

Ortega, Josué; Segal-Halevi, Erel

Working Paper

Obvious Manipulations in Cake-Cutting

QMS Research Paper, No. 2020/04

Provided in Cooperation with:

Queen's University Belfast, Queen's Business School

Suggested Citation: Ortega, Josué; Segal-Halevi, Erel (2020) : Obvious Manipulations in Cake-Cutting, QMS Research Paper, No. 2020/04, Queen's University Belfast, Queen's Management School, Belfast,
<https://doi.org/10.2139/ssrn.3504487>

This Version is available at:

<https://hdl.handle.net/10419/271235>

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.

Obvious Manipulations in Cake-Cutting

JOSUÉ ORTEGA, Queen’s University Belfast, UK.

EREL SEGAL-HALEVI, Ariel University, Israel.

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 (2019), 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 the literature are obviously manipulable, including the original moving knife mechanism and some other variants of it. Not-obviously manipulable mechanisms (including leftmost leaves) are significantly less often manipulated in practice than other obviously manipulable mechanisms.

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 research for finding mechanisms that produce fair allocations without giving agents incentives to misrepresent their preferences over the cake. Unfortunately, besides some restricted preference domains, this research has so far produced mainly negative results, showing a strong tension between fairness and incentive properties (see Section 2).

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 an adaptation of the moving knife mechanism [Dubins and Spanier 1961], in which all agents cut the cake *simultaneously* 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 mechanisms are manipulated less frequently than other cake-cutting mechanisms in practice, as shown by [Kyropoulou et al. 2019] in a recent lab experiment.

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].

2.1 Computational models

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 *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.

Direct revelation mechanisms for cake-cutting have also been studied. However, such mechanisms cannot be practically implemented for the entire preference domain, since a general valuation function might require uncountably many values to represent. Therefore, direct mechanisms have been studied mainly for restricted classes of preferences. *Piecewise-uniform valuations* mean that each agent has a “desired subset” of the cake; the agent’s value-density is positive and identical inside the desired subset and zero outside. *Piecewise-constant valuations* generalize piecewise-uniform valuations by allowing each agent to have several desired subsets where the agent’s value-density is constant inside each subset but different in different subsets. *Piecewise-linear valuations* generalize piecewise-constant valuations by allowing the agent’s value-density to change linearly inside each desired subset.

Most papers assume that the cake is *all-good*, i.e., all agents value every piece of cake as weakly positive. Some papers assume that the cake is *all-bad*, i.e., all agents value every piece of cake as weakly negative (such a cake is often called a “chore”). Recently, some papers have studied a more general model in which the cake is *mixed*, i.e., each piece of cake may be positive for some agents and negative for others [Avvakumov and Karasev 2019; Aziz et al. 2019; Bogomolnaia et al. 2019; Meunier and Zerbib 2019; Segal-Halevi 2018].

2.2 Randomized truthful-in-expectation mechanisms

The first truthful mechanisms developed for cake-cutting were randomized direct-revelation mechanisms. Mossel and Tamuz [2010] describe a mechanism based on a *consensus division* — a division in which each agent values every piece at exactly $1/n$. Such a division always exists, it is proportional and envy-free, and it is simple to compute if the valuations are piecewise-linear. Once a consensus division is found, each agent is assigned a piece at random. Since the expected value of each agent is always $1/n$ regardless of the reports, an agent can never gain by misreporting. However, restricting the value to exactly $1/n$ goes against the spirit of finding “win-win” divisions. Therefore, Mossel and Tamuz [2010] describe a randomized mechanism which always finds a *super-proportional allocation* (an allocation giving each agent strictly more than $1/n$ of the total value) if and only if it exists. They prove that this super-proportionality cannot be guaranteed by any *deterministic* truthful mechanism.

