with a total value of $n + 1$ for the whole of $\mathcal{T}$.

Now let $\{T_i\}_{1 \leq i \leq n}$ be any allocation that assigns a connected piece to each agent. Then there are at least $n - t$ agents whose share includes no members of the cutset X, so that for each of these agents their share must be a subset of some single connected components $D_1, \ldots, D_{t+k-1}, D_{t+k}$. If for each of these agents their share is a subset of D_{t+k}, then as the value of D_{t+k} is only $n - t$, some *deprived* agent receives a share valued at 1 or less. If some agent receives a share that is a subset of one of the other components, then their value for their share is also 1 or less, and they are deprived. The remaining $n - 1$ agents share a value of at least n, so that at least one of them receives at least $1 + \frac{1}{n-1}$ and the deprived agent envies them by at least the amount $b = \frac{1}{n-1}$. □

Notice that for the first three examples below, the gap ≥ 2 cutset mentioned consists of the singular points of degree 3 or greater:

- The $\mathcal{Y}$ tangle has a gap ≥ 2 cutset of cardinality 1; connected EF allocations to n agents can fail for each $n \geq 2$.
- The friendly diamond tangle has a gap ≥ 2 cutset of cardinality 2, so connected EF allocations to n agents can fail for each $n \geq 3$. Cardinality 1 cutsets each have gap 1 or less, and we know that this tangle does guarantee connected EF allocations for 2 agents (because it is the continuous analogue of a graph having a bipolar numbering—see [4] and Sect. 5).
- The lips tangle has a gap ≥ 2 cutset of cardinality 3, so connected EF allocations to n agents can fail to exist for each $n \geq 4$. Cardinality 2 cutsets each have gap 1 or less, and we show that lips guarantees connected EF allocations for two agents (Sect. 5) and for three agents (Sect. 6).
- For the unit interval (and for every stringable tangle) cardinality t cutsets each have gap 1 or less, for every choice of t. These tangles guarantee connected EF allocations for every number of agents (via Theorem 11).

- A tangle that is itself disconnected (so, not actually a tangle, according to the definition we have given here) has a gap ≥ 2 cutset of cardinality 0. It does not guarantee connected allocations to even a single agent.

Definition 14 For a tangle $\mathcal{T}$, let $\Sigma_3(\mathcal{T})$ denote the set of singular points of $\mathcal{T}$ having degree 3 or greater, $\sigma_3(\mathcal{T})$ denote $|\Sigma_3(\mathcal{T})|$, and $\epsilon(\mathcal{T})$ denote the number of edges of $\mathcal{T}$.

Lemma 15 *For any tangle $\mathcal{T}$ satisfying $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq 2$, $\Sigma_3(\mathcal{T})$ is a cutset of gap ≥ 2.*

Proof We claim that deleting all points in $\Sigma_3(\mathcal{T})$ disconnects any tangle $\mathcal{T}$ (stringable or not) into its constituent edges. It is clear that any edge connecting two singular points in $\Sigma_3(\mathcal{T})$ is disconnected from everything else, and this is also evident if the edge connects one singular point in $\Sigma_3(\mathcal{T})$ with another having degree 1; of course, if the edge connects two degree 1 singular points, then $\mathcal{T}$ is necessarily the unit interval, and the claim holds trivially.

Lemma 16 *For any tangle $\mathcal{T}$ other than the six basic stringable ones, $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq 2$.*

The proof of Lemma 16, via degree sequences, is postponed to Sect. 4. Note that the Tangle Boundedness Theorem (Theorem 11) now follows immediately from Lemmas 13, 15, and 16.

As we have seen, most tangles are not stringable, and the Tangle Boundedness Theorem (b) paints only part of the picture. Below, we investigate this matter further.

Definition 17 The *gap threshold* of a tangle is the smallest integer t for which a gap ≥ 2 cutset of cardinality t exists; if no gap ≥ 2 cutset exists of any size, then the gap threshold is ∞.

The gap threshold for the $\mathcal{Y}$ tangle is 1, for the friendly diamond is 2, and for the lips is 3; for these cases the set of singular points of degree 3 or 4 is the relevant cutset; deleting these points leaves 3 disconnected pieces for the $\mathcal{Y}$, 4 for the friendly diamond, and 5 for lips. The *Gap* ≥ 2 *Lemma* now tells us that the friendly diamond does not guarantee connected envy-free allocations for three agents or more, while the lips tangle does not guarantee connected envy-free allocations for four or more. Stringables, of course, have gap threshold ∞. Moreover each non-stringable tangle has a finite gap threshold, which serves as an upper bound n_0 for Theorem 11. But does the gap threshold provide the *least* upper bound? Sometimes it does:

Remark 18 For each tangle $\mathcal{T}$ appearing in Fig. 1, $\mathcal{T}$ guarantees envy-free connected allocations for n agents whenever n is less than or equal to $\mathcal{T}$'s gap threshold.

This observation summarizes several *positive* results. In Sect. 5, we characterize those tangles that guarantee EF allocations of connected shares for two agents, in terms of the existence of a *bipolar ordering*. The characterization and its proof are parallel to the closely related characterization for graphs and two agents in [4], and

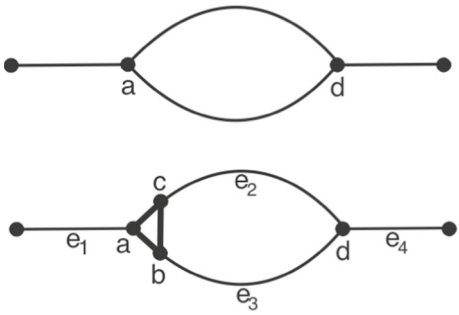

Fig. 3 The δ-diamond tangle (below) with the friendly diamond (above)

thus represents a transfer from the discrete to the continuous realm. It follows that the friendly diamond tangle and the lips tangle $\mathcal{L}$ each guarantee connected, envy-free allocations for two agents. Moreover, in Sect. 6, we show that $\mathcal{L}$ guarantees connected and envy-free allocations for three agents. This result, taken along with those described in the previous paragraph, the mentioned values for some specific gap thresholds, and the assertions about stringable tangles given before the statement of Theorem 11, verify Remark 18.

Remark 18 does not refer to all tangles, however. Figure 3 shows the friendly diamond tangle, along with a variant—the *δ-diamond tangle* $\mathcal{T}_{\delta\diamond}$—which replaces the single degree 3 singular point a of the friendly diamond with triangle abc, a subtangle of $\mathcal{T}_{\delta\diamond}$. It is easy to check that, unlike the original friendly diamond, $\mathcal{T}_{\delta\diamond}$ has no gap ≥ 2 cutset of cardinality 2, while $\{a, b, c\}$ provides one of cardinality 3, so that $\mathcal{T}_{\delta\diamond}$'s gap threshold is 3. In a moment, we will argue that $\mathcal{T}_{\delta\diamond}$ fails to guarantee connected EF allocations for three agents.

Definition 19 The *agent bound* of a non-stringable tangle $\mathcal{T}$ is the largest integer n such that envy-free and connected allocations of $\mathcal{T}$ are guaranteed for n agents with monotone continuous valuations.[15]

In fact, $\mathcal{T}_{\delta\diamond}$'s *agent bound* is 2, which is less than its gap threshold. The fact that connected EF allocations of $\mathcal{T}_{\delta\diamond}$ are guaranteed for two agents follows from Theorem 27 of Sect. 5. To see that such allocations may fail to exist for three agents, consider $\triangle abc$, a closed subset and subtangle of $\mathcal{T}_{\delta\diamond}$ consisting of the three points a, b, and c, together with the three small tangle edges joining these three points. Let $X_1 = \triangle abc$ and $X_2 = \{d\}$. We will argue that $X_{\delta\diamond} = \{X_1, X_2\}$ behaves like a gap ≥ 2 cutset of cardinality 2; it can block envy-free connected allocations of the δ-diamond tangle in the same way that $X_\diamond = \{a, d\}$ does for the original friendly diamond.

Proposition 20 *The δ-diamond tangle $\mathcal{T}_{\delta\diamond}$ does not guarantee envy-free connected allocations for three or more agents with identical edgewise constant valuations.*

Proof Let $X_{\delta\diamond} = \{X_1, X_2\}$ be defined as in the previous paragraph. An allocation of $\mathcal{T}_{\delta\diamond}$ may award to some agent i a piece T_i that includes a majority (either two or three)

[15] The definition does not say, "guaranteed for *n or fewer* agents." In fact, we do not know if the following "population monotonicity" holds: if a tangle guarantees envy-freeness for n agents, then it guarantees envy-freeness for $n - 1$ agents.

of the three points a, b, and c; in this case we will say that T_i *dominates* X_1. Clearly, at most one such piece can dominate X_1. Thus, with a connected allocation to three or more agents, there must be at least one agent k whose piece T_k neither dominates X_1 nor contains d as a member. This T_k must be a connected subset of the "leftovers" $\mathcal{L}_k = \mathcal{T}_{\delta\diamond} \setminus \cup_{i \neq k} T_i$ obtained by excising the other two pieces. But the four edges e_1, e_2, e_3, and e_4 that do not belong to $\triangle abc$ are disconnected in $\mathcal{L}_k$. If a path $\mathcal{P}$ in $\mathcal{T}_{\delta\diamond}$ connects a point from one of these four edges to a point in a different one of the four, then it passes either through d or through at least two of the three points a, b, and c. Thus no such path can be contained in T_k, and so T_k contains elements from at most one of these four edges.

The rest of the argument proceeds as in the proof of the Gap ≥ 2 Lemma 13, treating the edges e_1, e_2, e_3, and e_4 as if they were the disconnected components D_j in that earlier proof. In the case of exactly three agents, this would mean assigning identical valuations that spread one unit of value along each e_j, with some negligible (or zero) value assigned to the edges of $\triangle abc$. □

The situation seems to call for some suitably generalized version of the cutset definition, such as the following:

Definition 21 A *generalized gap* ≥ 2 *cutset of cardinality* t is a finite set $X = \{X_1, X_2, \ldots, X_t\}$ of pairwise disjoint connected subtangles[16] of a tangle $\mathcal{T}$, for which the set $\mathcal{T} \setminus (\cup_{i=1}^{m} X_i)$ contains $t + 2$ or more disconnected components $C_1, C_2, \ldots, C_t, C_{t+1}, C_{t+2}, \ldots$, and these components satisfy:

(a) For each i and j the intersection $X_i \cap \overline{C_j}$ of X_i with the closure (in $\mathcal{T}$) of C_j is either empty, or contains a single point $s_{i,j}$, known as the *contact point* for X_i and C_j.
(b) For each i, X_i is either

- a "type I" member, containing a single point of $\mathcal{T}$,
- or a "type II" member: a set with more than one element, satisfying that exactly three components C_j share a contact point with X_i, and that these contact points are distinct.

Indeed, one can show that $X_{\delta\diamond}$ satisfies this definition, and then lift Proposition 20 and its proof to obtain a corresponding *generalized* gap ≥ 2 lemma. Moreover, for the special case $t = 1$, Definition 21 nicely reduces to the definition of *trident*, as employed in Theorem 29, Sect. 5. Suppose we define the *generalized gap threshold* to be the smallest integer t for which a generalized gap ≥ 2 cutset of cardinality t exists, and that this threshold has value 2 for some tangle $\mathcal{S}$ (as it does for $\mathcal{S} = \mathcal{T}_{\delta\diamond}$). Then Theorem 29—which characterizes exactly which tangles guarantee connected

[16] Note that a subtangle of $\mathcal{T}$ is necessarily closed in the topology of $\mathcal{T}$, so that deleting all points from a generalized cutset leaves connected components that are open in $\mathcal{T}$. For any generalization of the Gap ≥ 2 Lemma to the context of Definition 21, we require open components in order to be sure that monotone continuous valuations can assign nonzero values to these components (or values greater than that assigned to $\emptyset$—see footnote 10). Without that requirement, a component might consist of single point, and any continuous monotone valuation must assign to a single point the same value as that assigned to the empty set (because "nonatomic" follows from "continuous").

EF allocations for two agents—tells us that the generalized gap threshold coincides with the agent bound (the largest n for which connected EF allocations are guaranteed to exist for n agents having monotone continuous valuations over $\mathcal{S}$).

The situation is more complex, however, when the generalized gap threshold exceeds 2 in value. Other tangles exist, for which the generalized gap threshold is 3, yet arguments similar to those in Proposition 20 nonetheless show the agent bound to be only 2. These examples suggest how to further modify Definition 21 to incorporate these newer and even more general cutsets. However, we have then repeatedly been able to construct even more subtle examples that escape the latest round of generalization. Our goal had been to identify some most general definition of cutset—one for which the corresponding gap threshold achieves the agent bound in the general case—because this would point towards a characterization of the agent bound for an arbitrary tangle. But that goal is not yet within reach, and will have to serve as the subject of a future inquiry.

4 Degree sequences for tangles

Recall that one key step in the proof of the Tangle Boundedness Theorem was Lemma 16, which asserts every tangle $\mathcal{T}$ is either one of the six basic stringable tangles, or satisfies the inequality:

$$\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq 2, \tag{1}$$

where $\epsilon(\mathcal{T})$ is the number of $\mathcal{T}$'s edges and $\sigma_3(\mathcal{T})$ is the number of singular points of degree three or higher. The proof of this lemma, provided below, uses *degree sequences* as topological invariants for tangles. These are almost the same as degree sequences for graphs (which are well-known invariants in the graph context). Note that in order to avoid circularity, throughout this section we cannot assume that the only stringable tangles are the six basic ones (from Theorem 11).

Definition 22 If $\mathcal{T}$ is a tangle and $i \geq 1$ is any integer, let d_i be the number of singular points of degree i (So that d_2 is, perforce, 0).[17] The *degree sequence* $\overrightarrow{deg}(\mathcal{T})$ *of a tangle* is the vector $\langle d_1, d_2, \ldots, d_r \rangle$, where $d_r > 0$ and $d_s = 0$ for all $s > r$.

