

Resource-monotonicity and Population-monotonicity in Connected Cake-cutting

Erel Segal-Halevi^a, Balázs Sziklai^{b,c}

^aBar-Ilan University, Ramat-Gan 5290002, Israel. erelsgl@gmail.com

^b'Momentum' Game Theory Research Group, Centre for Economic and Regional Studies, Hungarian Academy of Sciences. H-1112 Budapest Budaörsi út 45., Email: sziklai.balazs@krtk.mta.hu

^cCorvinus University of Budapest, Department of Operations Research and Actuarial Sciences, H-1093 Budapest Fővám tér 8.

Abstract

In the classic cake-cutting problem (Steinhaus, 1948), a heterogeneous resource has to be divided among n agents with different valuations in a *proportional* way — giving each agent a piece with a value of at least $1/n$ of the total. In many applications, such as dividing a land-estate or a time-interval, it is also important that the pieces are *connected*. We propose two additional requirements: resource-monotonicity (RM) and population-monotonicity (PM). When either the cake or the set of agents changes and the cake is re-divided using the same rule, the utility of all remaining agents must change in the same direction. Classic cake-cutting protocols are neither RM nor PM. Moreover, we prove that no Pareto-optimal proportional division rule can be either RM or PM. Motivated by this negative result, we search for division rules that are *weakly-Pareto-optimal* — no other division is strictly better for all agents. We present two such rules. The *relative-equitable* rule, which assigns the maximum possible relative value equal for all agents, is proportional and PM. The so-called *rightmost mark* rule, which is an improved version of the Cut and Choose protocol, is proportional and RM for two agents.

Keywords: fair division, cake-cutting, resource-monotonicity, population-monotonicity, connected utilities

1. Introduction

Monotonicity axioms have been extensively studied with respect to cooperative game theory (Calleja et al., 2012), political representation (Balinski and Young, 1982), computer resource allocation (Ghodsi et al., 2011) and many other fair division problems (Moulin, 2004, chapters 3 6 7), (Thomson, 2011, chapter 7).

These axioms express the idea of solidarity among agents: whenever the environment changes in a way that requires the re-allocation of resources, the welfare of all agents not responsible for the change should be affected in the same direction — either they should all be made at least as well off as they were initially, or they should all be made at most as well off. This is the so called replacement principle which was formulated by Thomson (1997).

Two common monotonicity axioms are *resource monotonicity* (RM) and *population monotonicity* (PM). Resource-monotonicity, sometimes known as aggregate monotonicity, requires that when new resources are added, and the same division rule is used consistently, the utility of all agents should weakly increase. Population-monotonicity is concerned with changes in the number of participants. It requires that when someone leaves the division process and abandons his share, the utility of the remaining participants should weakly increase. Conversely, when a new agent joins the process, all existing participants should participate in supporting the new agent, thus their utility should weakly decrease.

Monotonicity axioms are sometimes conceived as more important than other, more basic fairness axioms. A prominent example is the practical problem of *apportionment*: there is a parliament with a fixed number of seats and administrative regions with different number of voters. The seats have to be distributed among the regions in such way that the resulting allotment ensures proportional representation. The solution originally employed in the USA congress was the Hamiltonian rule, which guaranteed proportional representation. However, it was found that this

rule exhibits the *Alabama paradox* — increasing the number of seats would have rendered state Alabama with less seats. In other words, the rule violates resource-monotonicity. Later, it was found that this rule also exhibits the *new state paradox* (the “*Oklahoma paradox*”) — the addition of Oklahoma to the USA would have rendered state Maine with more seats. In other words, the rule also violates population-monotonicity. These violations pressed legislators to adopt a new apportionment method. The currently used method, the so called Huntington-Hill method, fails to uphold Hare-quota, a basic guarantee of proportionality, but it satisfies the monotonicity axioms (Balinski and Young, 1975; Caulfield, 2008).

The present paper studies these two monotonicity requirements in the framework of the *fair cake-cutting* problem (Steinhaus, 1948), where a single heterogeneous resource - such as land or time - has to be divided fairly. We approach the problem from the classic point of view, when each agent is interested in getting a connected piece. Connected utility functions make sense in many applications. E.g., when dividing land, a large connected piece can be used for building a house, but a collection of small disconnected patches of land is virtually useless. Similarly when departments dispute over the availability of a conference room, each of them is interested in reserving the room for a contiguous period which is free of disruption. Another example is a long TV ad, which needs to be aired in one piece. These examples show that assuming connected utilities in some cases is a reasonable restriction, and indeed, many cake-cutting papers explicitly assume that each agent must be allocated a connected piece.

1.1. Results

We survey many traditional cake-cutting protocols and show that they do not satisfy either of the monotonicity axioms. In particular, all methods based on the Cut and Choose scheme violate both resource-monotonicity and population-monotonicity. This motivates a search for division rules that are both fair in the conventional sense and monotonic. We conducted this search under two different assumptions regarding the agents’ utility functions, which are equally common in the cake-cutting literature.

In both models, each agent has a value measure defined over the cake. In the *additive* model, the utility of a piece of cake is just the value measure of that piece; the geometry of the piece has no importance. In the *connected* model, the utility of a piece of cake is the value of the most valuable connected component of the piece.

Our results for the additive model can be found in another manuscript (Sziklai and Segal-Halevi, 2015). These results were mainly positive: we found several Pareto-optimal proportional division rules that satisfy one or both monotonicity axioms. In particular, the *Nash-optimal* rule, maximizing the product of values, is envy-free (hence also proportional), resource-monotonic and population-monotonic.

The present paper studies the connected model. Here, the situation is not so positive. Each of the monotonicity properties is incompatible with proportionality and Pareto-optimality. That is, no Pareto-optimal proportional division rule can be either resource- or population-monotonic. Thus the fair divider has to choose between Pareto-optimality and monotonicity. While from an economics perspective Pareto-optimality is crucial, public opinion may not always agree. In some cases people are willing to sacrifice efficiency to get fairness (Herreiner and Puppe, 2009). As a compromise, we suggest several division rules which are proportional and *weakly*-Pareto-optimal (no other allocation is strictly preferred by all agents; see e.g. Varian (1974)) while satisfying one of the monotonicity axioms. The *max-equitable-connected* rules, which give equally-valuable pieces to each agent while maximizing this value, are both population monotonic. There are two such rules: the rule equalizing the relative values (normalized such that the entire cake value is 1) is proportional but not resource-monotonic, and the rule equalizing the absolute (not normalized) values is resource-monotonic but not proportional. Additionally, we present a proportional and resource-monotonic division protocol for two agents. It is an open question whether there exists a weakly-Pareto-optimal, proportional and resource-monotonic rule for n agents.

The equitable rule belongs to the cardinal welfarism framework (cf. chapter 3 of Moulin (2004)). It relies on inter-agent utility comparison, and makes sense if and when such a comparison is feasible. For example, suppose the agents are firms, each of which wants to use the land to a pre-specified purpose (e.g. one firm plans to dig for oil, another firm wants to build housing complexes, etc.). Then, economic models can be used to estimate the monetary utility of each firm for each piece of land and the estimates can be used to calculate equitable divisions (see the conclusion of Chambers (2005a) for further discussion of the additive utility model).

The paper is organized as follows. Section 2 reviews the related literature. Section 3 formally presents the cake-cutting problem and the monotonicity axioms. Section 4 examines classic cake-cutting protocols and shows that they are not monotonic. Sections 5 and 6 present our negative and positive results. Section 7 concludes and presents a table summarizing the various rules' properties.

2. Related Work

The cake-cutting problem originates from the work of the Polish mathematician Hugo Steinhaus and his students Banach and Knaster (Steinhaus, 1948). Their primary concern was how to divide the cake in a fair way. Since then, game theorists analyzed the strategic issues related to cake-cutting, while computer scientists were focusing mainly on how to implement solutions, i.e. the computational complexity of cake-cutting protocols.

Many economists regard land division as an important application of division procedures (e.g. Berliant and Raa, 1988; Berliant et al., 1992; Legut et al., 1994; Chambers, 2005b; Dall'Aglio and Maccheroni, 2009; Hüsseinov, 2011; Nicolò et al., 2012)). Hence, they note the importance of imposing some geometric constraints on the pieces allotted to the agents. Connectivity is the most well-studied constraint.

As we already noted in the introduction there is a vast literature on monotonicity related issues. To our knowledge our paper is the first that explicitly defines RM and PM for the cake cutting setting. However, there are a few other axioms which bear resemblance to these two.