Brânzei and Miltersen [2015] translate the above results to indirect-revelation mechanisms in the Robertson–Webb query model. They show that an approximate consensus division (in which each agent values every piece between $1/n - \epsilon$ and $1/n + \epsilon$) can be found by $O(n^2/\epsilon)$ queries; this leads to a randomized query-based mechanism that is truthful in expectation and approximately-proportional. On the other hand, they show that, under a “hungry players” assumption (all agents value every part of the cake positively), every *deterministic* truthful Robertson–Webb mechanism

assigns no cake to at least one agent.¹ Their result builds on a weaker result by [Kurokawa et al. 2013] regarding mechanisms that have a bounded number of messages. Obviously, such a mechanism cannot be considered fair.

Aziz and Ye [2014] designed a different randomized mechanism for *piecewise-constant* valuations, called “Constrained Serial Dictatorship”, which is truthful in expectation, proportional, and satisfies a property called unanimity: if each agent’s most preferred $1/n$ length of the cake is disjoint from other agents, then each agent gets their most preferred $1/n$ length of the cake. This is a kind of “super-proportionality” that is not satisfied by the mechanisms based on consensus division.

Another randomized mechanism is implied (though not explicitly stated) by Cole et al. [2013]. They present a *partial allocation mechanism* (PAM) for homogeneous divisible goods. PAM initially finds a *max-product allocation* — an allocation maximizing the product of agents’ utilities. Then it calculates, for each agent i , a ratio f_i , defined as the product of all agents when i is present, divided by the max-product of all agents when i is absent. Then it allocates to each agent i a fraction f_i of each resource allocated to i in the original max-product allocation. They prove that $f_i \geq 1/e \approx 0.368$, and thus their mechanism guarantees each agent at least 0.368 of his/her utility in the max-product allocation. PAM can be adapted to cake-cutting by interpreting f_i as a probability: with probability f_i , agent i receives his/her piece in the max-product cake-allocation (such an allocation always exists; see Segal-Halevi and Sziklai [2019]); otherwise i receives nothing. This mechanism guarantees envy-freeness (ex-ante), but does not guarantee proportionality (neither ex-ante nor ex-post); it might also discard a large part of the cake.

In contrast to the above papers, we focus on deterministic mechanisms, that guarantee fairness and never discard a part of the cake.

2.3 Deterministic truthful mechanisms

All deterministic truthful mechanisms that we are aware of only work in the direct-revelation model and for the very restricted domain of *piecewise-uniform* valuations. In this line, Chen et al. [2013] provide a deterministic mechanism that is strategy-proof, envy-free and Pareto-efficient. Later, Li et al. [2015] proved that their mechanism remains truthful even in the presence of externalities (when agents derive small benefits from the welfare of others), as long as the magnitude of externalities is sufficiently small. On the negative side, their mechanism assumes free-disposal, since it may have to discard pieces of cake in order to attain truthfulness. In particular, it does not work for a negative cake (“chore”) which must be fully allocated.

Alijani et al. [2017] provide several deterministic truthful mechanisms, but only for a sub-class of piecewise-uniform valuations in which each agent’s desired subset is a single interval.

Bei et al. [2018] provide a mechanism that does not need to dispose any part of the cake. However, it works only for two agents with piecewise-uniform valuations, or for n agents where each agent’s desired subset is a single interval. An advantage of their mechanism is that it works with a negative cake too (they do not mention whether their mechanism works with a mixed cake).

For *piecewise-constant valuations*, most results regarding deterministic truthful mechanisms are negative: No such mechanism is both proportional and Pareto-optimal [Aziz and Ye 2014]; No such mechanism is approximately-proportional and non-wasteful, or approximately-proportional and connected [Menon and Larson 2017], or envy-free and *position-oblivious* (the allocation of a

¹Their dictatorship theorem assumes that a Robertson–Webb protocol may cut the cake only in cut-points determined by agents’ replies to queries; see subsection 3.2 for a formal definition. Without this assumption, there is a trivial strategyproof mechanism that simply ignores the agents’ actions and gives each agent a fixed non-empty piece of cake. Such a mechanism clearly does not achieve proportionality.

cake-part is based only on the agents' valuations of that part, and not on its relative position on the cake) [Bei et al. 2017].

