

The Price of Connectivity in Fair Division

Xiaohui Bei,¹ Ayumi Igarashi,² Xinhang Lu,¹ Warut Suksompong³

¹ School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore

² Principles of Informatics Research Division, National Institute of Informatics, Japan

³ School of Computing, National University of Singapore, Singapore

Abstract

We study the allocation of indivisible goods that form an undirected graph and quantify the loss of fairness when we impose a constraint that each agent must receive a connected subgraph. Our focus is on the well-studied fairness notion of maximin share fairness. We introduce the *price of connectivity* to capture the largest gap between the graph-specific and the unconstrained maximin share, and derive bounds on this quantity which are tight for large classes of graphs in the case of two agents and for paths and stars in the general case. For instance, with two agents we show that for biconnected graphs it is possible to obtain at least $3/4$ of the maximin share with connected allocations, while for the remaining graphs the guarantee is at most $1/2$. Our work demonstrates several applications of graph-theoretic tools and concepts to fair division problems.

1 Introduction

We consider a classical resource allocation setting where a set of goods are to be allocated among interested agents. Our goal is to find an allocation that is fair to all agents. This problem has been addressed in a large body of literature on *fair division* (Brams and Taylor 1996a; Moulin 2003), which has found applications ranging from divorce settlement (Brams and Taylor 1996b) to credit assignment (de Clippel, Moulin, and Tideman 2008). One of the most prominent fairness notions in the literature is *proportionality*. An allocation is said to be proportional if every agent receives value at least $1/n$ of her value for the entire set of goods, where n denotes the number of agents.

Our focus in this paper is on the setting where we allocate *indivisible* goods. This pertains to the allocation of houses, cars, artworks, electronics, and many other common items. When goods are indivisible, a proportional allocation does not always exist—consider two agents trying to divide a single valuable good. As a result, relaxations of proportionality have been studied. A relaxation that has received considerable attention in recent years is *maximin share fairness*—the *maximin share (MMS)* of an agent is the largest value that the agent can guarantee for herself if she is allowed to divide the goods into n parts and always receives the worst part (Budish 2011). An allocation that gives every agent her maximin

share—said to satisfy maximin share fairness—does not always exist for additive utilities when there are at least three agents, but a constant multiplicative approximation can be obtained (Kurokawa, Procaccia, and Wang 2018).

Perhaps the best-known fair division protocol is the *cut-and-choose protocol*, which can be used to allocate a *divisible* good between two agents. In this protocol, the first agent divides the good into two equal parts (this is possible because the good is divisible), and the second agent chooses the part that she prefers. The cut-and-choose protocol has a direct analogue in the indivisible goods setting: since an equal partition may no longer exist, the first agent now divides the goods into two parts that are as equal as possible in her view. The resulting allocation is guaranteed to satisfy maximin share fairness. However, this guarantee relies crucially on the assumption that any allocation of the goods to the two agents can be chosen—in reality, there are often constraints on the allocations that we desire. A common type of constraints is captured by a model of Bouveret et al. (2017), where the goods are vertices of a connected undirected graph and each agent must be allocated a connected subgraph. For instance, the goods could represent offices in a university building that we wish to divide between research groups, and it is desirable for each group to receive a connected set of offices in order to facilitate communication within the group. In this paper we therefore ask the following question:

To what extent do maximin share fairness guarantees hold when connectivity constraints are imposed, and how does the answer depend on the underlying graph? Put differently, what is the price in terms of fairness that we have to pay if we desire connectivity?

1.1 Our Contributions

In this paper, we study maximin share fairness for agents with additive utilities. We define the *price of connectivity (PoC)* of a graph to be the largest gap between the maximin share defined over all possible partitions and the *graph maximin share (G-MMS)*, which is defined over all partitions that respect the connectivity constraints of the graph. For any graph and any number of agents, it follows from the definitions that if the PoC is α and one can give each agent β times her G-MMS, then it is also possible to guarantee all agents a β/α fraction of their MMS. Moreover, in cases where giving every agent their full G-MMS is possible (i.e., $\beta = 1$), we

Class of graphs	$n = 2$	$n \geq 2$
Paths	1 if $m = 2$ 2 if $m \geq 3$ (Thm. 3.1)	1 if $m < n$ $m - n + 1$ if $n \leq m < 2n - 1$ (Thm. 4.7) n if $m \geq 2n - 1$
Stars	$m - 1$ (Thm. 3.1)	$m - n + 1$ (Thm. 4.3)
Vertex connectivity 1	k (Thm. 3.1, see caption)	$\leq m - n + 1$ (Thm. 4.1)
Vertex connectivity 2	$\frac{4}{3}$ (Cor. 3.6)	$\leq m - n + 1$ (Thm. 4.1)
Vertex connectivity ≥ 3	$\leq \frac{4}{3}$ (Thm. 3.4)	$\leq m - n + 1$ (Thm. 4.1)

Table 1: Summary of our PoC bounds, where n and m denote the number of agents and goods, respectively. For $n = 2$ and graphs with vertex connectivity 1 (which include all trees), the parameter k denotes the maximum number of components that can result from deleting a single vertex from the graph.

observe in **Section 2** that the resulting factor $1/\alpha$ is tight—in other words, the PoC is the reciprocal of the optimal MMS approximation that can be achieved. Since it is known from prior work that $\beta = 1$ for two agents and arbitrary graphs as well as for any number of agents and trees, our PoC notion precisely captures the best possible MMS guarantee in these cases. Hence, determining the PoC, whose definition only involves a single utility function, allows us to identify the optimal MMS guarantee for agents with possibly different utility functions.