Proposition 23 *For any tangle* $\mathcal{T}$*:*

(a) If $\overrightarrow{deg}(\mathcal{T}) = \langle\,\rangle$*, then* $\mathcal{T}$ *must be the circle,*
(b) If $\overrightarrow{deg}(\mathcal{T}) = \langle 2 \rangle$*, then* $\mathcal{T}$ *must be the unit closed interval,*
(c) If $\overrightarrow{deg}(\mathcal{T}) = \langle 1, 0, 1 \rangle$*, then* $\mathcal{T}$ *must be the lollipop,*
(d) If $\overrightarrow{deg}(\mathcal{T}) = \langle 0, 0, 0, 1 \rangle$*, then* $\mathcal{T}$ *must be the Fig. 8, and*
(e) If $\overrightarrow{deg}(\mathcal{T}) = \langle 0, 0, 2 \rangle$*, then* $\mathcal{T}$ *must either be the handcuffs or the theta.*

Proof Take a segment of any tangle edge in some tangle $\mathcal{T}$; extending it in both directions must lead to two distinct singular points, or to a common singular point of degree at least 3, unless the edge loops back on itself. Thus degree sequences of $\langle\,\rangle$

[17] We ruled out singular points of degree two by fiat, in Definition 5. Note that *all* tangles have the same number of degree 2 points—continuum many—so counting them does not help distinguish among tangles.

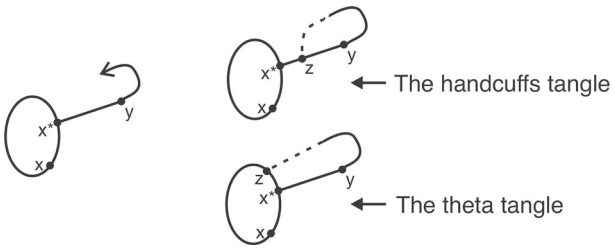

Fig. 4 Illustration of the proof of Proposition 23

or $\langle 2 \rangle$ can only arise when $\mathcal{T}$ has a single such edge, so that $\mathcal{T}$ must be the circle or unit closed interval, respectively. Degree sequences of $\langle 1, 0, 1 \rangle$ or $\langle 0, 0, 0, 1 \rangle$ can only arise from exactly two tangle edges; if the singular point joining them is degree 3, we are left with the lollipop, while degree 4 leads to the Fig. 8.

Turning to part (e), if there exists a continuous function $F : [0, 1] \to \mathcal{T}$ that is injective on $(0, 1)$, and satisfies $F(0) = F(1)$ then we will refer to the image of $[0, 1]$ in $\mathcal{T}$ as a *cycle* of tangle $\mathcal{T}$. We claim that if $\mathcal{T}$ has no degree 1 singular points then it must contain a cycle. Start a path at any point in $\mathcal{T}$ and continue it as far as possible without allowing the path to visit any point a second time. Such a path must get "stuck" at some point z, for one of two reasons. Perhaps there is one way in to z and no ways out—in that case z is a singular point of degree 1, but $\mathcal{T}$ has none of those. Otherwise, it must be that z has been visited earlier by the path, in which case the portion $\mathcal{C}$ of the path between the first and second visit to z must be a cycle.

Now assume that $\overrightarrow{deg}(\mathcal{T}) = \langle 0, 0, 2 \rangle$, and let $\mathcal{C} \subseteq \mathcal{T}$ be a cycle. Clearly $\mathcal{T}$ must contain points that do not belong to $\mathcal{C}$. As $\mathcal{T}$ is connected, there is some path $\mathcal{P} \subseteq \mathcal{T}$ that starts at some point $x \in \mathcal{C}$ and ends at some point $y \in \mathcal{T} \backslash \mathcal{C}$ (see Fig. 4 left). The cycle $\mathcal{C}$ is closed as a subset of $\mathcal{T}$, so when we trace path $\mathcal{P}$ from x there will be some last point $x^\star$ along $\mathcal{P}$ that lies in $\mathcal{C}$. Remove the portion of the path between x and $x^\star$ so that the new path $\mathcal{P}^\star$ now begins at $x^\star$. Now $x^\star$ will be the only point of $\mathcal{C}$ along $\mathcal{P}^\star$ and will be a singular point of degree 3 or greater. Assume that $\mathcal{P}^\star$ does not yet visit any point twice, and now extend $\mathcal{P}^\star$ in $\mathcal{T}$ until it either revisits some point $z \in \mathcal{P}^\star$ for a second time (Fig. 4, upper right), or hits a point $z \in \mathcal{C}$ (Fig. 4, lower right) noting that $z \neq x^\star$ (otherwise $x^\star$ is a singular point of degree 4 or higher). After this extension $\mathcal{P}^\star \cup \mathcal{C}$ is a tangle, with two singular points $x^\star$ and z, each of degree three. If $z \in \mathcal{C}$, then $\mathcal{P}^\star \cup \mathcal{C}$ is Θ; otherwise $\mathcal{P}^\star \cup \mathcal{C}$ is handcuffs. Finally, if there were any additional points of $\mathcal{T}$ outside $\mathcal{P}^\star \cup \mathcal{C}$ they would be connected to $\mathcal{P}^\star \cup \mathcal{C}$ via some path that adds additional singular points, or changes the degree of $x^\star$ or of z, so no such additional points exist. □

We are now ready to prove Lemma 16, stating that for any tangle $\mathcal{T}$ other than the six basic stringable ones, $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq 2$.

Proof of Lemma 16 For graphs, we know the number of edges is half the sum of the degrees of the vertices, and it is easy to see that this holds for tangles, as well,[18]

[18] The circle tangle is the unique exception, as its single edge does not contribute to the degree of any vertex.

with *singular points* replacing *vertices*. Thus, from the degree sequence of a tangle $\mathcal{T}$ we can derive both the number $\epsilon(\mathcal{T})$ of edges and the number $\sigma_3(\mathcal{T})$ of singular points of degree three or higher. The relationship also allows us to eliminate degree sequences that are *parity blocked* because they imply an odd sum. The rest of the proof breaks all parity unblocked degree sequences of tangles into ten cases (grouped in 4 categories), with each case either falling under Proposition 23 (thus leading only to a basic stringable tangle) or shown via a very short calculation (as sketched earlier in this paragraph) to imply $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq 2$.[19] Note that the parameters r, s, and t below are assumed to be nonnegative integers.

Case 1 $\mathcal{T}$ has no singular point of degree 3 or higher.

- $\overrightarrow{deg}(\mathcal{T}) = \langle\rangle$: See Proposition 23.
- $\overrightarrow{deg}(\mathcal{T}) = \langle r\rangle$: We must have $r \leq 2$, lest $\mathcal{T}$ be disconnected; for $r = 2$ see Proposition 23.

Case 2 $\mathcal{T}$ has at least one singular point of degree 3 and no singular point of degree 4 or higher.

- $\overrightarrow{deg}(\mathcal{T}) = \langle 0, 0, 2\rangle$ or $\langle 1, 0, 1\rangle$: See Proposition 23.
- $\overrightarrow{deg}(\mathcal{T}) = \langle t, 0, 2\rangle, t \geq 2$: Then $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq \frac{8}{2} - 2 = 2$
- $\overrightarrow{deg}(\mathcal{T}) = \langle t, 0, 2r\rangle, r \geq 2$: Here $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq \frac{6r}{2} - 2r = r \geq 2$
- $\overrightarrow{deg}(\mathcal{T}) = \langle t, 0, 1\rangle, t \geq 3$: In this case $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq \frac{t+3}{2} - 1 \geq 3 - 1 = 2$.
- $\overrightarrow{deg}(\mathcal{T}) = \langle t, 0, 2r+1\rangle, t \geq 1, r \geq 1$: Then $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq \frac{3(2r+1)+1}{2} - (2r + 1) = r + 1 \geq 2$.

Case 3 $\mathcal{T}$ has at least one singular point of degree 4 and no singular point of degree 5 or higher.

- $\overrightarrow{deg}(\mathcal{T}) = \langle 0, 0, 0, 1\rangle$: See Proposition 23.
- $\overrightarrow{deg}(\mathcal{T}) = \langle t, 0, r, s\rangle, s \geq 1, t+r+s \geq 2$ (note that when $s \geq 1$ and $t+r+s \leq 1$, then we are back to the case above): Then $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq \lceil \frac{t+3[r+(s-1)]+4}{2} - (r + s)\rceil = \lceil \frac{t+r+s+1}{2}\rceil \geq 2$.

Case 4 $\mathcal{T}$ has $r + 1 \geq 1$ singular points of degree 3 or greater, at least one of which is degree 5 or greater: Then $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T}) \geq \lceil \frac{5+3r}{2} - (r + 1)\rceil = \lceil \frac{3+r}{2}\rceil \geq 2$. □

5 Envy-free allocations of tangles for two agents

In Theorem 11, we have established that stringable tangles are the only ones that guarantee EF for any number of agents. What if we aim to guarantee EF for a particular number of agents, instead of *any* number of agents? Our goal of this section is to characterize the family of tangles guaranteeing EF for two agents. Specifically, we

[19] The referee suggested an alternative approach that avoids cases. Observe that the quantity $\epsilon(\mathcal{T}) - \sigma_3(\mathcal{T})$ is given by the dot product $\overrightarrow{deg}(\mathcal{T}) \cdot \langle \frac{1}{2}, 0, \frac{1}{2}, 1, \frac{3}{2}, 2, \ldots, \frac{k-2}{2}, \ldots\rangle$. This dot product is a nonnegative integer and can only equal 0 or 1 for the degree sequences listed in Proposition 23; for all others, its value is at least 2.

show that a tangle $\mathcal{T}$ guarantees connected envy-free allocation if and only if it has no generalized gap ≥ 2 cutset of cardinality 1. This characterization theorem for tangles uses the notion of a *bipolar ordering* of a tangle; it is analogous to, and builds upon, the result of Bilò et al. [4] (restated, below, as Theorem 27).

We first review the definitions of a *separating vertex*, a *non-separable graph*, a *block*, and a *block tree* of a multi-graph. As Bilò et al. only consider the class of simple graphs, their definition of a block is slightly different from that for multi-graphs that allow for loops and parallel edges in [5]. Consider a multi-graph G. A *decomposition* of $G = (V, E)$ is a family $\mathcal{F}$ of edge disjoint subgraphs $F = (V_F, E_F)$ of G such that the union of the edges $\bigcup_{F \in \mathcal{F}} E_F$ is E. A *separation* of a connected multi-graph G is a decomposition of the graph into two non-empty connected subgraphs G_1 and G_2 that have only one vertex in common, which is said to be a *separating vertex* of the graph. A multi-graph G is called *non-separable* if it is connected and has no separating vertices; otherwise, it is called *separable*. A *block* of a multi-graph is a subgraph that is non-separable and is maximal with respect to this property.

Lemma 24 (Bondy and Murty [5] (Proposition 5.3)) *Let G be a multi-graph. Then:*

- *any two blocks of G have at most one vertex in common,*
- *the blocks of G form a decomposition of G,*
- *each cycle of G is contained in a block of G.*

For a connected multi-graph G, consider a bipartite graph $B(G)$ with bipartition $(\mathcal{B}, S)$, where $\mathcal{B}$ is the set of blocks of G and S the set of separating vertices of G; a block B and a separating vertex v are adjacent in $B(G)$ if and only if v belongs to B. The block graph $B(G)$ does not contain a cycle by the above lemma, and is therefore referred to as the *block tree* of G. Bilò et al. [4] characterize the graphs that guarantee connected, EF1 allocations for two agents, using the notions of a *bipolar numbering* and a *trident*.

Definition 25 A *bipolar numbering* of a connected multi-graph G is an enumeration $v_1, v_2, \ldots, v_k$ of G's vertices satisfying any of the following three (equivalent) conditions:

- the subgraph induced by any initial or final segment of the numbering is connected in G
- for each j with $1 < j \leq k$, v_j is adjacent to some v_i with $1 \leq i < j$ and for each ℓ with $1 \leq \ell < k$, v_ℓ is adjacent to some v_h with $\ell < h \leq k$
- for each j with $1 < j \leq k$, G contains a path $v_{i_1}, v_{i_2}, \ldots, v_{i_r}$ from v_1 to v_j satisfying $1 = i_1 < i_2 < \cdots < i_r = j$ and for each ℓ with $1 \leq \ell < k$, G contains a path $v_{i_1}, v_{i_2}, \ldots, v_{i_s}$ from v_k to v_ℓ satisfying $k = i_1 > i_2 > \cdots > i_s = \ell$.[20]

A *trident* is the graph analogue of a generalized gap ≥ 2 cutset of cardinality 1.

Definition 26 A *trident* is a substructure of a connected graph G consisting of connected subgraphs C, P_1, P_2, P_3 of G such that (i) each P_i contains exactly one *contact*

[20] The standard definition does not mention this third condition. It can be shown to follow from the second condition via an easy inductive argument.

vertex s_i in common with C, (ii) each P_i contains at least two vertices, (iii) [the "double door" condition] for each $i \neq j$ any path connecting a vertex in P_i with a vertex in P_j passes through both s_i and s_j, and (iv):

(I) [type I trident] either C consists of a single vertex s (whence $s = s_1 = s_2 = s_3$), or
(II) [type II trident] the three contact vertices are distinct.

Theorem 27 (Bilò et al. [4]) *The following conditions are equivalent for every simple connected graph G:*

(a) G admits a bipolar numbering.
(b) G guarantees EF1 for two agents with monotone valuations.
(c) G guarantees EF1 for two agents with identical, additive, binary valuations[21].
(d) G does not contain a trident.
(e) The block tree $B(G)$ is a path.

The proof of Theorem 27 appears in [4], and consists of a cycle of implications:

$$(a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (d) \Rightarrow (e) \Rightarrow (a). \tag{2}$$

The notion corresponding to a bipolar numbering for tangles is as follows:

Definition 28 A *bipolar ordering* for a tangle $\mathcal{T}$ is a linear order relation $<$ on the elements of $\mathcal{T}$ for which every initial segment of $\mathcal{T}$ is (path) connected as a subset of $\mathcal{T}$, as is every final segment. Here a subset S of $\mathcal{T}$ is an *initial segment* under $<$ if for all $x, y \in \mathcal{T}$ with $x < y$, if $y \in S$ then $x \in S$. It is a *final segment* if for all $x, y \in \mathcal{T}$ with $x < y$, if $x \in S$ then $y \in S$.

Of the tangles appearing in Fig. 1, only the Y fails to have a bipolar ordering. The next result is an analogous characterization to (2) for tangles and two agents.

