

Ortega, Josué; Segal-Halevi, Erel

Working Paper

Obvious manipulations in cake-cutting

ZEW Discussion Papers, No. 19-056

Provided in Cooperation with:

ZEW - Leibniz Centre for European Economic Research

Suggested Citation: Ortega, Josué; Segal-Halevi, Erel (2019) : Obvious manipulations in cake-cutting, ZEW Discussion Papers, No. 19-056, ZEW - Leibniz-Zentrum für Europäische Wirtschaftsforschung, Mannheim

This Version is available at:

<https://hdl.handle.net/10419/208371>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

// NO.19-056 | 12/2019

DISCUSSION PAPER

// JOSUE ORTEGA AND
EREL SEGAL-HALEVI

Obvious Manipulations in Cake-Cutting

Obvious Manipulations in Cake-Cutting

Josué Ortega^{a,b}, Erel Segal-Halevi^c

^a*Queen's Management School, Queen's University Belfast, UK.*

^b*Center for European Economic Research, Mannheim, Germany.*

^c*Department of Computer Science, Ariel University, Israel.*

Abstract

In cake-cutting, strategy-proofness is a very costly requirement in terms of fairness: for $n = 2$ it implies a dictatorial allocation, whereas for $n \geq 3$ it requires that one agent receives no cake. We show that a weaker version of this property recently suggested by Troyan and Morrill, called *not-obvious manipulability*, is compatible with the strong fairness property of *proportionality*, which guarantees that each agent receives $1/n$ of the cake. Both properties are satisfied by the *leftmost leaves* mechanism, an adaptation of the Dubins - Spanier moving knife procedure. Most other classical proportional mechanisms in literature are obviously manipulable, including the original moving knife mechanism. Not-obvious manipulability explains why leftmost leaves is manipulated less often in practice than other proportional mechanisms.

Keywords: cake-cutting, not-obvious manipulability, prior-free mechanism design.

JEL Codes: D63, D82.

Email addresses: j.ortega@qub.ac.uk, erelsgl@gmail.com.

We thank Maria Kyropoulou, Thayer Morrill, Hervé Moulin and Gabriel Ziegler for their helpful comments and Sarah Fox, Jacopo Gambato and Fabian Spühler for proof-reading this paper. Any errors are our own.

October 17, 2019

1. Introduction

The division of a single good among several agents who value different parts of it distinctly is one of the oldest fair division problems, going as far back as the division of land between Abram and Lot (Genesis 13). Since its formalization as the cake-cutting problem (Steinhaus, 1948), this research question has inspired a large interdisciplinary literature which has proposed mechanisms that produce fair allocations without giving agents incentives to misrepresent their preferences over the cake. Unfortunately, Brânzei and Miltersen (2015) show that any deterministic strategy-proof mechanism suggests a dictatorial allocation for $n = 2$, and that one agent must receive no cake at all for $n \geq 3$, documenting a strong tension between fairness and incentives properties.

Nevertheless, recent results in applied mechanism design have shown that, even if mechanisms can be manipulated in theory, they are not always manipulated in practice. Some manipulations are more likely to be observed than others, particularly those which are salient or require less computation. Based on this observation, Troyan and Morrill (2019) have proposed a weaker version of strategy-proofness for direct mechanisms, called *not-obvious manipulability* (NOM). They define a manipulation as obvious if it yields a higher utility than truth-telling in either the best- or worst-case scenarios. A mechanism is NOM if it admits no obvious manipulation. Their notion of NOM is a compelling one, since it does not require prior beliefs about other agents' types, and compares mechanisms only based on two scenarios which are particularly salient and which require less cognitive effort to compute. They show that NOM accurately predicts the level of manipulability that different mechanisms experience in practice in school choice and auctions.

In this paper, we provide a natural extension of NOM to indirect mechanisms, and show that the stark conflict between fairness and truth-telling in cake-cutting disappears if we weaken strategy-proofness to NOM. In particular, NOM is compatible with the strong fairness property of *proportionality*, which guarantees each agent $1/n$ of the cake. Both properties are satisfied by a discrete adaptation of the moving knife mechanism (Dubins and Spanier, 1961), in which all agents cut the cake and the agent with the smallest cut receives all the cake to the left of his cut and leaves. This procedure is also procedurally fair and easy to implement in practice. NOM is violated by most other classical proportional mechanisms, even by the original Dubins-Spanier procedure, which shows that theoretically equivalent mechanisms may have different “obvious” incentive properties for boundedly rational agents. NOM partially explains why leftmost leaves is manipulated less frequently than other cake-cutting mechanisms in practice (Kyropoulou et al., 2019).

2. Related Literature

The cake-cutting problem has been studied for decades, given its numerous applications to the division of land, inheritances, and cloud computing (Brams and Taylor, 1996; Moulin, 2004). Most of the cake-cutting literature studies indirect revelation mechanisms, in which agents are asked to reveal their valuation over the cake via specific messages such as cuts and evaluations. In particular, the computer science literature focuses on the so called *Robertson–Webb mechanisms* (Robertson and Webb, 1998), in which agents can only use two types of messages: either an agent *cuts* a piece of the cake having a specific value, or he *evaluates* an existing piece by revealing his utility for it. Most well-known mechanisms in the literature, such as cut and choose, can be expressed as a combination of these two operations.