With this relationship in hand, we proceed to determine the PoC of various graphs; our results are summarized in Table 1. In the two-agent case (**Section 3**), we show that the PoC is related to the *vertex connectivity* of the graph, i.e., the minimum number of vertices whose deletion disconnects the graph. For graphs with connectivity exactly 1, including all trees, we show that the PoC is equal to the maximum number of connected components that result from deleting one vertex. As a consequence, the PoC is at least 2 for any graph in this class. On the other hand, we show an upper bound of $4/3$ for all graphs with connectivity at least 2—this bound is tight for all graphs with connectivity exactly 2 and, perhaps surprisingly, for certain graphs with connectivity up to 5. In addition, we pose an intriguing conjecture that the PoC of any graph with connectivity at least 2 is closely related to its “linkedness”—the two-agent case would be completely solved if the conjecture holds—and verify our conjecture when the graph is a complete graph with an arbitrary matching removed.

For any number of agents (**Section 4**), we establish a general upper bound of $m - n + 1$ on the PoC (where m and n denote the number of goods and agents, respectively), and show that this implies the existence of a connected allocation that gives every agent at least a $1/(m - n + 1)$ fraction of her MMS with respect to any graph. We also derive the exact PoC for paths and stars. Notably, in order to establish the PoC for paths, we introduce a new relaxation of proportionality that we call the *indivisible proportional share (IPS) property*. This notion strengthens a number of relaxations of proportionality in the literature while maintaining guaranteed existence, so we believe that it may be of independent interest as well.

From a technical point of view, our work makes extensive use of tools and concepts from graph theory, including vertex connectivity, linkedness, ear decomposition, and bipolar ordering. We believe that establishing these connections enriches the growing literature and lays the groundwork for fruitful collaborations between researchers across the two well-established fields. Furthermore, we remark that with the exception of Theorem 3.11, all of our guarantees are constructive. In particular, we exhibit polynomial-time algorithms that produce allocations satisfying the guarantees.

1.2 Related Work

The paper most closely related to ours is the one by Bouveret et al. (2017) that we already mentioned. They showed that for any number of agents with additive utilities, there always exists an allocation that gives every agent her maximin share when the graph is a tree, but not necessarily when the graph is a cycle. It is important to note that their maximin share notion corresponds to our G-MMS notion and is defined based on the graph, with only connected allocations with respect to that graph taken into account in an agent’s calculation. As an example of a consequence, even though a cycle permits strictly more connected allocations than a path, it offers less guarantee in terms of the G-MMS. Our approach of considering approximations to the (complete-graph) MMS allows us to directly compare the guarantees that can be obtained for different graphs.

Besides Bouveret et al. (2017), a number of other authors have recently studied fairness under connectivity constraints. Bilò et al. (2019) considered the same model with respect to relaxations of another important fairness property, envy-freeness. Lonc and Truszczynski (2018) investigated maximin share fairness in the case of cycles, also using the G-MMS notion, while Suksompong (2019) focused on paths and provided approximations of envy-freeness, proportionality, as well as another fairness notion called equitability. Igarashi and Peters (2019) examined fairness in conjunction with the economic efficiency notion of Pareto optimality, and Bouveret, Cechlárová, and Lesca (2019) studied the problem of *chore division*, in which all items yield disutility to the agents.

2 Preliminaries

Let $N = \{1, 2, \dots, n\}$ denote the set of agents, and $M = \{1, 2, \dots, m\}$ the set of goods. There is a bijection between the goods in M and the m vertices of a connected undirected graph G ; we will refer to goods and vertices interchangeably. A *bundle* is a subset of goods, and an *allocation* is a partition of M into n bundles (M_1, \dots, M_n) such that agent i receives bundle M_i . A bundle is called *connected* if the goods in it form a connected subgraph of G , and an allocation or a partition is *connected* if all of its bundles are connected. We assume in this paper that allocations are required to be connected.

Each agent i has a nonnegative utility $u_i(M')$ for each bundle $M' \subseteq M$, where we assume without loss of generality that $u_i(\emptyset) = 0$ for all i . For a good $g \in M$, we will use $u_i(\{g\})$ and $u_i(g)$ interchangeably. We assume that utilities are *additive*, i.e., $u(M') = \sum_{g \in M'} u(g)$ for all $M' \subseteq M$; this assumption is commonly made in the fair division literature, especially when studying maximin share fairness (Bouveret et al. 2017; Kurokawa, Procaccia, and Wang 2018; Lonc and Trzuszczynski 2018; Gourvès and Monnot 2019). An *instance* consists of the goods, their underlying graph, the agents, and their utilities for the goods.

We are ready to define maximin share fairness.

Definition 2.1. Given a graph G , an additive utility function u , and the number of agents n , the *graph maximin share (G-MMS)* for G, u, n is defined as

$$\text{G-MMS}(G, u, n) := \max_{(M_1, \dots, M_n)} \min_{i=1, \dots, n} u(M_i),$$

where the maximum is taken over all partitions (M_1, \dots, M_n) that are connected with respect to G . The *maximin share (MMS)* for u, n is defined as

$$\text{MMS}(u, n) := \text{G-MMS}(CG, u, n),$$

where CG denotes the complete graph over the goods. When the parameters are clear from the context, we will refer to the graph maximin share and the maximin share simply as G-MMS and MMS, respectively. A partition for which the maximum is attained is called a *G-MMS partition* (resp., *MMS partition*).

