MAGYAR TUDOMÁNYOS AKADÉMIA
Közgazdaság- és Regionális Tudományi Kutatóközpont

MTA KRTK KTI

Centre for Economic and Regional Studies
HUNGARIAN ACADEMY OF SCIENCES

MŰHELYTANULMÁNYOK DISCUSSION PAPERS

MT-DP – 2015/52

Resource-monotonicity and Population-monotonicity in Cake-cutting

EREL SEGAL-HALEVI – BALÁZS SZIKLAI

INSTITUTE OF ECONOMICS, CENTRE FOR ECONOMIC AND REGIONAL STUDIES,
HUNGARIAN ACADEMY OF SCIENCES BUDAPEST, 2015

Discussion papers
MT-DP – 2015/52

Institute of Economics, Centre for Economic and Regional Studies,
Hungarian Academy of Sciences

Resource-monotonicity and Population-monotonicity in Cake-cutting

Authors:

Balázs Sziklai
junior research fellow
Momentum Game Theory Research Group
Institute of Economics - Centre for Economic and Regional Studies
Hungarian Academy of Sciences
e-mail: sziklai.balazs@krtk.mta.hu

Erel Segal-Halevi
Bar-Ilan University, Ramat-Gan 5290002, Israel
erelsgl@gmail.com

October 2015

ISBN 978-615-5594-18-2
ISSN 1785 377X

Resource-monotonicity and Population-monotonicity in Cake-cutting

Erel Segal-Halevi – Balázs Sziklai

Abstract

We study the monotonicity properties of solutions in the classic problem of fair cake-cutting – dividing a single heterogeneous resource among agents with subjective utilities. Resource- and population-monotonicity relate to scenarios where the cake, or the number of participants who divide the cake, changes. It is required that the utility of all participants change in the same direction: either all of them are better-off (if there is more to share) or all are worse-off (if there is less to share). We formally introduce these concepts to the cake-cutting setting and present a meticulous axiomatic analysis. We show that classical cake-cutting protocols, like the Cut and Choose, Banach-Knaster, Dubins–Spanier and many other fail to be monotonic. We also show that, when the allotted pieces must be contiguous, proportionality and Pareto-optimality are incompatible with each of the monotonicity axioms. We provide a resource-monotonic protocol for two players and show the existence of rules that satisfy various combinations of contiguousness, proportionality, Pareto-optimality and the two monotonicity axioms.

Keywords: resource-monotonicity, population-monotonicity, cake-cutting, leximin divisions, equitable divisions

JEL classification: D63

Erőforrás-monotonitás és népesség-monotonitás a tortaszeletelésben

Erel Segal-Halevi – Sziklai Balázs

Összefoglaló

A megoldások monotonitási tulajdonságait vizsgáljuk a klasszikus tortaszeletelési probléma kapcsán – azaz egy heterogén erőforrást osztunk szét szubjektív hasznossági függvényekkel rendelkező szereplők között. Erőforrás- és népesség-monotonitás olyan esetekben értelmezhető, amelyekben a torta mérete, vagy a tortát elosztó egyének száma megváltozik. Ekkor a résztvevők hasznosságának egy irányban kell változnia: mindenki gyengén jól jár (ha megnőtt a torta mérete), vagy mindenki gyengén rosszabbul jár (ha kevesebb a szétosztható erőforrás). Ezeket a fogalmakat a tortaszeletelési problémák osztályára formálisan is definiáljuk, és a különböző axiómák viszonyát aprólékosan megvizsgáljuk. Megmutatjuk, hogy a klasszikus tortaszelető eljárások, mint a Feloszt–Választ, a Banach–Knaster, a Dubins–Spanier, illetve számos más eljárás is megsérti ezeket a monotonitási tulajdonságokat. Belátjuk továbbá, hogy ha a szeleteknek egybefüggőknek kell lenniük, akkor az arányosság és a Pareto-hatékonyság követelménye kizárja a megoldás monotonitását. Ismertetünk egy kétszereplős erőforrás-monoton eljárást, illetve olyan megoldásokra mutatunk példákat, amelyek az egybefüggőség, az arányosság, a Pareto-hatékonyság és a két monotonitási axióma különböző kombinációit elégítik ki.

Tárgyszavak: erőforrás-monotonitás, népesség-monotonitás, tortaszeletelés, leximin felosztás, egyenlő felosztás

JEL kód: D63

Resource-monotonicity and Population-monotonicity in Cake-cutting

Erel Segal-Halevi

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

Balázs Sziklai

Centre for Economic and Regional Studies, Hungarian Academy of Sciences, E-mail: sziklai.balazs@krtk.mta.hu

October 19, 2015

Abstract

We study the monotonicity properties of solutions in the classic problem of fair cake-cutting - dividing a single heterogeneous resource among agents with subjective utilities. Resource- and population-monotonicity relate to scenarios where the cake, or the number of participants who divide the cake, changes. It is required that the utility of all participants change in the same direction: either all of them are better-off (if there is more to share) or all are worse-off (if there is less to share). We formally introduce these concepts to the cake-cutting setting and present a meticulous axiomatic analysis. We show that classical cake-cutting protocols, like the Cut and Choose, Banach-Knaster, Dubins-Spanier and many other fail to be monotonic. We also show that, when the allotted pieces must be contiguous, proportionality and Pareto-optimality are incompatible with each of the monotonicity axioms. We provide a resource-monotonic protocol for two players and show the existence of rules that satisfy various combinations of contiguousness, proportionality, Pareto-optimality and the two monotonicity axioms.

1 Introduction

Two siblings Ann and Miriam inherit a piece of land from their father. Just as soon as they agree on how to share the land, they learn that their father recently bought another strip of land. They re-divide the whole estate. However Miriam ends up with a more valuable piece of land while Ann is worse off. Ann feels that they should try another division protocol.

After a little reasoning Ann and Miriam manage to divide the estate. Just as soon as they settle the issue, they discover that they have a little step-brother. They decide to divide the land threeways by using the same division rule as before. However Ann ends up with a more valuable piece of land while Miriam is worse off. Miriam feels that they should try another division protocol.

These two examples motivate the so called *resource-monotonicity (RM)* and *population-monotonicity (PM)* axioms. Resource-monotonicity, sometimes known as aggregate monotonicity, requires that when new resources are added, and the same division rule is used consistently, the welfare of the participants should weakly increase. Population-monotonicity is concerned with changes in the number of participants. No one should profit from the arrival of a new agent, when more people share the same resource, and everyone should be weakly better off if someone leaves.

In both cases fairness stems from solidarity. If there is less to share some agents will necessarily lose. Hence it is only fair to require that the loss is (weakly) shared by all agents. On the other hand economic growth might be objected if the benefits are not shared by everyone[1].

We study the two monotonicity requirements in the framework of the classic *fair cake-cutting* problem (Steinhaus, 1948), where a single heterogeneous resource - such as land or time - has to be divided fairly. Cake-cutting protocols can be applied in inheritance cases and divorce settlements. They can be also used to divide broadcast time of advertisements and priority access time for customers of an Internet service provider (Caragiannis et al., 2011; Procaccia, 2015).