In contrast to the above papers, we aim for mechanisms that work for arbitrary valuations, and do not dispose a part of the cake. Moreover, our mechanisms can also be used, *mutatis mutandis*, to divide bad cakes (those in which every piece yields dis-utility to agents but the whole cake must be allocated) and even to divide cakes with a mixture of good and bad parts.

2.4 Relaxations of truthfulness

Since truthfulness is unattainable with general preferences, several authors have studied relaxations of it.

Menon and Larson [2017] study *ϵ -strategyproof mechanisms* — allocation mechanisms in which an agent can only increase his/her utility by a ϵ (a fraction of the total utility) by reporting false preferences, compared to reporting true preferences. They prove that the Even-Paz algorithm [Even and Paz 1984] is ϵ -truthful for a constant ϵ that becomes worse as the number of agents grows. They present a variant in which the approximation factor is better (smaller). [Kyropoulou et al. 2019] obtain bounds on the amount of extra utility that agents can guarantee by lying in other proportional cake-cutting mechanisms.

Ianovski [2012] presents a mechanism for piecewise-uniform valuations, that maximizes the social welfare when the agents are truthful, and even when the agents are strategic, it has Nash equilibria that are Pareto-efficient and envy-free.

Bei et al. [2017] present mechanisms in which truth-telling is a *Nash equilibrium* (rather than a dominant-strategy equilibrium). In particular, they prove that, if each agent is replicated k times, then in any envy-free mechanism, truth-telling converges to a Nash equilibrium when k approaches infinity. If the pieces are allowed to be disconnected, then there is a mechanism in which truth-telling is a Nash equilibrium already when $k \geq 2$.

In this paper we explore a NOM which is a different relaxation of strategy-proofness. In contrast to the other relaxations, which are based on a quantitative constant or on replicating agents, the NOM relaxation is based on a natural qualitative assumption on the behavior of bounded-rationality agents. 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.

2.5 Other approaches

A new research avenue, recently proposed by [Bogomolnaia et al. 2019] has taken a different approach to fairness properties in cake-cutting. They provide utility levels that agents can guarantee by truthful behavior, no matter what their opponents do (or that can be implemented in protective equilibrium [Barberá and Dutta 1982]). They propose a generalized divide-and-choose mechanism that guarantees to each agent the utility associated to the best piece of cake in the worst cake partition.

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 set function that maps each piece to a real number.

The cakes we consider in most of the paper are *all-good*, in the sense that the utility of every piece is weakly-positive for each agent. These cakes are by the far the most studied in the literature. Section 7 extends the results to an *all-bad* cake, in which the utility of every piece is weakly-negative for each agent. Extending the results to a *mixed* cake, which may have both positive and negative parts, remains an open problem.

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 $u_i([x, z]; \theta_i) = \lambda u_i([x, y]; \theta_i)$.²

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 .³ 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) *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 previously made cut points or 0 or 1.
- (2) *Cut query*: the query $\text{CUT}(i; x, \alpha)$ asks agent i for an arbitrary $y \in [0, 1]$ such that $u_i([x, y]; \theta_i) = \alpha$; where $\alpha \in \mathbb{R}$ and x is an existing cut point or 0 or 1. The point y becomes a cut point.

All the mechanisms that we consider in this paper can be expressed as a Robertson–Webb extensive form.

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.⁴ 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 .

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 $E(\theta) = \{e^h(\theta) | h \in H\}$. This is, $E(\theta)$ is the set of all edges corresponding to truthful reports. An alternative set of edges in which all agents except i answer all queries

²Both 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.

³An arborescence is a directed, rooted tree in which all edges point away from the root.

⁴This 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 $[0, 0.3)$.