The main result relating fairness and incentive properties in cake-cutting was put forward by Brânzei and Miltersen (2015). They show that with three or more agents, every strategy-proof Robertson–Webb mechanism assigns no cake to at least one agent. Their result builds on a weaker result by Kurokawa et al. (2013) regarding mechanisms that have a bounded number of messages.

To allow for some degree of fairness and truth-telling, the literature has considered four research avenues.

The first one is to use *randomized mechanisms*. For example, Mossel and Tamuz (2010) show that truth-telling in expectation and approximate proportionality can be obtained with probabilistic mechanisms. In this paper we consider deterministic mechanisms only.

The second is to restrict the set of possible valuations over the cake. In this line, Chen et al. (2013) provide a complex deterministic mechanism that is both strategy-proof and proportional for a restricted class of utilities called *piecewise linear*. Their mechanism may waste pieces of cake which remain unassigned, something that never occurs with the proportional and NOM mechanism we identify.

A third research avenue studies allocation mechanisms in which an agent can only increase his utility by ϵ by misrepresenting his preferences, compared to the utility he obtains from truth-telling. Menon and Larson (2017) and Kyropoulou et al. (2019) obtain bounds on the amount of extra utility that agents can guarantee by lying in proportional cake-cutting mechanisms.

In this paper, we explore a fourth way to escape the conflict between fairness and truth-telling, which is relaxing the strategy-proofness definition to the one of NOM. NOM is a stronger version of an existing weak truth-telling property in the cake-cutting literature called *maximin strategy-proofness*, proposed by Brams et al. (2006, 2008). We discuss the relationship between these two concepts in detail in the next section.

3. Model

3.1. Outcomes, agents and types

In a general mechanism design problem, there is a set of possible outcomes \mathcal{X} , a set of agents N , and a set of possible agent-types $\Theta = \prod_{i \in N} \Theta_i$. In the particular case of cake cutting, there is an interval $[0, 1]$ called the *cake*. A union of subintervals of $[0, 1]$ is called a *piece* of the cake. The outcome-set \mathcal{X} is the set of *allocations* — ordered partitions of $[0, 1]$ into n disjoint pieces. $X \in \mathcal{X}$ denotes an arbitrary allocation and X_i denotes the piece allocated to agent i .

The type-set Θ is the set of possible *utility functions*, where a utility function is a non-atomic measure — an additive function that maps each piece to a non-negative number. The utility of the entire cake is normalized to 1. The utility of agent i with type θ_i is denoted by u_i . That u_i is non-atomic allows us to ignore the boundaries of intervals. Another implication of nonatomicity is that u_i is divisible, i.e. for every subinterval $[x, y]$ and $0 \leq \lambda \leq 1$, there exists a point $z \in [x, y]$ such that $u_i([x, z]; \theta_i) = \lambda u_i([x, y]; \theta_i)$.¹

3.2. Extensive forms and mechanisms

An *extensive form* is an arborescence A that consists of a set of labelled nodes H and a set of directed edges E .² The root node is h_0 . Each terminal node is labelled with an allocation of the cake. Each non-terminal node h is labelled with a non-empty subset of agents $N(h)$ who have to answer a query about their type. $N(h)$ is said to be the set of players *active at* h .

In a general extensive form, the query may be arbitrary, for example asking agents to fully reveal their type. In a *Robertson–Webb extensive form*, only two types of queries are allowed:

1. *Cut query*: the query $\text{cut}(i; x, \alpha)$ asks agent i the minimum point $y \in [0, 1]$ such that $u_i([x, y]; \theta_i) = \alpha$; where $\alpha \in [0, 1]$ and x is an existing cut point or 0 or 1. The point y becomes a cut point.
2. *Eval query*: the query $\text{eval}(i; x, y)$ asks agent i for its value for the interval $[x, y]$, that is, $\text{eval}(i; x, y) = u_i([x, y]; \theta_i)$ where x, y are a previously made cut points or 0 or 1.

¹Both u_i and θ_i are equivalent — the type of an agent is the agent’s utility, so our notation is slightly redundant. Nevertheless, we use it to make the comparison with [Troyan and Morrill \(2019\)](#) straightforward.

²An arborescence is a directed, rooted tree in which all edges point away from the root.

All agents must reply to the queries and their answers must be dynamically consistent, meaning that there must be a possible type for which such a sequence of answers is truthful.³ However, agents' answers may be untruthful. For every possible combination of agents' answers to the queries at node h , there is an edge from h to another node h' . Thus, answers to queries and edges have a one-to-one relationship. Note that, since the valuations are real numbers, there are uncountably many edges emanating from each node. Figure 1 illustrates a small subset of an extensive form A .




Figure 1: An example of an extensive form representing cut and choose.

We make the following two standard informational assumptions (Moore and Repullo, 1988). First, at each node h all agents know the entire history of the play. Second, if more than one agent is active at node h , they answer their corresponding queries simultaneously.