Experimental studies show that people value certain fairness criteria more than others. Herreiner and Puppe (2009) demonstrated that people are willing to sacrifice Pareto-efficiency in order to reach an envy free allocation. To our knowledge no study was ever conducted to unfold the relationship between monotonicity and efficiency or proportionality. However some indirect evidence points toward that monotonicity of the solution is in some cases as

[1]A recurring theme in recent social movements like the "99%" and "Occupy Wall Street" is that although the current economic system indeed brings economic growth, it does only to a small fraction of the population, while the other 99% becomes poorer.

important as proportionality.

In the apportionment problem there is a parliament with a fixed number of seats and administrative regions with different number of voters. The question is – assuming that electoral districts are of equal size within each of the administrative regions – how to distribute the seats among the administrative regions to minimize the differences between the voters influences. The problem is analogous to cake-cutting where the cake corresponds to the parliamentary seats which have to be distributed. During the 1880 US census C.W. Seaton, a Chief Clerk of the Census Office, noted that an enlargement of the House of Representatives from 299 to 300 would result in a loss of seat for State Alabama. This anomaly together with the later discovered population and new state paradoxes pressed the legislators to adopt newer and newer apportionment rules. The currently used seat distribution method is free from such anomalies, however it does not satisfy the so called Hare-quota, a basic guarantee of proportionality (Balinski and Young, 1975).

Here we present a systematic analysis of the relationship of resource and population-monotonicity and classical axioms such as Pareto-efficiency and proportionality. To the best of our knowledge, this is the first paper that studies these questions in a cake-cutting setting[2].

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 resource-monotonicity. Moreover, we show that, if each agent must receive a contiguous piece, then each of the monotonicity properties is incompatible with Pareto-optimality.

On the other hand, we present several division rules that satisfy some - but not all - of the properties desired from a cake division:

- The *rightmost-mark* protocol for two agents, which is resource-monotonic with contiguous pieces;
- The *max-equitable-connected* rule for n agents, which is population-monotonic with contiguous pieces;
- Two variants of the *leximin* rule with disconnected pieces, which are Pareto-optimal and population-monotonic, in addition one of them satisfies proportionality while the other is resource-monotonic.

[2]Although the idea itself is not new. Berliant et al. (1992) already mentioned in the conclusion that "...there are a number of important issues that should be tackled next pertaining, in particular, to the existence of selections from the no-envy solution satisfying additional properties, Examples are monotonicity with respect to the amount to be divided (all agents should benefit from such an increase), and with respect to changes in the number of claimants (all agents initially present should lose in such circumstances)."

The equitable and leximin rules belong to the cardinal welfarism framework (cf. chapter 3 of Moulin (2004)). They rely on inter-agent utility comparison, and make 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 leximin/equitable divisions.

Although these two rules have several drawbacks, such as high computational complexity and strategic issues that make them hard to implement in real life, their mere existence shows that it is possible to combine monotonicity with certain fairness axioms.

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 classical cake-cutting protocols and shows that they are not monotonic. Section 5 generalizes these examples and proves that monotonicity is impossible to attain in combination with contiguousness, Pareto-optimality and proportionality. Section 6 shows that, if one of these requirements is relaxed, monotonic division rules exist. The rules are formally defined and their properties are proved. 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.

Monotonicity issues have been extensively studied with respect to cooperative game theory (Calleja et al., 2012), political representation (Balinski and Young, 1982) and other fair division problems: Chapter 6 of (Moulin, 2004) and Chapter 7 of Thomson (2011) provide extensive reviews of monotonicity axioms.

Thomson (1997) defines the *replacement principle*, which requires that, whenever any change happens in the environment, the welfare of all agents not responsible for the change should be affected in the same direction - 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 most general way of expressing the idea of *solidarity* among agents. The PM and RM axioms are special cases of this principle.

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.

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 (2005) 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. 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.

3 Model

3.1 Cake-cutting

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

- $N = \{1, 2, \ldots, n\}$ 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, which is assumed to be an interval in $\mathbb{R}$. By convention the left endpoint of this interval is the origin, that is $C = [0, c]$ for some real number c. We call a Lebesgue-measureable subset of C a *slice*. A slice that is allotted to an agent is called a *piece*.

- $\hat{v}_i^C$ is the value measure of agent i, which is a real-valued function on the slices of C. When it does not cause confusion we will omit C from the upper index, and simply write $\hat{v}_i$.

As the term "measure" implies, the value measures of all agents are countably additive. Moreover, we assume that the value measures are non-negative, monotonic, bounded and non-atomic. That is, $\hat{v}_i$ assigns a non-negative, but finite number to each slice of C. The value measures of a union of disjoint slices is just the sum of the values of the slices. Note that non-negativity and additivity already imply monotonicity, i.e. if S, T are two slices and $S \subseteq T$ then $\hat{v}_i(S) \leq \hat{v}_i(T)$, for all $i \in N$. Finally, non-atomicity means that any slice with a Lebesgue-measure of zero has value zero. In particular any point on the interval is worthless for the agents. These assumptions are standard in the cake-cutting literature.

Our model diverges from the standard cake-cutting setup is 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 higher or lower. Hence, we differentiate between *absolute* and *relative* value measures:

- The absolute value measure of the entire cake, $\hat{v}_i(C)$, 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^C and defined by: $v_i^C(S) = \hat{v}_i^C(S)/\hat{v}_i^C(C)$. Again, for convenience's sake we will write simply $v_i(S)$ when there is only one cake.

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. As we noted before cake-cutting problems can be studied from a strategic angle. Here, however, we will not analyze the strategic behaviour of the agents.

The aim is to divide the cake into n pairwise-disjoint slices. A *division rule* is a function that takes a cake-cutting problem as input and returns a *division* $X = (X_1, \ldots, X_n)$, or a set of divisions. Note that a division does not necessarily compose a partition of C (i.e. free disposal is assumed). A division rule R is called *essentially single-valued* if $X, Y \in R(\Gamma)$ implies that $\hat{v}_i(X_i) = \hat{v}_i(Y_i)$ for all $i \in N$. That is, even if R returns a set of divisions, all agents are indifferent between these divisions because they have exactly the same value-vector.

The classical requirements of fair cake-cutting are the following.

- A division X is called *proportional* (PROP) if each agent gets at least $1/n$ fraction of the cake according to his own evaluation, i.e. $v_i(X_i) \geq 1/n$ for all $i \in N$. Note that the definition uses relative values.
- A division 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.
- A division X is called *contiguous* (CON) if each agent receives a connected piece (i.e. each X_i is an interval).
- A division X is called *whole* (WH) if the entire cake is divided, $\cup_{i=1}^{n} X_i = C$.

Often, we are interested only in contiguous divisions. In these cases, Pareto-optimality might not make sense as the augmenting division might not consist of contiguous pieces.

- We call a division *Pareto-optimal for Contiguous Allocations* (POCA) if there is no division made up of *contiguous* pieces which is weakly better for all agents and strictly better for at least one agent.

A division rule is called *proportional* it returns only proportional divisions. The same applies to Pareto-optimality, contiguousness and wholeness.