It follows from the definition that

$$\text{G-MMS}(G, u, n) \leq \text{MMS}(u, n) \leq \frac{u(M)}{n}$$

for all G, u, n , and $\text{G-MMS}(G_1, u, n) \leq \text{G-MMS}(G_2, u, n)$ if G_1 is a subgraph of G_2 . Moreover, $\text{G-MMS}(G, u, n) = \text{MMS}(u, n) = 0$ if $m < n$.

Next, we define the price of connectivity.

Definition 2.2. Given a graph G and the number of agents n , the *price of connectivity (PoC)* of G for n agents is defined as

$$\sup_u \frac{\text{MMS}(u, n)}{\text{G-MMS}(G, u, n)},$$

where the supremum is taken over all possible additive utility functions u .¹ We denote the PoC of a graph G for n agents by $\text{PoC}(G, n)$.

¹We interpret $\frac{0}{0}$ in this context to be equal to 1.

By definition of the PoC, we have

$$\text{PoC}(G, n) \cdot \text{G-MMS}(G, u, n) \geq \text{MMS}(u, n) \quad (1)$$

for any G, u, n , and the factor $\text{PoC}(G, n)$ cannot be replaced by any smaller factor. When G and n are clear from the context, we will refer to $\text{PoC}(G, n)$ simply as PoC. Note that the PoC is always at least 1, and is exactly 1 for complete graphs of any size. Moreover, the PoC is 1 if $m \leq n$.

Suppose that for some graph G and number of agents n , there always exists a connected allocation that gives each agent at least β times her G-MMS. By (1), this allocation also gives each agent at least $\beta/\text{PoC}(G, n)$ times her MMS. Prior work has established that $\beta = 1$ when $n = 2$ and G is arbitrary (Lonc and Trzuszczynski 2018, Cor. 2), as well as when G is a tree and n is arbitrary (Bouveret et al. 2017, Thm. 5.4). Hence, in these cases, we can guarantee each agent at least $1/\text{PoC}(G, n)$ times her MMS. The factor $1/\text{PoC}(G, n)$ is also the best possible. To see this, consider n agents with the same utility function u . From the definition of G-MMS, any connected allocation gives some agent a value of at most $\text{G-MMS}(G, u, n)$. By considering u such that $\text{G-MMS}(G, u, n)$ is arbitrarily close to $\text{MMS}(u, n)/\text{PoC}(G, n)$, this agent receives arbitrarily close to $1/\text{PoC}(G, n)$ times her MMS. To summarize, we have the following proposition.

Proposition 2.3. *Let n be any positive integer and G be any graph. If $n = 2$ (and G is arbitrary), or if G is a tree (and n is arbitrary), then there always exists a connected allocation that gives each agent at least $1/\text{PoC}(G, n)$ times her MMS. Moreover, the factor $1/\text{PoC}(G, n)$ is tight in both cases.*

Proposition 2.3 implies that if there are two agents or G is a tree, in order to determine the optimal MMS approximation for agents with possibly different utilities, it suffices to determine the value $\text{PoC}(G, n)$, which only concerns a single utility function. The rest of this paper is devoted to finding (or obtaining bounds on) the PoC in these cases.

3 Two Agents

We first focus on the setting of two agents. Before we proceed to our results, we remark here that this setting is fundamental in fair division. Indeed, a number of fair division applications including divorce settlements, inheritance division, and international border disputes often fall into this setting, and numerous prominent works in the field deal exclusively with the two-agent case (Brams and Fishburn 2000; Brams, Kilgour, and Klamler 2014; Kilgour and Vetschera 2018).² In addition, as we will see, under connectivity constraints the setting with two agents is already surprisingly rich and gives rise to several challenging questions.

Recall that all graphs we consider are assumed to be connected. The *vertex connectivity* (or simply *connectivity*) of a graph G is the minimum number of vertices whose deletion disconnects G . A graph with vertex connectivity at least k is said to be *k-connected*. By definition, every connected graph is 1-connected. A 2-connected graph is also called *biconnected*. A *bipolar ordering* (also called *bipolar numbering*)

²See also (Plaut and Roughgarden 2020b) for further discussion on the importance of the two-agent setting.

of a graph is an ordering of its vertices such that every prefix and every suffix of the ordering forms a connected subgraph.

We begin by establishing the PoC for all graphs with connectivity 1.

Theorem 3.1. *Let G be a graph with connectivity exactly 1, and let $k \geq 2$ be the maximum number of connected components that can result from deleting a single vertex of G . Then $PoC(G, 2) = k$.*

Proof. First, we show that the PoC of G is at least k . Let v be a vertex of G whose deletion results in k components. Consider a utility function with value k for v , value 1 for an arbitrary vertex in each of the k components, and value 0 for all other vertices. The MMS is k . In any connected bipartition, the part that does not contain v is a subset of one of the k components, so this part has value at most 1. Hence the PoC is at least k .

Next, we show that the PoC of G is at most k . Take an arbitrary utility function u , and assume without loss of generality that $u(M) = 1$. Since $MMS(u, 2) \leq u(M)/2 = 1/2$, the desired claim follows if there is a connected bipartition such that both parts have value at least $1/(2k)$. Assume that no such bipartition exists.