Theorem 29 *Let $\mathcal{T}$ be a tangle other than the circle. Let $G_{\mathcal{T}}^{\dagger}$ be the graph in the topological class $\mathcal{G}(\mathcal{T})$ obtained by starting with $G_{\mathcal{T}}$ (see Definition 7) and then inserting one additional degree 2 subdivision vertex*[22] *along each edge. Then the following are equivalent:*

(1) $G_{\mathcal{T}}^{\dagger}$ admits a bipolar numbering.
(2) $\mathcal{T}$ admits a bipolar ordering.
(3) $\mathcal{T}$ guarantees connected EF allocations for two agents with monotone continuous valuations.
(3') $\mathcal{T}$ guarantees connected EF allocations for two agents with identical edgewise constant valuations with at most one non-zero density.
(4) $\mathcal{T}$ has no generalized gap ≥ 2 cutset of cardinality 1.
(5) The block graph $B(G_{\mathcal{T}})$ is a path.

[21] A valuation function $v_i : 2^O \to \mathbb{R}_+$ is *additive* if $v_i(X) = \sum_{o \in X} v_i(\{o\})$ for each $X \subseteq O$. An additive valuation $v_i : 2^O \to \mathbb{R}_+$ is *binary* if $v_i(o) \in \{1, 0\}$ for every item $o \in O$.

[22] "At least one" would also work, but would be less convenient for the argument.

(6) The block graph $B(G^{\dagger}_{\mathcal{T}})$ is a path (equivalently, $G^{\dagger}_{\mathcal{T}}$ does not admit a trident).

The proof consists of a modified cycle:

$$(1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (3') \Rightarrow (4) \Rightarrow (5) \Rightarrow (6) \Rightarrow (1), \tag{3}$$

The desired characterization for tangles and two agents is an immediate corollary.

Corollary 30 (Tangle Characterization for two agents) *For any tangle $\mathcal{T}$*[23] *the following are equivalent:*

(i) $\mathcal{T}$ admits a bipolar ordering.
(ii) $\mathcal{T}$ guarantees connected EF allocations for two agents with monotone continuous valuations.
(iii) $\mathcal{T}$ has no generalized gap ≥ 2 cutset of cardinality 1.

Proof of Theorem 29 *((1) $\Rightarrow$ (2))* Our main goal is to convert a bipolar numbering of graph $G^{\dagger}_{\mathcal{T}}$ into a bipolar ordering of tangle $\mathcal{T}$. Because $\mathcal{T}$ includes actual points along the tangle edges, these points need to be inserted into $G^{\dagger}_{\mathcal{T}}$'s vertex ordering. But the bipolar numbering induces an orientation of each edge $e = \{a, b\}$ of $G^{\dagger}_{\mathcal{T}}$: if a appears before b in the numbering then we orient e from a (source) to b (sink). If we ignore, momentarily, the added degree 2 subdivision vertices of $G^{\dagger}_{\mathcal{T}}$ (so that graph edges correspond to tangle edges), we can now obtain a bipolar ordering of $\mathcal{T}$ by inserting the points along any tangle edge after the source a for that edge and before the sink b, with the ordering of points within the edge determined by the ordering of real numbers in the corresponding copy of [0, 1], having previously identified 0 with the edge's source and 1 with its sink.

More precisely (and with the degree 2 subdivision vertices now accounted for), let $v_1, v_2, \ldots, v_k$ be a bipolar numbering for $G^{\dagger}_{\mathcal{T}}$. For each r with $1 < r \leq k$ let $\#r$ denote the number of edges that link v_r to some vertex appearing earlier in the bipolar numbering, let $e_{1,r}, e_{2,r}, \ldots, e_{\#(r),r}$ enumerate those edges, in any order, and let $v_{s(i,r)}$ be the vertex linked to v_r by $e_{i,r}$ (so that $s(i, r) < r$). In what follows, adding a $\star$ to the notation for a vertex (respectively, for an edge) of $G^{\dagger}_{\mathcal{T}}$ converts it into the corresponding point (respectively, corresponding half of a tangle edge) in $\mathcal{T}$.

If $v_{s(i,r)}$ corresponds to one of the singular points $v^{\star}_{s(i,r)}$ of $\mathcal{T}$ then v_r must be one of the added subdivision vertices, and edge $e_{i,r}$ in the graph corresponds to a half tangle edge $e^{\star}_{i,r}$; without loss of generality we may assume it is the $[0, \frac{1}{2}]$ half of a copy of [0, 1], with $v^{\star}_{s(i,r)}$ identified with the 0 of this copy. Then, $e^{\star}_{i,r}$ will designate the interior $(0, \frac{1}{2})$ of this half copy and $v^{\star}_r$ the point $\frac{1}{2}$ of this copy.

If, instead, $v_{s(i,r)}$ is one of the added subdivision points of $G^{\dagger}_{\mathcal{T}}$ then v_r must correspond to one of the singular points $v^{\star}_r$ of $\mathcal{T}$. Edge $e_{i,r}$ in the graph again corresponds to a half tangle edge $e^{\star}_{i,r}$, and without loss of generality we may assume it is the $[\frac{1}{2}, 1]$ half of a copy of [0, 1], with $v^{\star}_{s(i,r)}$ set equal to the $\frac{1}{2}$ of this copy. In this case, $e^{\star}_{i,r}$ will designate the interior $(\frac{1}{2}, 1)$ of this half copy.

[23] Corollary 30 clearly holds for the circle, even though Theorem 29 does not apply to it.

Our bipolar ordering for $\mathcal{T}$ has then the following template:

$$\mathbf{v_1}^\star < e^\star_{1,2} < e^\star_{2,2} < \cdots < e^\star_{\#(2),2} < \mathbf{v_2}^\star < \cdots < \mathbf{v_3}^\star < \cdots < e^\star_{\#(k),k} < \mathbf{v_k}^\star. \quad (4)$$

Our bipolar ordering $<_\mathcal{T}$ extends this template by ordering points within any $e^\star_{i,j}$ according to the ordering of real numbers within the interval $(0, \frac{1}{2})$ or $(\frac{1}{2}, 1)$, depending on which of these intervals has previously been identified with $e^\star_{i,j}$, and respecting the orientation of $e^\star_{i,j}$ as specified earlier.

To see that initial segments of form $Y = \{x \in \mathcal{T} \mid x <_\mathcal{T} a\}$ or $Y = \{x \in \mathcal{T} \mid x \leq_\mathcal{T} a\}$ are path-connected in $\mathcal{T}$, it suffices to show that each point $b \in Y$ is path-connected to the first element $v^\star_1$ in the bipolar ordering.

Case 1 If $b = v^\star_j$ for some j, use the bipolar numbering of $G^\dagger_\mathcal{T}$ to choose a path $v_{i_1}, v_{i_2}, \ldots, v_{i_r}$ from v_1 to v_j in $G^\dagger_\mathcal{T}$ satisfying $1 = i_1 < i_2 < \cdots < i_r = j$. Then, the points $v^\star_{i_1}, v^\star_{i_2}, \ldots, v^\star_{i_r}$ all lie in Y, as do all of the half tangle edges that link any $v^\star_{i_t}$ to $v^\star_{i_{t-1}}$ for $2 \leq t \leq r$. The desired path in $\mathcal{T}$ from $v^\star_1$ to b contains, as its elements, the points $v^\star_{i_1}, v^\star_{i_2}, \ldots, v^\star_{i_r}$ along with all elements of these particular half tangle edges.

Case 2 If $b \neq v^\star_j$ for any j, then b belongs to the half tangle edge $e^\star_{i,r}$ from $v^\star_j$ to $v^\star_r$ for some j, r with $1 \leq j < r \leq k$. Let P be the path in $\mathcal{T}$ from $v^\star_1$ to $v^\star_j$, as constructed in the previous case. To get the desired path for this case, add to P the points x in $e^\star_{i,r}$ satisfying $x \leq_\mathcal{T} b$.

The argument for final segments (of form $Y = \{x \in \mathcal{T} \mid x >_\mathcal{T} a\}$ or $Y = \{x \in \mathcal{T} \mid x \geq_\mathcal{T} a\}$) is similar, and we skip the details.

[(2) $\Rightarrow$ (3)] We apply "cut and choose." Given a bipolar ordering $<_\mathcal{T}$ on $\mathcal{T}$ and a point $x \in \mathcal{T}$, let L_x denote the initial segment $\{y \in \mathcal{T} \mid y <_\mathcal{T} x\}$. Consider $S = \{x \in \mathcal{T} \mid \nu_1(L_x) < \nu_1(\mathcal{T} \backslash L_x)\}$, where ν_1 denotes the monotone continuous valuation of agent 1. Note that S is an initial segment of $\mathcal{T}$ under $<_\mathcal{T}$, and $\mathcal{T} \backslash S$ is a final segment, so both of these sets are connected. Now, we will show the following claim:

Claim: For agent 1, the initial segment S and its complement are equally valuable[24], namely, $\nu_1(S) = \nu_1(\mathcal{T} \backslash S)$.

From the claim, for "cut and choose" agent 1 specifies an initial segment S [worth, to her, the same as its complement], agent 2 chooses either S or $\mathcal{T} \setminus S$ [whichever has the higher value for her], and 1 gets the remaining piece.

To prove the claim, first construct a sequence $\{s_n\}_{n \in \mathbb{N}}$ of members of S such that the corresponding sequence of sets $S_n = L_{s_n}$ has union S. This can be done because $\mathcal{T}$ has a countable dense subset (inherited from the rational numbers within the copies of [0, 1] used to build $\mathcal{T}$). Then $d(S_n, S) \to 0$ as $n \to \infty$ so by continuity $\nu_1(S_n) \to \nu_1(S)$. Similarly, $\nu_1(\mathcal{T} \backslash S_n) \to \nu_1(\mathcal{T} \backslash S)$. As $\nu_1(S_n) < \nu_1(\mathcal{T} \backslash S_n)$ holds for each n, we conclude that $\nu_1(S) \leq \nu_1(\mathcal{T} \backslash S)$.

[24] If this construction sounds somewhat overcomplex, keep in mind that there may not exist a point $x \in \mathcal{T}$ with $v_1(L_x) = v_1(\mathcal{T} \backslash L_x)$. For example, consider the friendly diamond on the top of Fig. 3, and its bipolar ordering that goes through the left edge first, then point a, then the top edge (not including d), then the bottom edge (not including a), then point d, and then the right edge. If the valuation v spreads 1 unit of value evenly over each of the four edges, then there does not exist a point x such that $v_1(L_x) = v_1(\mathcal{T} \setminus L_x)$.

For the reverse inequality, construct a sequence $\{q_n\}_{n\in\mathbb{N}}$ of members of $\mathcal{T}\backslash S$ such that the corresponding sequence of sets $Q_n = L_{q_n}$ satisfies that $d(Q_n, S) \to 0$ as $n \to \infty$. By continuity, $\nu_1(Q_n) \to \nu_1(S)$. Note that $d(\mathcal{T}\backslash Q_n, \mathcal{T}\backslash S) \to 0$ as $n \to \infty$, so $\nu_1(\mathcal{T}\backslash Q_n) \to \nu_1(\mathcal{T}\backslash S)$. As $\nu_1(Q_n) \geq \nu_1(\mathcal{T}\backslash Q_n)$ for each n, we conclude that $\nu_1(S) \geq \nu_1(\mathcal{T} \setminus S)$.

[(3) $\Rightarrow$ (3′)] This is immediate since edgewise constant valuations are monotone continuous.

[(3′) $\Rightarrow$ (4)] (via the contrapositive) Let $X = \{X_1\}$ be a generalized gap ≥ 2 cutset of cardinality 1. We will produce an identical edgewise constant valuation with at most one non-zero density such that no connected EF allocation exists for two agents. If X_1 is a type I member then X is an ordinary gap ≥ 2 cutset of cardinality 1, so the result follows immediately from the Gap ≥ 2 Lemma. Assume X_1 is type II, with the three components sharing contact points with X_1 being C_1, C_2, C_3. For each $i = 1, 2, 3$ spread 1 unit of value uniformly over one of C_i's tangle edges, with value 0 spread over all other edges. Now let $\{T_1, T_2\}$ be any allocation of $\mathcal{T}$ that assigns connected pieces to two agents. Without loss of generality, assume that T_1 includes at most one of the three contact points $s_{1,1}, s_{1,2}, s_{1,3}$. Then T_1 intersects at most one of the three components C_1, C_2, C_3, because any path from one of these components to a different one must pass through at least two of the contact points. Thus the identical value of T_1 is at most 1, leaving agent 1 envious of agent 2, whose share T_2 has value at least 2.

[(4) $\Rightarrow$ (5)] We will show the contrapositive. Assume that the block graph $B(G_{\mathcal{T}})$ is not a path. Then, $B(G_{\mathcal{T}})$ contains a vertex s with degree at least three. We have the following two cases: either s is a separating vertex belonging to three distinct blocks of $G_{\mathcal{T}}$, or s is a block containing three distinct separating vertices of $G_{\mathcal{T}}$.

First, suppose that the vertex s is a single separating vertex v of $G_{\mathcal{T}}$ belonging to three distinct blocks B_i ($i = 1, 2, 3$). It is easy to see that v corresponds to a singular point $v^\star$ of $\mathcal{T}$. In this case, we argue that $\{\{v^\star\}\}$ is a generalized gap ≥ 2 cutset of cardinality 1 for $\mathcal{T}$. Condition (a) of Definition 21 is immediately satisfied, and we need only to show that deleting $v^\star$ from $\mathcal{T}$ leaves at least three connected components. We know that each of the blocks B_i ($i = 1, 2, 3$) contains at least one distinct edge e_i, and shares v as the unique common vertex in $G_{\mathcal{T}}$. Thus, the subtangles corresponding to each B_i ($i = 1, 2, 3$) are disconnected from each other after $v^\star$ is deleted from $\mathcal{T}$.

Second, suppose that the vertex s corresponds to a block B of $G_{\mathcal{T}}$. Then, B is a connected subgraph of $G_{\mathcal{T}}$ and contains at least three distinct separating vertices s_1, s_2, s_3, each of which belongs to a distinct block B_i ($i = 1, 2, 3$). In this case, B corresponds to a subtangle X of $\mathcal{T}$ and s_1, s_2, s_3 correspond to the singular points $s_1^\star, s_2^\star, s_3^\star$ of $\mathcal{T}$, respectively. Deleting X from $\mathcal{T}$ leaves at least three disconnected components corresponding to B_1, B_2, and B_3, but may have several additional components. Let X' denote the union of X with these additional components. Then the definition of a block tree implies that $\{X'\}$ satisfies all requirements for being a generalized gap ≥ 2 cutset of cardinality 1 for $\mathcal{T}$.