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 let $(\hat{v}_i^C)_{i \in N}$, $(\hat{v}_i^{C'})_{i \in N}$ be two value measures such that $(\hat{v}_i^{C'})_{i \in N}$ coincides with $(\hat{v}_i^C)_{i \in N}$ on the Lebesgue-measurable subsets of C. The cake-cutting problem $\Gamma' = (N, C', (\hat{v}_i^{C'})_{i \in N})$ is called a *cake-enlargement* of the problem $\Gamma = (N, C, (\hat{v}_i^C)_{i \in N})$.

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 6.3) are valid 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.

The definition of resource-monotonicity in Thomson (2011) requires *all* divisions in $R(\Gamma)$ 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 idea 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 principle even if the rule suggests many non-monotonic allocations as well. This distinction will be important later on in case of set valued solutions. For essentially single-valued solutions, however, the two approaches are equivalent. Note that the definition of RM cares only about absolute values. The relative value of an agent, $v_i(X_i)$, is allowed to decrease.

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.

For essentially-single-valued solutions, downwards resource (or population) monotonicity implies upwards resource (or population) monotonicity

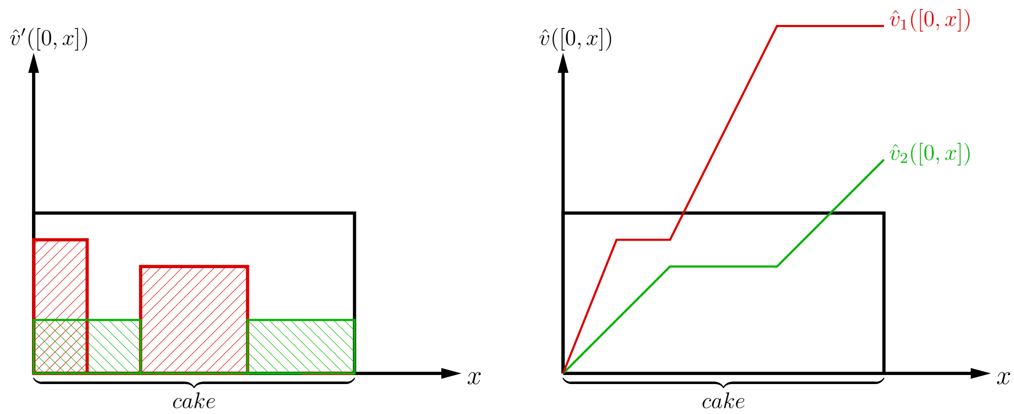

Figure 1: Piecewise homogeneous cake with two players. The derivative of $\hat{v}([0,x])$ indicates the density of the utility. Note that the value measures are not normalized, hence $\hat{v}_1([0,c]) \neq \hat{v}_2([0,c])$.

and vice versa. Set valued solutions, however, may satisfy only one direction of these axioms.

4 Monotonicity and classical 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 all do not satisfy these axioms. A detailed explanation for most of these can be found in (Brams and Taylor, 1996). Admittedly this list is incomplete, still it illustrates our point well.

All the counterexamples below feature *piecewise homogeneous* cakes. Such cakes 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 piecewise linear (see Figure 1). We stress that our division rules, presented in the next section, hold for arbitrary cakes - not only piecewise-homogeneous. The use of piecewise-homogeneous cakes is only for counter-examples.

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

First let us examine why Cut and Choose is not resource-monotonic. In the next couple of examples the ▼ sign over a column indicates the enlargement.

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

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

Table 1: Colored cells indicate the pieces of the agents.

Consider the cake described in Table 1. While the extra piece is not present, Agent A (Alice) cuts the cake after the second slice, allowing Agent B (Bob) to choose the piece worth 5 for him. However if we enlarge the cake, Alice cuts after the third slice and no matter which piece Bob chooses, he loses utility. This simple counterexample implies that the Banach-Knaster, Dubins-Spanier, Even-Paz and the Fink methods are not resource-monotonic either, as they all produce the same division on the above cake as the Cut and Choose.[3]

Note that the fact that we enlarge the cake on the right-hand side is not an issue here. As the next example shows the Dubins-Spanier method violates resource-monotonicity even if we reverse its order, i.e. we move the knives and execute the cuts from right to left.

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

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

Before the enlargement, Bob got the rightmost slice and Carl got the middle slice, leaving Alice with a piece worth 4 for her. With the extra piece present, Alice picks the rightmost slice[4]. Similar reasoning shows that the reverse Even-Paz is not resource-monotonic either.

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

[4] Here we rely on the assumption that the value measures are private information, otherwise Alice would be prompted to act strategically.

Finally, the Selfridge-Conway protocol - the only envy-free method we survey here - is not resource-monotonic either. Consider the following cake.

$\hat{v}_A$	4	2	2	4	6
$\hat{v}_B$	4	2	2	4	0
$\hat{v}_C$	0	0	2	4	0

						▼	▼
$\hat{v}_A$	4	2	2	4	6	6	12
$\hat{v}_B$	4	2	2	4	0	0	0
$\hat{v}_C$	0	0	2	4	0	0	0

Alice divides the cake into three equal parts: the two leftmost slices, the third and fourth slices, and the rightmost slice. The two most valuable parts for Bob are worth the same so he passes. Then the agents choose in the order Carl, Bob, Alice and each obtain 6 units of utility. However if the cake is enlarged, Bob finds only one part of the cake valuable, so he puts that aside. Carl takes the uncut part (which is worthless for him). Then Bob cuts the remaining part into three equal parts. The agents take one slice each in the order Carl, Alice, Bob. This way, Carl ends up with a piece worth only 4 for him.

Population-monotonicity is not applicable to protocols with fixed number of agents, such as Cut and Choose and Selfridge-Conway. For symmetry reasons, reversing the direction of how a protocol processes the cake has no consequence on whether the protocol violates population-monotonicity or not. In the next couple of examples the cells of the leaving player are colored gray.

First we show that the Banach-Knaster method is not population-monotonic. In the following table ε is a very small positive real number.

$\hat{v}_A$	2	2	2	1	1
$\hat{v}_B$	0	$2+\varepsilon$	2	1	1
$\hat{v}_C$	0	0	0	6	6
$\hat{v}_D$	$1+\varepsilon$	0	1	1	1

$\hat{v}_A$	2	2	2	1	1
$\hat{v}_B$	0	$2+\varepsilon$	2	1	1
$\hat{v}_C$	0	0	0	6	6
$\hat{v}_D$	$1+\varepsilon$	0	1	1	1

Suppose the order of the agents is: Alice, Bob, Carl, David. While Alice is present, she starts the process and she cuts after the first slice. Since David believes that this slice is slightly larger than 1/4 fraction of the cake, he will take it. Next, Alice cuts after the second slice. Bob thinks it is slightly bigger than his proportional share, so he diminishes ε part of it. No one else after him trims the slice, hence he takes it. Only Alice and Carl remain. They apply Cut and Choose. Since Alice cuts, Carl can take the piece worth 12 for him.

If Alice leaves the scene, Bob starts the process. He cuts after the second slice and he takes it too. The rest is divided with Cut and Choose and Carl loses utility. Note that Carl would lose utility even if David were to cut.