Pick a spanning tree T of G , and let v be an arbitrary vertex. The removal of v results in a number of subtrees of T ; clearly, at most one of these subtrees can have value more than $1/2$. If such a subtree exists, we move from v towards the adjacent vertex in that subtree and repeat the procedure with the new center vertex. Note that we will never traverse back an edge—otherwise there are two disjoint subtrees with value more than $1/2$ each, contradicting $u(M) = 1$. Since the tree is finite, we eventually reach a vertex v such that all subtrees T_1, \dots, T_r resulting from the removal of v have value at most $1/2$ each.

Since T_i and $T \setminus T_i$ are both connected for every i , by our earlier assumption, each of the subtrees T_1, \dots, T_r has value less than $1/(2k)$. Recall that in the original graph G , removing v can result in at most k components. This means that if $r > k$, the r subtrees must be connected by some edges not belonging to T . If subtrees T_i and T_j are connected by such an edge, we can merge T_i and T_j into one component. Note that $T_i \cup T_j$ has value less than $1/(2k) + 1/(2k) = 1/k \leq 1/2$, so since $T_i \cup T_j$ and $T \setminus (T_i \cup T_j)$ are both connected, $T_i \cup T_j$ must again have value less than $1/(2k)$. Our procedure can be repeated until the components can no longer be merged, at which point we are left with at most k components. Each of these components has value less than $1/(2k)$, which implies that v has value more than $1 - k/(2k) = 1/2$. In this case, a bipartition with v as one part is an MMS partition, so $MMS(u, 2) = 1 - u(v)$. On the other hand, at least one of the (at most) k components has value at least $(1 - u(v))/k$, which is $1/k$ of the MMS. We can take a connected bipartition with such a component as one part and obtain the desired result. \square

We remark that the proof of Theorem 3.1 also yields a polynomial-time algorithm for computing a bipartition such that both parts have value at least $1/k$ of the MMS. To compute an allocation between two agents such that both agents

receive $1/k$ of their MMS, we simply let the first agent compute a desirable bipartition, and let the second agent choose the part that she prefers. Since $MMS(u, 2) \leq u(M)/2$, the second agent is always satisfied.

Before we move on to results about graphs with higher connectivity, we show the following lemma, which will help simplify our subsequent proofs. The lemma implies that in order to prove an upper bound on the PoC in the case of two agents, it suffices to establish the bound for utility functions such that in an MMS partition, the two parts are of equal value. The proof of this lemma, along with all other omitted proofs, can be found in the full version of this paper (Bei et al. 2019).

Lemma 3.2. *For $n = 2$ and any graph G , the PoC remains the same if instead of taking the supremum in Definition 2.2*

$$\sup_u \frac{MMS(u, 2)}{G \text{-}MMS(G, u, 2)}$$

over all utility functions u , we only take the supremum over all utility functions u such that in any MMS partition according to u , the two parts are of equal value.

Next, we consider biconnected graphs, i.e., graphs with connectivity at least 2. We show that the PoC is at most $4/3$ for all such graphs. For this result, we will use a property of biconnected graphs which we state in the following proposition. An *open ear decomposition* of a graph consists of a cycle as the first *ear* and a sequence of paths as subsequent ears such that in each path, the first and last vertices (which must be different) belong to previous ears while the remaining vertices do not.

Proposition 3.3 (Whitney (1932a,b)). *In a biconnected graph with at least three vertices, any two vertices belong to a common cycle, and there exists an open ear decomposition. Moreover, we may choose any cycle in the graph as the first ear.³*

Theorem 3.4. *Let G be a biconnected graph. Then $PoC(G, 2) \leq 4/3$.*

Proof. The case $m \leq 2$ is trivial since $n = 2$ and the PoC is 1 in this case, so consider $m \geq 3$. Take an arbitrary utility function u , and assume without loss of generality that $u(M) = 1$. By Lemma 3.2, we may also assume that $MMS(u, 2) = 1/2$. Call a good *heavy* if it has value strictly more than $1/4$. Since there can be at most one heavy good in each part of an MMS partition, there are at most two heavy goods in total. Pick goods g_1 and g_2 so that together they include all of the heavy goods. By Proposition 3.3, there is a cycle in G containing g_1 and g_2 , and an open ear decomposition with this cycle as the first ear.

We will construct a bipolar ordering of the vertices that begins with g_1 and ends with g_2 . Assume that the first ear is a cycle with vertex order $g_1, h_1, \dots, h_i, g_2, h_{i+1}, \dots, h_j$. We arrange these vertices as

$$g_1, h_1, h_2, \dots, h_i, h_j, h_{j-1}, \dots, h_{i+1}, g_2.$$

³There is also a linear-time algorithm for computing an open ear decomposition with an arbitrary cycle as the first ear (Schmidt 2013).

For each subsequent ear, suppose that the two vertices belonging to previous ears are h and h' , where h appears before h' in the current ordering. We insert the remaining vertices on the path from h to h' into the ordering directly after h , following the same order as in the path. One can check (for example, by induction on the number of ears) that the resulting ordering is a bipolar ordering beginning with g_1 and ending with g_2 .

Consider first the case where $\max\{u(g_1), u(g_2)\} > 1/2$; assume without loss of generality that $u(g_1) > 1/2$. In this case, $\text{MMS}(u, 2) = 1 - u(g_1) < 1/2$, contradicting the assumption that $\text{MMS}(u, 2) = 1/2$.