[(5) $\Rightarrow$ (6)] It suffices to prove the following claim, stating that the operation of subdividing a graph preserves the underlying block graph structure.

Lemma 31 *For a multi-graph G and arbitrary edge e, if the block graph $B(G)$ of G is a path, then the block graph $B(G')$ of the graph G' that results from subdividing e (i.e., insert a single degree 2 vertex into e) is also a path.*

Proof of Lemma 31 Let G be a multi-graph and e be an edge of G. Suppose that $B(G)$ is a path. Let G' be the graph that results from subdividing e. If e itself, along with its two vertices, forms a block in G, then it is immediate that $B(G')$ continues to be a path. Thus, suppose that e belongs to a block B containing at least three vertices in G. Let B' be the subgraph of G' whose edges consist of the original edges in B together with the new subdivided edges e_1, e_2, i.e., $B' = (B \backslash \{e\}) \cup \{e_1, e_2\}$. It suffices to show that B' is a block of G'.

Recall that a graph is non-separable if and only if any two of its edges lie on a common cycle [21] (see also Theorem 5.2 of [5]). Thus, it suffices to show that

- any two edges g, h in B' lie on a common cycle (i.e., B' is non-separable).
- if two edges do not belong to a common cycle in G, they do not lie on a common cycle in G' (i.e., B' is maximal as a non-separable subgraph of G').

To see that B' is non-separable, let g', h' be a pair of edges in B'. We claim that there is a cycle C' containing g', h' in B'. Let g, h be the original edges corresponding to g', h' in B, respectively; namely, $g = e$ (respectively, $h = e$) if $g' = e_i$ (respectively, $h' = e_i$) for $i = 1, 2$ and $g = g'$ (respectively, $h = h'$) otherwise. We show that there is a cycle C containing the edges g, h in B. Indeed, if $g \neq e$ or $h \neq e$, such cycle does exist since B is non-separable; if $g = h = e$, then there is a cycle C that contains g, h in B since B has at least three vertices. Now we use C to show the existence of a cycle containing g', h' in B'. Clearly, if C does not include e, then C is a cycle containing g', h' in B'. If C includes e, we obtain a cycle C' containing g', h' in B' by replacing the single edge e with the two edges e_1 and e_2. Hence, B' is non-separable.

To see maximality, suppose toward a contradiction that two edges $g, h \neq e$ do not lie on a common cycle in G but they belong to a common cycle C' in G'. Then, C' needs to include one of the new subdivided edges e_1 and e_2. Further, the new subdivided vertex incident to both e_1 and e_2 has degree 2 in G' and thus C' includes both of the new subdivided edges e_1 and e_2. Thus, C' corresponds to a cycle C in G that contains g, h, a contradiction. A similar argument shows that if the edge e and edge $g \neq e$ do not lie on a common cycle in G, then for each e_i $(i = 1, 2)$, e_i and g do not lie on a common cycle in G'. Thus, B' remains non-separable and maximal with respect to this property in G'.

Applying the above lemma repeatedly on the edges in $G_{\mathcal{T}}$, we get that $B(G^{\dagger}_{\mathcal{T}})$ is a path if $B(G_{\mathcal{T}})$ is a path.

[(6) $\Rightarrow$ (1)] Suppose that the block graph $B(G^{\dagger}_{\mathcal{T}})$ is a path. We will show that $G^{\dagger}_{\mathcal{T}}$ admits a bipolar numbering.

Let $(B_1, s_1, B_2, s_2, \ldots, s_{k-1}, B_k)$ be the block graph $B(G^{\dagger}_{\mathcal{T}})$ where

- each B_i for $i \in [k]$ is a block of $G^{\dagger}_{\mathcal{T}}$,
- each s_i for $i \in [k-1]$ is a separating vertex of $G^{\dagger}_{\mathcal{T}}$, and
- each B_i for $i \in [k-1]$ is adjacent to s_i in $B(G^{\dagger}_{\mathcal{T}})$.

Recall that by construction of $G^{\dagger}_{\mathcal{T}}$, B_1 cannot be a single loop, since we have inserted an additional degree 2 subdivision vertex along every edge of $G_{\mathcal{T}}$. Thus, B_1 has a vertex, say s_0, distinct from s_1. Similarly, B_k has a vertex, say s_k, distinct from s_{k-1}.

It is known that for any non-separable graph H and any pair of two distinct vertices x, y of H, there is a bipolar numbering $v_1, v_2, \ldots, v_\ell$ of the vertices of H such that $v_1 = x$ and $v_\ell = y$ (see Exercise 5.3.4 in [5]). Exploiting this fact, we create a bipolar numbering P_i from s_{i-1} to s_i of each B_i for $i \in [k]$. We concatenate these $(P_1, P_2, \ldots, P_k)$ in the order of indices, by adding P_{i+1} to the end of P_i for $i \in [k-1]$. It can be easily checked that this results in a desired bipolar numbering of $G^{\dagger}_{\mathcal{T}}$. □

6 Lips guarantees EF allocations for three agents

For more than two agents, we have no general result characterizing which tangles guarantee existence of connected EF allocations. Here, we show in Theorem 32 that the Lips tangle guarantees connected envy-free allocation for three agents, implying that stringable tangles are not the only ones that guarantee EF for three agents.

Our proof elaborates Stromquist's moving knife argument [18] for a connected EF allocation of [0, 1]. Recall that in Stromquist's original version, a referee places a single sword at 0 and slowly moves it to the right towards 1, while each of the three agents continually re-position their individual knives so as to divide the portion to the right of the sword into two pieces of equal value (in their own view). Cuts will eventually be made by the sword and the median knife. The instant an agent sees that the piece L^* to the left of the sword has value at least as great as the other two (as defined by the median knife), she shouts cut and receives L^*. The remaining two pieces are now assigned to the others (with the agent whose knife is weakly to the right of the median knife getting the rightmost piece), and the resulting allocation is envy-free.

Now, consider a lips tangle, with three singular points (arranged in Fig. 1 from left to right as a, b, and c) and five edges. One approach might be to pretend that the top two a–b edges, as well as the top two b–c edges, are almost pinched together, and then apply Stromquist's moving knife protocol as if the lips tangle were a circle. However, if two cuts were made—each cutting through both pinched edges—we might lose connectivity. Thus our actual modification of Stromquist's argument is divided into three stages, changing directions of swords and knives when threatened by this sort of disconnection.

Theorem 32 *For three agents with monotone continuous valuations over the lips tangle $\mathcal{L}$, an EF allocation with connected shares always exists.*

Proof *(INITIAL STAGE)* Two swords, labeled L and R in Fig. 5a, both begin at the center of the edge $a - c$ and maintain a radial orientation; L moves in a clockwise direction (hence, initially to the left) while R moves counterclockwise, in such a way that L will reach a at the same moment that R reaches c.

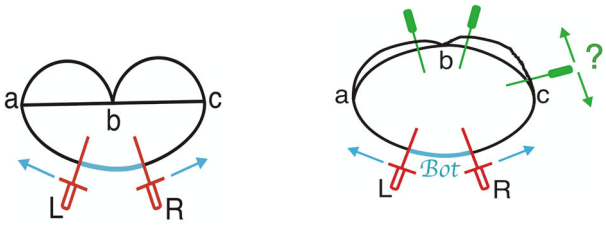

(a) Two swords move

(b) Three knives also move ... though who knows which directions they move in.

Fig. 5 Initial Stage

However, during the initial stage we assume they have not yet reached a and c. The open-ended portion of $a - c$ lying *strictly* between L and R is the *central piece* of the partition; the part of $\mathcal{L}$ outside the central piece is the *complement*.

We imagine the two $a-b$ edges as almost pinched together, with the two $b-c$ edges similarly pinched, so that $\mathcal{L}$ looks roughly like a circle, as in Fig. 5b. As the two swords rotate, each of the three agents continually adjusts the position of her knife (which is always held radially to the "circle") so as to divide the complement into two parts of equal value in her eyes. If, during the initial stage, any one of these knives should cut the complement at any location, it would divide the complement into two connected pieces. Any agent can call "cut" during the initial stage if she sees the central piece (labeled *Bot*) as exactly equal in value to the larger of the two pieces formed when the complement is cut by the median knife. If this happens then the procedure ends, with cuts made at L, R, and the median knife. The caller is awarded the central portion, and the pieces of the complement are awarded to the other two agents as they would be in Stromquist's original procedure. If the two swords reach a and c without anyone having called "cut," we move to the second stage, which depends on the position of the median knife at the end of the initial one.

[SECOND STAGE] It may happen that the two swords reach a and c without anyone having called "cut," as in Fig. 6a. In this case, the swords pause at these locations *with the points a and c still considered to be members of the complement*, while we identify the current location of the median knife. Without loss of generality assume this knife is either at b or counterclockwise of b (as in Fig. 6a). In this case, the left sword now remains fixed at a while the right one continues its relentless counter-clockwise rotation, further enlarging the central piece as in Fig. 6b. Notice that in this case, the median knife will rotate counter-clockwise (weakly, perhaps), so that during this second stage, it never passes from strictly on one side of b to strictly on the other side. It may be that some agent calls "cut" before the right sword reaches b. In this case the procedure terminates, with cuts at L, R, and the median knife. Again, it is clear that the pieces are connected,[25] and that when the pieces are assigned to agents in the standard way, the allocation is EF.

[25] One must be careful, in case the median knife is exactly at b, to assign b itself to the piece on the clockwise side of b.

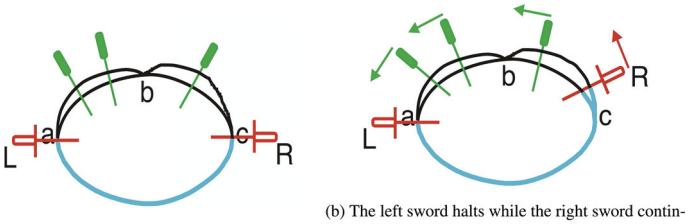

(a) The swords may reach a and c

(b) The left sword halts while the right sword continues to rotate

Fig. 6 Second Stage

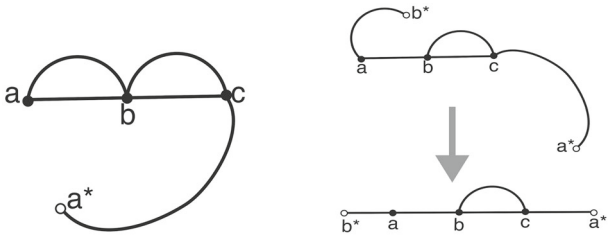

(a) The first disconnection

(b) The second disconnection

Fig. 7 Third Stage

[THIRD STAGE] It may happen that in the second stage the right sword reaches b and still no one has called "cut." At this point, the complement is the closed loop consisting of the two $a - b$ edges (including the points a and b themselves). It must be that all three agents have a strictly higher valuation for this loop than they have for the rest of $\mathcal{L}$; otherwise, someone would have called "cut" in an earlier stage. With this *loop value guarantee* in hand, we abandon the previous approach and launch a different version of Stromquist, as the third stage of the procedure. It is easier to picture this version by first introducing two disconnections—or "breaks"—that convert $\mathcal{L}$ to $\mathcal{L}^\star$ (though one could instead work with the original topology on $\mathcal{L}$ by slightly altering the following description).[26] There is no harm in switching to $\mathcal{L}^\star$, because any shares connected in $\mathcal{L}^\star$ remain connected in $\mathcal{L}$.

First disconnect at a, yielding Fig. 7a. Such a disconnection can be done in two ways, depending on which side of the break remains attached to a; Fig. 7a indicates that a remains attached to the $a - b$ loop by showing a as a solid disk on the $a - b$ loop and as an open disk labeled $a^\star$ on the $a - c$ edge. Second, disconnect at b, unrolling the $a - b$ loop, as in Fig. 7b (note that b itself remains attached to the $b - c$ loop).

We will treat $\mathcal{L}^\star$ as if it were an open line segment $(b^\star, a^\star)$ by squeezing the two $b - c$ edges almost together. We proceed to use a standard Stromquist on this open segment, with the sword starting on the left (the $b^\star$ end). We know that someone calls "cut" before the sword reaches b, thanks to the loop value guarantee that applies whenever stage three is reached. The part of $\mathcal{L}^\star$ to the right of the sword cut will consist of two connected pieces once it is cut by the median knife. □

[26] As topological spaces, $\mathcal{L}$ and $\mathcal{L}^\star$ share the same set of points, but the collection of open sets for $\mathcal{L}^\star$ is expanded from that for $\mathcal{L}$.

7 Discrete lips guarantee EF1 for three agents

We turn our attention to the positive transfer from tangles to graphs: a discrete version of the previous moving knife algorithm returns a contiguous EF1_{outer} allocation for any graph G in the topological class $\mathcal{G}(\mathcal{L})$ associated with the lips tangle $\mathcal{L}$.

Theorem 33 *Every graph in the topological class $\mathcal{G}(\mathcal{L})$ of the lips tangle $\mathcal{L}$ guarantees $EF1_{\text{outer}}$ allocations for three agents who have monotone valuations.*

In conjunction with the continuous ideas from Sect. 6, the proof uses a slight modification of the discrete Stromquist moving-knife algorithm developed in [4] for path graphs. We will refer to this modification as *the discrete moving-knife algorithm* $\mathcal{A}_{discrete}$. The proof of Theorem 33 applies $\mathcal{A}_{discrete}$ several times, in stages. Theorem 34 establishes that $\mathcal{A}_{discrete}$ achieves what is required of it, in the proof of Theorem 33. The detailed description of $\mathcal{A}_{discrete}$ and the proof of Theorem 34 appear in Appendix A.

While both continuous and discrete procedures are similar in spirit, the discrete version differs in fundamental ways from the version of the continuous procedure for Lips in Sect. 6. Indeed, we may run into a trouble if we simply follow the same strategy as in the continuous case: For example, if at the end of the first stage, the median lies at b and the two swords lie at a and c, moving one of the swords above may disconnect some piece between the swords and the median, since the swords can only move in discrete steps. Informally, instead of moving the sword over the top of the lips in the latter stages, we will ensure that the sword only hovers over the path at every stage.