For the Dubins-Spanier moving-knife method, consider the following cake.

$\hat{v}_A$	2	2	1	1	2
$\hat{v}_B$	0	0	3	1	8
$\hat{v}_C$	0	0	1	1	2
$\hat{v}_D$	2	2	0	0	8

$\hat{v}_A$	2	2	1	1	2
$\hat{v}_B$	0	0	3	1	8
$\hat{v}_C$	0	0	1	1	2
$\hat{v}_D$	2	2	0	0	8

If we apply the Dubins-Spanier moving knife methods, then Alice is the first to stop the algorithm and takes the first slice. Since the next section of the cake is not that valuable for David, Bob and Carl will be the next two cutters. Thus David receives a piece worth 8 for him. However if Alice leaves the process, then David will be the first to stop the knife, taking a piece of value 4.

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, 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.

Finally we show that the Even-Paz method is not population-monotonic.

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

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

Since there are three of them, the agents divide the above cake into left and right part with ratio 1 : 2. Since the leftmost mark belongs to Alice, she gets the leftmost piece. Bob and Carl continue to cut the rest of the cake with ratio 1 : 1 and obtain pieces worth 3 and 4 to them respectively. If Bob leaves, then Carl and Alice divide the cake with ratio 1 : 1. The leftmost mark belongs to Carl so he gets the leftmost piece which is worth only 3 for him.

5 Impossibility Results

In this short section we present two impossibility results. Both of them relate to cake-cutting with contiguous pieces.

5.1 CON+PROP+POCA+RM are incompatible

The following cake shows that, when connectivity is required, any division rule which is proportional and Pareto-optimal cannot be resource-monotonic, even when there are only two agents.

Example 5.1. Consider the following cake, where the enlargement is marked by the ▼ sign:

$\hat{v}_A$	8	0	2
$\hat{v}_B$	0	5	5

				▼
$\hat{v}_A$	8	0	2	8
$\hat{v}_B$	0	5	5	0

In the smaller cake, any rule that is CON+PROP+POCA must cut after the first slice. No other division is possible. Hence, Alice gets 8 and Bob gets 10. In the larger cake, by proportionality we must give Alice at least 9. Regardless of how it is done, Bob gets less than 10.

5.2 CON+PROP+POCA+PM are incompatible

The following cake shows that, when connectivity is required, any algorithm which is proportional and Pareto-optimal cannot be population-monotonic.

$\hat{v}_A$	0	5	0	2	0	0
$\hat{v}_B$	0	0	0	0	7	0
$\hat{v}_C$	2	0	2	0	0	3

Example 5.2. Consider the following valuations, where the value of the entire cake is 7 for all agents:

We first prove that the agents must receive pieces in the following order, from left to right: Alice, Bob, Carl. Note that by PROP, each agent must receive a value of at least 7/3.

- If Carl receives the left piece, then by PROP the piece must also touch its middle "2" slice. But then, only a value of 2 is left for Alice. By a similar argument, Bob cannot receive the left piece. Hence, the leftmost piece must be given to Alice.
- If Carl receives the middle piece, then by PROP it must touch its rightmost "3" slice. But this leaves no value for Bob. Hence, the middle piece must be given to Bob (and it must contain Bob's "7" slice).
- this leaves Carl with the rightmost slice and a value of 3. By POCA, Alice must receive all the cake to the left of the "7" slice of Bob. This means that her value 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 entails that Alice receives a value of at most 5.

Note that without POCA there is no contradiction. The initial value of Alice could be less than 7. Indeed, in the next section we show that CON+PROP+PM can be satisfied if Pareto-optimality is not required.

6 Existence Results

Let us remind the reader that v_i denotes the relative value measure. The value $v_i(X)$ can be thought of the percentage of utility Agent i realizes (compared to the utility of the whole cake) if he receives piece X.

6.1 Equitable divisions

We present a population-monotonic division rule which is based on the notion of *equitable cake divisions.*

Definition 6.1. (a) A cake division X is called *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 *equitable value* of the division.

(b) An equitable division is called *max-equitable* if its equitable value is maximal among all equitable divisions.

(c) An equitable division with contiguous pieces is called *max-equitable-connected* if its equitable value is maximal among all equitable divisions with contiguous pieces.

Max-equitable-connected divisions for two agents were studied by Jones (2002). The generalization for n agents was mentioned by Brams et al. (2006), but only for a special case in which the valuations of all agents are *pairwise-absolutely-continuous*. This means that all valuations have the same support, i.e. if a piece of cake has positive value to one agent then it has positive value for all agents. This is a strong assumption which we do not make here. We now show that a max-equitable division exists even with general valuations.

6.1.1 Exactly-proportional divisions

A special case of equitable solutions is where each agent receives exactly his proportional share. As we will see such divisions are both resource- and population-monotonic, but not Pareto-optimal.

Definition 6.2. A division X is called *exactly-proportional* if it gives every agent a relative value of exactly $\frac{1}{n}$. Formally: $v_i(X_i) = \frac{1}{n}$.

Every division rule that returns only exactly-proportional divisions obviously satisfies PROP. Moreover, it satisfies RM+PM:

- 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.

Exactly-proportional divisions exist both with connected and disconnected pieces, but they are not very efficient. When the pieces must be connected, it is not always possible to divide the entire cake in an exactly-proportional way.

Example 6.1. Suppose every agent wants a distinct sub-interval of the cake:

$\hat{v}_A$	2	0
$\hat{v}_B$	0	2

An exactly-proportional division must give each agent a value of exactly 1, yet every whole connected division gives one of them a value of 2.

An exactly-proportional division can be attained only if we allow to discard some of the cake (i.e. we give up WH). We define the *exactly-proportional-connected* division rule as the following variant of the Dubins-Spanier procedure:

- Every agent marks a point x_i such that $\hat{v}_i([0, x_i]) = \frac{1}{n}$.
- 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 exactly-proportional-connected division rule is not WH since there is a remainder at the rightmost end of the cake. However, it is CON, RM and PM.

When the pieces do not have to be connected, it is possible to divide the entire cake in an exactly-proportional way. This is a corollary of the Dubins-Spanier theorem (Dubins and Spanier, 1961). The exactly-proportional division rule without connectivity is WH, PROP, RM and PM. Moreover, there is a moving-knife procedure that actually finds an exactly-proportional division; see Austin (1982) and Brams and Taylor (1996).

However, the exactly-proportional rule is not Pareto-optimal. This is shown by Example 6.1, in which the exactly-proportional rule gives both agents a value of 1 while the Pareto-optimal division gives them 2. More efficient division rules are considered in the following subsections.

6.1.2 Max-equitable-connected divisions

In this subsection we show that there exists a max-equitable-connected solution, which satisfies CON, PROP, WH and PM.

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

(b) A 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 division in which Agent 1 receives the leftmost piece, Agent 3 receives the middle piece and Agent 2 receives the rightmost piece.