The indirect mechanism M corresponding to extensive form A takes as input the answers to each query in A and returns the allocation obtained at the corresponding terminal node. This is, the input to the mechanism consists of a path from the root node to a terminal node labelled with an allocation. At node h , the edge corresponding to a truthful answer by all agents in $N(h)$ with a type profile θ is denoted by $e^h(\theta)$, and $E(\theta) = \{e \in E : e^h(\theta) \text{ for some } h \in H\}$. This is, $E(\theta)$ is the set of all edges corresponding to truthful reports. The set $E(\theta)$ is an edge cover, i.e. is a set of edges incident

³This is a standard requirement (Brânzei and Miltersen, 2015). For example, if agent i is asked the query $\text{eval}(i; 0.3, 1)$ and replies a value of 0, then if later asked the query $\text{cut}(i; 0, 0.5)$ his answer must be in the interval $[0, 0.3)$.

to every node in A .⁴ An alternative edge cover in which all agents except i answer all queries truthfully, and agent i answers the queries as if he was of type θ'_i , is denoted by $E(\theta'_i, \theta_{-i})$. By our assumption of dynamically consistent answers, each possible set of untruthful answers by each agent is associated to a possible type in Θ_i . $M_i(E(\theta_i, \theta_{-i}))$ denotes agent i 's individual allocation in mechanism M when agents' answers correspond to the set of edges $E(\theta_i, \theta_{-i})$. A special case of our definition is a direct revelation mechanism, when there is a unique non-terminal node — the root h_0 —, $N(h_0) = N$ and the query for each agent is to fully reveal his type.

A mechanism M is called *proportional* if it guarantees to a truthful agent a utility of $1/n$ regardless of what every other agent reports, i.e. $u_i(M_i(E(\theta_i, \theta'_{-i}); \theta_i)) \geq 1/n$ for all $i \in N$, for all $\theta_i \in \Theta_i$, and all $\theta'_{-i} \in \Theta_{-i}$.

3.3. Manipulations

A *manipulation* by agent i is a set of (partially) untruthful answers that satisfy the consistency requirement, i.e., correspond to some type $\theta'_i \neq \theta_i$.

A manipulation is called *profitable* (for mechanism M and type θ_i) if $u_i(M_i(E(\theta'_i, \theta_{-i})); \theta_i) > u_i(M_i(E(\theta_i, \theta_{-i})); \theta_i)$. If some type of some agent has a profitable manipulation, then mechanism M is called *manipulable*.

A mechanism is called *strategy-proof* (SP) if it is not manipulable, that is, $u_i(M_i(E(\theta_i, \theta_{-i})); \theta_i) \geq u_i(M_i(E(\theta'_i, \theta_{-i})); \theta_i)$ for all $i \in N$, for all $\theta_i, \theta'_i \in \Theta_i$, and all $\theta_{-i} \in \Theta_{-i}$. This definition of SP is very demanding: it requires a truthful report from *every* agent, *every* time he is asked to answer a query.⁵ Even if just one type of an agent has an incentive to give a non-truthful answer to a single query, the mechanism is no longer strategy-proof.

Troyan and Morrill (2019) suggest a weaker version of SP for direct-revelation mechanisms, which only compares the best and worst case scenarios from both truthful and untruthful behavior. We extend their definition to indirect mechanisms as follows. A mechanism M is *not-obviously manipulable* (NOM) if, for any profitable manipulation of agent i , corresponding to pretending being a type θ'_i , the following two conditions hold:⁶

$$\inf_{\theta_{-i}} u_i(M_i(E(\theta'_i, \theta_{-i})); \theta_i) \leq \inf_{\theta_{-i}} u_i(M_i(E(\theta_i, \theta_{-i})); \theta_i) \quad (1)$$

$$\sup_{\theta_{-i}} u_i(M_i(E(\theta'_i, \theta_{-i})); \theta_i) \leq \sup_{\theta_{-i}} u_i(M_i(E(\theta_i, \theta_{-i})); \theta_i) \quad (2)$$

⁴Edge covers represent a complete contingent plan of answers.

⁵Strategy-proofness is a notion more commonly used for direct-revelation mechanisms. This extension to indirect mechanisms follows the definition of Kurokawa et al. (2013).

⁶They present this definition using maximum and minimum, which may not exist with a continuous cake; instead we consider the supremum and infimum.

If any of the previous two conditions do not hold for some θ'_i , then θ'_i is said to be an *obvious manipulation* for agent i with type θ_i ; and the mechanism M is *obviously manipulable*. In other words, a manipulation is obvious if it either makes the agent better off in the worst-case, or if it makes him better off in the best-case. This is a strengthening of the *maximin strategy-proofness* defined by [Brams et al. \(2006, 2008\)](#), who only impose condition (1). Brams, Jones and Klamler write: “We assume that players try to maximize the minimum-value pieces (maximin pieces) that they can guarantee for themselves, regardless of what the other players do. In this sense, the players are risk-averse and never strategically announce false measures if it does not guarantee them more-valued pieces”.

Both relaxations of strategy-proofness have the advantages that agents do not require beliefs about other agents’ actions, and that comparing best- and worst-cases scenarios requires less cognitive effort than comparing expected values using an arbitrary distribution over agents’ types. However, maximin strategy-proofness is a mild property that is satisfied by a very large class of mechanisms.⁷ On the other hand, NOM is a property that most classical proportional mechanisms in the literature fail, with the *leftmost leaves* mechanism being a remarkable exception (Theorem 2).