Before we proceed, we present Theorem 34, stating that given a path graph, $\mathcal{A}_{discrete}$ returns an EF1_{outer} allocation of the specific form. To formally present the theorem, we need the notion of a *median lumpy tie* introduced in [4], which serves as the discrete analogue of the middle knife. Let $P = (v_1, v_2, \ldots, v_m)$ be an enumeration of vertices. We define $P(v_s, v_t)$ to be $(v_s, v_{s+1}, \ldots, v_t)$ when $s \leq t$ and $\emptyset$ when $s > t$. We define $L(v_r) = P(v_1, v_{r-1})$ to be the sub-enumeration of vertices strictly left of v_r and $R(v_r) = P(v_{r+1}, v_m)$ to be the sub-enumeration of vertices strictly right of v_r. For an agent i, we say that v_r $(1 \leq r \leq m)$ is a *lumpy tie* over P for i if r is an index such that

$$v_i(L(v_r) \cup \{v_r\}) \geq v_i(R(v_r)) \quad \text{and} \quad v_i(R(v_r) \cup \{v_r\}) \geq v_i(L(v_r)).$$

A lumpy tie always exists. For example, a vertex v_r with the smallest index r such that $v_i(L(v_r) \cup \{v_r\}) \geq v_i(R(v_r))$ is a lumpy tie [4].

Suppose that there are three agents, i, j, k. A vertex v_r is a *median lumpy tie* over the sequence P of vertices if there exist vertices v_a, v_b, v_c that are lumpy ties for i, j, k, respectively, such that index r is the median of $\{a, b, c\}$. Note that the above definition of a lumpy tie is slightly different from that in [4] in that lumpy ties for an agent may not be uniquely determined. When agents' valuations are monotone, multiple lumpy tie vertices will appear consecutively.

Theorem 34 uses a different version of EF1: For an allocation $A = \{A_i\}_{1 \leq i \leq n}$ and set X of items, we say that agent i *does not envy* j *up to* X if the envy of i towards j can be eliminated by hiding items in X, i.e., $v_i(A_i) \geq v_i(A_j \setminus X)$. Allocation A is *envy-free up to* X if for any pair of agents $i, j \in [n]$, i does not envy j up to X. The discrete

moving knife algorithm $\mathcal{A}_{discrete}$ returns such an allocation A with $|X \cap A_i| \leq 1$ for each agent $i = 1, 2, \ldots, n$.

Theorem 34 (proved in Appendix 1) *Assume that the discrete moving-knife algorithm $\mathcal{A}_{discrete}$ for three agents with monotone valuations is applied to the disjoint union of a finite set I of vertices— called the* initial endowment of the left bundle*—and an enumeration $P(v_1, v_m) = \{v_1, v_2, \ldots, v_m\}$ of additional vertices, along with a pair ℓ, r of indices with initial values $\ell = 0$, $r = r_0$ such that v_{r_0} is a median lumpy tie over $P(v_1, v_m)$ and no agent strictly prefers I to both $P(v_1, v_{r-1})$ and $P(v_{r+1}, v_m)$. Then the algorithm increments r at certain stages and ℓ at certain other stages, and returns an EF1 allocation A of $I \cup \{v_1, v_2, \ldots, v_m\}$ that has the following property:*

(i) *A is EF up to $\{v_\ell, v_r\}$ and partitions $I \cup P$ into either $(I \cup P(v_1, v_\ell), P(v_{\ell+1}, v_r), P(v_{r+1}, v_m))$ or $(I \cup P(v_1, v_\ell), P(v_{\ell+1}, v_{r-1}), P(v_r, v_m))$; or*
(ii) *A is EF up to $\{v_{\ell+1}, v_r\}$, and partitions $I \cup P$ into $(I \cup P(v_1, v_\ell), P(v_{\ell+1}, v_{r-1}), P(v_r, v_m))$.*

The algorithm terminates before or at the time when ℓ becomes m. If the algorithm does not terminate just after increasing ℓ by one, no agent weakly prefers $I \cup P(v_1, v_\ell)$ to both $P(v_{\ell+1}, v_{r-1})$ and $P(v_{r+1}, v_m)$. Every time just after the algorithm increases ℓ by one, v_r is a median lumpy tie over $P(v_{\ell+1}, v_m)$. Further, if it returns an allocation just after increasing ℓ by one, the resulting allocation is of the form (i). *Otherwise, the resulting allocation is of the form* (ii).

Notice that in case of (i), the hidden items v_ℓ and v_r belong to the left bundle $I \cup P(v_1, v_\ell)$ and to either the middle bundle $P(v_{\ell+1}, v_r)$, or the right bundle $P(v_r, v_m)$; in case of (ii), such items belong to the middle bundle $P(v_{\ell+1}, v_{r-1})$ and right bundle $P(v_r, v_m)$. It is immediate to see that by the above theorem, the resulting allocation of $\mathcal{A}_{discrete}$ is contiguous and EF1_{outer} if $I = \emptyset$ and the given enumeration P is a path.

Next, we will repeatedly apply Theorem 34 to show that any discrete version of the lips tangle guarantees EF1_{outer} for three agents. Our algorithm recursively identifies a decomposition (X, H') of the current subgraph H, called a *handle decomposition*, and applies the discrete moving knife algorithm over H with the initial endowment I being the vertices in $V(G) \setminus V(H)$ over which the sword has already hovered. Formally, we say that a graph H admits a *handle decomposition* (X, H') if H' is a subgraph of H, X is a path in H where no vertex of X lie in H', and $V(H) = V(H') \cup V(X)$; we call X a *handle* of H'. Here, we denote by $V(H)$ the vertex set of a graph H.

The following lemma states that the algorithm results in EF1_{outer} as long as the enumeration of vertices is "almost" a path and the sword stops "soon enough." See Fig. 8 for an illustration. In what follows, for two enumerations X and Y of vertices, $X \cdot Y$ will denote the concatenation of X and Y.

Lemma 35 *Suppose that a subgraph H of graph G admits a handle decomposition (X, H') for which H' admits a bipolar numbering Y over H', where the end vertex of X is adjacent to the first vertex of Y. Let $P = X \cdot Y = (v_1, v_2, \ldots, v_m)$. Suppose that $\mathcal{A}_{discrete}$ applies to a set of items $I = V(G) \setminus V(H)$ as an initial endowment, ordering P, and v_r where*

(i) I is contiguous, and if $I \neq \emptyset$, some vertex of I is adjacent to v_1; and

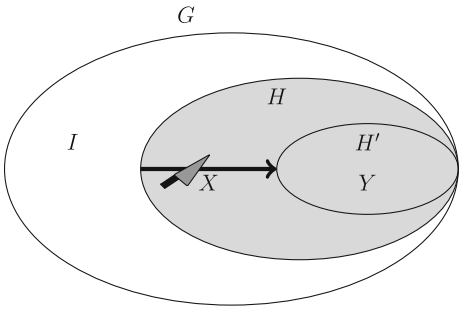

Fig. 8 Illustration of Lemma 35. The initial endowment I corresponds to the vertices in G outside H. The subgraph H admits a handle decomposition (X, H') with Y being a bipolar numbering over H'

(ii) v_r *is a median lumpy tie over* P*; and*
(iii) no agent strictly prefers I *to both* $P(v_1, v_{r-1})$ *and* $P(v_{r+1}, v_m)$.

If the algorithm stops before or at the time when v_ℓ *becomes the last vertex of* X*, then the resulting allocation is contiguous* $EF1_{\text{outer}}$*. Otherwise, no agent weakly prefers* $I \cup V(X)$ *to both* $P(v_{\ell^*+1}, v_{r^*-1})$ *and* $P(v_{r^*+1}, v_m)$ *where* ℓ^* *and* r^* *are the values of indices* ℓ *and* r*, respectively, at the time when* v_ℓ *becomes the last vertex of* X*, and so* v_{r^*} *is a median lumpy tie over* $P(v_{\ell^*+1}, v_m)$.

Proof It is not difficult to see that by definition, the enumeration P that results from concatenating X and Y is a bipolar numbering over the original subgraph H. We let $X = (u_1, u_2, \ldots, u_k)$. If the algorithm $\mathcal{A}_{discrete}$ stops before or during the time when $v_\ell = u_{k-1}$, the resulting allocation is contiguous and EF up to $\{v_\ell, v_r\}$, or $\{v_{\ell+1}, v_r\}$, by Theorem 34. The resulting bundles of the form:

- $(I \cup P(v_1, v_\ell), P(v_{\ell+1}, v_r), P(v_{r+1}, v_m))$; or
- $(I \cup P(v_1, v_\ell), P(v_{\ell+1}, v_{r-1}), P(v_r, v_m))$

remain contiguous even after removing $\{v_\ell, v_r\}$, or $\{v_{\ell+1}, v_r\}$ since X is a path and Y is a bipolar numbering over H'. The same argument applies to the case when the algorithm $\mathcal{A}_{discrete}$ stops at the time when v_ℓ becomes u_k; the difference is that the resulting allocation is EF up to $\{v_\ell, v_r\}$ so that the middle bundle of the form $P(v_{\ell+1}, v_{r-1})$ or $P(v_{\ell+1}, v_r)$ is guaranteed to be contiguous even after removing these items. The second statement immediately follows from Theorem 34. □

Applying Lemma 35 in three stages (each initiated if the algorithm fails to halt in the previous), we will prove that contiguous $EF1_{outer}$ allocations always exist for discrete versions of lips.

Proof of Theorem 33 Let G be a graph that belongs to the topological class $\mathcal{G}(\mathcal{L})$ of the lips tangle. Define vertices $a, b, c, a_1, b_1, b_2, b_3, c_1, c_2$ as in Fig. 9.

Note that when some of the vertices a_1, b_1, b_2, c_1, and c_2 do not exist, the problem reduces to that over a path and hence Theorem 3 applies. For example, when b_1 does not exist, we can remove the top left edge between a and b and the reduced graph H is a graph in the topological class $\mathcal{G}(\Theta)$ of the theta tangle Θ, where H admits a Hamiltonian path (for example, a path from the middle left neighbor b_3 of b through b to c_1). It is clear that any $EF1_{outer}$ connected allocation of H is an $EF1_{outer}$ connected

Fig. 9 A lips graph $G \in \mathcal{G}(\mathcal{L})$

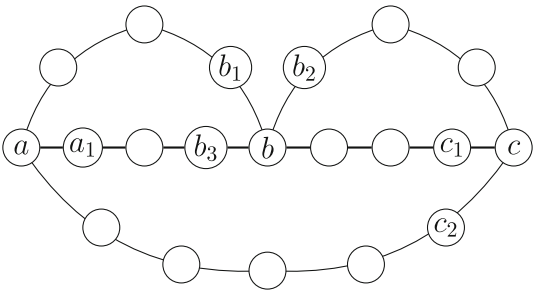

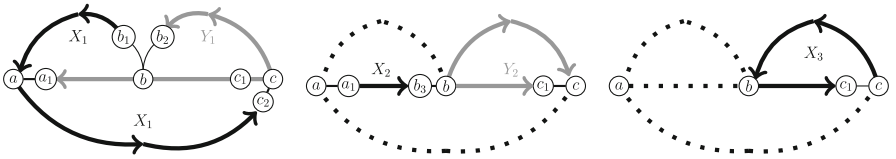

(a) Initial stage over $X_1 \cdot Y_1$ (b) Second stage over $X_2 \cdot Y_2$ (c) Third stage over $X_3 \cdot Y_3$

Fig. 10 Discrete moving-knife algorithm. The black arrows correspond to the handles X_i and the gray arrows correspond to the bipolar orderings Y_i over G_{i+1}. The dotted regions correspond to the set of vertices in I at the beginning of each stage

allocation of G. A similar argument holds when one of a_1, b_2, and c_1 does not exist. Thus, in the sequel we assume that a_1, b_1, b_2, and c_1 exist. If there is no vertex in the middle left path between a_1 and b, we let $b_3 = a_1$.

Now, we create the following sequence of subgraphs G_{i+1} together with their handles X_i for $i = 1, 2, 3$, where the first vertex of each X_i is adjacent to some vertex of an earlier X_j ($j < i$). See Fig. 10. We let $G_1 = G$.

- Let X_1 be a path that starts from b_1, goes through the top-left path from b_1 to a, then through the bottom path from a to c_2. Let G_2 be the subgraph of G induced by the other vertices. Y_1 is a bipolar numbering over G_2 that starts from c, goes through the top-right path from c to b_2, then through the middle path from c_1 to the right neighbor a_1 of a.
- Let X_2 be the middle left path from a_1 to the vertex b_3. Let G_3 be the top-right cycle of the graph G. Y_2 corresponds to the reverse sub-enumeration of Y_1 restricted to the vertices in G_3, meaning that Y_2 first goes through the middle right path from b to c_1 and then through the top-right path from b_2 to c.
- Let X_3 be a path over the vertices in the top-right cycle that starts from c and ends with the middle left neighbor c_1 of c. Similarly, X_3' is a path over the vertices in the top-right cycle that starts from b and ends with the top-right neighbor b_2 of b. Let Y_3 be the empty path and G_4 be the empty graph.

The subgraph G_i of G admits a handle decomposition (X_i, G_{i+1}) for $i = 1, 2, 3$. Further, the first vertex of Y_i is adjacent to the last vertex of X_i for $i = 1, 2$.

The following procedure iteratively applies the discrete moving-knife algorithm $\mathcal{A}_{discrete}$ while keeping the bundle that is left of the sword contiguous. At each stage, I denotes the initial endowment and P denotes the sequence of vertices to which $\mathcal{A}_{discrete}$ is applied.

[INITIAL STAGE] Initialize the initial endowment $I = \emptyset$, the sequence $P = X_1 \cdot Y_1$, and the indices $\ell = 0$, and r such that v_r is any median lumpy tie over P. Apply the discrete moving-knife algorithm $\mathcal{A}_{discrete}$ to I, P, ℓ, and r. If the algorithm terminates before or at the time when v_ℓ becomes the last vertex of X_1, return the resulting allocation. If not, go to the second stage.

[SECOND STAGE] Initialize $I = V(X_1)$, $P = X_2 \cdot Y_2$, $\ell = 0$, and v_r to the same vertex as the one in the last step of the initial stage. Apply the discrete moving-knife algorithm $\mathcal{A}_{discrete}$ to I, P, ℓ, and r. If the algorithm terminates before or at the time when v_ℓ becomes the last vertex of X_2, return the resulting allocation. If not, go to the third stage.