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

Proof. We present here a moving-knife procedure that finds the desired partition. A formal proof, using the famous Borsuk-Ulam theorem, is given in Appendix A. Without loss of generality we may assume that π is the order $1, \ldots, n$. The procedure starts with the agents holding their knives at the leftmost end of the cake. There is a large screen where the current 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 allocated to agent i is the piece between knife of Agent i and the knife of Agent $i-1$ (or for $i = 1$ the leftmost end of the cake). We assume that the valuations are common knowledge. As the value increases on the screen each agent moves his knife to the right until one of the following two things happen:

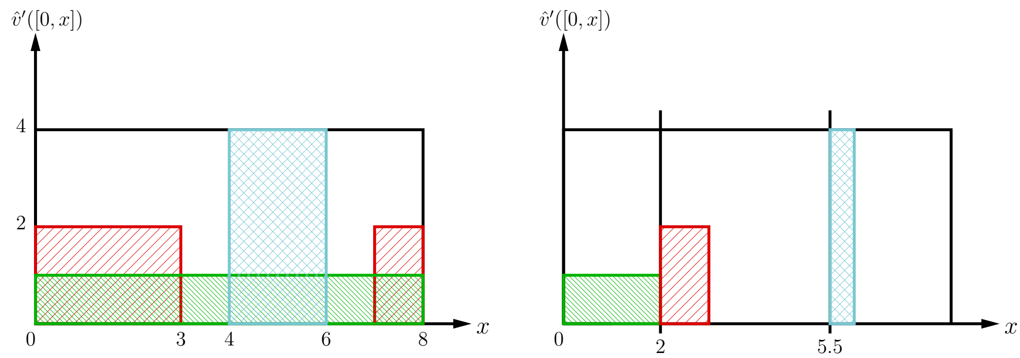

Figure 2: The moving knife procedure described in Lemma 6.1 applied in a cake-division with three agents: Green (densely striped), Red (sparsely striped), Blue (grid).

(a) The rightmost knife reaches the end of the cake.

(b) The knife of an agent reaches an interval for which the value measure of that particular agent is 0.

In the first case the procedure stops and we have obtained an equitable partition of the entire cake. In the second case we stop the knives temporarily, freeze the value on the screen and the procedure enters the second phase. Let j be the rightmost agent whose piece is adjacent to an interval which has a zero measure. Every agent who comes after j in the predefined order can 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. We continue till either (a) holds, or another agent reaches an interval of zero measure, or Agent j reaches the end of his zero measure interval. In the first case the procedure stops, in the second case we continue the second phase with this new agent, in the third case we resume the first phase. Since the rightmost knife is moving continuously and monotonically to the right, eventually it will reach the end of the cake[5]. □

Example 6.2. We demonstrate the somewhat informal description of the above moving-knife procedure on an example. 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. Suppose they agreed on using the above procedure with the Green, Red, Blue order. Immediately at the beginning, we are at option (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. As the knives move to the right, the agents increase the value of their pieces until they reach a utility of 2. 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 an interval which has no value for him. 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.

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

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

Proof. We prove that $v^X_{min} \le v^Y$, the proof for $v^Y \le v^X_{max}$ is analogous. Assume w.l.o.g. that π is the ordering $1, \ldots, n$. Assume by contradiction that $v^Y < v^X_{min}$. 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.

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

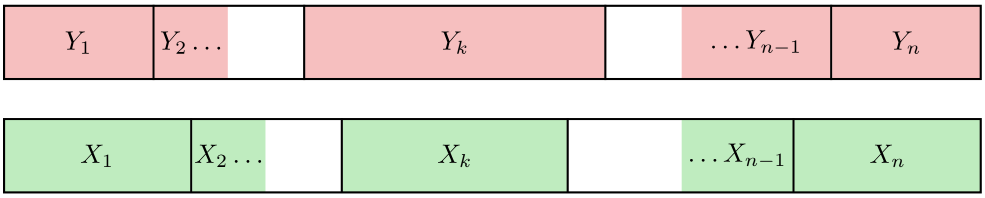

Figure 3: Proof of Lemma 6.2.

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 $v^Y < v^X_{min}$. □

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 equitable value in all equitable π-partitions.*

Proof. Assume that there are two 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 $v^X = v^Y$. □

Lemma 6.4. *There exists an agent-ordering π for which the equitable value v^π is at least $\frac{1}{n}$.*

Proof. Let X be a contiguous proportional division of the cake (by Steinhaus (1948) such a division always exists). Because X is proportional, $v^X_{min} \geq \frac{1}{n}$. Hence, by Lemma 6.2, the value of an equitable division in the same ordering as X is at least $\frac{1}{n}$. □

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. *A max-equitable-connected division always exists.*

We define the *max-equitable-connected* division rule as the rule that returns all max-equitable-connected divisions of the cake.

Theorem 6.1. *The max-equitable-connected division rule is contiguous, proportional, whole, and population-monotonic.*

Proof. Contiguousness and wholeness hold by design, while proportionality due to Lemma 6.4.

We now prove PM. Note that the max-equitable-connected rule is essentially single-valued, so it is sufficient to prove downwards-PM. Let X be a max-equitable-connected division 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 whole and contiguous division for $n-1$ agents, in which the smallest relative 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 probably larger). By Lemma 6.2, the maximum equitable value in the new situation is at least v^X. Hence, in the max-equitable-connected division for $n-1$ agents, the relative value of all agents is weakly larger than in the previous division. The value of the entire cake has not changed, so the absolute value of all remaining agents is weakly larger too. □

Remark 6.1. Thomson (2011) discusses the *egalitarian-equivalent* division rule for resource allocation and proves that it is resource-monotonic. The egalitarian-equivalent rule is similar in spirit to the max-equitable-connected rule (although the setting is different - multiple homogeneous resources versus a single heterogeneous resource); hence it is not surprising that the max-equitable rule satisfies RM.

The max-equitable-connected division rule has drawbacks as well. It does not satisfy POCA (cf. Example 5.2) except in cases when the value measures are pairwise-absolutely-continuous, i.e. have the same support (Brams et al., 2006). As the next example shows it also violates RM.

Example 6.3. Consider the following cake.

$\hat{v}_A$	5	5	1	1
$\hat{v}_B$	1	1	5	5

					▼	▼
$\hat{v}_A$	5	5	1	1	5	5
$\hat{v}_B$	1	1	5	5	1	1

In the smaller cake, the unique max-equitable-connected division gives the two leftmost slices to Alice and the two rightmost slices to Bob. The equitable value is 10/12. However in the larger cake the unique max-equitable-connected division is attained by cutting exactly in the middle, decreasing the equitable value to 1/2. While Alice gains from the division and her (absolute) value increases to 11, Bob loses since his absolute value drops to 7.

6.2 Leximin-optimal divisions: PROP+PO+PM

Lexicographic optimization as a solution concept is used in various fields of fair division (Moulin, 2004). Dubins and Spanier (1961) were the first to employ it in a cake-cutting setting. Later, Dall'Aglio (2001); Dall'Aglio and Hill (2003); Dall'Aglio and Di Luca (2012) presented various properties and approximation algorithms for finding such leximin divisions.