Arzi et al. (2011) study the "dumping paradox" in cake-cutting. They show that, in some cakes, discarding a part of the cake improves the total social welfare of any envy-free division. This implies that enlarging the cake might decrease the total social welfare. This is related to resource-monotonicity; the difference is that in our case we are interested in the welfare of the individual agents and not in the total social welfare.

Chambers (2005a) studies a related cake-cutting axiom called "division independence": if the cake is divided into sub-plots and each sub-plot is divided according to a rule, then the outcome should be identical to dividing the original land using the same rule. He proves that the only rule which satisfies Pareto-optimality and division independence is the utilitarian-optimal rule - the rule which maximizes the sum of the agents' utilities. The rule is only feasible when the utilities are additive (with no connectivity constraints). Unfortunately, this rule does not satisfy fairness axioms such as proportionality.

Walsh (2011) studies the problem of "online cake-cutting", in which agents arrive and depart during the process of dividing the cake. He shows how to adapt classic procedures like cut-and-choose and the Dubins-Spanier in order to satisfy online variants of the fairness axioms. Monotonicity properties are not studied, although the problem is similar in spirit to the concept of population-monotonicity.

Finally, we mention that the consistency axiom (cf. Young (1987) or Thomson (2012)) is related to population-monotonicity, but it is fundamentally different as in that case the leaving agents take their fair shares with them.

2.1. Equitable divisions

The “heroes” of the present paper are the equitable division rules. The equitability condition in cake-cutting is much less studied than other properties such as proportionality and envy-freeness. Some notable exceptions are presented below.

The first proof to the *existence* of an equitable division is implied by the seminal work of Dubins and Spanier (1961). The piece allocated to each agent can be an arbitrary member of a σ -algebra, i.e, not necessarily connected. The result of Alon (1987) implies the existence of an equitable division with a limited number of cuts, but still not necessarily connected. Max-relative-equitable divisions without the connectivity requirement were studied extensively by Dall’Aglio (2001) (he calls such divisions *equi-optimal*). Max-relative-equitable divisions with the connectivity requirement for two agents were studied by Jones (2002). The generalization for n agents was mentioned by Brams et al. (2006, 2007). They related the problem of equitable-connected cake-cutting to a set of integral equations, but did not prove they are solvable. The latter point was discussed by Mawet et al. (2010) for the special case when the valuations are piecewise-constant and everywhere-positive. Aumann and Dombb (2010) proved that equitable-connected allocations exist for general valuations. Cechlárová et al. (2013) extended this result and proved that an equitable-connected division exists for any ordering of the agents, and for at least one ordering it is also proportional.

The *computability* of equitable allocations is discussed by several recent works. Cechlárová and Pillárová (2012) proved that there is no finite discrete procedure for finding an allocation that is equitable, connected and proportional. Procaccia and Wang (2017) showed that this impossibility holds even without the connectivity and proportionality requirements. On the positive side, Cechlárová and Pillárová (2011, 2012) provided discrete procedures that attain ϵ -*equitable connected divisions* — divisions in which the difference between the value of every two agents is at most ϵ . Independently and contemporaneously to our work, Brânzei and Nisan (2017) presented a moving-knife procedure for equitable cake-cutting, for the special case in which all players are “hungry” (i.e, all valuations are strictly positive).

The main contribution of the present paper to the literature on equitable division is in showing its advantages over other, more famous cake-cutting procedures. In particular, we show that it is population monotonic, and can be made either proportional or resource-monotonic depending on whether relative or absolute values are used.

A secondary contribution is a moving-knives procedure for finding an equitable-connected division in any ordering of the agents, which is applicable for general valuations (not only strictly positive). This does not contradict the impossibility results mentioned above, since a moving-knife procedure is continuous rather than discrete.

3. Model

3.1. Cake-cutting

A cake-cutting problem is a triple $\Gamma(N, C, (\widehat{v}_i)_{i \in N})$ where:

- $N = \{1, 2, \dots, n\}$ denotes the set of agents who participate in the cake-cutting process. In examples with a small number of agents, we often refer to them by names (Alice, Bob, Carl...).
- C is the cake. For simplicity we assume that C is a interval, $C = [0, c]$ for some real number c . We call a Borel subset of C a *slice*.
- \widehat{v}_i is the value measure of agent i . It is a finite real-valued function defined on the Borel subsets of $[0, \infty)$.

As the term “measure” implies, the value measures of all agents are countably additive: the value measures of a union of disjoint slices is the sum of the values of the slices. Moreover, we assume that the value measures are non-negative and bounded. That is, \widehat{v}_i assigns a non-negative, but finite number to each slice of C . We also assume

that the value measures are absolutely-continuous with respect to Lebesgue measure: this means that a slice with zero length has zero value to anyone. Therefore it is unimportant to specify which agent gets the endpoints of an interval, since the endpoints have zero value. All these assumptions are standard in the cake-cutting literature.

Our model diverges from the standard cake-cutting setup in that we do not require the value measures to be normalized. That is, the value of the entire cake is not necessarily the same for all agents. This is important because we examine scenarios where the cake changes, so the cake value might become larger or smaller. Hence, we differentiate between *absolute* and *relative* value measures:

- The **absolute** value measure of the entire cake, $\widehat{v}_i(C)$, can be any positive value and it can be different for different agents.
- The **relative** value of the entire cake is 1 for all agents. Relative value measures are denoted by v_i and defined by: $v_i(S) := \widehat{v}_i(S)/\widehat{v}_i(C)$.

It is also common to assume that value measures are private information of the agents. This question leads us to whether agents are honest about their preferences. Cake-cutting problems can be studied from a strategic angle, however, the results are mostly negative. For example, in any deterministic discrete strategy-proof protocol, there always exists an agent that gets the empty piece (Brânzei and Miltersen, 2015). Here, we will not analyze the strategic behavior of the agents but assume they act truthfully.

The *utility* of an agent is based on its value measure. In the present paper we assume that:

$$\widehat{u}_i(X) = \sup_{I \subset X} \widehat{v}_i(I) \qquad u_i(X) = \sup_{I \subset X} v_i(I)$$

where the supremum is over all connected intervals I that are subsets of X . That is, an agent can only use a single connected piece.

The aim is to divide the cake into n pairwise-disjoint slices. A *division rule* is a correspondence that takes a cake-cutting problem as input and returns a *division* $X = (X_1, \dots, X_n)$, or a set of divisions. Note that a division does not necessarily compose a partition of C (i.e. free disposal is assumed).

Since all agents have connected utilities, we can assume without loss of generality that each agent receives a connected piece, i.e, for all i , X_i is an interval. Under this assumption, $\widehat{u}_i(X_i) = \widehat{v}_i(X_i)$ and $u_i(X_i) = v_i(X_i)$ for all i , so from now on we will use only \widehat{v}_i and v_i .

A division rule R is called *essentially single-valued (ESV)* if $X, Y \in R(\Gamma)$ implies that for all $i \in N$, $\widehat{v}_i(X_i) = \widehat{v}_i(Y_i)$. That is, even if R returns a set of divisions, all agents are indifferent between these divisions.

The classic requirements of fair cake-cutting are the following. A division X is called:

- *Pareto-optimal* (PO) if there is no other division which is weakly better for all agents and strictly better for at least one agent.
- *Weakly-Pareto-optimal* (WPO) if there is no other division which is strictly better for all agents.
- *Proportional* (PROP) if each agent gets at least $1/n$ fraction of the cake according to his own evaluation, i.e. for all $i \in N$, $v_i(X_i) \geq 1/n$. Note that the definition uses relative values.
- *Envy-free* (EF) if each agent gets a piece which is weakly better, for that agent, than all other pieces: for all $i, j \in N$, $v_i(X_i) \geq v_i(X_j)$. Note that here it is irrelevant whether absolute or relative values are used. Note also that PO+EF imply PROP.

A division rule is called *Pareto-optimal* (PO) if it returns only PO divisions. The same applies to WPO, PROP and EF.

3.2. Monotonicity

We now define the two monotonicity properties. In the introduction we defined them informally for the special case in which the division rule returns a single division. Our formal definition is more general and applicable to rules that may return a set of divisions.

Definition 3.1. Let N be a fixed set of agents, $C = [0, c]$, $C' = [0, c']$ two cakes where $c < c'$, and $(\hat{v}_i)_{i \in N}$ value measures on $[0, \infty)$. The cake-cutting problem $\Gamma' = (N, C', (\hat{v}_i)_{i \in N})$ is called a *cake-enlargement* of the problem $\Gamma = (N, C, (\hat{v}_i)_{i \in N})$.