[THIRD STAGE] Initialize $I = V(X_1) \cup V(X_2)$. Initialize r so that v_r remains the same as the one in the last step of the second stage.

(3–1) If v_r appears strictly after c_1 in the numbering Y_2 (equivalently, v_r is a top-right vertex), then apply the discrete moving-knife algorithm $\mathcal{A}_{discrete}$ to the initial endowment I, the sequence $P = X_3$ from c to c_1, $\ell = 0$, and r.

(3–2) If v_r appears strictly before or at c_1 in the numbering Y_2 (equivalently, v_r is a middle right vertex), then apply the discrete moving-knife algorithm $\mathcal{A}_{discrete}$ to the initial endowment I, the sequence $P = X_3'$ from b to b_2, $\ell = 0$, and r.

The algorithm for the third stage will terminate before or at the time when v_ℓ becomes the last vertex of each of the sequences X_3 and X_3' by Theorem 34. Note that the last case distinction is necessary so as to ensure, as shown below, that the sequence P of the third stage is a path for which v_r is a median lumpy tie.

Now, we show that at the beginning of each stage, the three conditions in Lemma 35 are satisfied:

(i) I is contiguous, and if $I \neq \emptyset$, some vertex of I is adjacent to v_1; and
(ii) v_r is a median lumpy tie over P; and
(iii) no agent strictly prefers I to both $P(v_1, v_{r-1})$ and $P(v_{r+1}, v_m)$,

which, together with Lemma 35, then implies that the resulting allocation is contiguous $\mathrm{EF1}_{outer}$.

Consider the beginning of the initial stage. The conditions (i) and (ii) hold, according to the construction of I, P, and v_r. The condition (iii) holds as well, by monotonicity and as $I = \emptyset$.

Consider the beginning of the second stage. The condition (i) holds, according to the construction of I and P. The conditions (ii) and (iii) also hold, by the fact that that the algorithm failed to terminate during the previous stage, and by Lemma 35.

Consider the beginning of the third stage. Again, (i) holds, according to the construction of I and P. Indeed, in case (3–1), $v_1 = c$ and the vertex just below c (i.e., c_2) is adjacent to c and belongs to I. In case (3–2), $v_1 = b$ and the left neighbor of b (i.e., b_3) belongs to I. Regarding (ii), at the beginning of the third stage, v_r is a median lumpy tie over Y_2 by Lemma 35. If v_r appears strictly after c_1 in the numbering Y_2, then it is easy to see that v_r is a median lumpy tie over X_3, since at least two agents weakly prefer the top-right path from c to v_r, to the other part in G_3, and at least two agents weakly prefer the union of the middle path from b to c_1 and the top right path

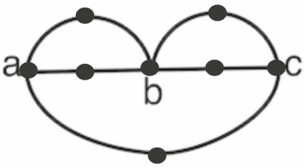

Fig. 11 The graph $L^\star$ that results from subdividing $G_{\mathcal{L}}$

from b to v_r, to the rest in G_3. Similarly, if v_r appears strictly before or at c_1 in the numbering Y_2, v_r is a median lumpy tie over X_3'. Condition (iii) holds as in the second stage. □

8 Negative transfer from tangles to graphs

The Tangle Boundedness Theorem of Sect. 3 told us that every non-stringable tangle has an agent bound n_0; connected envy-free allocations sometimes fail to exist, whenever the number of agents exceeds n_0 (even for agents with identical additive valuations). Our goal here is to prove the Negative Transfer Principle, which carries that result over to a corresponding one for contiguous EFk allocations of graphs.

Recall that any tangle $\mathcal{T}$ can be thought of as a drawing of a graph $G_{\mathcal{T}}$, with the singular points and tangle edges of $\mathcal{T}$ becoming the vertices and edges, respectively, of $G_{\mathcal{T}}$; the topological class of $\mathcal{T}$ then arises from from $G_{\mathcal{T}}$ by inserting arbitrarily many subdivision vertices along the edges of $G_{\mathcal{T}}$. We were led to this definition by considering the statement of problem [♣] in Sect. 1 – how could we add structure to the notion of "infinite class of graphs" in a way that seems appropriate for the setting and also allows progress on the problem?

Note that any stringable tangle has a topological class consisting entirely of traceable graphs; the injective map from $(0, 1)$ to $\mathcal{T}$ (from Definition 10) induces the Hamiltonian path for any graph in $\mathcal{G}(\mathcal{T})$. What about the topological class of a non-stringable tangle?

Example 1 Consider the class $\mathcal{G}(\mathcal{L})$ of the lips tangle. It contains some traceable graphs; in fact, $G_{\mathcal{L}}$ itself is traceable. However, if we insert a single subdivision vertex along each edge of $G_{\mathcal{L}}$ the resulting graph $L^\star$ is no longer traceable (see Fig. 11), and adding additional subdivision vertices can never reverse this. Yet, even with thousands of such additional vertices inserted, $G_{\mathcal{L}}$ would remain traceable if there were a single original edge that received none of the added vertices.

Part of this result generalizes to any non-stringable tangle $\mathcal{T}$; once at least one subdivision vertex has been added to each edge of $G_{\mathcal{T}}$, the resulting graph is not traceable. (If it were traceable then the Hamiltonian path would induce the onto function f: $[0, 1] \to \mathcal{T}$ witnessing $\mathcal{T}$'s stringability.[27]) One might say that $G_{\mathcal{T}}$ becomes non-traceable once "enough" subdivision vertices are inserted, but as the above example shows, "enough" cannot be expressed in terms of a single number alone.

Thus stringable tangles have traceable topological graph classes, and non-stringable tangles have non-traceable topological graph classes, but the latter assertion requires

[27] If an edge of $G_{\mathcal{T}}$ had no subdivision vertices added, then the induced function f could avoid covering the corresponding tangle edge.

an “enough" qualifier. The Negative Transfer Principle itself relies on similar notion of “enough": for any non-stringable tangle $\mathcal{T}$, there is a bound n_0 such that once enough subdivision vertices are inserted along edges of $G_{\mathcal{T}}$ the resulting graph no longer guarantees connected EF1 allocations for more than n_0 agents; once even more are inserted, EF2 can fail for more than n_0 etc. up to EFk for any fixed k.

Example 1 (continued) The graph $L^\star$ already has enough subdivision vertices to defeat the guaranteed existence of contiguous EF1 allocations for four agents. To see this, consider identical, additive valuations for the four agents that assign value 1 to each vertex. Some *deprived* agent does not receive any of the three vertices labeled a, b, and c, so their connected share has at most one vertex. The remaining 7 vertices are divided among the other three agents, so at least one agent has at least 3 vertices, leaving the deprived agent envying her by at least two vertices. If we instead added 28 subdivision vertices to each edge of $G_{\mathcal{L}}$, a similar argument shows that connected EF10 allocations can fail to exist.

In what follows we formalize “enough" in a way that further strengthens our precise statement of Negative Transfer. There are a variety of contexts in which it is natural to consider a property that holds for “almost all" members of a base set S; in Measure Theory a measure 1 subset X of the closed interval $[0, 1]$ is considered to be *almost all* of $[0, 1]$; in Number Theory a subset of the natural numbers $\mathbb{N}$ might be considered to be *almost all* of $\mathbb{N}$ if its complement in $\mathbb{N}$ is finite. A different notion of *almost all* for $\mathbb{N}$ is “has a limiting density of 1." These examples all share certain properties that have come to be viewed as necessary and sufficient to justify use of the phrase “almost all of S," and these properties have been formalized in the following definition:

Definition 36 (Non-principal Filter) A collection $\mathfrak{F}$ of subsets of a set S is a *non-principal filter over* S if

- $S \in \mathfrak{F}$ and no finite or empty set is in $\mathfrak{F}$,
- $Y \in \mathfrak{F}$ whenever $X \subseteq Y \subseteq S$ and $X \in \mathfrak{F}$,
- $X \cap Y \in \mathfrak{F}$ whenever $X \in \mathfrak{F}$ and $Y \in \mathfrak{F}$, and
- $S \setminus X \notin \mathfrak{F}$, whenever $X \in \mathfrak{F}$.[28]

In particular, the measure 1 sets form a non-principal filter over $[0, 1]$, the co-finite sets form a non-principal filter over $\mathbb{N}$, and the sets of limiting density 1 form a different non-principal filter over $\mathbb{N}$ (which contains the previous filter).

We next define a mathematically natural non-principal filter $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$ over the set of graphs in $\mathcal{G}(\mathcal{T})$, with the property that when $\mathcal{T}$ is not stringable, the set of graphs in $\mathcal{G}(\mathcal{T})$ that do not guarantee connected EFk allocations to n agents belongs to $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$ whenever n exceeds the agent bound arising from $\mathcal{T}$ via Lemma 1. When referring to membership in the particular filter $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$ we will use “eventually all" in place of the more generic “almost all," by way of acknowledging that $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$'s definition is based on final segments in the *subdivision pre-ordering*. The precise definitions now follow.

Definition 37 (The Subdivision Pre-ordering on Graphs) For graphs H and G, $H \succeq G$ holds if H is isomorphic to some subdivision of G.[29]

[28] This last condition is implied by the others, and is included for emphasis.

[29] If H and G are isomorphic graphs that are not identical, then $H \succeq G$ and $G \succeq H$ both hold, which is why $\succeq$ is a pre-order and not a partial order. Footnote 30 has more to say on this.

Proposition 38 *When restricted to any topological class $\mathcal{G}(\mathcal{T})$ the subdivision order is directed*[30]*—that is, for any $H_1, H_2 \in \mathcal{G}(\mathcal{T})$, there exists a $K \in \mathcal{G}(\mathcal{T})$ with $K \succeq H_1$ and $K \succeq H_2$.*

Proof To construct K, add more vertices to each edge than were added by H_1 or H_2. □

The notions introduced in the next definition are only useful when the underlying partial order is directed:

Definition 39 (Final Segments) Let $H \in \mathcal{G}(\mathcal{T})$ be any member of a topological class. Then $\widehat{\mathbb{H}}$ denotes the set of all $K \in \mathcal{G}(\mathcal{T})$ such that $K \succeq H$; any set $\widehat{\mathbb{H}}$ that can be formed in this way from an $H \in \mathcal{G}(\mathcal{T})$ is called a *final segment* of $\mathcal{G}(\mathcal{T})$. A subset $\mathcal{F} \subseteq \mathcal{G}(\mathcal{T})$ of a topological class is *final in* $\mathcal{G}(\mathcal{T})$ if $\widehat{\mathbb{H}} \subseteq \mathcal{F}$ holds for some final segment $\widehat{\mathbb{H}}$ of $\mathcal{G}(\mathcal{T})$, and $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$ denotes the collection of all subsets of $\mathcal{G}(\mathcal{T})$ that are final in $\mathcal{G}(\mathcal{T})$.

Proposition 40 (Final Segment Filter) *Let $\mathcal{G}(\mathcal{T})$ be a topological class of graphs. Then $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$ is a non-principal filter over $\mathcal{G}(\mathcal{T})$.*

Proof The first two properties of a non-principal filter follow immediately. For the third, assume that $\mathcal{F}_1, \mathcal{F}_2 \in \mathfrak{F}_{\mathcal{G}(\mathcal{T})}$. Choose H_1, H_2 with $\widehat{\mathbb{H}_1} \subseteq \mathcal{F}_1$ and $\widehat{\mathbb{H}_2} \subseteq \mathcal{F}_2$. As $\succeq$ is directed on $\mathcal{G}(\mathcal{T})$, we can find a $K \in \mathcal{G}(\mathcal{T})$ with $K \succeq H_1$ and $K \succeq H_2$. Then $\widehat{\mathbb{K}} \subseteq \mathcal{F}_1 \cap \mathcal{F}_2$, whence $\mathcal{F}_1 \cap \mathcal{F}_2 \in \mathfrak{F}_{\mathcal{G}(\mathcal{T})}$, as desired. □

Definition 41 (Eventually All) Suppose that **P** is some property of graphs, and that the set of graphs in the topological class $\mathcal{G}(\mathcal{T})$ that have property **P** belongs to the final segment filter $\mathfrak{F}_{\mathcal{G}(\mathcal{T})}$. In this case, we will say that *eventually all of the members of $\mathcal{G}(\mathcal{T})$ are* **P**.

Example 1 (conclusion) For example, "property **P** holds for eventually all graphs in $\mathcal{G}(\mathcal{L})$" means that there exists some subdivision $L^\dagger$ of $G_\mathcal{L}$ such that **P** holds for all graphs isomorphic to a subdivision of $L^\dagger$.

It follows that if eventually all members of $\mathcal{G}(\mathcal{T})$ are **P** and eventually all members of $\mathcal{G}(\mathcal{T})$ are **Q**, then eventually all members of $\mathcal{G}(\mathcal{T})$ are (simultaneously) **P** *and* **Q**. Similarly, if eventually all members of $\mathcal{G}(\mathcal{T})$ are **P** it cannot also be the case that eventually all members of $\mathcal{G}(\mathcal{T})$ are $\neg$**P**. More generally, the collection of properties **P** that hold for eventually all members of a topological class form a *logically consistent* and *deductively closed "final segment theory"* (which will not be logically *complete*).

Proposition 42 *Let $\mathcal{T}$ be a tangle. Then:*

1. *eventually all members of* any *topological class $\mathcal{G}(\mathcal{T})$ are simple graphs,*
2. *if $\mathcal{T}$ is stringable, then all graphs in $\mathcal{G}(\mathcal{T})$ are traceable, and*

[30] The fact that the order is directed when restricted to a topological class relies on including the phrase "is isomorphic to" in Definition 37; alternately, we could have defined a topological class initially to be a set of isomorphism classes. Ordinarily one sweeps these sorts of subtleties under the rug, but here there are some interesting consequences. For example, one can show that there can be more than one *minimal* K with $K \succeq H_1$ and $K \succeq H_2$; in particular, the directed set may fail to be a lattice.

3. if $\mathcal{T}$ is non-stringable, then eventually all graphs in $\mathcal{G}(\mathcal{T})$ are non-traceable.