The leximin idea is simple: First we narrow down the solution set to divisions that maximize the utility of the poorest agent. Then among those solutions that satisfy this criterion, we select those which maximize the utility of the second poorest agent. We repeat this process with the third, fourth, etc. poorest agent. To introduce this concept formally first we define lexicographical ordering of real vectors.

We say that vector $y \in \mathbb{R}^n$ is *lexicographically greater* than $x \in \mathbb{R}^n$ (denoted by $x \prec y$) if $x \neq y$ and there exists a number $1 \leq j \leq n$ such that $x_i = y_i$ if $i < j$ and $x_j < y_j$.

Definition 6.4. (a) For a cake division X, define the *relative-leximin-welfare vector* as a vector of length n which contains the relative values of the agents under division X in a non-decreasing order.

(b) A cake division Y is said to be *relative-leximin-better* than X if the relative-leximin-welfare vector of Y is lexicographically greater than the relative-leximin-welfare vector of X.

(c) A cake division X is called *relative-leximin-optimal* if no other division is relative-leximin-better than X.

(d) We define the *relative-leximin division rule* as the rule that returns all relative-leximin-optimal divisions of the cake.

The terms *absolute-leximin-welfare vector*, *absolute-leximin-welfare better* and *absolute-leximin-welfare optimal* and the *absolute-leximin division rule* are defined analogously.

Example 6.4. In the following cake, the absolute-leximin division rule splits the leftmost slice between Alice and Bob, giving each of them 6. The rightmost slice is given to Carl. Thus, the absolute-leximin-optimal vector is

(6,6,9). The relative value vector of the above division is $(6/12, 6/12, 9/30)$.

$\hat{v}_A$	12	0
$\hat{v}_B$	12	0
$\hat{v}_C$	21	9

The corresponding relative-leximin welfare vector is $(9/30, 6/12, 6/12) = (3/10, 1/2, 1/2)$, which is not optimal. The relative-leximin rule divides the leftmost slice between all three agents, giving 1/6 to Carl (absolute value 3.5) and 5/12 to Alice and Bob (absolute value 5). The rightmost slice is given to Carl. The relative-leximin-optimal vector is $(5/12, 5/12, 5/12)$.

Dubins and Spanier (1961) use the compactness of the space of value-matrices to prove that leximin-optimal cake divisions always exist. The proof is equally valid for absolute and relative leximin-optimal divisions. Hence, both the absolute- and the relative- leximin division rules are well-defined.

We now prove some properties which are common to both absolute- and relative- leximin rules.

Lemma 6.5. *The absolute- and the relative-leximin division rules are both Pareto-optimal.*

Proof. By definition, if a division Y Pareto-dominates a division X, then both the absolute and the relative leximin-welfare vectors of Y are larger than those of X. □

Given a division X, we say that agent i is *relatively-poorer* than agent j if $v_i(X_i) < v_j(X_j)$ and *relatively-richer* if $v_i(X_i) > v_j(X_j)$. We say that i is *weakly-relatively-poorer* than j if $v_i(X_i) \leq v_j(X_j)$ and *weakly-relatively-richer* if $v_i(X_i) \geq v_j(X_j)$. The terms *(weakly-)absolutely-poorer* and *(weakly-)absolutely-richer* are defined analogously based on the absolute values $\hat{v}_i, \hat{v}_j$.

Lemma 6.6. *For every absolute/relative-leximin-optimal division X, if agent j is absolutely/relatively poorer than agent i then agent j believes that the piece of agent i is worthless: $\hat{v}_j(X_i) = v_j(X_i) = 0$.*

Proof. If this were not the case, then we could take a small bit of X_i and give it to agent j, thus achieving an absolute/relative-leximin-better division. But this contradicts the leximin-optimality of X. □

In Example 6.4, in case of the absolute-leximin-optimal division, Carl is absolutely-richer than Alice and Bob, and indeed his share is worthless for both of them.

Suppose X is an old division and Y is a new division of the same cake. We say that an agent i *conceded a slice* to agent j if there is a slice that belonged to agent i in X and belongs to agent j in Y (in other words, $X_i \cap Y_j$ has a positive Lebesgue-measure). If X and Y are both absolute/relative-leximin-optimal, then by Pareto-optimality, $X_i \cap Y_j$ has positive value to both who concedes the slice (i) and the recipient (j). Hence, we have the following corollary of Lemma 6.6:

Corollary 6.2. *Let X and Y be two leximin-optimal divisions. If, when switching from X to Y, agent i conceded a slice to agent j, then in division X, agent i is weakly-poorer than j, and in division Y, agent i is weakly-richer than j.*

Lemma 6.7. *The absolute and the relative leximin division rules are essentially single-valued.*

Proof. The proof is the same for the absolute and the relative leximin rules. Hence, we omit the adjectives "absolute" and "relative" during the proof.

Suppose X and Y are two different leximin-optimal divisions. By switching from X to Y, some agents may have lost value; call these agents *unlucky*. Similarly, call the agents who gained value *lucky*. We are going to prove that there are neither unlucky nor lucky agents; this implies that all agents have exactly the same value in both divisions.

Suppose by contradiction that agent i is unlucky. Then, he must have conceded a slice to at least one other agent, say j. By Corollary 6.2, i is weakly-poorer than j in X and weakly-richer in Y. But this means that j is also unlucky. So, all slices conceded by unlucky agents are held by other unlucky agents. Suppose the unlucky agents take back all the slices that they conceded. This has no effect on the lucky agents, but strictly increases the value of the unlucky agents, since they now have at least the value that they had in X. But this contradicts the optimality of Y. Hence, there are no unlucky agents. If Y had a lucky agent without having any unlucky agent that would contradict the optimality of X.

Since X and Y were arbitrary leximin-optimal divisions it follows that all leximin-optimal divisions have the same value vector. □

Lemma 6.8. *The absolute and the relative-leximin division rules are population-monotonic.*

Proof. Again the proof is the same for the absolute and the relative leximin rules. Thanks to Lemma 6.7 it is sufficient to prove downwards-PM.

Let X be a leximin-optimal division of C. Suppose that agent i leaves and abandons his share X_i. So X is now a division of $C \setminus X_i$ among the agents $N \setminus \{i\}$. If X_i is divided arbitrarily among $N \setminus \{i\}$, the result is X^+, a division of C which is weakly leximin-better than X. Let Y be a leximin-optimal division of C among $N \setminus \{i\}$. Y is weakly leximin-better than X^+ and hence weakly leximin-better than X.

The rest of the proof is very similar to the proof of Lemma 6.7. Call the agents that lost value in the move from X to Y, unlucky. Suppose by contradiction that agent i is unlucky. Then he must have conceded a slice to an agent j, who must also be unlucky. All slices conceded by unlucky agents, are held by other unlucky agents. If those took back all the slices that they conceded, then all of them would be strictly better off while the lucky ones would remain unaffected. This contradicts the optimality of Y. Hence there are no unlucky agents and PM is proved.