By definition the cake is always enlarged on the right hand side. This might be critical for some protocols. For instance in the Dubins-Spanier moving knife protocol the cake is processed from left to right (Dubins and Spanier, 1961). However, most of our results (except that of Subsection 6.4) are valid whenever $C \subset C'$, regardless of whether the cake is enlarged from the left, right or middle.

Definition 3.2. (a) A division rule R is called *upwards resource-monotonic*, if for all pairs (Γ, Γ') , where Γ' is a cake-enlargement of Γ , for every division $X \in R(\Gamma)$ there *exists* a division $Y \in R(\Gamma')$ such that $\hat{v}_i(Y_i) \geq \hat{v}_i(X_i)$ for all $i \in N$ (i.e all agents are weakly better-off in the new division).

(b) A division rule R is called *downwards resource-monotonic*, if for all pairs (Γ', Γ) , where Γ' is a cake-enlargement of Γ , for every division $Y \in R(\Gamma')$ there *exists* a division $X \in R(\Gamma)$ such that $\hat{v}_i(X_i) \leq \hat{v}_i(Y_i)$ for all $i \in N$ (i.e all agents are weakly worse-off in the new division).

(c) A division rule is *resource-monotonic* (RM), if it is both upwards and downwards resource-monotonic.

Definition 3.3. Let C be a fixed cake, N and N' two sets of agents such that $N \supset N'$ and $(\hat{v}_i)_{i \in N}$ their value measures. The cake-cutting problem $\Gamma' = (N', C, (\hat{v}_i)_{i \in N'})$ is called a *population-reduction* of the problem $\Gamma = (N, C, (\hat{v}_i)_{i \in N})$.

Definition 3.4. (a) A division rule R is called *upwards population-monotonic*, if for all pairs (Γ', Γ) such that Γ' is a population-reduction of Γ , for every division $Y \in R(\Gamma')$ there exists a division $X \in R(\Gamma)$ such that $\hat{v}_i(X_i) \leq \hat{v}_i(Y_i)$ for all $i \in N'$ (all the original agents are weakly worse-off in the new division).

(b) A division rule R is called *downwards population-monotonic*, if for all pairs (Γ, Γ') such that Γ' is a population-reduction of Γ , for every division $X \in R(\Gamma)$ there exists a division $Y \in R(\Gamma')$ such that $\hat{v}_i(Y_i) \geq \hat{v}_i(X_i)$ for all $i \in N'$ (all remaining agents are weakly better-off in the new division).

(c) A division rule is *population-monotonic* (PM), if it is both upwards and downwards population-monotonic.

Remark 3.1. As usual in the literature, the monotonicity axioms care only about absolute values. In other words, it is not considered a violation of RM if the relative value of an agent decreases when the cake grows.

Remark 3.2. For essentially-single-valued solutions, downwards resource (or population) monotonicity implies upwards resource (or population) monotonicity and vice versa. Set valued solutions, however, may satisfy only one direction of these axioms.

Remark 3.3. The monotonicity axioms in Thomson (2011) require that *all* divisions in $R(\Gamma)$ have to be weakly better/worse than all divisions in $R(\Gamma')$. In contrast, our definition only requires that there *exists* such a division. This is closer to the definition of aggregate monotonicity, which originates from cooperative game theory (Peleg and Sudhölter, 2007). The rationale is that even if a set-valued solution is used, only a single allocation will be implemented. Hence, the divider can be faithful to the monotonicity principles even if the rule suggests many non-monotonic allocations as well.

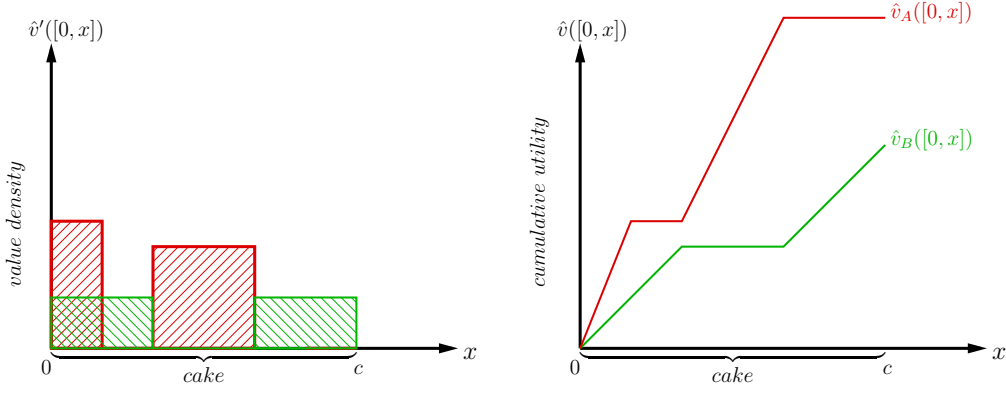


Figure 1: Piecewise homogeneous cake with two players. The derivative of $\hat{v}([0, x])$ indicates the density of the value. Note that the value measures are not normalized, hence $\hat{v}_A([0, c]) \neq \hat{v}_B([0, c])$.

Because our monotonicity requirements are weaker, any impossibility result in our model is valid in Thomson’s model, too. This is not true in general for positive results; however, the specific positive results in the present paper are all based on essentially-single-valued rules. Hence, by the previous remark, they are valid in Thomson’s model, too.

4. Monotonicity of classic cake-cutting protocols

Although resource- and population-monotonicity are well established axioms in various fields of fair division, the cake-cutting literature has not adopted these ideas so far. Moreover, classical division methods like the Banach-Knaster (Steinhaus, 1948), Cut and Choose, Dubins-Spanier (Dubins and Spanier, 1961), Even-Paz (Even and Paz, 1984), Fink (Fink, 1964) or the Selfridge-Conway protocol do not satisfy these axioms. A detailed explanation for most of these can be found in (Brams and Taylor, 1996). For completeness, our survey includes procedures that return disconnected pieces.

All the counterexamples below feature *piecewise homogeneous* cakes. These are finite unions of disjoint intervals, such that on each interval the value densities of all agents are constant (although different agents may evaluate the same piece differently). In such cases, the function $\hat{v}_i([0, x])$ – which displays the value (for agent i) of the piece which lies left to the point $x \in \mathbb{R}$ – is a piecewise-linear function (see Figure 1). Piecewise-homogeneous cakes are interesting on their own (see e.g. Aziz and Ye (2014)), however, we stress that our results, hold for arbitrary cakes - not only for piecewise-homogeneous ones.

Piecewise homogeneous cakes can be represented by a simple table containing the value densities of the agents on the different slices. For example the cake in Figure 1 has the following representation.

\hat{v}_A	2.5	0	2	2	0	0
\hat{v}_B	1	1	0	0	1	1

4.1. Resource-monotonicity

First let us examine the Cut and Choose protocol for two agents. We define the *cut-and-choose rule* as the rule in which one pre-specified agent (say, Alice) cuts the cake into two pieces equal in her eyes, the other agent (Bob) picks the piece that he prefers, and the first agent receives the remaining piece. The following example shows that this rule is not resource-monotonic. In the examples below, the \blacktriangledown sign over a column indicates the enlargement, and the colored cells in an agent’s row indicate the agent’s piece.

▼

\hat{v}_A	1	1	1	1
\hat{v}_B	1	1	3	3

\hat{v}_A	1	1	1	1	2
\hat{v}_B	1	1	3	3	2

When the extra piece is not present (left), Alice cuts the cake after the second slice, allowing Bob to choose the piece worth 6 for him. However, when the cake is enlarged, Alice cuts after the third slice and Bob's utility drops to 5.

This example implies that the Banach-Knaster, Dubins-Spanier, Even-Paz and the Fink methods are not resource-monotonic either, as they all produce the same divisions on the above cake as Cut and Choose.¹

Finally we examine the Selfridge-Conway envy-free protocol for three agents. This protocol has a pre-specified *cutter* (who cuts the cake to three equal pieces) and a pre-specified *trimmer* (who trims his best piece to make it equal to his second-best piece). W.l.o.g, we define the *Selfridge-Conway-rule* as the rule in which Alice is the cutter and Bob is the trimmer. The following example shows that this rule is not RM.

▼

\hat{v}_A	4	2	2	4	4	2
\hat{v}_B	5	2	3	4	1	1
\hat{v}_C	1	2	4	4	1	1

\hat{v}_A	4	2	2	4	4	2	6
\hat{v}_B	5	2	3	4	1	1	1
\hat{v}_C	1	2	4	4	1	1	1

In the smaller cake (left), Alice cuts the cake into three parts worth 6 to her, each made of two adjacent slices. The two most valuable parts for Bob are worth the same (7) so he passes. Then the agents choose in the order Carl, Bob, Alice. Carl's utility is 8.