Before we proceed to present some examples, we clarify a difference between the notion of indirect mechanisms that we use (standard in the computer science literature) and the definition in economics ([Moore and Repullo, 1988](#)). In economics, each non-terminal node is labelled not with a query, but with an action. Thus, to define leftmost leaves one could just specify that, at each period t , an agent’s possible actions are to cut a cake at any point, and the one who provides the smallest cut exits with the leftmost part. There is no question where to cut: agents simply cut the cake wherever they want. But, if one does not specify the query asked to each player, it is not at all clear what truth-telling behavior is. Where should an agent with uniform valuation over the cake cut when dividing a cake against 3 agents if he is not asked a specific question? The computer science definition emphasizes that the mechanism designer can make a mechanism hard to manipulate by cleverly choosing which queries to ask at each node in the extensive form. Indeed, if we slightly modify the queries asked in *leftmost leaves*, the mechanism becomes obviously manipulable, as we show after the proof of Theorem 2.

⁷[Chen et al. \(2013\)](#) call maximin SP a “strikingly weak notion of truthfulness”.

4. Results

We first show, using an example, that many classic cake-cutting procedures are obviously manipulable.

Theorem 1. *The cake-cutting mechanisms known as cut-and-choose, cut-middle, Banach-Knaster last diminisher and Dubins-Spanier moving knife are all obviously manipulable.*

Proof. For simplicity, consider a cake-cutting problem with piecewise uniform valuations, i.e. agents either like or dislike certain intervals, and each desirable interval of the same length has the same value. One agent, called Blue, has valuations as in Figure 2.




Figure 2: The preferences of a (blue) agent over the cake.




Figure 3: Obvious manipulations for agent blue (black arrow), supporting preferences for the other agent (red), final allocation received by blue agent (green).

In the well-known *cut-and-choose* mechanism, Blue (the cutter) is asked $\text{cut}(\text{blue}; 0, 1/2)$. Truthful behavior requires him to cut at 0.5, guaranteeing a utility of 0.5 in all cases. If Blue chooses a profitable manipulation instead, say to cut at 0.4, the best case is that the other agent chooses the left piece of the cake, leaving to Blue a utility of 0.75. Thus, in inequality (2), the supremum at the left-hand side is at least 0.75 while the supremum at the right-hand side is 0.5. Cut-and-choose is therefore obviously manipulable.

In the *cut-middle* mechanism, both agents are asked $\text{cut}(i; 0, 1/2) = x_i$ simultaneously, and the cake is divided at $\frac{x_1 + x_2}{2}$, with each agent obtaining the part of the cake which contains his cut. If Blue is truthful and cuts at 0.5, the best case is that the other agent cuts at ϵ (or $1 - \epsilon$) and thus the cut point becomes $\frac{0.5 + \epsilon}{2} \approx 0.25$, so Blue receives 0.75 utility. Nevertheless, if Blue chooses a profitable manipulation such as cutting at ϵ , the best that could happen is that the other agent cuts at $\eta < \epsilon$, and thus Blue receives a

utility almost equal to 1. In inequality (2), the supremum at the left-hand side is 1 and the supremum at the right-hand side is 0.75.

We conclude that cut-middle, too, is obviously manipulable.

Cut-and-choose and cut-middle are mechanisms for dividing a cake among two agents. We now turn to mechanisms that can be used to divide cake among two or more agents.

First, consider the *Banach-Knaster last diminisher* mechanism (Steinhaus, 1948), in which agents are assigned a fixed order. The first agent is asked the point $x_1 = \text{cut}(1; 0, 1/n)$, and is tentatively assigned the piece $[0, x_1]$. Then, the second agent has an option to “diminish” this piece: he is asked the point $x_2 = \text{cut}(2; 0, 1/n)$, and if $x_2 < x_1$, then the previous tentative assignment is revoked, and agent 2 is now tentatively assigned the piece $[0, x_2]$. This goes on up to agent n . Then, the tentative assignment becomes final: some agent k receives the piece $[0, x_k]$ and the other agents recursively divide the remaining cake $[x_k, 1]$.

This procedure is obviously manipulable even with two agents. For example, consider Figure 2(a). Suppose agent 1 answers truthfully $x_1 = 0.5$. If agent 2 answers truthfully $x_2 = 0.2$, then he gets a value of exactly 0.5. A profitable manipulation for agent 2 is to answer $x'_2 = x_1 - \epsilon$: it yields a utility of 1. This is true in both the best- and worst-case scenarios, which are the same in this case.

We turn to the *Dubins-Spanier moving knife* (Dubins and Spanier, 1961), in which a knife moves from 0 to 1 until the first agent stops it at the point x for which $u_i([0, x]; \theta_i) = 0.5$, receiving himself the interval $[0, x]$. The best that can happen to a truthful agent is that the other agent cuts the cake at ϵ , so that he receives all the remaining cake which gives him almost 1 utility. However, once the moving knife has reached point x , a truthful agent should stop the knife, implying that he gets 0.5 with certainty, whereas a profitable manipulation would yield a higher utility in the best-case scenario (one example is to stop the knife at $1 - \epsilon$). Thus, Dubins-Spanier moving knife is also obviously manipulable. \square

We now present a slight variation of the Dubins-Spanier moving knife mechanism, in which, in each period, all agents are asked the cut query simultaneously rather than sequentially. We call this variant **leftmost leaves**. The leftmost leaves mechanism works as follows:

- In period 1 all agents are asked $\text{cut}(i; 0, 1/n)$. The agent who cuts the cake at the smallest point (denoted x^1) leaves with the interval $[0, x^1]$. In case of a tie, the agent with the smallest index of all those who cut at x^1 leaves with $[0, x^1]$.