□

Corollary 6.3. *The absolute leximin division rule is resource-monotonic.*

Proof. The same proof works as in case of the previous lemma. The cake enlargement can be treated as a piece that was acquired from an agent who left the scene. □

Note that we can not extend Corollary 6.3 to the relative leximin rule. The reason is simple: cake enlargements leave the absolute values of the pieces intact, but might alter their relative values. In particular it may happen that an agent – say j – lost utility although he received the same piece in both cakes, i.e. $X_j = Y_j$, but $v_j^C(X_j) > v_j^{C'}(Y_j)$, where C' denotes the enlarged cake. Also it is possible that the relative-leximin-optimal division of the enlarged cake is lexicographically smaller than the relative-leximin-optimal division of the smaller cake. The following subsections are devoted to uncovering the main differences between the absolute- and relative- leximin rules.

6.2.1 Absolute-leximin rule: PO+RM+PM

Theorem 6.2. *The absolute-leximin division rule is Pareto-optimal, population-monotonic and resource-monotonic.*

Proof. PO follows from Lemma 6.5, PM from Lemma 6.8 and RM from Corollary 6.3. □

Unfortunately, the absolute-leximin rule is not PROP. For instance, in Example 6.4, the absolute-leximin-optimal division gives Carl a value of 9, which is only $\frac{3}{10}$ of his total cake value.

6.2.2 Relative-leximin rule: PROP+PO+PM

Theorem 6.3. *The relative-leximin division rule is proportional, Pareto-optimal and population-monotonic.*

Proof. PROP holds because proportional divisions exist. The relative-leximin-welfare of a proportional division is at least $(\frac{1}{n}, \dots, \frac{1}{n})$, hence the optimal relative-leximin-welfare vector must be at least $(\frac{1}{n}, \dots, \frac{1}{n})$. PO follows from Lemma 6.5 and PM from Lemma 6.8. □

Example 6.5. The following cake shows that the relative-leximin rule is not RM. The largest value that can be given to both Alice and Bob is 9.

						▼
$\hat{v}_A$	9	9	0	0	0	0
$\hat{v}_B$	9	9	0	0	0	0
$\hat{v}_C$	4	4	10	0	0	18
$\hat{v}_D$	4	4	0	10	0	18
$\hat{v}_E$	4	4	0	0	10	18

Hence, in the smaller cake, the optimal relative-leximin-welfare vector is $(9/18, 9/18, 10/18, 10/18, 10/18)$. It is attained by halving the two leftmost slices between Alice and Bob, and giving the three slices at their right to Carl, David and Eve, in that order.

In the larger cake, the largest value that can be given to Carl, David and Eve from the additional slice at the right is 6. So the largest value that can be given to them from the four rightmost slices is 16. However, the total cake value has doubled for them. Hence, if Alice and Bob keep their share of 9, the relative-leximin-welfare vector changes to $16/36, 16/36, 16/36, 9/18, 9/18$.

Note that Alice and Bob are now relatively-richer than Carl, David and Eve, and their pieces have a positive value for them. Thus, by Lemma 6.6, the division in which Alice and Bob keep their current shares cannot be relative-leximin-optimal.

6.2.3 Relative-leximin-connected rule: CON+PROP+POCA but not PM+RM

For completeness we mention the *relative-leximin-connected* rule, which returns the divisions that are relative-leximin-optimal among all contiguous divisions. It is clearly CON+PROP+POCA. By Example 5.2, it cannot be PM. Indeed, in that example, the initial leximin-welfare vector is (3,7,7) with A receiving a value of 7, but after agent B leaves the leximin-welfare vector becomes (5,5) and A's value drops down to 5. Similarly, by Example 5.1 it cannot be RM.

Remark 6.2. Instead of leximin principle, we could have used another welfare-optimization criterion such as the *utilitarian* criterion - maximizing the sum of utilities. The absolute- and relative-utilitarian rules are both Pareto-optimal and population-monotonic, and absolute-utilitarian is also resource-monotonic. However, in contrast to our leximin rules, they are not essentially-single-valued and not proportional. The merit of the leximin rules is the combination of efficiency (PO) and fairness (PROP) with (at least some) monotonicity properties.

6.3 The rightmost-mark rule: CON+PROP+WH+RM

In this section we present an resource-monotonic procedure that produces a proportional and contiguous division of the whole cake for 2 agents.

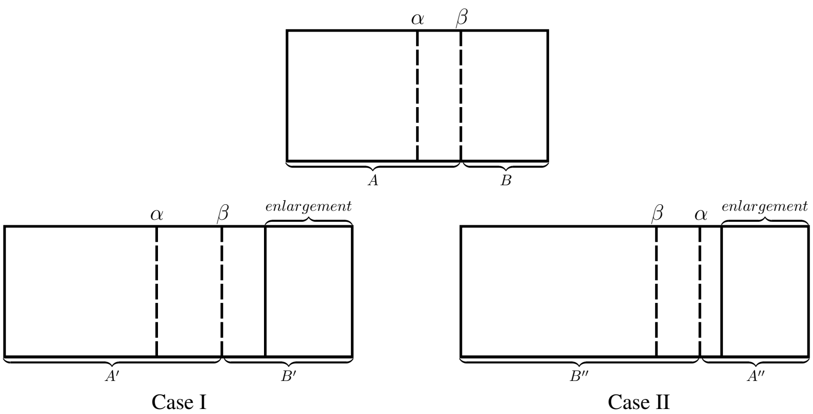

Figure 4: Illustration of the rightmost-mark division rule. Alice's cut mark is denoted by α, while Bob's cut mark by β.

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 their own valuation.
- 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, the rightmost-mark procedure is contiguous, whole, proportional and resource-monotonic.*

Proof. CON, WH and PROP are obvious. 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 more than 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 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 A" contains A, her original piece. Thus, she is not worse off. Similarly, Bob evaluates A and B the same, and he received B", thus he is not worse off either. □

There are several ways to generalize the rightmost mark rule. The reverse Dubins-Spanier and reverse Even-Paz methods coincide with it for two agents. Unfortunately as we have shown in Section 4 neither of these protocols are resource-monotonic.

Remark 6.3. The rightmost-mark procedure is CON+PROP+WH+RM. By Example 5.1, it cannot be Pareto-optimal. Indeed, in that example Bob makes the rightmost mark, which falls between the 2nd and 3rd slice. Bob receives the third slice which is worth 5 for him and Alice receives the other two slices which are worth 8 to her. But this is not PO because Bob could get the 2nd slice and increase his value to 10 without decreasing Alice's utility.

7 Conclusion and Future Work

We studied monotonicity properties in combination with the classical fairness axioms of proportionality, contiguousness and Pareto-optimality. Table 2 summarizes the properties of the various division rules.