In the larger cake (right), Alice cuts three parts worth 8 to her, made of 3, 2 and 2 slices. The leftmost part is most valuable for Bob, so he trims it to make it equal to the middle part. Carl takes the uncut part, which is worth 5 for him. Now, Carl divides the remainder to 3 equal pieces and the agents choose in order: Bob, Alice, Carl. Carl receives a piece worth at most 2, so his total utility is at most 7.

4.2. Population-monotonicity

Population-monotonicity is not applicable to protocols with fixed number of agents, such as Cut and Choose and Selfridge-Conway. The following example shows that the Dubins-Spanier moving-knife protocol is not PM. In the next couple of examples the cells of the leaving player are colored gray.

\hat{v}_A	20	1	1	1	10	27
\hat{v}_B	1	20	10	28	1	1
\hat{v}_C	1	1	18	10	29	1

\hat{v}_A	20	1	1	1	10	27
\hat{v}_B	1	20	11	28	1	1
\hat{v}_C	1	1	18	10	29	1

When all three agents are present, Alice is the first to stop the knife and get a piece. In the second round, Bob stops the knife and gets a piece, and finally Carl receives the reminder which is worth for him 40. However, if Bob is not present then Carl will be the first to stop the knife and his value will be 30.

The Even-Paz method and the Banach-Knaster protocol (when agents are ordered Alice-Bob-Carl) produce the same allocations. Thus, none of these three methods is population-monotonic.

Consider now the Fink procedure. This procedure was specifically designed with upwards-population-monotonicity in mind: when a new agent joins an existing division, he takes a proportional share from each of the existing agents,

¹A more recent protocol, the Recursive Cut and Choose, proposed by Tasnádi (2003), violates resource-monotonicity for the same reason.

so all existing agents are weakly worse off; they all participate in supporting the new agent, which is what PM is all about. However, the Fink procedure is not downwards-PM, as the following example shows:

\hat{v}_A	2	2	2	2	1	1	2
\hat{v}_B	0	0	0	4	2	2	4
\hat{v}_C	0	0	0	2	1	1	2

\hat{v}_A	2	2	2	2	1	1	2
\hat{v}_B	0	0	0	4	2	2	4
\hat{v}_C	0	0	0	2	1	1	2

Suppose that initially Alice and Bob use Cut and Choose and Bob is the chooser. He is able to salvage the whole cake according to his own evaluation. Now they divide their pieces into three equal parts, and Carl gets to choose one slice from each of them. Hence, Bob ends up with a piece worth at least 8 for him. But if Alice leaves, then Bob and Carl have to redivide the cake using Cut and Choose. Then, no matter who cuts, Bob ends up with only 6.

The above example seemingly contradicts our claim that upwards-PM implies downwards-PM and vice versa for single valued solutions. However, there is a subtle difference here. The Fink procedure is based on a predefined order of the agents, and it is only upwards monotonic if the new agent is the last in the order. An alternative explanation is to treat the Fink rule as a set-valued rule, which returns $n!$ possible allocations, for all $n!$ possible orderings of the agents. Under this definition, the Fink rule is upwards-PM, but not downwards-PM as shown in the example.

5. Negative Results

When the agents do not care about connectivity, the ideal division rule (at least in terms of fairness) is the *Nash-optimal rule*, which maximizes the product of utilities: it is RM, PM, PO and EF, hence also proportional (Sziklai and Segal-Halevi, 2015). Moreover, for every Nash-optimal allocation there exists a price-vector that is a competitive-equilibrium from equal-incomes (CEEI). Therefore, it makes sense to ask whether Nash-optimality and/or CEEI have all these desirable properties with connectivity too. Unfortunately, the answer is no.

Consider first the Nash-optimal rule and the following cake:

\hat{v}_A	2	2	2	2	2	2
\hat{v}_B	1	1	4	4	1	1

Without connectivity, it gives the two central slices to Bob and the four peripheral slices to Alice. The Nash-welfare is $8 \cdot 8 = 64$. The allocation is EF and PROP. Moreover, it is supported in a competitive-equilibrium from equal-incomes, in which the price of a central slice is 2 and the price of a peripheral slice is 1 and the income of both agents is 4.

In contrast, with connectivity, the Nash-optimal rule is not proportional. To see this, observe that in both of the connected proportional divisions, Alice and Bob each get three slices and a value of 6, so the Nash-welfare is 36. However, when Bob gets four slices and Alice two slices, the Nash-welfare is 40. Hence neither of the two possible proportional allocations is Nash-optimal.

Moreover, with connectivity, a CEEI allocation might not exist at all. Recall that any CEEI allocation is both PO and EF; in following cake, no PO+EF allocation exists:

Example 5.1. [No connected allocation is both PO and EF]:

\hat{v}_A	2	0	3	0	2	0	0
\hat{v}_B	0	0	0	0	0	7	0
\hat{v}_C	0	2	0	2	0	0	3

EF requires to give Bob a part of his 7 slice; PO then requires to give him his entire 7 slice. Carl's piece can then be either at Bob's left or at Bob's right:

- If Carl is at Bob's left and his utility is 2, then the allocation is not PO since Carl can be moved to the rightmost slice and get a utility of 3 without harming any other agent.
- If Carl is at Bob's left and his utility is more than 2, then Alice envies him since he holds her 3 slice while here utility is at most 2.
- If Carl is at Bob's right, then by PO the entire left is given to Alice, but then Carl envies her.

□

Since both the Nash-optimal and the CEEI rules fail in the presence of connectivity requirements, we have to look for different rules. But first we show that, with connectivity requirements, Pareto-optimality and proportionality are incompatible with resource-monotonicity:

Theorem 5.1. *When there are two or more agents with connected utilities, any division rule which is proportional and Pareto-optimal cannot be resource-monotonic.*

Proof. Consider the following cake, where the enlargement is marked by the ▼ sign:

\hat{v}_A	6	0	1	1
\hat{v}_B	0	4	2	2

▼

\hat{v}_A	6	0	1	1	6
\hat{v}_B	0	4	2	2	0

In the smaller cake, any PROP+PO rule must give the leftmost slice to Alice and the rest of the cake to Bob. Hence, Alice's utility is 6 and Bob's utility is 8. In the larger cake, any PROP rule must give Alice a utility of at least 7. This leaves Bob a utility of at most 6. □

Moreover, Pareto-optimality and proportionality are incompatible with population-monotonicity, too:

Theorem 5.2. *When agents have connected utilities, any division rule which is proportional and Pareto-optimal cannot be population-monotonic.*

Proof. Consider again the cake of Example 5.1:

\hat{v}_A	2	0	3	0	2	0	0
\hat{v}_B	0	0	0	0	0	7	0
\hat{v}_C	0	2	0	2	0	0	3

Note that all agents value the entire cake as 7, so by PROP each agent must receive a connected piece with a value of at least $2\frac{2}{3}$. Bob's piece must be in his 7 slice (by PROP) and must contain all this slice and nothing more (by PO).

Carl's piece can be either left or right of Bob's 7 slice. If Carl's piece is at Bob's left, then by PROP it must contain the two 2 slices of Carl. This leaves Alice a utility of at most 2, which violates PROP. So Carl's piece must be the Bob's right and Carl's utility is 3.

This leaves the entire region at Bob's left to Alice. By PO, she receives this entire region and her utility is 7.

Suppose Bob leaves. Now $n = 2$, so Carl must get a value of at least $7/2 = 3.5$, so his piece must touch his middle "2" slice. But this leaves Alice a utility of at most 5. □

Despite the crucial importance of Pareto-optimality in economics, in our case it is problematic: it is incompatible with the stronger fairness criterion of envy-freeness. Even with the weaker criterion of proportionality, it is incompatible with any of the monotonicity axioms.

If we believe that a division rule is fair only if it satisfies both proportionality and monotonicity, we must compromise on efficiency. Some possible compromises are presented in the following section.

6. Positive Results

6.1. *Exactly-proportional rule: PROP+RM+PM*

Our first division rule is both resource- and population-monotonic, but very inefficient. We present it merely as a benchmark for comparison with the more advanced rules that come later.

Definition 6.1. A division X is called *exactly-proportional* if it gives every agent a relative value of exactly $1/n$. Formally: $\forall i \in N : v_i(X_i) = 1/n$.