- In period 2, every remaining agent i is asked $\text{cut}(i; x^1, \frac{u_i([x^1, 1]; \theta_i)}{n-1})$. The agent who submits the smallest point (denoted x^2) leaves with the interval $(x^1, x^2]$. In case of a tie, the agent with the smallest index of all those who cut at x^2 leaves with $(x^1, x^2]$.
- In period t , all remaining agents are asked $\text{cut}(i; x^{t-1}, \frac{u_i([x^{t-1}, 1]; \theta_i)}{n-t+1})$. The procedure continues until only one agent remains. That agent receives $(x^{n-1}, 1]$.

Despite leftmost leaves being equivalent (in a sense in which we describe below) to the Dubins-Spanier moving knife mechanism, they differ in terms of obvious manipulability.

Theorem 2. *The leftmost leaves mechanism is proportional and not-obviously-manipulable.*

Before presenting the proof, we present a few remarks about this result and its possible extensions.

First, note that leftmost leaves is theoretically equivalent to Dubins-Spanier moving knife mechanism, in that when applied to truthful agents with the same types, both mechanisms always suggest the same allocation. How can mechanisms that are theoretically equivalent, such as the two we just presented, rank differently in terms of incentives? This idea goes back to Li (2017), who shows that two equivalent mechanisms, such as the ascending auction and the second price auction, in which bidding truthfully is a weakly dominant strategy, are different in terms of incentive properties for boundedly rational agents. The intuition in both results is similar: both in the second price auction and in leftmost leaves, agents have no restriction in the prior about their opponents' types when they reveal their type through either their bids or their cuts; whereas in both the ascending auction and the moving-knife procedure, the fact that the knife or the clock has reached some point tells the agents' something about their opponents' types, and thus modifies what to expect in the best- and worst-case scenarios.

Second, the leftmost leaves mechanism satisfies several other desiderata that make it a good candidate to divide goods in practice. Is not only proportional and NOM, but also procedurally fair (up to tie-breaks) (Crawford, 1977; Nicolò and Yu, 2008), since agents' identities do not affect the allocation produced. It also generates an assignment of a connected piece of cake for each agent, a desirable property for applications such as the division of land (Segal-Halevi et al., 2017).

Third, a follow-up question to Theorem 2 is whether leftmost leaves is the only proportional and NOM mechanism in cake-cutting. The answer is

no, as leftmost leaves can be slightly modified in several ways retaining both properties. One of such modifications is to start cutting the cake from the right instead of from the left.⁸ Another less trivial one is a modification of the protocol of [Even and Paz \(1984\)](#), which works exactly the same as leftmost leaves for $n = 2$ but requires fewer queries for larger values of n (the proof is an almost verbatim copy of the proof of Theorem 2, so we omit it).⁹ Obtaining a characterization of all NOM and proportional mechanisms remains an interesting, albeit challenging, open question.

5. Proof of Theorem 2

We denote the leftmost leaves mechanism by M^l . We show that M^l is not-obviously manipulable. Since M^l is an anonymous mechanism in which the identity of the agents does not play a role, it is necessary to check conditions (1) and (2) only for one agent, denoted by i .

First, we show that no manipulation yields a higher utility in the worst-case scenario. We use the following lemma.

Lemma 1. *At period t ,*

$$\inf_{\theta_{-i}} u_i(M^l(E(\theta_i, \theta_{-i})); \theta_i) = \frac{u_i([x^{t-1}, 1]; \theta_i)}{n - t + 1}$$

Proof. If x_i^t is the smallest cut at period t , then the result is immediate. Otherwise,

$$u_i([x^{t-1}, x^t]; \theta_i) \leq \frac{u_i([x^{t-1}, 1]; \theta_i)}{n - t + 1}$$

A direct implication is that the remainder of the cake $[x^t, 1]$ must be

⁸*Rightmost leaves* is the only mechanism to divide cake among two agents that is weakly Pareto optimal, proportional and resource monotonic under specific restrictions on agents' utilities ([Segal-Halevi and Sziklai, 2018](#)).

⁹The “leftmost leaves Even-Paz” mechanism works as follows (for clarity let n be a power of 2). Given a cake $[y, z]$, all agents choose cuts x_i such that $u_i([y, x_i]; \theta_i) = u_i([y, z]; \theta_i)/2$. Let x^* be the median of the n cuts, rounded to the nearest smallest integer. Then the procedure breaks the cake-cutting problem into two: all agents who choose cuts $x_i \leq x^*$ are to divide the cake $[y, x^*]$, whereas all agents who chose cuts above x^* are to divide the cake $[x^*, z]$. Each half is divided recursively among the $n/2$ partners assigned to it. When the procedure is called with a singleton set of agents $\{i\}$ and an interval I it assigns $X_i = I$. For example, if $n = 4$, agents cut the cake in two equivalent pieces and the cake is cut at the second smallest cut. Then the two agents with the smallest (largest) cuts play *leftmost leaves* on the left (right) side of the cake.

worth at least $\frac{n-t}{n-t+1}$ of $u_i((x^{t-1}, 1); \theta_i)$, i.e.