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

truthfully, and agent i answers the queries as if he were 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. In such a mechanism, there is a unique non-terminal node, which is the root h_0 . The set of active agents there is $N(h_0) = N$, 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 at least $1/n$ the total value, regardless of what every other agent reports at any node of the corresponding extensive form. Formally, for all $i \in N$, all $\theta_i \in \Theta_i$ and all $\theta'_{-i} \in \Theta_{-i}$:

$$u_i(M_i(E(\theta_i, \theta'_{-i}); \theta_i)) \geq u_i([0, 1]; \theta_i)/n.$$

3.3 Manipulations

A *manipulation* by agent i to the mechanism M is a set of (partially) untruthful answers along the corresponding extensive form 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 $u_i(M_i(E(\theta'_i, \theta_{-i})); \theta_i) > u_i(M_i(E(\theta_i, \theta_{-i})); \theta_i)$ for at least one θ_{-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*

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].

manipulable (NOM) if, for any profitable manipulation of agent i , corresponding to pretending being a type θ'_i , the following two conditions hold:⁶

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

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

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. We *emphasize* that NOM depends on the particular extensive form of the mechanism because it requires that, at each node in which an agent is active, answering a query truthfully leads to a weakly larger payoff than telling any lie (consistent with previous answers) in both the best- and worst-case scenarios. Such requirement is present in our definition via $E(\theta)$ and $E(\theta_i, \theta_{-i})$.

NOM is a strengthening of *maximin strategy-proofness*, as defined by Brams, Jones and Klamler [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 4.1).

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”, for example, 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 of a mechanism, using queries, emphasizes that there is a “true” answer to each query. Therefore the concept of truthfulness is meaningful. 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 4.1.

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

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

4 RESULTS

4.1 Obviously-manipulable mechanisms

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

The example considers a cake-cutting problem with positive piecewise uniform valuations, i.e. agents either like or dislike certain intervals, each desirable interval of the same length has the same positive value, and each undesirable interval has a value of 0. One agent, called Blue, has valuations as in Figure 2.




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




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

We now consider several classic mechanisms in turn.

Cut and choose. In this well-known mechanism, Blue (the cutter) is asked CUT (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.

Cut middle. In this mechanism, both agents are asked 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.

Banach-Knaster (Last Diminisher). In this mechanism [Steinhaus 1948], 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 3(a). Suppose agent 1 answers truthfully $x_1 = 0.5$. If agent 2 answers truthfully $x_2 = 0.2$, then agent 2’s value is 0.5. A profitable manipulation for agent 2 is to answer $x'_2 = x_1 - \epsilon$ for some small $\epsilon > 0$; it yields a value of 1. This is true in both the best- and worst-case scenarios, which are the same in this case.

Dubins-Spanier. In this procedure [Dubins and Spanier 1961], a knife moves continuously over the cake from left to right, the first agent who shouts “stop!” stops the knife and receives the piece to its left, and the process repeats with the remaining agents. Since this procedure is continuous, it does not fit our discrete model of a mechanism. Below we describe two common adaptations of this procedure to a discrete mechanism, both of which are obviously-manipulable.

First, consider the adaptation described by Procaccia [2016]. Here, in period 1, all agents i are simultaneously asked $\text{CUT}(i; 0, \frac{u_i([0,1]; \theta_i)}{n})$, and the agent who replied with the smallest number (say x^1) receives the piece $[0, x^1]$. Then, in each period $t = 2, \dots, n-1$, all agents i are simultaneously asked $\text{CUT}(i; x^{t-1}, \frac{u_i([0,1]; \theta_i)}{n})$, the agent who replied with the smallest number (x^t) receives the piece $[x^{t-1}, x^t]$. Finally, agent n receives the piece $[x^{n-1}, 1]$.