Proof For part 1, insert at least two subdivision vertices along each original loop edge (edge from a vertex back to itself) and insert at least one subdivision vertex along each original non-loop edge. This guarantees that multiple edges joining the same two vertices no longer exist, while loop edges that join a vertex to itself also are banished. The proofs of parts 2 and 3 were sketched informally in the paragraphs just before and after Example 1. □

With our notion of *eventually all* in hand, we are finally ready to present our main tool for converting results about tangles to results about graphs. In particular, the Negative Transfer Principle converts a negative result for a tangle into a negative result for eventually all graphs in the corresponding topological class.

Theorem 43 (Negative Transfer Principle) *Suppose the tangle $\mathcal{T}$ does not guarantee envy-free, connected allocations for arbitrarily many agents with monotone continuous valuations. Then $\mathcal{T}$'s topological class $\mathcal{G}(\mathcal{T})$ has an agent bound n_0 of the following kind: for each positive integer k, and each $n \geq n_0$,* eventually all *graphs in the class fail to guarantee contiguous EFk allocations for n agents, even in the special case wherein agent valuations are identical, non-negative, and additive.*

Proof Suppose there exists a greatest integer $n^\dagger$ such that the tangle $\mathcal{T}$ guarantees envy-free, connected allocations for $n^\dagger$ or fewer agents with monotone continuous valuations. Then $\mathcal{T}$ must be non-stringable, and thus has a finite gap threshold t. Set $n_0 = t$. Then for any $n > n_0$ our proof of Lemma 13 provides an identical additive valuation ν over $\mathcal{T}$ for all n agents, for which there is a strictly positive lower bound $b = \frac{1}{n-1}$ on envy; this means that for every allocation $\{T_i\}_{1 \leq i \leq n}$ of $\mathcal{T}$ into connected shares there exist agents i and j such that $\nu(T_j) \geq b + \nu(T_i)$, so that agent i envies j's share "by at least b."

Fix any integer $k \geq 1$. Let $\{e_r\}_{1 \leq r \leq \epsilon(\mathcal{T})}$ enumerate $\mathcal{T}$'s edges. For each edge, let J_r be an integer that is both greater than k and sufficiently large to make $\frac{\nu(e_r)}{J_r} < \frac{b}{k}$. Construct a graph H in $\mathcal{G}(\mathcal{T})$ that has J_r subdivision vertices on the edge of $G_\mathcal{T}$ corresponding to e_r, for each r. Consider the following identical valuations for agents over the vertices of H: assign value $\frac{\nu(e_r)}{J_r}$ to each subdivision vertex along $[0, 1]_r$ and value 0 to each vertex of H whose degree is not 2. Note that this discrete distribution mimics the continuous distribution ν that each agent had over $\mathcal{T}$, replacing the identical value $\nu(e_r)$ that each agent placed on e_r with a large number of vertices of identical value, placed along this same edge, whose values sum to the original value $\nu(e_r)$. This number J_r of vertices is high, so that the individual value of any vertex is low—so low that the total value of any k vertices to any agent is less than b.

We now claim there is no contiguous EFk allocation of H for these additive valuations of vertices. If such an allocation $\{A_i\}_{1 \leq i \leq n}$ existed, we could use it to obtain a connected allocation $\{T_i\}_{1 \leq i \leq n}$ of $\mathcal{T}$ such that for each j the identical value $\nu(T_j)$ to each agent i is the same as the value they placed on A_j. But then the greatest envy any agent has for another's share is at most the equivalent of the total value of k vertices of H, which is less than b, a contradiction. The same argument works for any graph G in the final segment $\widehat{\mathbb{H}}$ (for example by giving value 0 to any additional subdivision vertices in G). □

Remark 44 Because the proof factors through Lemma 1, it yields a bound n_0 for the class $\mathcal{G}(\mathcal{T})$ of graphs that can be higher than the bound $n^\dagger$ for the tangle. It is not clear whether the argument could be extended to guarantee a bound for $\mathcal{G}(\mathcal{T})$ as low as that for $\mathcal{T}$—we could no longer assume that valuations witnessing the tangle bound were identical, or additive. In that case, one might have a higher bound $n^{\dagger\dagger}$ for $\mathcal{T}$ that applies only to identical, additive, and non-atomic valuations. A referee points out that we could then argue that this same upper bound $n^{\dagger\dagger}$ applies to $\mathcal{G}(\mathcal{T})$, for agents with additive monotone valuations. Such an extension would no longer factor through Lemma 1, so it would seem to require an additional compactness argument to obtain a lower bound $b > 0$ on the minimal amount of envy for the tangle.

Our main negative result on EFk fair division of non-traceable graphs now follows as a consequence of the Tangle Boundedness Theorem and the Negative Transfer Principle.

Theorem 45 *Let $\mathcal{G}(\mathcal{T})$ be any topological class containing at least one non-traceable graph. Then there exists a positive integer n_0 such that for each $n > n_0$, and $k \geq 1$, eventually all graphs in $\mathcal{G}(\mathcal{T})$ fail to guarantee EFk for n agents with monotone valuations, even if those valuations are identical and additive.*

Proof If $\mathcal{G}(\mathcal{T})$ contains at least one non-traceable graph then by Proposition 42, $\mathcal{T}$ is not a stringable tangle, whence Theorem 11 applies, providing an integer n_0 such that $\mathcal{T}$ fails to guarantee connected EF allocations for n agents whenever $n > n_0$. Theorem 45 now follows immediately from the Negative Transfer Principle (Theorem 43). □

Note that Theorem 45 directly addresses the question of whether the restriction to traceable graphs was actually needed for the positive results about three or more agents in [4]. The answer is affirmative for results that hold for arbitrarily many agents and apply to topological classes.

9 Conclusions: envy-freeness for tangles and graphs

We have provided partial resolutions to problem [♣], restated here:

[♣] *"For the case of three or more agents, it is a challenging open problem to find an infinite class of non-traceable graphs that guarantee EF1."*

Does such a class exist? We have provided two opposing answers, depending on how one interprets "three or more agents." Restating these answers informally (in particular, with "class" interpreted as "topological class"):

- *No*: for any $k \geq 1$, the only infinite classes that guarantee EFk for an arbitrary number of agents are classes of traceable graphs.
- *Yes*: such an infinite class of non-traceable graphs exists, for the case of 3 agents.

The precisely phrased versions of these answers appeared earlier in the form of Theorems 43 and 33. Note that our negative result (Theorem 45) holds even with very simple valuations (agents have identical, additive valuations), while the positive

result (Theorem 33) holds for a very broad class of valuations. Note, as well (and as earlier pointed out) that eventually all graphs in the topological class $\mathcal{G}(\mathcal{L})$ of lips are non-traceable, providing the desired infinite class of non-traceable graphs.

For *positive* transfers from the continuous realm to the discrete one, such as Theorem 33, we know of no wholesale method analogous to Theorem 43. Instead, the process operates at the retail level; one takes a proof that connected EF allocations are always possible (for [0, 1], for example) and elaborates the argument to guarantee contiguous EFk allocations for a corresponding topological class of graphs (the path graphs, for example). The moving knife argument of Stromquist [18] guarantees connected envy-free allocations of [0, 1] for three agents, for example, and in [4] this argument was adapted to the discrete setting, to guarantee contiguous EF1_{outer} allocations of any path graph for three agents.[31]

Important questions left open by our investigations here include:

1. Does there exist a single (finite) graph that is not traceable, yet guarantees contiguous EF1_{outer} allocations for arbitrarily many agents with "reasonable" valuations?
2. Can the methods used for the lips tangle be extended to characterize the class of tangles that guarantee connected and envy-free allocations for three agents?
3. Does some generalization of the gap ≥ 2 cutset notion provide a plausible characterization of the agent bound for an arbitrary tangle?

More generally, our work establishes an intriguing connection between the fields of cake-cutting and graph theory, so that exploring further connection would be a promising research direction. In particular, the notions we develop in this paper, such as gap ≥ 2 cutsets and their generalizations, are seemingly new and may be interesting on their own from a graph-theoretic viewpoint. For example, would an approach similar to that taken in the Tangle Boundedness Theorem have any success at analyzing the class of traceable graphs? What is the relationship between our results and any conditions known to be necessary for traceability in graph theory?

Acknowledgements We were fortunate to have referees who were particularly thorough; their comments led to significant improvements in the manuscript. We thank Jiehua Chen, Dominik Peters, Manuel Sorge, and Walter Stromquist for discussions that contributed substantially to our understanding of the ideas presented here.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

[31] As a second example of this retail process, [4] also took the Sperner's lemma proof that connected envy-free allocations of [0, 1] for n agents are guaranteed (stated earlier in this introduction as Theorem 1) as the inspiration for a proof that contiguous EF2_{outer} allocations of any path graph for n agents are guaranteed. None of these adaptations were routine; new ideas and additional details were required.

Appendix: Proof of Theorem 34

Here we define the algorithm $\mathcal{A}_{discrete}$ and prove Theorem 34 of Sect. 7, which states that the algorithm produces an EF1 allocation of the desired form, given an initial endowment and a sequence of vertices. $\mathcal{A}_{discrete}$ is a slight modification of the version in [4], reformulated to better suit its role here in Sect. 7.

At a high level, in the discrete moving-knife algorithm of [4], the referee's sword moves from left to right over both vertices and edges, while the agents' knives identify the median *lumpy tie* for the part lying to right of the sword. The algorithm in [4] maintains three bundles L, M, and R; intuitively, these bundles are the discrete analogues of the left, middle, and right pieces cut by the sword and the median knife, respectively, in Stromquist's original continuous version [18].

In comparison with the original (continuous) version, the additional complexity is that the sword has to stop at one vertex and wait until the knives reach their new lumpy ties by moving step-wise from the previous median lumpy ties. This algorithm produces a contiguous allocation that is envy-free up to one of the two boundary items.

We start by introducing a few auxiliary definitions. For a given subsequence $P(v_s, v_t)$ of $P = (v_1, v_2, \ldots, v_m)$ with $s \leq t$ and a vertex v_r, we say that agent i is a *left agent* (respectively, *right agent*) over $P(v_s, v_t)$ with respect to v_r if every lumpy tie for i appears strictly before (respectively, strictly after) v_r. We say that agent i is a *middle agent* over $P(v_s, v_t)$ with respect to v_r if v_r is a lumpy tie for i.

Suppose that v_r is a median lumpy tie over the subsequence $P(v_s, v_t)$, and let i be an agent. Then by the definitions of lumpy tie and left and right agents, we have that

$$\begin{array}{llll} v_i(L(v_r)) \geq v_i(R(v_r) \cup \{v_r\}) & \text{if} & i & \text{is a left agent with respect to } v_r, \quad \text{and} \\ v_i(R(v_r)) \geq v_i(L(v_r) \cup \{v_r\}) & \text{if} & i & \text{is a right agent with respect to } v_r. \end{array} \tag{5}$$

Further, if v_r is a median lumpy tie over $P(v_s, v_t)$, at most one agent can be a left agent; similarly, at most one agent can be a right agent, and at least one agent is a middle agent.

Given a vertex v_r of the subsequence $P(v_s, v_t)$, and a two-agent set S, $Lumpy(S, v_r, P(v_s, v_t))$ identifies agents i and k in S so that i's leftmost lumpy tie appears before or at k's leftmost lumpy tie (breaking ties arbitrarily), and returns the allocation of the items in $P(v_s, v_t)$ to S such that

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

We show that $Lumpy(S, v_r, P(v_s, v_t))$ returns an EF1 allocation when v_r is a median lumpy tie.

Lemma 46 (Median Lumpy Ties Lemma (Lemma 9 in [4])) *Suppose that there are three agents. Let $S = \{i, k\}$ be a pair of agents, and v_r be a median lumpy tie over the subsequence $P(v_s, v_t)$. Then $Lumpy(S, v_r, P(v_s, v_t))$ is an EF allocation of $P(v_s, v_t)$ to S up to $\{v_r\}$. Furthermore, each agent in S receives a bundle weakly better than the two bundles $L(v_r)$ and $R(v_r)$.*

Proof Let $A = \{A_i, A_k\}$ be an allocation returned by $Lumpy(S, v_r, P(v_s, v_t))$. If i is a left agent, neither of these agents envies the other by (5) and by the definition of a lumpy tie. If i is a middle agent, she does not envy agent k by the definition of a lumpy tie and k does not envy i up to $\{v_r\}$ by construction. □

Before we proceed to prove Theorem 34, we observe the following auxiliary lemmas.

Lemma 47 *If $\mathcal{A}_{discrete}$ terminates at Step 2, there is a shouter s who is a middle agent over $P(v_{\ell+1}, v_m)$ with respect to v_r. Further,*

$$v_s(\{v_{\ell+1}\} \cup M) > v_s(L) \geq v_s(M), \text{ and} \tag{6}$$

$$v_s(\{v_r\} \cup R) \geq v_s(\{v_{\ell+1}\} \cup M) > v_s(R). \tag{7}$$

Proof Consider any agent i who shouted in Step 2. Then, i strictly prefers $\{v_{\ell+1}\} \cup M$ or R to L because i did not shout at the previous step. In either case, we have that

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

where the second inequality holds because i shouted in Step 2. Thus, i strictly prefers $\{v_{\ell+1}\} \cup M$ to R. Thus, the leftmost item of R cannot be a lumpy tie for each of the two shouters, and so every lumpy tie over $P(v_{\ell+1}, v_m)$ for the shouters appears strictly before or at v_r. Since there are two agents who shouted and at most one agent can be the left agent, at least one shouter s is a middle agent. The inequalities (6) and (7) hold by the above argument and by the fact that s is a middle agent. □

Discrete moving-knife algorithm $\mathcal{A}_{discrete}$ for a set N of three agents over the disjoint union of a finite set I of vertices and an enumeration $P(v_1, v_m) = \{v_1, v_2, \ldots, v_m\}$ of additional vertices, along with a pair ℓ, r of indices with initial values $\ell = 0, r = r_0$ such that v_{r_0} is a median lumpy tie over $P(v_1, v_m)$ and no agent strictly prefers I to both $P(v_{\ell+1}, v_{r_0-1})$ and $P(v_{r_0+1}, v_m)$:
An agent i is a *shouter* if i weakly prefers L to both M and R.

Step 0. Set $L = I$, $M = P(v_{\ell+1}, v_{r_0-1})$, and $R = P(v_{r_0+1}, v_m)$.

Step 1. If no agent shouts, go to Step 2. Otherwise, allocate L as follows: If there is a shouter s_{left} who shouted at the previous step of 2 (b) or 3(b), s_{left} receives L; else, give L to an arbitrary shouter s_{left}. Allocate the remaining items according to $Lumpy(N \setminus \{s_{\text{left}}\}, v_r, P(v_{\ell+1}, v_m))$.