$$u_i([x^t, 1]; \theta_i) \geq \frac{n-t}{n-t+1} u_i([x^{t-1}, 1]; \theta_i)$$

Dividing both sides by $n-t$

$$\frac{u_i([x^t, 1]; \theta_i)}{n-t} \geq \frac{u_i([x^{t-1}, 1]; \theta_i)}{n-t+1}$$

Note that the left-hand side of the previous expression is the utility that the truthful agent i would receive if he cut the cake at the smallest point in period $t+1$. If his cut was not the smallest at period $t+1$, a recursive formulation shows that he would receive a share of the cake that he values even more in period $t+2$. Thus, the worst that can happen to a truthful agent in period t is to obtain a utility of $\frac{u_i([x^{t-1}, 1]; \theta_i)}{n-t+1}$. \square

Setting $t = 1$ in Lemma 1 shows that leftmost leaves is proportional:

$$\inf_{\theta_{-i}} u_i(M^l(E(\theta_i, \theta_{-i})); \theta_i) = \frac{u_i([0, 1]; \theta_i)}{n}$$

Now we show that any manipulation at period t , $\hat{x}_i^t \neq x_i^t$, yields a utility weakly smaller than $\frac{u_i([x^{t-1}, 1]; \theta_i)}{n-t+1}$ in the worst-case scenario.

If $\hat{x}_i^t < x_i^t$, in the worst-case scenario \hat{x}_i^t would be the smallest cut, and agent i would therefore receive the allocation $[x^{t-1}, \hat{x}_i^t]$, which by construction yields a weakly lower utility than the allocation $[x^{t-1}, x_i^t]$.

In the other case, if $\hat{x}_i^t > x_i^t$, in the worst-case scenario the smallest cut in period t , x^t , would be such that $x_i^t < x^t < \hat{x}_i^t$. Thus, the rest of the cake $[x^t, 1]$ is such that

$$u_i([x^t, 1]; \theta_i) < \frac{n-t}{n-t+1} u_i([x^{t-1}, 1]; \theta_i)$$

In period $t+1$, agent i cannot choose a manipulation $\hat{x}_i^{t+1} \leq x_i^{t+1}$, as per the previous step he would receive a smaller utility, so agent i must choose a manipulation $\hat{x}_i^{t+1} > x_i^{t+1}$.

But in the worst-case scenario, the smallest cut in period $t+1$ is x^{t+1} such that $x_i^{t+1} < x^{t+1} < \hat{x}_i^{t+1}$, so that the rest of the cake $[x^{t+1}, 1]$ is worth less than

$$u_i([x^{t+1}, 1]; \theta_i) < \frac{n-t-1}{n-t} u_i([x^t, 1]; \theta_i)$$

Following this argument, we see that agent i 's only alternative is to receive

the last piece of cake $[x^{n-1}, 1]$, which is, by construction

$$u_i([x^{n-1}, 1]; \theta_i) < u_i([x^{n-2}, x_i^{n-1}]; \theta_i) < \dots < u_i([x^{t-1}, x_i^t]; \theta_i) = \frac{u_i([x^{t-1}, 1]; \theta_i)}{n - t + 1}$$

This concludes the proof that no manipulation yields a higher utility than truth-telling in the worst-case scenario, so inequality (1) is satisfied.

Next, we show that no manipulation yields a higher utility in the best-case scenario. We use the following lemma.

Lemma 2. *At period t ,*

$$\sup_{\theta_{-i}} u_i(M^l(E(\theta_i, \theta_{-i})); \theta_i) = u_i([x^{t-1}, 1]; \theta_i)$$

This is, agent i , by being truthful in period t , can expect (in the best-case scenario) to obtain the whole cake available in period t .

Proof. Let us remember that truthful agent i cuts the cake at a point x_i^t such that $u_i([x^{t-1}, 1]; \theta_i) = \frac{u_i([x^{t-1}, 1]; \theta_i)}{n-t+1}$. In the best-case scenario, the smallest cut at period t is $x^t < x_i^t$ such that $x^t = x^{t-1} + \epsilon$. This guarantees that $u_i([x^{t-1}, 1]; \theta_i) - \epsilon' = u_i([x^t, 1]; \theta_i)$ for an arbitrarily small ϵ' , i.e. no valuable piece of cake for agent i has been allocated. That we can find an ϵ so small comes directly from the standard assumption that the utilities are divisible.

The best-case scenario is that all further smallest cake cuts x^{t+1}, x^{t+2}, \dots are epsilon increments of x^t , such that in the end agent i receives $X_i = [x^{n-1}, 1]$, which gives him a utility arbitrarily close to $u_i([x^{t-1}, 1])$, i.e. the utility of eating all of the remaining cake available. \square

Since this is the maximum utility attainable, inequality (2) in the NOM definition is satisfied. This concludes the proof that no manipulation gives a higher utility to a truthful agent in the best-case scenario.

We conclude that no manipulation is better than truth-telling in either the best or the worst-case scenario, thus no manipulation is obvious and leftmost leaves is NOM.

Remark 1. A minor modification of leftmost leaves, sometimes used in the literature (Procaccia, 2016), destroys the NOM of the procedure. In this modification, at each period t , agents are asked cut($i; x^{t-1}, 1/n$).