Assume now that $\max\{u(g_1), u(g_2)\} \leq 1/2$, and recall that $u(g) \leq 1/4$ for all $g \notin \{g_1, g_2\}$. Since $\text{MMS}(u, 2) = 1/2$, it suffices to find a connected bipartition such that both parts have value at least $3/8$. Let $S = \{g_1\}$, so $u(S) \leq 1/2$. We add one good at a time to S following the bipolar ordering until $u(S) \geq 1/2$. Since $u(g_2) \leq 1/2$, we stop (not necessarily directly) before we add g_2 . Moreover, since each good besides g_1 and g_2 has value at most $1/4$, at some point during this process we must have $3/8 \leq u(S) \leq 5/8$. In the bipartition with S as one part, both parts are connected and have value at least $3/8$, completing the proof. \square

Unlike for Theorem 3.1, the proof of Theorem 3.4 does not directly lead to a polynomial-time algorithm for computing an allocation such that both agents receive at least $3/4$ of their MMS. The problematic step is when we apply Lemma 3.2, since computing the MMS value is NP-hard by a straightforward reduction from the partition problem. Woeginger (1997) showed that a PTAS for the problem exists—using his PTAS, we can obtain a $(3/4 - \epsilon)$ -approximation algorithm that runs in polynomial time for any constant $\epsilon > 0$. Nevertheless, we show in the full version of this paper that by building upon the proof of Theorem 3.4, we can also achieve a polynomial-time $3/4$ -approximation algorithm (Bei et al. 2019).

In light of Theorems 3.1 and 3.4, it is tempting to believe that for graphs with connectivity 3 or higher, the PoC is strictly less than $4/3$. Perhaps surprisingly, this is not the case: a counterexample is the wheel graph shown in Figure 1, which has connectivity 3. In the instance shown in the figure, the MMS is 4 while the G-MMS is 3, so the PoC of the graph is at least $4/3$ (and by Theorem 3.4, exactly $4/3$). The key point of this example is that the graph cannot be partitioned into two connected subgraphs in such a way that one subgraph contains the vertices with value 1 and 3, while the other subgraph contains the two vertices with value 2. This observation allows us to generalize the counterexample. A graph is said to be *2-linked* if for any two disjoint pairs of vertices (a, b) and (c, d) , there exist two vertex-disjoint paths, one from a to b and the other from c to d .

Proposition 3.5. *Let G be a graph that is not 2-linked. Then $\text{PoC}(G, 2) \geq 4/3$.*

Proof. Suppose that G is not 2-linked, and let (a, b) and (c, d) be disjoint pairs of vertices such that there do not exist two disjoint paths, one from a to b and the other from c to d . Consider a utility function u such that $u(a) = u(b) = 2$,

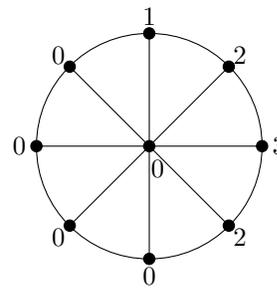


Figure 1: An instance showing that the PoC of a wheel graph is at least $4/3$.

$u(c) = 3$, $u(d) = 1$, and $u(g) = 0$ for every other vertex g . We have $\text{MMS}(u, 2) = 4$. On the other hand, the graph cannot be partitioned into two connected subgraphs in such a way that one subgraph contains a and b while the other subgraph contains c and d —indeed, such a partition would give rise to two disjoint paths that cannot exist by our assumption. This means that $\text{G-MMS}(G, u, 2) \leq 3$. Hence $\text{PoC}(G, 2) \geq 4/3$. \square

Every graph with connectivity at most 2 is not 2-linked,⁴ and Figure 1 shows an example of a 3-connected graph that also does not satisfy the property. In fact, Mészáros (2015) constructed a 5-connected graph that still fails to be 2-linked!⁵ Combining these facts with Theorem 3.4 yields the following corollaries:

Corollary 3.6. *Let G be a graph with connectivity 2. We have $\text{PoC}(G, 2) = 4/3$.*

Corollary 3.7. *There exists a graph G with connectivity 5 such that $\text{PoC}(G, 2) = 4/3$.*

While we have not been able to precisely determine the PoC for all graphs with connectivity 3 or above, we present a conjecture that, if settled in the affirmative, would complete the picture for the two-agent case. Before we can describe the conjecture, we need the following generalization of 2-linkedness (Mészáros 2015):

Definition 3.8. Given positive integers a, b , a graph G is said to be (a, b) -linked if for any disjoint set of vertices M_1, M_2 with $|M_1| = a$ and $|M_2| = b$, there exist disjoint connected subgraphs G_1, G_2 of G such that M_i is contained in G_i for $i = 1, 2$.

For example, $(2, 1)$ -linkedness is equivalent to biconnectivity, while $(2, 2)$ -linked graphs correspond to what we have so far called 2-linked graphs. The new definition allows us to extend the lower bound from Proposition 3.5.

Proposition 3.9. *Let k be a positive integer, and let G be a graph that is not $(2, k)$ -linked. Then $\text{PoC}(G, 2) \geq 2k/(2k - 1)$.*

⁴Indeed, given such a graph, let a, b be two vertices whose removal disconnects the graph, and let c, d be vertices from distinct components in the resulting graph. Then any path between c and d must go through either a or b .