To see that this variant is obviously manipulable, consider the case of a truthful agent i , who has uniform preferences over the whole $[0, 1]$ interval and has to cut the cake against 4 other agents. Her first cut is $x^1_i = 0.2$, which guarantees her a utility of at least 0.2. Suppose the lowest cut (submitted by someone else) was $x^1 = 0.1$. So in period $t = 2$, the remaining cake is $[0.1, 1]$, and agent i is asked $\text{CUT}(i; 0.1, 0.2)$ to which the truthful answer is 0.3. However, there is a manipulation that is guaranteed to yield i a higher utility in the worst-case scenario. If she cuts the cake instead at $\hat{x}^2_i = 0.325$ (i.e. the point at which $u_i([0.1, \hat{x}^2_i]; \theta_i) = 1/4$), 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}^t_i such that $x^t_i = \frac{([x^{t-1}, 1])}{n-t+1}$ for all subsequent t , she could make sure to receive a utility of at least 0.225 (see Lemma 1 below), which is larger than the worst-case scenario utility received by being truthful, 0.2.

Second, consider the adaptation (similar to the one described by Brânzei and Nisan [2017]), in which a knife moves from 0 to 1 in discrete increments of size δ , which can be an arbitrarily small positive number. After each increment, every agent is asked whether $u_i([0, x]; \theta_i) \geq 1/n$, where x is the current knife location. Whenever one or more agents answer “yes”, the knife is stopped, and one of these agents (with the smallest index) receives the interval $[0, x]$. The remaining agents then divide the rest of the cake in the same fashion.

This variant is obviously manipulable even with two agents. For example, consider Figure 3(a) and suppose $\delta = 10^{-k}$, for some integer $k \geq 1$. When the knife arrives at $x = 0.2$, the truthful action of agent 2 is to say “yes” and stop the knife. A profitable manipulation for agent 2 is to say “no”, and keep saying “no” until $x > 0.4$ (since agent 1 also says “no” in this case). This manipulation guarantees agent 2 a value of 1.

4.2 A non-obviously-manipulable mechanism

We now present a third adaptation of the Dubins–Spanier mechanism, in which, in each period, all agents are asked the cut query simultaneously. We call this variant **leftmost leaves**. It is presented as Algorithm 1.

Algorithm 1 Leftmost-leaves algorithm for an all-good cake.

Require: A cake $[0, 1]$, and n agents who have a weakly-positive value for every piece of cake.

Ensure: A proportional allocation of the entire cake.

- 1: In period 1, each agent i is asked $\text{CUT} \left(i; 0, \frac{u_i([0,1];\theta_i)}{n} \right)$. 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 $[0, x^1]$.
 - 2: In period 2, every remaining agent i is asked $\text{CUT} \left(i; x^1, \frac{u_i([x^1,1];\theta_i)}{n-1} \right)$. 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 $(x^1, x^2]$.
 - 3: Similarly, in period t , all remaining agents are asked $\text{CUT} \left(i; x^{t-1}, \frac{u_i([x^{t-1},1];\theta_i)}{n-t+1} \right)$, and the agent who submits the smallest point (denoted x^t) leaves with the interval $(x^{t-1}, t^t]$. 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 4.1. *The leftmost leaves mechanism is proportional and not-obviously-manipulable.*

Before presenting the proof, let us present a few remarks.

First, the leftmost-leaves mechanism differs subtly than the first adaptation of Dubins–Spanier presented in subsection 4.1 (following Procaccia [2016]). The difference is that in the former adaptation, the agents in period t are asked $\text{CUT} \left(i; x^{t-1}, \frac{u_i([0,1];\theta_i)}{n} \right)$, while in leftmost-leaves, they are asked $\text{CUT} \left(i; x^{t-1}, \frac{u_i([0,1];\theta_i)}{n-t+1} \right)$. In the example of the previous subsection, in period $t = 2$, agent i was asked $\text{CUT} (i; 0.1, 0.2)$, while in leftmost-leaves the query is $\text{CUT} (i; 0.1, 0.25)$. This apparently small change makes a big difference in terms of strategic properties: the former adaptation is obviously-manipulable while the latter is not.

Second, 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 yield the same allocation (up to the arbitrarily small increment δ). 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. A dynamic implementation is thus helpful for strategyproof mechanisms and harmful for manipulable ones.

Third, the leftmost-leaves mechanism satisfies several other desiderata that make it a good candidate to divide a cake in practice. Besides being proportional and NOM, it is 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 time or land.