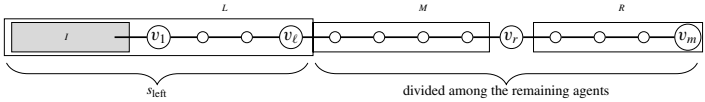

Step 2. Delete the left-most point of the middle bundle, i.e., set $M = \{v_{\ell+2}, v_{\ell+3}, \ldots, v_{r-1}\}$; then consider the following cases.

(a) If at least two agents shout, we show that there is a shouter s who is a middle agent over $P(v_{\ell+1}, v_m)$ with respect to v_r (see Lemma 47). Then, allocate L to a shouter s_{left} distinct from s. Let the agent c, who is distinct from s and s_{left}, choose his preferred bundle among $\{v_{\ell+1}\} \cup M$ and $\{v_r\} \cup R$. Agent s receives the other bundle.
(b) If the number of shouters is smaller than two and v_r is a median lumpy tie over $P(v_{\ell+2}, v_m)$, go to Step 4.
(c) If the number of shouters is smaller than two and and v_r is not a median lumpy tie over $P(v_{\ell+2}, v_m)$, go to Step 3.

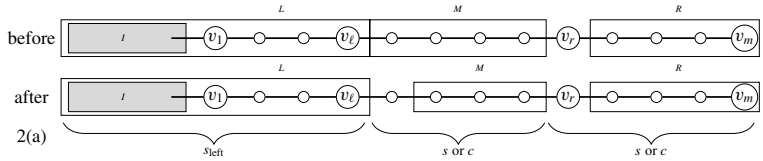

Step 3. Set $r = r+1$, $M = \{v_{\ell+2}, v_{\ell+3}, \ldots, v_{r-1}\}$, and $R = \{v_{r+1}, v_{r+2}, \ldots, v_m\}$; then, consider the following cases (a)–(c).

(a) If at least two agents shout, find a shouter s who did not shout at the previous step. If there is a shouter s_{left} who shouted at the previous step, s_{left} receives L; else, give L to an arbitrary shouter s_{left} distinct from s. The agent c, distinct from s and s_{left}, chooses 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.
(b) If v_r is a median lumpy tie over $P(v_{\ell+2}, v_m)$ but at most one agent shouts, go to Step 4.
(c) Otherwise v_r is not a median lumpy tie over $P(v_{\ell+2}, v_m)$: Repeat Step 3.

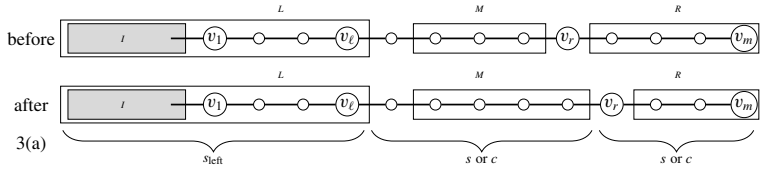

Step 4. Add an additional item to L, i.e., set $\ell = \ell + 1$ and $L = L \cup \{v_1, v_2, \ldots, v_\ell\}$. Go to Step 1.

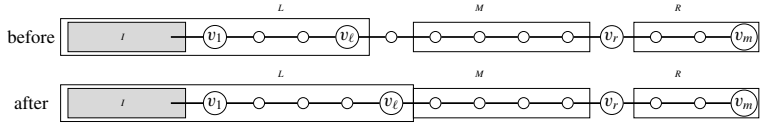

Lemma 48 *If $\mathcal{A}_{discrete}$ terminates at Step 3(a), and i is a shouter who did not shout in the previous step, then*

$$v_i(\{v_r\} \cup R) > v_i(L) \geq v_i(M). \tag{8}$$

Proof Observe that at the beginning of Step 3 just before Step 3(a), the algorithm increased r by 1, adding the previous v_r to the middle bundle M and deleting the leftmost item of the previous right bundle $\{v_r\} \cup R$. Thus, since i is not a shouter at the previous step and i is a shouter at Step 3(a), we have that

$$\max\{v_i(M \setminus \{v_{r-1}\}), v_i(\{v_r\} \cup R)\} > v_i(L) \geq \max\{v_i(M), v_i(R)\},$$

which implies the desired inequality (8). □

Now, we are ready to prove Theorem 34.

Proof of Theorem 34 To see that the algorithm is well-defined: by Lemma 47, there is a shouter s who is a middle agent over $P(v_{\ell+1}, v_m)$ when at least two agents shout in Step 2. The algorithm terminates and returns an allocation before or at the time when ℓ becomes m, because when ℓ becomes m, each of the middle and right bundles becomes empty and every agent shouts thanks to monotonicity. By construction, the algorithm also ensures that

- every time it increases ℓ by 1 in Step 4, v_r is a median lumpy tie over $P(v_{\ell+1}, v_m)$, and
- if it does not terminate at Step 1, no agent weakly prefers $L = P(v_1, v_\ell)$ to both $R = P(v_{\ell+1}, v_{r-1})$ and $M = P(v_{r+1}, v_m)$.

We will show that the resulting allocation A is

- EF up to $\{v_\ell, v_r\}$ if the algorithm terminates at Step 1.
- EF up to $\{v_{\ell+1}, v_r\}$ if the algorithm terminates at Step 2(a), or Step 3(a).

To see this, we will consider each of the three steps at which the algorithm might terminate.

Step 1. Suppose that the algorithm terminates at Step 1. We claim that each agent does not envy others up to $\{v_\ell\}$ or $\{v_r\}$.

- Since s_{left} is a shouter, agent s_{left} obtains L and does not envy the other agents up to $\{v_r\}$.
- Agent i who is not a shouter does not envy s_{left} because i strictly prefers either M or R to L, and thus by Lemma 46 receives a bundle preferred to L.
  Agent i does not envy the other agent $j \neq s_{\text{left}}$ up to $\{v_r\}$ by Lemma 46.
- We claim that any shouter agent $i \neq s_{\text{left}}$ does not envy s_{left} up to $\{v_\ell\}$. If this is the first time when Step 1 is implemented, i weakly prefers M or R to L by the initial assumption of the algorithm, and by Lemma 46 i receives a bundle weakly preferred to M or R; thus, i does not envy s_{left}. Now consider the case when the last step just before is Step 4. Recall that Step 4 follows an instance of Step 2(b) or Step 3(b), and by definition of these steps, at most one agent shouted at this step.

Thus, by the choice of s_{left}, i did not shout at the previous step of 2(b) or 3(b), and hence i strictly prefers either M or R to the previous left bundle $L \setminus \{v_\ell\}$ of Step 2(b) or Step 3(b). By Lemma 46, agent i receives a bundle weakly preferred to M or R. Thus, i does not envy s_{left} up to $\{v_\ell\}$.
Moreover, according to Lemma 46, agent i does not envy the other agent $j \neq s_{\text{left}}$ up to $\{v_r\}$.

Step 2(a). Suppose that the algorithm terminates at Step 2(a). We claim that each agent does not envy others up to $\{v_{\ell+1}\}$ or $\{v_r\}$.

- Agent s_{left} does not envy the agent receiving the bundle $\{v_{\ell+1}\} \cup M$ up to $\{v_{\ell+1}\}$ because s_{left} shouted and receives L. Similarly, s_{left} does not envy the agent who receives bundle $\{v_r\} \cup R$ up to $\{v_r\}$.
- Agent c does not envy agent s_{left}; since Step 2 follows an instance of Step 1, agent c strictly prefers $\{v_{\ell+1}\} \cup M$ or R to L. Moreover, c gets his preferred bundle among $\{v_{\ell+1}\} \cup M$ and $\{v_r\} \cup R$. Thus, c does not envy s_{left}.
  Similarly, agent c does not envy agent s since c obtains his preferred bundle among $\{v_{\ell+1}\} \cup M$ and $\{v_r\} \cup R$.
- Agent s is a middle agent, and by Lemma 47, we have that

$$\begin{aligned} &v_s(\{v_{\ell+1}\} \cup M) > v_s(L) \geq v_s(M), \textit{and} \\ &v_s(\{v_r\} \cup R) \geq v_s(\{v_{\ell+1}\} \cup M) > v_s(R). \end{aligned}$$

  Thus, if s receives bundle $\{v_{\ell+1}\} \cup M$, then she does not envy s_{left}, and does not envy c up to $\{v_r\}$. If s receives bundle $\{v_r\} \cup R$, then she does not envy the other agents.

Step 3(a). Suppose that the algorithm terminates at Step 3(a). We claim that each agent does not envy others up to $\{v_{\ell+1}\}$ or $\{v_r\}$.

- Agent s_{left} does not envy the agent who gets the bundle $\{v_{\ell+1}\} \cup M$ up to $\{v_{\ell+1}\}$ because s_{left} shouted and receives L. Similarly, s_{left} does not envy the agent who receives bundle $\{v_r\} \cup R$ up to $\{v_r\}$.
- Agent c gets his preferred bundle among $\{v_{\ell+1}\} \cup M$ and $\{v_r\} \cup R$, and hence does not envy agent s who receives the other bundle. Also, we claim that agent c does not envy s_{left}. If c is not a shouter, she strictly prefers either M or R to L and hence does not envy s_{left}. If c is a shouter, by the choice of c, she did not shout in the previous step; applying the inequality (8) in Lemma 48, we have that

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

  meaning that c does not envy s_{left}.
- Agent s does not envy others up to $\{v_{\ell+1}\}$ or $\{v_r\}$. First, consider the case when agent s receives bundle $\{v_r\} \cup R$. Since s is a shouter who did not shout in the previous step, we can apply (8) in Lemma 48 and obtain

$$v_s(\{v_r\} \cup R) > v_s(L) \geq v_s(M).$$

Thus, s does not envy s_{left} who receives L and does not envy c who receives $\{v_{\ell+1}\} \cup M$ up to $\{v_{\ell+1}\}$. Second, consider the case when agent s receives bundle $\{v_{\ell+1}\} \cup M$. This means that c strictly prefers $\{v_r\} \cup R$ over $\{v_{\ell+1}\} \cup M$ by the tie breaking rule. Recall that the initial part of Step 3 before Step 3(a) incremented r by 1 since the vertex v_{r-1} is not a median lumpy tie over $P(v_{\ell+2}, v_m)$. Thus, c cannot be a left or middle agent over $P(v_{\ell+1}, v_m)$ with respect to v_{r^*} where v_{r^*} denotes the previous median lumpy tie over $P(v_{\ell+1}, v_m)$ in Step 2 just before the algorithm enters this iteration of Step 3. Note that $r^* \leq r - 1$. Hence, s is either a left or middle agent over $P(v_{\ell+1}, v_m)$ with respect to v_{r^*}, and by definition of a median lumpy tie,

$$\begin{aligned} v_s(\{v_{\ell+1}\} \cup M) &\geq v_s(\{v_{\ell+1}, v_{\ell+2}, \ldots, v_{r^*}) \\ &\geq v_s(\{v_{r^*+1}, v_{r^*+2}, \ldots, v_m\}) \\ &\geq v_s(\{v_r\} \cup R). \end{aligned}$$

Combining this with (8), we conclude that s does not envy the other agents.

We conclude that the allocation returned by any of the steps satisfies the desired properties. □

References

1. Barbanel, J.B., Brams, S.: Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. Math. Soc. Sci. **48**(3), 251–269 (2004)
2. Bei, X., Igarashi, A., Lu, X., Suksompong, W.: The price of connectivity in fair division. SIAM J. Discret. Math. **36**(2), 1156–1186 (2022)
3. Bei, X., Suksompong, W.: Dividing a graphical cake. Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), pp. 5159–5166 (2021)
4. Bilò, V., Caragiannis, I., Flammini, M., Igarashi, A., Monaco, G., Peters, D., Vinci, C., Zwicker, W.S.: Almost envy-free allocations with connected bundles. Games Econom. Behav. **131**, 197–221 (2022)
5. Bondy, J.A., Murty, U.S.R.: Graph Theory, 1st edn. Springer, Berlin (2008)
6. Bouveret, S., Cechlárová, K., Elkind, E., Igarashi, A., Peters, D.: Fair division of a graph. Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), pp. 135–141 (2017)
7. Brânzei, S., Nisan, N.: The query complexity of cake cutting. arXiv:1705.02946v3 (2017)
8. Brams, S., Taylor, A., Zwicker, W.S.: A moving-knife solution to the four-person envy-free cake-division problem. Proc. Am. Math. Soc. **125**, 547–554 (1997)
9. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. J. Polit. Econ. **119**, 1061–1103 (2011)
10. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. ACM Trans. Econ. Comput. **7**(3), 1–32 (2019)
11. Deligkas, A., Eiben, E., Ganian, R., Hamm, T., Ordyniak, S.: The parameterized complexity of connected fair division. In Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), pp. 139–145 (2021)
12. Deligkas, A., Eiben, E., Ganian, R., Hamm, T., Ordyniak, S.: The complexity of envy-free graph cutting. In Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI), pp. 237–243 (2022)
13. Elkind, E., Segal-Halevi, E., Suksompong, W.: Graphical cake cutting via maximin share. Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), pp. 161–167 (2021)
14. Greco, G., Scarcello, F.: The complexity of computing maximin share allocations on graphs. In Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI), 2006–2013 (2020)

15. Igarashi, A., Peters, D.: Pareto-optimal allocation of indivisible goods with connectivity constraints. In Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI), pp. 2045–2052 (2019)
16. Igarashi, A.: How to cut a discrete cake fairly. arXiv:2209.01348 (2022)
17. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In Proceedings of the 5th ACM conference on Electronic commerce (ACM-EC), pp. 125–131 (2004)
18. Stromquist, W.: How to cut a cake fairly. Am. Math. Mon. **87**(8), 640–644 (1980)
19. Su, F.E.: Rental harmony: Sperner's lemma in fair division. Am. Math. Mon. **106**(10), 930–942 (1999)
20. Truszczynski, M., Lonc, Z.: Maximin share allocations on cycles. J. Artif. Intell. Res. **69**, 613–655 (2020)
21. Whitney, H.: Non-separable and planar graphs. Trans. Am. Math. Soc. **34**(2), 339–362 (1932)
22. Woodall, D.R.: Dividing a cake fairly. J. Math. Anal. Appl. **78**(1), 233–247 (1980)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.