The *exactly-proportional rule* returns an exactly-proportional division; such a division can be found, for example, using the following variant of the Banach-Knaster procedure (Steinhaus, 1948):

- Every agent marks a point x_i such that $\widehat{v}_i([0, x_i]) = \frac{1}{n}$.
- The procedure selects the leftmost point x_{\min} (breaking ties arbitrarily) and gives $[0, x_{\min}]$ to the agent that made that mark.
- The remaining agents divide the cake recursively in the same way (keeping the fraction $1/n$ fixed).
- The cake that remains after the n -th step is discarded.

Theorem 6.1. *The exactly-proportional rule is proportional and resource-monotonic and population-monotonic, but not weakly-Pareto-optimal.*

Proof. PROP is obvious by definition.

RM holds because when the cake grows/shrinks, all agents receive the same fraction of a larger/smaller whole.

PM holds because when an agent leaves/joins, the remaining agents receive a larger/smaller fraction of the same whole.

The following cake shows that the rule is not WPO:

\widehat{v}_A	2	0
\widehat{v}_B	0	2

An exactly-proportional division must give each agent a utility of exactly 1, yet it is possible to give each agent a utility of 2. □

In essence, the exactly-proportional rule tells the agents “keep your happiness at the minimum proportional level of exactly $1/n$, so that when new resources become available, you can only become happier”. This guarantees PROP and RM and PM, but it is very inefficient.

6.2. *Relative-equitable rule: WPO+PM+PROP*

We present a population-monotonic division rule based on the notion of *equitable cake divisions*. The idea of an equitable division is that all agents are equally happy — each agent receives a piece with the same personal value.

Definition 6.2. (a) A cake division X is called *relative-equitable* if all agents receive exactly the same relative value. Formally: $v_i(X_i) = v_j(X_j)$ for all $i, j \in N$. This value is called the *relative-equitable value* of the division.

(b) A relative-equitable division is called *max-relative-equitable* if its relative-equitable value is weakly larger than of all relative-equitable divisions.

Definition 6.3. (a) An *agent-ordering*, denoted by π , is a permutation on the set of agents N .

(b) A connected partition of the cake into n intervals is called a π -*partition* if the intervals are assigned to the n agents in the order specified by π .

For example, a $\overline{132}$ -partition is a connected partition in which Agent 1 receives the leftmost piece, Agent 3 receives the middle piece and Agent 2 receives the rightmost piece.

Independently of ours, some of the following lemmata were proved by Cechlárová et al. (2013). We mention each lemma that was previously proved. For completeness, we provide alternative proofs.

Lemma 6.1. *For every agent-ordering π , there exists a relative-equitable π -partition.*

Proof. A proof using generalized-inverse functions is given in Theorem 5 of Cechlárová et al. (2013).

A new and shorter proof, using the Borsuk-Ulam theorem, is given in Appendix A. □

Below, we present a moving-knife procedure that finds an equitable-connected division for any ordering of the agents. Note that an equitable-connected division cannot be found by a discrete procedure (Cechlárová and Pillárová, 2012), so a moving-knife procedure is a natural alternative.

Without loss of generality, assume that π is the order $1, \dots, n$. The procedure starts with the agents holding their knives at the leftmost end of the cake. There is a large screen where the current relative-equitable value is displayed, which is zero at the beginning. During the procedure the positions of the knives determine a division of the cake: the piece allotted to Agent i is the piece between the knife of Agent i and the knife of Agent $i - 1$ (or for $i = 1$ the leftmost end of the cake). The value on the screen increases continuously and all agents moves their knives to the right such that the value of each piece matches the value on the screen. This goes on until one of the following two things happen:

- (a) The rightmost knife reaches the end of the cake.
- (b) The knife of an agent reaches the leftmost endpoint of an interval in which the value density of that particular agent is 0.

In case (a), the procedure stops and we have obtained an relative-equitable partition of the entire cake. The relative-equitable value is the value on the screen.

In case (b), the value on the screen is frozen temporarily and the procedure enters its second phase. Let j be the rightmost agent whose piece is adjacent to a zero-value interval. So every agent who comes after j in the predefined order can strictly increase the value of his piece by moving his knife to the right. We ask all agents starting from j to move their knives to the right such that their value remains constant. This goes on until either (a) holds, or another agent $k > j$ reaches an interval of zero measure so (b) holds for that agent, or Agent j reaches the end of his zero-measure interval. In the first case the procedure stops, in the second case the procedure continues at second phase with agent k , in the third case the procedure goes back to the first phase.

Since the rightmost knife is moving continuously and monotonically to the right, eventually it reaches the end of the cake and an equitable division is found.² □

We demonstrate the somewhat informal description of the above moving-knife procedure on an example.

Example 6.1. Consider the piecewise homogeneous cake depicted in Figure 2. Three agents: Green, Red and Blue seek an equitable division of the cake, which has total value of 8 for each of them (this is a special case where the same relative value indicates the same absolute value). They agree on using the above procedure with the order Green, Red, Blue. Immediately at the beginning, we are at case (b) because Blue's knife is at a zero-value region. Thus, we enter phase 2 and Blue's knife moves to $x = 4$. Then we return to phase 1.

²Here we implicitly use the assumption that the value measures are bounded. If $\widehat{v}_i([0, c])$ were infinite then the rightmost knife could move to the right indefinitely without reaching the end of the cake by slowing down.




Figure 2: The moving knife procedure described in Lemma 6.1 applied in a cake-division with three agents: Green (densely striped), Red (sparsely striped), Blue (grid). **Left:** the agents' valuations. **Right:** the resulting equitable division




Figure 3: Proof of Lemma 6.2.

As the knives move to the right, the agents increase the value of their pieces until they reach a relative value of $2/8$. At that moment Green's knife rests at $x = 2$, Red's knife at $x = 3$ and Blue's knife at $x = 4.5$. The value displayed at the screen becomes fixed at this point since Red reached a zero-value interval. As Red gradually increases his piece, Blue moves his knife to the right making sure his value does not change. This continues until Blue himself reaches $x = 6$, which is the start of an interval of zero value. Red stops his knife at $x = 5.5$, but Blue continues until he reaches the right end of the cake. The resulting division is relative-equitable with value $2/8$.

Let X be a certain division of a cake. We denote the smallest relative value obtained by an agent by v_{min}^X and the largest by v_{max}^X , such that for all $i = 1, \dots, n$: $v_{min}^X \leq v_i(X_i) \leq v_{max}^X$. Note that $v_{min}^X = v_{max}^X$ if and only if X is a relative-equitable division.

Lemma 6.2. *Let π be an agent-ordering and X a π -partition of a cake. Let Y be a relative-equitable π -partition of the same cake, having a relative-equitable value v^Y . Then: $v_{min}^X \leq v^Y \leq v_{max}^X$.*

Proof. The proof that $v_{min}^X \leq v^Y$ is in Lemma 7 of Cechlárová et al. (2013); the proof that $v^Y \leq v_{max}^X$ (which is analogous) is in Lemma 1 of Cechlárová and Pillárová (2012). We provide an alternative, graphic proof that that $v_{min}^X \leq v^Y$ (see Figure 3).

Assume w.l.o.g. that π is the ordering $1, \dots, n$. Assume by contradiction that $v^Y < v_{min}^X$. In particular, this means that Agent 1 receives a smaller value in partition Y than in partition X , that is, $v_1(Y_1) < v_1(X_1)$. Hence, the cut-point between pieces Y_1 and Y_2 is to the left of the cut-point between pieces X_1 and X_2 .

The same is true for the n -th agent: $v_n(Y_n) < v_n(X_n)$. Hence the cut-point between pieces Y_{n-1} and Y_n is to the right of the cut-point between pieces X_{n-1} and X_n . Because the leftmost cut-point moved to the left and the rightmost cut-point moved to the right, there must be a pair of adjacent cut-points such that the left one moved to the left and the right one moved to the right (see Figure 3). Hence, there must be an index k , such that:

- The left boundary of piece Y_k is to the left of the left boundary of X_k , and
- The right boundary of Y_k is to the right of the right boundary of X_k .

This means that $Y_k \supset X_k$ which in turn implies that $v_k(Y_k) \geq v_k(X_k)$. This contradicts our assumption that $v^Y < v_{min}^X$. \square

For every agent-ordering π , there may be many different equitable π -partitions. However, all these divisions have the same equitable value:

Lemma 6.3. *For every agent-ordering π , there is a unique value v^π which is the relative-equitable value in all relative-equitable π -partitions.*

Proof. This simple corollary is also proved in Corollary 2 of Cechlárová and Pillárová (2012).