5 PROOF OF THEOREM 4.1

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 arbitrary agent, denoted below by i .

First, we show that no manipulation yields a higher utility in the worst-case scenario. We use the following lemma, where for convenience we define $x^0 = 0$.

Lemma 1. *For all $t \geq 1$, at period t , the worst possible utility of a truthful agent i is*

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

PROOF. If x_i^t is chosen as the smallest cut at period t , then i receives the piece $[x^{t-1}, x^t] = [x^{t-1}, x_i^t]$ and the result is immediate.

Otherwise, $x^t \leq x_i^t$. Since we assume that the cake is all-good, this implies that

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

so the piece given away $[x^{t-1}, x^t]$ is worth for i at most $\frac{u_i([x^{t-1}, 1]; \theta_i)}{n - t + 1}$.

Therefore, the remainder of the cake $[x^t, 1]$ must be 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 his cut was chosen as the smallest cut in period $t + 1$. If his cut was not the smallest at period $t + 1$, an inductive argument 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 for agent i a utility weakly smaller than $\frac{u_i([x^{t-1}, 1]; \theta_i)}{n - t + 1}$ in the worst-case scenario.

Suppose that, from period t onwards, agent i manipulates in some way, while all other $n - t$ agents behave exactly as a truthful agent i would behave (i.e., they answer all queries as if their utility function is u_i). By Lemma 1, each such agent j must receive some piece X_j such that $u_i(X_j; \theta_i) \geq \frac{1}{n-t+1} u_i([x^{t-1}, 1]; \theta_i)$. The total value of the pieces given to these $n - t$ agents is at least $\frac{n-t}{n-t+1} u_i([x^{t-1}, 1]; \theta_i)$. Hence, the value remaining for agent i is at most $\frac{1}{n-t+1} u_i([x^{t-1}, 1]; \theta_i)$.

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. *For all $t \geq 1$, at period t ,*

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

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. The supremum is obviously at most $u_i([x^{t-1}, 1]; \theta_i)$, since the agent cannot get more than the whole remaining cake. To prove that the supremum is at least $u_i([x^{t-1}, 1]; \theta_i)$, it is sufficient to prove that, for every $\epsilon > 0$, there exists some θ_{-i} such that $u_i(M^{\text{ll}}(E(\theta_i, \theta_{-i})); \theta_i) \geq u_i([x^{t-1}, 1]; \theta_i) - \epsilon$.

Let us remember that a truthful agent i cuts the cake at a point $x_i^t > x^{t-1}$ such that

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

Given some $\epsilon > 0$, let θ_{-i} be such that the smallest cut at each period $t' \geq t$, denoted $x^{t'}$, is to the left of the cut of agent i , but only slightly to the right of the previous cut. Formally:

$$\begin{aligned} x^{t'-1} &< x^{t'} < x_i^{t'} \\ u_i([x^{t'-1}, x^{t'}]; \theta_i) &\leq \epsilon/n \end{aligned}$$

Such $x^{t'}$ values exist due to the standard assumption that the utilities are divisible (u_i is a continuous function).

With such a sequence of cuts, agent i remains the last agent, and receives the piece $X_i = [x^{n-1}, 1]$. Since the pieces given to each of the other agents are worth for i at most ϵ/n , the remaining piece is worth for him at least $u_i([x^{t-1}, 1]; \theta_i) - \epsilon$. \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.

6 ALTERNATIVE MECHANISMS

A natural follow-up question to Theorem 4.1 is whether leftmost leaves is the only proportional and NOM mechanism in cake-cutting. The answer is no, as shown below.