⁵On the other hand, a 6-connected graph is always 2-linked (Jung 1970).

Our conjecture is that for biconnected graphs, the PoC is exactly captured by $(2, k)$ -linkedness:

Conjecture 3.10. *Let $k \geq 2$ be an integer, and let G be a graph that is $(2, k - 1)$ -linked but not $(2, k)$ -linked. Then $PoC(G, 2) = 2k/(2k - 1)$.*

The case $k = 2$ of Conjecture 3.10 holds by Corollary 3.6. We demonstrate next that the conjecture also holds for ‘almost-complete’ graphs, i.e., for complete graphs with a nonempty matching removed. These graphs have minimum degree $m - 2$, where m is the number of vertices (i.e., goods). We show that the PoC of these graphs is always exactly $(2m - 4)/(2m - 5)$, with the only exception being the graph L_5 that results from removing two disjoint edges from the complete graph K_5 (Figure 2). The graph L_5 is not 2-linked, so Proposition 3.5 (or alternatively, the utilities in Figure 2) implies that its PoC is at least $4/3$ instead of $6/5$. In fact, since the graph has connectivity 3, Theorem 3.4 tells us that its PoC is exactly $4/3$.

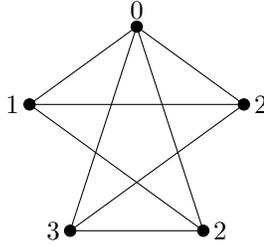


Figure 2: Graph L_5 and utilities showing that its PoC is at least $4/3$.

Theorem 3.11. *Let G be a graph that results from removing a nonempty matching from a complete graph with at least three vertices, and assume that G is different from L_5 . Then $PoC(G, 2) = (2m - 4)/(2m - 5)$.*

One can check that any such graph G is $(2, m - 3)$ -linked but not $(2, m - 2)$ -linked, so Theorem 3.11 confirms Conjecture 3.10 for this class of graphs.

4 Any Number of Agents

We move on to the general setting where the goods are divided among an arbitrary number of agents. In this setting, it is no longer true that the PoC alone captures the MMS approximation that can be guaranteed to the agents—this is evident in the case of a complete graph, where the PoC is 1 by definition, but an allocation that gives all agents their full MMS does not always exist (Kurokawa, Procaccia, and Wang 2018). At first glance, it may seem conceivable that certain graphs do not admit any useful MMS approximation. However, we provide a non-trivial guarantee for arbitrary graphs that depends only on the number of agents and goods and, in particular, not on the utilities (Theorem 4.2). We begin by establishing a general upper bound on the PoC.

Theorem 4.1. *For any graph G and any number of agents n , we have $PoC(G, n) \leq m - n + 1$.*

As we will see in Theorems 4.3 and 4.7, the bound $m - n + 1$ is tight for sufficiently short paths and all stars. We now give a maximin share guarantee for arbitrary graphs.

Theorem 4.2. *For any graph G and any number of agents n , there exists a connected allocation that gives each agent at least $1/(m - n + 1)$ of her MMS.*

Next, we derive tight bounds on the PoC in the cases of paths and stars for any number of agents. By Proposition 2.3, this also yields the optimal MMS approximation for each of these cases. We begin with stars.

Theorem 4.3. *Let $n \geq 2$ and let G be a star. Then*

$$PoC(G, n) = \begin{cases} m - n + 1 & \text{if } m \geq n; \\ 1 & \text{if } m < n. \end{cases}$$

Proof. The following simple fact about MMS will be useful.

Lemma 4.4. *Let $m \geq n$, and let $M' \subseteq M$ be an arbitrary set of at least $m - n + 1$ goods. For an agent with utility function u , we have $u(M') \geq MMS(u, n)$.*

Proof. Observe that in any partition of the vertices into n parts, at least one of the parts is contained in M' . In particular, this holds for an MMS partition. It follows that $MMS(u, n) \leq u(M')$, as claimed. \square

We proceed to the proof of Theorem 4.3. If $m < n$ the PoC is 1, so assume that $m \geq n$. We first show that the PoC is at least $m - n + 1$. Consider a utility function u with value $m - n + 1$ for the center vertex and for $n - 2$ of the leaves, and value 1 for each of the remaining $m - n + 1$ leaves. We have $MMS(u, n) = m - n + 1$. In any connected partition into n parts, at least $n - 1$ parts contain a single leaf. This means that at least one of these parts contains a single leaf with value 1. Hence the PoC is at least $m - n + 1$.

Next, we show that the PoC is at most $m - n + 1$. Take an arbitrary utility function u , let v^* be the center vertex, and let v_1, v_2, \dots, v_{n-1} be the leaves with the highest value where $u(v_1) \geq \dots \geq u(v_{n-1})$. Consider a connected partition Π with each of these $n - 1$ vertices as a part, and the remaining $m - n + 1$ vertices as the last part.