Assume that there are two relative-equitable π -partitions: X with equitable value v^X and Y with equitable value v^Y . By Lemma 6.2, $v^X \leq v^Y \leq v^X$. Hence $v^X = v^Y$. \square

A straightforward corollary of the above lemmata is that the orderings can be sorted by their equitable value. Since for n agents there are finitely many different orderings the following holds.

Corollary 6.1. *There exist max-relative-equitable divisions with connected pieces.*

Define the *relative-equitable rule* as the rule that returns all connected max-relative-equitable divisions of the cake. By Lemma 6.3, this rule is essentially-single-valued.

Lemma 6.4. *The relative-equitable division rule is weakly-Pareto-optimal.*

Proof. Let Y be a max-relative-equitable division with equitable value v^Y . Suppose by contradiction that there is a division X in which the utility of all agents is strictly more than v^Y . Let π be the agent ordering in X . By Lemma 6.2, the relative-equitable-value of the relative-equitable division in ordering π is at least $v_{min}^X > v^Y$. But this contradicts the maximality of Y . \square

Lemma 6.5. *The relative-equitable division rule is population-monotonic.*

Proof. Since the rule is essentially-single-valued, it is sufficient to prove downwards-PM.

Let X be a max-relative-equitable for n agents with equitable value v^X . Suppose that an agent $i \in N$ abandons his share. Give agent i 's piece to an agent that holds an adjacent piece, e.g. to agent $i + 1$. Call the resulting division Y . We obtained a connected division for $n - 1$ agents, in which the smallest value enjoyed by an agent is at least v^X (indeed, the value of all agents except $i + 1$ is exactly v^X , and the value of agent $i + 1$ is at least as large). By Lemma 6.2, the maximum equitable value in the new situation is at least v^X . Hence, in the max-relative-equitable for $n - 1$ agents, the value of all agents is at least as large as in the previous division. \square

Lemma 6.6. *The relative-equitable division rule is proportional.*

Proof. This is also proved in Corollary 1 of Cechlárová and Pillárová (2012).

Let X be any connected proportional division of the cake (by Steinhaus (1948) such a division always exists). Because X is proportional, $v_{min}^X \geq 1/n$. Hence, by Lemma 6.2, the value of a relative-equitable division in the same ordering as X is at least $1/n$. Hence, the maximum relative-equitable value is at least $1/n$. \square

Unfortunately, the relative-equitable rule is not resource-monotonic.

Example 6.2. Consider the following cake, where M is a large constant, $M \gg 2$:

								▼	▼
\hat{v}_A	M	M	1	1					
\hat{v}_B	1	1	M	M					

In the smaller cake, the unique max-relative-equitable division gives the two leftmost slices to Alice and the two rightmost slices to Bob. The relative-equitable value is $M/(M+2)$. However, in the larger cake the unique max-relative-equitable division is attained by cutting exactly in the middle, decreasing the relative-equitable value to $1/2$. While Alice gains from the division and her (absolute) value increases by 1, Bob loses since his absolute value drops from $2M$ to $M+2$.

The following theorem summarizes the properties of the relative-equitable rule:

Theorem 6.2. *The relative-equitable rule is weakly-Pareto-optimal and population-monotonic and proportional, but not resource-monotonic.*

6.3. Absolute-equitable rule: $WPO+PM+RM$

As shown by Example 6.2, the relative-equitable rule is not RM since the relative value of some agents is made smaller when the cake becomes larger. This may imply that, if we use *absolute* instead of relative values, we can get resource-monotonicity.

Fortunately, almost all definitions, examples, procedures, lemmata and proofs from the previous subsection can easily be adapted to absolute values by just replacing “relative” with “absolute”; the only exception is Lemma 6.6.

Theorem 6.3. *The absolute-equitable rule is weakly-Pareto-optimal and population-monotonic and resource-monotonic, but not proportional.*

Proof. WPO holds by Lemma 6.4 and PM by Lemma 6.5, replacing “relative” by “absolute”.

The proof of RM is essentially the same as the proof of Lemma 6.5: the cake enlargement can be treated as a piece that was acquired from an agent who left the scene.³

To see that the rule is not PROP, suppose that Alice values the entire cake as 1 and Bob values the entire cake as $M \gg 2$. Then, any absolute-equitable division must give Bob at most $1/M \ll 1/2$ of his value. \square

6.4. Rightmost-mark rule: $WPO+PROP+RM$

In this section we present a resource-monotonic procedure that produces an envy-free (hence proportional) division of the whole cake for 2 agents, giving each agent a connected piece.

The procedure is called the *rightmost-mark rule* and consists of the following steps.

- Ask both agents to make a mark which cuts the cake in half according to their own valuation. If more than one point satisfies this criterion, i.e. the middle of the cake is worthless to one of the agents, take the rightmost such point.
- Cut the cake at the rightmost mark and give the slice on the right to the agent who made the mark.
- The remaining part is given to the other agent.

Theorem 6.4. *For two agents with connected utilities, the rightmost-mark procedure is envy-free, proportional, weakly Pareto-optimal and resource-monotonic.*

³Note that this argument is not true for the relative-equitable-connected rule. When the cake grows, while the absolute value of all agents weakly increases, the relative value of some agents may decrease. Hence, the relative-equitable value in the enlarged cake might be smaller than in the original cake, and this may make some agents worse-off. See Example 6.2.




Figure 4: Illustration of the rightmost-mark division rule. Alice's cut mark is denoted by α , while Bob's cut mark by β .

Proof. EF and PROP are obvious.

To prove WPO, suppose w.l.o.g. that Bob made the rightmost mark. Suppose we want to give Bob a piece worth strictly more than his current utility of $1/2$. This can be done in two ways. One way is to keep Bob at the right side and move the division line leftwards; this necessarily does not increase Alice's utility. The other way is to switch between Alice and Bob. But then, if Bob's utility is to be improved, he must receive at least Alice's current share (which is worth for him $1/2$). This leaves at most $1/2$ to Alice. Hence, there is no division in which the utilities of both agents are strictly higher.

We now prove that the procedure is RM. Since the rule is single-valued, it is sufficient to prove upwards-RM.

Suppose w.l.o.g. that Bob made the rightmost mark on the smaller cake (see Fig. 4). Thus, Bob obtained the piece marked with B, which is worth exactly half for him. Alice received the part marked with A, which is worth at least half for her. When the cake is enlarged two cases are possible: The order of the cut marks made by the agents remains the same or gets reversed. In the first case, Bob still receives the rightmost cake (marked with B'). Since it is still worth for him half of the cake, and since the cake is enlarged, he is not worse off. Neither is Alice, who receives a piece that contains her original share.

In the second case, Alice receives the rightmost piece A''. Note that she believes that the pieces A'' and B'' represent the same value, and B'' contains A, her original piece. Thus, she is not worse off. Similarly, Bob evaluates A and B the same, and he received B'' which contains A, thus he is not worse off either. \square

Now we show that, in the special case in which the value-densities of both agents are strictly-positive, the rightmost-mark division rule is the *only* rule which is PROP+WPO+RM. First we need a lemma.

Lemma 6.7. *Suppose there are $n = 2$ agents, Alice and Bob, with strictly-positive valuations, and half-points h_A, h_B respectively. Then, any PROP+WPO division rule that allocates connected pieces must:*

- (1) *Cut the cake in the closed interval between h_A and h_B (that is, $[h_A, h_B]$ if $h_A \leq h_B$ or $[h_B, h_A]$ if $h_B \leq h_A$);*
- (2) *Allocate the leftmost piece to the agent with the leftmost half-point (that is, Alice if $h_A \leq h_B$ or Bob if $h_B \leq h_A$).*

Proof. If $h_A = h_B$ then obviously the only PROP allocation is to cut at that point and give a half to each agent.

W.l.o.g, we now assume that $h_A < h_B$. If the cake is cut at $x < h_A$, then the piece to the left of x is worth less than $1/2$ to both agents, so it cannot be given to any of them. Similarly, if the cake is cut at $x > h_B$, then the piece to the right of x is worth less than $1/2$ to both agents. Hence, the cake must be cut at $x \in [h_A, h_B]$.