To see that this variant is obviously manipulable, consider the case of an agent who has uniform preferences over the whole $[0, 1]$ interval and has to cut the cake against 4 other agents. Her first cut must be $x_i^1 = 0.2$, which guarantees a utility of 0.2. However, suppose the lowest cut (submitted by

someone else) was instead $x^1 = 0.1$. In period $t = 2$, the remaining cake is $[0.1, 1]$ and if agent i is truthful, she should cut the cake at $x_i^2 = 0.3$ so that $u_i([0.1, 0.3]; \theta_i) = 1/5 = 0.2$. However, there is a manipulation that is guaranteed to yield a higher utility in the worst-case scenario. If she cuts the cake instead at $\hat{x}_i^2 = 0.325$ (i.e. the point at which $u_i([0.1, \hat{x}_i^2]; \theta_i) = 1/4$), in the worst-case scenario, in which her cut is the lowest, she would guarantee herself a utility of $0.225 > 0.2$. If her cut was not the lowest, by continuing to cut the cake at the point \hat{x}_i^t such that $x_i^t = \frac{([x^{t-1}, 1])}{n-t+1}$ for all subsequent t , she could make sure to receive a utility of at least 0.225 (Lemma 1), which is larger than the worst-case scenario utility received by being truthful, 0.2 . We conclude that the modified version of leftmost leaves is obviously manipulable, and thus the choice of which query to make at each step of the mechanism is crucial to make the leftmost leaves procedure not-obviously manipulable.

6. Direct-revelation mechanisms

While leftmost leaves is proportional and connected, it does not satisfy other desirable properties such as *envy-freeness* (no agent prefers the piece of cake received by someone else over his own piece) and *Pareto-optimality* (no other allocation is better for one agent and not worse for the others).

In the Robertson-Webb model, we could not yet find NOM mechanisms satisfying these properties. For example, the classic mechanism of Selfridge-Conway for three agents (see [Brams and Taylor, 1996](#) for a detailed description) is envy-free, but it is obviously-manipulable. In this mechanism, the first agent cuts the cake into three pieces of equal worth. A truthful agent knows that one of those pieces will never belong to him, and thus he can achieve a maximum utility of at most 0.67 . However, a lying agent can cut the cake in one piece of value $1 - \epsilon$, and two pieces of almost no value at all. In the best case scenario, he will keep the most valued piece entirely, showing that the Selfridge-Conway procedure is obviously manipulable.

Interestingly, NOM is easier to achieve in the direct-revelation model.

Lemma 3. *Every direct-revelation mechanism that always returns proportional allocations satisfies inequality (1).*

Proof. By proportionality, a truthful agent always receives a utility of at least $1/n$. Consider now an untruthful agent i who reports a type $\theta'_i \neq \theta_i$ (equivalently, reports a utility function $u'_i \neq u_i$). Consider the case when all other $n - 1$ agents have a utility of u_i (the true utility of agent i). A proportional mechanism must give each of these $n - 1$ agents a piece with a value, by the function u_i , of at least $1/n$. Hence, the piece remaining for

agent i has a value, by the function u_i , of at most $1/n$. Hence, in inequality (1), the infimum is at most $1/n$ at the left and at least $1/n$ at the right, and the inequality holds. \square

Lemma 4. *Every direct-revelation mechanism that always returns Pareto-optimal allocations satisfies inequality (2).*

Proof. Consider the case when agent i is truthful, all other $n-1$ agents assign a positive value only to a tiny fraction of the cake, and assign a value of 0 to the rest of the cake. A Pareto-optimal mechanism must assign almost all the cake to agent i . Hence, in inequality (2), the supremum in the right-hand side equals 1 and the inequality holds. \square

Theorem 3. *There exists a NOM direct-revelation mechanism that finds envy-free and Pareto-optimal allocations.*

Proof. The *Nash-optimal mechanism* is a direct-revelation mechanism that, given n utility functions, selects an allocation that maximizes the product of utilities. Such an allocation is known to be Pareto-optimal and envy-free (Segal-Halevi and Sziklai, 2019), hence it is also proportional. Hence, by lemmas 3 and 4, the mechanism is NOM. \square

When the utility functions are *piecewise-constant*, the Nash optimal mechanism can be computed by an efficient algorithm described by Aziz and Ye (2014).

In contrast to the leftmost leaves rule, the Nash-optimal rule may return disconnected pieces. Moreover, it is known that *any* Pareto-optimal envy-free rule may have to return disconnected pieces (see Example 5.1 in Segal-Halevi and Sziklai (2018)). Since Pareto-optimality is crucial in the proof of Theorem 3, it remains an open question whether there exists a NOM mechanism that is both envy-free and connected.

A related question is whether there exists an algorithm that finds proportional, Pareto-optimal and connected allocations. If such an algorithm exists, then by Lemmas 3 and 4, it is NOM.

7. Conclusion and Experimental Evidence

Although it is impossible to cut a cake in a strategy-proof manner that is not completely unfair to some agent, we can divide a cake in a fair, proportional way that cannot be obviously manipulated using an easily implementable mechanism called leftmost leaves.