Let $A := M \setminus \{v^*, v_1, \dots, v_{n-1}\}$. By Lemma 4.4, $MMS(u, n) \leq u(A)$. Since there are $m - n + 1$ vertices in A and v_{n-1} is a vertex with the highest value, we have

$$u(v_{n-1}) \geq \frac{1}{m - n + 1} \cdot u(A) \geq \frac{1}{m - n + 1} \cdot MMS(u, n).$$

It follows that $u(v_i) \geq MMS(u, n)/(m - n + 1)$ for all $i = 1, 2, \dots, n - 1$, so the first $n - 1$ parts of Π have value at least $MMS(u, n)/(m - n + 1)$ each. The last part of Π is $B := M \setminus \{v_1, v_2, \dots, v_{n-1}\}$. By Lemma 4.4 again, we have $MMS(u, n) \leq u(B)$. This means that all parts of Π have value at least $MMS(u, n)/(m - n + 1)$, as desired. \square

To address the more involved case of paths, we introduce an approximation of proportionality that can be of interest even in the absence of connectivity considerations. Recall that an allocation is said to be *proportional* if it gives every agent at least her *proportional share*, which is defined as

$u(M)/n$. Even though a proportional allocation always exists for *divisible* goods, as we explained in the introduction, this is not the case for indivisible goods—our definition of *indivisible proportional share* therefore adapts proportionality to the setting of indivisible goods. In order to ensure a nontrivial approximation, we will need to hypothetically remove up to $n - 1$ goods from the entire bundle. When the number of goods is large, we will then be able to guarantee a $1/n$ fraction of the remaining value; however, for smaller numbers of goods, we can achieve a better approximation.

Definition 4.5. For positive integers n, m , define

$$\text{IPS}(n, m) = \begin{cases} \frac{1}{n} & \text{if } m \geq 2n - 1; \\ \frac{1}{m-n+1} & \text{if } n \leq m < 2n - 1; \\ 0 & \text{if } m < n. \end{cases}$$

Given n agents and m goods, a bundle A is said to satisfy the *indivisible proportional share (IPS) property* for an agent with utility function u if there exists a (possibly empty) set $B \subseteq M \setminus A$ with $|B| \leq n - 1$ such that

$$u(A) \geq \text{IPS}(n, m) \cdot u(M \setminus B).$$

An allocation is said to satisfy the IPS property if every agent receives a bundle that satisfies the IPS property. For brevity, we will refer to a bundle or allocation that satisfies the IPS property as being IPS.

We remark that IPS is a stronger property than $\text{PROP}^*(n-1)$ considered by Segal-Halevi and Suksompong (2019), which corresponds to taking $\text{IPS}(n, m) = 1/n$ for $m \geq n$ and 0 for $m < n$. It is also stronger than PROP1 considered by Conitzer, Freeman, and Shah (2017) and Aziz et al. (2019), as well as a proportionality relaxation studied by Suksompong (2019). Despite its strength, we show that an IPS allocation always exists. Moreover, we can obtain a connected IPS allocation if the graph is a path.

Proposition 4.6. *Let $n \geq 2$ and let G be a path. There exists a connected IPS allocation of the m goods to the n agents.*

Proof. If $m < n$, each agent needs utility 0 in an IPS allocation, so the claim holds trivially. Assume that $m \geq n$. Starting with an empty bundle, we process the goods along the path (say, from left to right) and add them one at a time to the current bundle until the bundle is IPS to at least one of the agents. We then allocate the bundle to one such agent, and repeat the procedure with the remaining goods and agents. Any leftover goods are allocated to the agent who receives the last bundle.

We claim that this procedure always results in an IPS allocation. Notice from Definition 4.5 that if a bundle is IPS for an agent, then so is any superset of the bundle. Hence it suffices to show that after $n - 1$ bundles are allocated, the last agent still finds the remaining bundle to be IPS. Assume without loss of generality that the bundles are allocated to agents $1, 2, \dots, n$ in this order, and let u be the utility function of agent n . The claim holds trivially if the empty bundle is IPS for agent n , so assume that it is not. For $1 \leq i \leq n-1$, let the bundle allocated to agent i be $M_i = X_i \cup Y_i$, where Y_i consists of the last good added to M_i (if M_i is nonempty),

and X_i consists of the remaining goods. Let $X = \cup_{i=1}^{n-1} X_i$ and $Y = \cup_{i=1}^{n-1} Y_i$. In particular, $|Y| \leq n - 1$.

Let M_n be the bundle allocated to agent n .

- Case 1: $m \geq 2n - 1$. By definition of the procedure, agent n does not find any of the bundles X_1, \dots, X_{n-1} to be IPS. In particular, noting that $Y \subseteq M \setminus X_i$ for each $1 \leq i \leq n - 1$ and taking $B = Y$ in Definition 4.5, we have $u(X_i) < \text{IPS}(n, m) \cdot u(M \setminus Y) = u(M \setminus Y)/n$ for all i . Hence,

$$\begin{aligned} u(M_n) &= u(M) - \sum_{i=1}^{n-1} u(X_i) - \sum_{i=1}^{n-1} u(Y_i) \\ &> u(M) - \frac{n-1}{n} \cdot u(M \setminus Y) - u(Y) \\ &= \frac{1}{n} \cdot u(M \setminus Y). \end{aligned}$$

Since $Y \subseteq M \setminus M_n$, bundle M_n is IPS for agent n .

- Case 2: $n \leq m \leq 2n - 1$. This case involves a more detailed analysis, which we defer to the full version of this paper (Bei et al. 2019). \square

Proposition 4.6 allows us to establish the PoC for paths, which we do next in Theorem 4.7. Conversely, the instances that we use to show the upper bound on the PoC in Theorem 4.7 also show that the factor $\text{IPS}(n, m)$ in the existence guarantee of Proposition 4.6 cannot be improved.