Since $h_A < h_B$, we must have either $h_A < x$ or $x < h_B$ (or both). In the former case, the piece to the right of x is worth less than $1/2$ to Alice; in the latter case, the piece to the left of x is worth less than $1/2$ to Bob. So in both cases, Alice must get the left and Bob must get the right. □

Theorem 6.5. *For $n = 2$ agents with strictly-positive valuations, any PROP+WPO+RM division rule that allocates connected pieces must allocate the rightmost agent a relative utility of exactly $1/2$.*

Proof. If $h_A = h_B$ then obviously both agents must receive exactly $1/2$. W.l.o.g, we now assume that $h_A < h_B$, so the rightmost agent is Bob. We normalize the agents' valuations to 2. So the cake looks like the following (where $a \in [0, 1]$ and $b \in [0, 1]$ are some constants):

	$[0, h_A)$	$[h_A, h_B)$	$[h_B, 1]$
\widehat{v}_A	1	a	$1 - a$
\widehat{v}_B	b	$1 - b$	1

We claim that a PROP+WPO+RM algorithm must give Bob a value of at most 1. Suppose by contradiction that Bob's value is $1 + 2d$, where $d > 0$ is a constant, $d \in (0, 1 - b)$. Now, the cake grows as follows:

▼

\widehat{v}_A	1	a	$1 - a$	$2a$
\widehat{v}_B	b	$1 - b$	1	d

In the extended cake, h_A moves rightwards and is located exactly between slices #2 and #3. h_B also moves slightly rightwards and is now located inside slice #3. By Lemma 6.7, the cake is cut at or to the right of the new h_A and Bob receives the rightmost piece. Hence, Bob's new value is at most $1 + d$ - in contradiction to RM. □

Theorem 6.5 implies that, when the value-densities are strictly-positive, the rightmost-mark rule is the only rule that satisfies PROP+WPO+RM with connected utilities.

7. Conclusion and Future Work

We studied monotonicity properties in combination with the classical axioms of proportionality and Pareto-optimality. Table 1 summarizes the properties of the various division rules. Most properties are proved in the paper body, except the WPO properties of the classic protocols, which are proved in Appendix B.

Each of our connected division rules satisfies three of the four properties {PROP, WPO, RM, PM}. Thus, the divider has to choose whether to give up proportionality (PROP) or efficiency (WPO) or give up one of the monotonicity properties. We are still missing a rule that satisfies PROP+WPO+RM for three or more agents, as well as a rule that satisfies PROP+WPO+RM+PM for two or more agents. Additionally, combining envy-freeness with monotonicity for three or more agents looks like a fairly challenging task.

In this paper we ignored strategic considerations and assumed that all agents truthfully report their valuations. An interesting future research topic is how to ensure monotonicity in truthful division procedures.

Finally, our procedure for equitable division uses moving knives and thus it is not discrete. Recently, Cechlárová and Pillárová (2011) and Cechlárová and Pillárová (2012) presented discrete procedures that attain approximately-equitable connected divisions. A division rule based on such procedures naturally attains approximate versions of proportionality and monotonicity. Further development of this idea is deferred to future work.

Connected rules	n	CON	EF	PROP	PO	WPO	RM	PM
exact-proportional	Any	Yes	No	Yes	No*	No	Yes	Yes
absolute-equitable	Any	Yes	No	No	No	Y.c.u.	Yes	Yes
relative-equitable	Any	Yes	No	Yes	No*	Y.c.u.	No	Yes
rightmost-mark	2	Yes	Yes	Yes	No*	Y.c.u.	Yes	No

Classic rules	n	CON	EF	PROP	PO	WPO	RM	PM
Banach-Knaster	Any	Yes	No	Yes	No	No	No	No
Cut and Choose	2	Yes	Yes	Yes	No	Y.c.u.	No	No
Dubins-Spanier	Any	Yes	No	Yes	No	No	No	No
Even-Paz	Any	Yes	No	Yes	No	No	No	No
Fink	Any	No	No	Yes	No	No	No	Upw
Selfridge-Conway	3	No	Yes	Yes	No	No	No	No

Table 1: Properties of division rules presented in this paper. **No** means that the property is not satisfied by the rule, while **No*** means that the property cannot be satisfied by *any* rule satisfying the other properties marked with **Yes** in the same line. In the WPO column **Y.c.u.** stands for 'Yes for connected utilities'. In case of the Fink rule, **Upw** indicates that – with some additional adjustments – the rule is upwards-PM.

8. Acknowledgments

The idea of this paper was born in the COST Summer School on Fair Division in Grenoble, 7/2015 (FairDiv-15). We are grateful to COST and the conference organizers for the wonderful opportunity to meet with fellow researchers from around the globe. In particular, we are grateful to Ioannis Caragiannis, Ulle Endriss and Christian Klamler for sharing their insights on cake-cutting. We are also thankful to Marcus Berliant, Shiri Alon-Eron, Christian Blatter and Ilan Nehama for their very helpful comments.

The authors acknowledge the support of the ‘Momentum’ Programme (LP-004/2010) of the Hungarian Academy of Sciences, the Pallas Athene Domus Scientiae Foundation, the OTKA grants K108383 and K109354, the ISF grant 1083/13, the Doctoral Fellowships of Excellence Program, the Wolfson Chair and the Mordecai and Monique Katz Graduate Fellowship Program at Bar-Ilan University. In addition Sziklai was supported by the ÚNKP-16-4-I. New National Excellence Program of the Ministry of Human Capacities.

A. Existence of equitable-connected divisions

The proof uses the *Borsuk-Ulam theorem*⁴. It is about functions defined on spheres. Define the sphere S^{n-1} as the set of points (x_1, \dots, x_n) satisfying: $|x_1| + \dots + |x_n| = 1$ (it is a sphere in the ℓ_1 metric).

Theorem (Borsuk-Ulam). *Let f_i , for $i = 1, \dots, n-1$, be real-valued functions of n variables, that are continuous on the sphere S^{n-1} .*

Then, there exists a point on the sphere, $X^ = (x_1, \dots, x_n) \in S^{n-1}$, such that for all i : $f_i(X^*) = f_i(-X^*)$.*

Assume that the cake is the interval $[0, 1]$. Each point $(x_1, \dots, x_n) \in S^{n-1}$ corresponds to a partition of the cake to n intervals, marked X_1, \dots, X_n , such that the length of interval X_i (the i -th interval from the left) is $|x_i|$. Note that each partition corresponds to many points which differ in the signs of some or all of the coordinates (this representation of the partition space was introduced by Alon and West (1986) and used e.g. by Simmons and Su (2003)).

⁴Independently and contemporaneously to our work, Chéze (2017) came up with a similar idea.

Suppose w.l.o.g. that the players are ordered from 1 to n , so that player i receives the piece X_i . For every point $X = (x_1, \dots, x_n) \in S^{n-1}$ and for every $i \in 1, \dots, n-1$, define the function $f_i(X)$ as follows:

$$f_i(X) = \text{sign}(x_i) \cdot v_i(X_i) - \text{sign}(x_{i+1}) \cdot v_{i+1}(X_{i+1})$$

Note that when $x_i = 0$, interval X_i is empty so $v_i(X_i) = 0$. Hence, the functions f_i are continuous on S^{n-1} .

Hence, by the Borsuk-Ulam theorem, there exists a point X^* on S^{n-1} such that for all i : $f_i(X^*) = f_i(-X^*)$. By definition of the f_i , the cake division that corresponds to X^* necessarily satisfies:

$$\text{sign}(x_i) \cdot v_i(X_i^*) = \text{sign}(x_{i+1}) \cdot v_{i+1}(X_{i+1}^*)$$

This is possible only if for all $i \in (1, \dots, n-1)$:

$$v_i(X_i^*) = v_{i+1}(X_{i+1}^*)$$

Hence the division X^* is equitable.

B. Weak Pareto-optimality of classic protocols

It is a well-known fact that the classic protocols that we discuss here are not Pareto-optimal. Now we show that – with the exception of Cut and Choose – they are not even weakly Pareto-optimal. In the following example the Banach-Knaster, Dubins-Spanier and Even-Paz protocols coincide.

\hat{v}_A	2	0	0	0	0	4
\hat{v}_B	2	3	1	1	5	0
\hat{v}_C	2	3	1	1	5	0

\hat{v}_A	2	0	0	0	0	4
\hat{v}_B	2	3	1	1	5	0
\hat{v}_C	2	3	1	1	5	0

Alice gets the first slice as it composes 1/3 of her cake value, while Bob and Carl perform a Cut and Choose on the rest of the cake. The table on the right presents an alternative allocation (still with connected pieces) which is strictly better for all agents.

The Fink method is not contiguous, hence we can obtain an improvement by composing the pieces from more slices.

\hat{v}_A	0	3	2	1
\hat{v}_B	2	1	2	$1 + \epsilon$

\hat{v}_A	0	3	2	1
\hat{v}_B	2	1	2	$1 + \epsilon$

For two agents the Fink method proceeds as the Cut and Choose: Alice cuts in the middle and Bob chooses the piece on the right. The second cake is a Pareto-improvement with four slices where every agent is strictly better off.