Troyan and Morrill’s notion of NOM not only allows us to escape the tradeoff between fairness and incentives in cake-cutting, but also helps us to

better understand real-life behavior when dividing an heterogeneous good. In the first lab experiment comparing cake-cutting mechanisms, [Kyropoulou et al. \(2019\)](#) report truthful behavior in NOM cake-cutting mechanisms in 44% of the cases, whereas the respective number for OM ones is of 31% (the difference is statistically significant with a p-value of 0.0000). In particular, leftmost leaves was significantly less manipulated than Banach-Knaster last diminisher when agents played against 2 opponents (difference of 29 percentage points, p-value of 0.0000) and 3 opponents (difference of 16 percentage points, p-value of 0.0000). The Even-Paz modification of leftmost leaves (which is NOM) was also significantly less manipulated than Banach-Knaster (difference of 35 percentage points, p-value of 0.0000).

Although in general NOM gives us testable predictions that map relatively well to observed behavior, it is intriguing that the Selfridge-Conway procedure also reports high rates of truth-telling, comparable to those of NOM mechanisms. Explaining this puzzling phenomenon remains an open problem which we leave for future research.

References

- AZIZ, H. AND C. YE (2014): “Cake cutting algorithms for piecewise constant and piecewise uniform valuations,” in *International Conference on Web and Internet Economics*, Springer, 1–14.
- BRAMS, S., M. JONES, AND C. KLAMLER (2006): “Better ways to cut a cake,” *Notices of the AMS*, 53, 1314–1321.
- BRAMS, S. AND A. TAYLOR (1996): *Fair division: From cake-cutting to dispute resolution*, Cambridge University Press.
- BRAMS, S. J., M. A. JONES, AND C. KLAMLER (2008): “Proportional pie-cutting,” *International Journal of Game Theory*, 36, 353–367.
- BRÂNZEI, S. AND P. B. MILTERSEN (2015): “A dictatorship theorem for cake cutting,” in *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- CHEN, Y., J. LAI, D. PARKES, AND A. PROCACCIA (2013): “Truth, justice, and cake cutting,” *Games and Economic Behavior*, 77, 284–297.
- CRAWFORD, V. P. (1977): “A game of fair division,” *Review of Economic Studies*, 44, 235–247.
- DUBINS, L. AND E. SPANIER (1961): “How to cut a cake fairly,” *American Mathematical Monthly*, 68, 1–17.
- EVEN, S. AND A. PAZ (1984): “A note on cake cutting,” *Discrete Applied Mathematics*, 7, 285 – 296.
- KUROKAWA, D., J. K. LAI, AND A. D. PROCACCIA (2013): “How to cut

- a cake before the party ends,” in *Twenty-Seventh AAAI Conference on Artificial Intelligence*.
- KYROPOULOU, M., J. ORTEGA, AND E. SEGAL-HALEVI (2019): “Fair cake-cutting in practice,” in *Proceedings of the 2019 ACM Conference on Economics and Computation*, ACM, 547–548.
- LI, S. (2017): “Obviously strategy-proof mechanisms,” *American Economic Review*, 107, 3257–87.
- MENON, V. AND K. LARSON (2017): “Deterministic, strategyproof, and fair cake cutting,” in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, AAAI Press, 352–358.
- MOORE, J. AND R. REPULLO (1988): “Subgame perfect implementation,” *Econometrica*, 1191–1220.
- MOSSEL, E. AND O. TAMUZ (2010): “Truthful fair division,” in *Algorithmic Game Theory*, ed. by S. Kontogiannis, E. Koutsoupias, and P. Spirakis, Springer, 288–299.
- MOULIN, H. (2004): *Fair division and collective welfare*, MIT press.
- NICOLÒ, A. AND Y. YU (2008): “Strategic divide and choose,” *Games and Economic Behavior*, 64, 268 – 289.
- PROCACCIA, A. (2016): “Cake cutting algorithms,” in *Handbook of Computational Social Choice*, ed. by F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, Cambridge University Press, 311–330.
- ROBERTSON, J. AND W. WEBB (1998): *Cake-cutting algorithms: Be fair if you can*, AK Peters/CRC Press.
- SEGAL-HALEVI, E., S. NITZAN, A. HASSIDIM, AND Y. AUMANN (2017): “Fair and square: Cake-cutting in two dimensions,” *Journal of Mathematical Economics*, 70, 1–28.
- SEGAL-HALEVI, E. AND B. R. SZIKLAI (2018): “Resource-monotonicity and population-monotonicity in connected cake-cutting,” *Mathematical Social Sciences*, 95, 19–30.
- (2019): “Monotonicity and competitive equilibrium in cake-cutting,” *Economic Theory*, 68, 363–401.
- STEINHAUS, H. (1948): “The problem of fair division,” *Econometrica*, 16, 101–104.
- TROYAN, P. AND T. MORRILL (2019): “Obvious manipulations,” in *Proceedings of the 2019 ACM Conference on Economics and Computation*, ACM, 865.



Download ZEW Discussion Papers from our ftp server:

<http://ftp.zew.de/pub/zew-docs/dp/>

or see:

<https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html>

<https://ideas.repec.org/s/zbw/zewdip.html>



IMPRINT

**ZEW – Leibniz-Zentrum für Europäische
Wirtschaftsforschung GmbH Mannheim**

ZEW – Leibniz Centre for European
Economic Research

L 7,1 · 68161 Mannheim · Germany

Phone +49 621 1235-01

info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.