In the connected case, each of our four rules satisfies two of the four properties {POCA,WH,RM,PM}. Thus, the fair divider has to choose between strong efficiency (POCA) without monotonicity, or strong monotonicity (RM+PM) without efficiency, or weak efficiency (WH) with only one monotonicity property (RM or PM). We are still looking for a rule that

Connected rules	n	CON	PROP	PO(CA)	WH	RM	PM	Sec.
exact-prop-conn.	Any	Yes	Yes	No*	No	Yes	Yes	6.1.1
max-equitable-conn.	Any	Yes	Yes	No*	Yes	No	Yes	6.1.2
rightmost-mark	2	Yes	Yes	No*	Yes	Yes	No	6.3
rel-leximin-conn.	Any	Yes	Yes	Yes	Yes	No*	No*	6.2.3

Disconnected rules	n	CON	PROP	PO(CA)	WH	RM	PM	Sec.
exact-prop.	Any	No	Yes	No	Yes	Yes	Yes	6.1.1
absolute-leximin	Any	No	No	Yes	Yes	Yes	Yes	6.2.1
relative-leximin	Any	No*	Yes	Yes	Yes	No	Yes	6.2.2

Classical rules	n	CON	PROP	PO(CA)	WH	RM	PM	Sec.
Banach-Knaster	Any	Yes	Yes	No	Yes	No	No	4
Cut and Choose	2	Yes	Yes	No	Yes	No	No	4
Dubins-Spanier	Any	Yes	Yes	No	Yes	No	No	4
Even-Paz	Any	Yes	Yes	No	Yes	No	No	4
Fink	Any	No	Yes	No	Yes	No	Upw	4
Selfridge-Conway	3	No	Yes	No	Yes	No	No	4

Table 2: 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 case of the Fink rule, Upw indicates that – with some additional adjustment – the rule is upwards-PM.

satisfies WH+RM for three or more agents, as well as a rule that satisfies WH+RM+PM for two or more agents.

In the disconnected case, each of the three rules satisfies three of the four properties {PROP,PO,RM,PM}. We are still looking for a rule that satisfies PROP+PO+RM+PM, or for a proof that such a rule is impossible.

Our work leaves a few open questions. The rules of the first two groups – with the exception of the rightmost mark rule – are difficult to implement, to say the least. We were mainly concerned with the theoretical possibilities: which combinations of axioms are feasible and which are not. We leave the actual design of protocols to future research.

Another interesting question is how the monotonicity axioms relate to envy-freeness, a fairness property as important and well-studied as proportionality.

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

Although we proved many negative results related to classical protocols, there is one promising lead. With appropriate adjustments the Fink rule turned out to be upwards population-monotonic. Thus, there may be set-valued rules, with desirable properties, which satisfy only one direction of the monotonicity axioms. In many cases the divider may not care about the other direction. For instance it is plausible that the agents are less sensitive about issues related to downwards-PM, as it is more likely that a new participant arrives than that there is someone who abandons his share and leaves the division process.

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.

The authors acknowledge the support of the 'Momentum' Programme (LP-004/2010) of the Hungarian Academy of Sciences and the OTKA grants K108383 and K109354, the ISF grant 1083/13, the Doctoral Fellowships of Excellence Program, the Adar Foundation, the Wolfson Chair and the Mordecai and Monique Katz Graduate Fellowship Program at Bar-Ilan University.

A Existence of max-equitable-connected divisions

Results about connected divisions use the *Borsuk-Ulam theorem.* It is about functions defined on spheres. Define the sphere S^{n-1} as the set of points $(x_1, \ldots, x_n)$ satisfying: $|x_1| + \cdots + |x_n| = 1$ (it is a sphere in the l_1 metric).

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

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

Assume that the cake is the interval $[0, 1]$. Each point $(x_1, \ldots, x_n) \in S^{n-1}$ corresponds to a partition of the cake to n intervals, marked $X_1, \ldots, X_n$, 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)).

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 $X = (x_1, \ldots, x_n) \in S^{n-1}$ and for every $i \in 1, \ldots, n-1$, 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 $i \in (1, \ldots, n-1)$:

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

Hence the division X^* is equitable.

References

Alon, N. and West, D. B. (1986). The Borsuk-Ulam theorem and bisection of necklaces. *Proceedings of the American Mathematical Society*, 98(4):623–628.

Arzi, O., Aumann, Y., and Dombb, Y. (2011). Throw one's cake - and eat it too. In Persiano, G., editor, *Algorithmic Game Theory*, volume 6982 of *Lecture Notes in Computer Science*, pages 69–80. Springer Berlin Heidelberg.

Austin, A. K. (1982). Sharing a Cake. *The Mathematical Gazette*, 66(437):212.

Balinski, M. and Young, H. P. (1975). The quota method of apportionment. *American Mathematical Monthly*, 82(7):701–730.

Balinski, M. and Young, H. P. (1982). *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press, New Haven.

Berliant, M., Thomson, W., and Dunz, K. (1992). On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216.

Brams, S. and Taylor, A. (1996). *Fair Division: From Cake Cutting to Dispute Resolution*. Cambridge University Press, Cambridge UK.

Brams, S. J., Jones, M. A., and Klamler, C. (2006). Better Ways to Cut a Cake. *Notices of the AMS*, 53(11):1314–1321.

Calleja, P., Rafels, C., and Tijs, S. (2012). Aggregate monotonic stable single-valued solutions for cooperative games. *International Journal of Game Theory*, 41(4):899–913.

Caragiannis, I., Lai, J. K., and Procaccia, A. D. (2011). Towards more expressive cake cutting. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*, IJCAI'11, pages 127–132. AAAI Press.

Chambers, C. P. (2005). Allocation rules for land division. *Journal of Economic Theory*, 121(2):236–258.

Dall'Aglio, M. (2001). The DubinsâĂŞSpanier optimization problem in fair division theory. *Journal of Computational and Applied Mathematics*, 130(1-2):17–40.

Dall'Aglio, M. and Di Luca, C. (2012). Finding maxmin allocations in cooperative and competitive fair division. arXiv preprint 1110.4241.

Dall'Aglio, M. and Hill, T. P. (2003). Maximin share and minimax envy in fair-division problems. *Journal of Mathematical Analysis and Applications*, 281(1):346–361.

Dubins, L. E. and Spanier, E. H. (1961). How to Cut A Cake Fairly. *The American Mathematical Monthly*, 68(1).

Even, S. and 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.

Herreiner, D. and Puppe, C. (2009). Envy Freeness in Experimental Fair Division Problems. *Theory and Decision*, 67(1):65–100.

Jones, M. A. (2002). Equitable, Envy-Free, and Efficient Cake Cutting for Two People and Its Application to Divisible Goods. *Mathematics Magazine*, 75(4):275–283.

Moulin, H. (2004). *Fair Division and Collective Welfare*. The MIT Press.

Peleg, B. and Sudhölter, P. (2007). *Introduction to the Theory of Cooperative Games*. Springer, Heidelberg.

Procaccia, A. D. (2015). Cake Cutting Algorithms. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press.

Simmons, F. W. and Su, F. E. (2003). Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical Social Sciences*, 45(1):15–25.

Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16(1):101–104.

Tasnádi, A. (2003). A new proportional procedure for the n-person cake-cutting problem. *Economics Bulletin*, 4(33):1–3.

Thomson, W. (1997). The Replacement Principle in Economies with Single-Peaked Preferences. *Journal of Economic Theory*, 76(1):145–168.

Thomson, W. (2011). Fair Allocation Rules. In *Handbook of Social Choice and Welfare*, pages 393–506. Elsevier BV.

Thomson, W. (2012). On The Axiomatics Of Resource Allocation: Interpreting The Consistency Principle. *Economics and Philosophy*, 28(03):385–421.

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(3):65–72.