Theorem 4.7. *Let $n \geq 2$ and let G be a path. Then*

$$\text{PoC}(G, n) = \begin{cases} n & \text{if } m \geq 2n - 1; \\ m - n + 1 & \text{if } n \leq m < 2n - 1; \\ 1 & \text{if } m < n. \end{cases}$$

5 Conclusion

In this paper, we have studied the fair allocation of indivisible goods under connectivity constraints and investigated maximin share fairness guarantees for various classes of graphs. In particular, we established a link between the graph-specific maximin share and the well-studied maximin share through our price of connectivity (PoC) notion, which allows us to directly compare the guarantees provided by different graphs. We presented a number of bounds on the PoC, several of which are tight, and left a tempting conjecture that would settle the two-agent case if it holds.

As is the case in most of the maximin share fairness literature, our results rely on the assumption that the agents' utility functions are additive. Maximin share fairness beyond additive utilities has been studied by Barman and Krishna Murthy (2017) and Ghodsi et al. (2018); for example, they showed that a constant approximation of the maximin share can be achieved for any number of agents with submodular utilities when the graph is complete. Since complementarity and substitutability occur in practice, it would be interesting to see how the graph-based approximations that we obtain in this paper change as we enlarge the class of utility functions considered. Indeed, as Plaut and Roughgarden (2020a) noted, there is a rich landscape of problems to explore in fair division with different classes of utility functions, and the graphical setting is likely to be no exception.

Acknowledgments

This work was partially supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (RG23/20), by the KAKENHI Grant-in-Aid for JSPS Fellows number 18J00997, by the European Research Council (ERC) under grant number 639945 (ACCORD), by an NUS Start-up Grant, and by JST, ACT-X.

References

- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2019. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 53–59.
- Barman, S.; and Krishna Murthy, S. K. 2017. Approximation algorithms for maximin fair division. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 647–664.
- Bei, X.; Igarashi, A.; Lu, X.; and Suksompong, W. 2019. The price of connectivity in fair division. *CoRR* abs/1908.05433.
- Bilò, V.; Caragiannis, I.; Flammini, M.; Igarashi, A.; Monaco, G.; Peters, D.; Vinci, C.; and Zwickler, W. S. 2019. Almost envy-free allocations with connected bundles. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, 14:1–14:21.
- Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 135–141.
- Bouveret, S.; Cechlárová, K.; and Lesca, J. 2019. Chore division on a graph. *Autonomous Agents and Multi-Agent Systems* 33(5): 540–563.
- Brams, S. J.; and Fishburn, P. C. 2000. Fair division of indivisible items between two people with identical preferences: Envy-freeness, Pareto-optimality, and equity. *Social Choice and Welfare* 17(2): 247–267.
- Brams, S. J.; Kilgour, D. M.; and Klamler, C. 2014. Two-person fair division of indivisible items: an efficient, envy-free algorithm. *Notices of the AMS* 61(2): 130–141.
- Brams, S. J.; and Taylor, A. D. 1996a. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brams, S. J.; and Taylor, A. D. 1996b. A procedure for divorce settlements. *Mediation Quarterly* 13(3): 191–205.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6): 1061–1103.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 629–646.
- de Clippel, G.; Moulin, H.; and Tideman, N. 2008. Impartial division of a dollar. *Journal of Economic Theory* 139: 176–191.
- Ghodsi, M.; Hajiaghayi, M. T.; Seddighin, M.; Seddighin, S.; and Yami, H. 2018. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 539–556.
- Gourvès, L.; and Monnot, J. 2019. On maximin share allocations in matroids. *Theoretical Computer Science* 754: 50–64.
- Igarashi, A.; and Peters, D. 2019. Pareto-optimal allocation of indivisible goods with connectivity constraints. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2045–2052.
- Jung, H. A. 1970. Eine Verallgemeinerung des n -fachen Zusammenhangs für Graphen. *Mathematische Annalen* 187(2): 95–103.
- Kilgour, D. M.; and Vetschera, R. 2018. Two-player fair division of indivisible items: Comparison of algorithms. *European Journal of Operational Research* 271(2): 620–631.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM* 64(2): 8:1–8:27.
- Lonc, Z.; and Truszczynski, M. 2018. Maximin share allocations on cycles. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 410–416.
- Mészáros, G. 2015. *Linkedness and path-pairability in the Cartesian product of graphs*. Ph.D. thesis, Central European University.
- Moulin, H. 2003. *Fair Division and Collective Welfare*. MIT Press.
- Plaut, B.; and Roughgarden, T. 2020a. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics* 34(2): 1039–1068.
- Plaut, B.; and Roughgarden, T. 2020b. Communication complexity of discrete fair division. *SIAM Journal on Computing* 49(1): 206–243.
- Schmidt, J. M. 2013. A simple test on 2-vertex- and 2-edge-connectivity. *Information Processing Letters* 113(7): 241–244.
- Segal-Halevi, E.; and Suksompong, W. 2019. Democratic fair allocation of indivisible goods. *Artificial Intelligence* 277: 103167.
- Suksompong, W. 2019. Fairly allocating contiguous blocks of indivisible items. *Discrete Applied Mathematics* 260: 227–236.
- Whitney, H. 1932a. Congruent graphs and the connectivity of graphs. *American Journal of Mathematics* 54(1): 150–168.
- Whitney, H. 1932b. Non-separable and planar graphs. *Transactions of the American Mathematical Society* 34(2): 339–362.
- Woeginger, G. J. 1997. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters* 20(4): 149–154.