The same example shows that the Cut and Choose is not WPO for additive utilities. However, it is a contiguous protocol, thus it makes sense to investigate whether we could improve it with connected pieces.

Lemma B.1. *The Cut and Choose method is weakly Pareto-optimal whenever the agents have connected utilities.*

Proof. Suppose Alice cuts at point $x \in [0, c]$ and Bob chooses. The cake is divided into the left piece (L) and the right piece (R). Without loss of generality we may assume that Alice receives the left piece. Alice's utility can be improved in only two ways: (a) Cut at some point $x' > x$ and give the left piece to Alice. But then Bob would receive $R' \subset R$, which obviously cannot be strictly better than R . (b) Cut at some point $x'' < x$ and give Alice the right piece $R'' \supset R$. Then Bob receives $L'' \subset L$. Since Bob preferred R to L he did not gain utility by this swap. So it is impossible to strictly increase the utility of Alice and Bob simultaneously. \square

Finally we show that the Selfridge-Conway does not satisfy WPO either. Consider the following cake.

\hat{v}_A	3	1	3	1	2	2
\hat{v}_B	1	3	1	3	1	2
\hat{v}_C	4	0	0	0	0	3

\hat{v}_A	3	1	3	1	2	2
\hat{v}_B	1	3	1	3	1	2
\hat{v}_C	4	0	0	0	0	3

Alice cuts the cake into three equal parts. Since the two most valuable slices have equal value for Bob he passes. The agents choose in the Carl-Bob-Alice order and obtain the pieces shown by the table on the left. However, as the table on the right shows, there is an allocation where every agent is strictly better off.

References

- Alon, N., 1987. Splitting necklaces. *Advances in Mathematics* 63, 247–253.
- Alon, N., West, D.B., 1986. The Borsuk-Ulam theorem and bisection of necklaces. *Proceedings of the American Mathematical Society* 98, 623–628.
- Arzi, O., Aumann, Y., Dombb, Y., 2011. Throw one’s cake - and eat it too, in: Persiano, G. (Ed.), *Algorithmic Game Theory*. Springer Berlin Heidelberg. volume 6982 of *Lecture Notes in Computer Science*, pp. 69–80.
- Aumann, Y., Dombb, Y., 2010. The Efficiency of Fair Division with Connected Pieces. *Web, Internet and Network Economics* 6484, 26–37.
- Aziz, H., Ye, C., 2014. Cake Cutting Algorithms for Piecewise Constant and Piecewise Uniform Valuations, in: Liu, T.Y., Qi, Q., Ye, Y. (Eds.), *Web and Internet Economics*. Springer International Publishing. volume 8877 of *Lecture Notes in Computer Science*, pp. 1–14.
- Balinski, M., Young, H.P., 1975. The quota method of apportionment. *American Mathematical Monthly* 82, 701–730.
- Balinski, M., Young, H.P., 1982. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press, New Haven.
- Berliant, M., Raa, T., 1988. A foundation of location theory: Consumer preferences and demand. *Journal of Economic Theory* 44, 336–353.
- Berliant, M., Thomson, W., Dunz, K., 1992. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics* 21, 201–216.
- Brams, S., Taylor, A., 1996. *Fair Division: From Cake Cutting to Dispute Resolution*. Cambridge University Press, Cambridge UK.
- Brams, S.J., Jones, M.A., Klamler, C., 2006. Better Ways to Cut a Cake. *Notices of the AMS* 53, 1314–1321.
- Brams, S.J., Jones, M.A., Klamler, C., 2007. Better Ways to Cut a Cake - Revisited. mimeo. Dagstuhl Seminar Proceedings 07261.
- Brânzei, S., Miltersen, P.B., 2015. A Dictatorship Theorem for Cake Cutting, in: *Proceedings of the 24th International Conference on Artificial Intelligence, AAAI Press*. pp. 482–488.
- Brânzei, S., Nisan, N., 2017. The Query Complexity of Cake Cutting. Preprint <https://arxiv.org/abs/1705.02946>.
- Calleja, P., Rafels, C., Tijs, S., 2012. Aggregate monotonic stable single-valued solutions for cooperative games. *International Journal of Game Theory* 41, 899–913.
- Caulfield, M.J., 2008. Apportioning representatives in the united states congress. *AMC* 10, 12.
- Cechlárová, K., Doboš, J., Pillárová, E., 2013. On the existence of equitable cake divisions. *Information Sciences* 228, 239–245.
- Cechlárová, K., Pillárová, E., 2011. A near equitable 2-person cake cutting algorithm. *Optimization* 61, 1321–1330.
- Cechlárová, K., Pillárová, E., 2012. On the computability of equitable divisions. *Discrete Optimization* 9, 249–257.
- Chambers, C.P., 2005a. Allocation rules for land division. *Journal of Economic Theory* 121, 236–258.
- Chambers, C.P., 2005b. Allocation rules for land division. *Journal of Economic Theory* 121, 236–258.
- Chéze, G., 2017. Existence of a simple and equitable fair division: A short proof. *Mathematical Social Sciences Short communication*.

- Dall’Aglio, M., 2001. The Dubins-Spanier optimization problem in fair division theory. *Journal of Computational and Applied Mathematics* 130, 17–40.
- Dall’Aglio, M., Maccheroni, F., 2009. Disputed lands. *Games and Economic Behavior* 66, 57–77.
- Dubins, L.E., Spanier, E.H., 1961. How to Cut A Cake Fairly. *The American Mathematical Monthly* 68.
- Even, S., Paz, A., 1984. A Note on Cake Cutting. *Discrete Applied Mathematics* 7, 285–296.
- Fink, A., 1964. A Note on the Fair Division Problem. *Mathematics Magazine* 37, 341–342.
- Ghods, A., Zaharia, M., Hindman, B., Konwinski, A., Shenker, S., Stoica, I., 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types, in: *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, USENIX Association, Berkeley, CA, USA. pp. 323–336.
- Herreiner, D., Puppe, C., 2009. Envy Freeness in Experimental Fair Division Problems. *Theory and Decision* 67, 65–100.
- Hüseinov, F., 2011. A theory of a heterogeneous divisible commodity exchange economy. *Journal of Mathematical Economics* 47, 54–59.
- Jones, M.A., 2002. Equitable, Envy-Free, and Efficient Cake Cutting for Two People and Its Application to Divisible Goods. *Mathematics Magazine* 75, 275–283.
- Legut, J., Potters, J.A.M., Tijs, S.H., 1994. Economies with Land - A Game Theoretical Approach. *Games and Economic Behavior* 6, 416–430.
- Mawet, S., Pereira, O., Petit, C., 2010. Equitable cake cutting without mediator, in: *5th Benelux Workshop on Information and System Security*.
- Moulin, H., 2004. *Fair Division and Collective Welfare*. The MIT Press.
- Nicolò, A., Perea, Roberti, P., 2012. Equal opportunity equivalence in land division. *SERIEs - Journal of the Spanish Economic Association* 3, 133–142.
- Peleg, B., Sudhölter, P., 2007. *Introduction to the Theory of Cooperative Games*. Springer, Heidelberg.
- Procaccia, A., Wang, J., 2017. A Lower Bound for Equitable Cake Cutting. Working Paper W2/2017. Carnegie Mellon University.
- Simmons, F.W., Su, F.E., 2003. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical Social Sciences* 45, 15–25.
- Steinhaus, H., 1948. The problem of fair division. *Econometrica* 16, 101–104.
- Sziklai, B., Segal-Halevi, E., 2015. Resource-monotonicity and Population-monotonicity in Cake-cutting. [arXivpreprint1510.05229](#).
- Tasnádi, A., 2003. A new proportional procedure for the n-person cake-cutting problem. *Economics Bulletin* 4, 1–3.
- Thomson, W., 1997. The Replacement Principle in Economies with Single-Peaked Preferences. *Journal of Economic Theory* 76, 145–168.
- Thomson, W., 2011. Fair Allocation Rules, in: *Handbook of Social Choice and Welfare*. Elsevier BV, pp. 393–506.
- Thomson, W., 2012. On The Axiomatics Of Resource Allocation: Interpreting The Consistency Principle. *Economics and Philosophy* 28, 385–421.
- Varian, H.R., 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* .

Walsh, T., 2011. Online Cake Cutting. *Algorithmic Decision Theory* 6992, 292–305.

Young, H., 1987. On dividing an amount according to individual claims or liabilities. *Mathematics of Operations Research* 12, 65–72.