6.1 Query-based mechanisms

Leftmost-leaves can be slightly modified in several ways retaining both NOM and proportionality. One such modification is to start cutting the cake from the right instead of from the left.⁸

Another less trivial one is an adaptation of the protocol of [Even and Paz 1984], which works the same as leftmost-leaves for $n = 2$, but requires fewer queries for larger values of n . The adaptation works as follows (for simplicity we present it for n that is a power of 2). Given a cake $[y, z]$, all agents choose cuts x_i such that $u_i([y, x_i]; \theta_i) = u_i([y, z]; \theta_i)/2$. Order the cuts in increasing order, and let x^* be the cut number $n/2$ from the left. Then the procedure breaks the cake-cutting problem

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].

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. Finally, when the procedure is called with two agents, they divide the remaining cake among them using leftmost-leaves. 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.

The proof that this variant of Even–Paz is NOM is similar to the proof of Theorem 4.1: at each point you are guaranteed the entire cake in the best case, and deviating from telling the truth may lead to a worse than proportional outcome.

Obtaining a characterization of all NOM and proportional mechanisms remains an interesting, albeit challenging, open question.

6.2 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 6.1. *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 [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 6.1, it remains an open question whether there exists a NOM mechanism that is both connected and envy-free (for three or more agents).

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

7 DIVIDING A CAKE WITH NEGATIVE PARTS

In this section we extend the model and allow agents to have negative utilities. We have managed to extend the leftmost-leaves mechanism to an all-bad cake, in which all utilities are negative. However, so far we have not managed to extend it to a mixed cake.

7.1 All-bad cakes

When all utilities are weakly-negative, the *rightmost-leaves* mechanism (Algorithm 2) is proportional and NOM. The main difference from the leftmost-leaves mechanism is that, at each step, the *largest* cut-point is selected instead of the smallest.

Algorithm 2 Rightmost-leaves algorithm for an all-bad cake.

Require: A cake $[0, 1]$, and n agents who have a weakly-positive value for every piece of cake.

Ensure: A proportional allocation of the entire cake.

- 1: In period 1, each agent i is asked $\text{CUT} \left(i; 0, \frac{u_i([0,1];\theta_i)}{n} \right)$. The agent who cuts the cake at the *largest* 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 $[0, x^1]$.
 - 2: In period 2, every remaining agent i is asked $\text{CUT} \left(i; x^1, \frac{u_i([x^1,1];\theta_i)}{n-1} \right)$. The agent who submits the largest 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 $(x^1, x^2]$.
 - 3: Similarly, in period t , all remaining agents are asked $\text{CUT} \left(i; x^{t-1}, \frac{u_i([x^{t-1},1];\theta_i)}{n-t+1} \right)$, and the agent who submits the largest point (denoted x^t) leaves with the interval $(x^{t-1}, x^t]$. The procedure continues until only one agent remains. That agent receives $(x^{n-1}, 1]$.
-

Lemma 1 is true as-is for the rightmost-leaves mechanism when the cake is all-bad: for all $t \geq 1$, at period t , the smallest possible utility of a truthful agent i is $\frac{u_i([x^{t-1},1];\theta_i)}{n-t+1}$ (note that this utility is now a negative number). The proof is very similar. The piece given away at each step t is at least as *large* as the piece that agent i values at $\frac{u_i([x^{t-1},1];\theta_i)}{n-t+1}$, and the cake is all-bad, so the value of the piece given away is still *at most* $\frac{u_i([x^{t-1},1];\theta_i)}{n-t+1}$. Thus, rightmost-leaves is proportional, and inequality (1) is satisfied.

Instead of Lemma 2, we have

Lemma 5. For all $t \geq 1$, at period t ,

$$\sup_{\theta_{-i}} u_i(M^{rl}(E(\theta_i, \theta_{-i})); \theta_i) = 0$$

This is, agent i , by being truthful in period t , can expect (in the best-case scenario) to obtain no bad cake at all.

PROOF. The supremum is obviously at most 0 since the cake is all-bad. To prove that the supremum is at least 0, it is sufficient to prove that, for every $\epsilon > 0$, there exists some θ_{-i} such that $u_i(M^{rl}(E(\theta_i, \theta_{-i})); \theta_i) \geq -\epsilon$.

Let us remember that a truthful agent i cuts the cake at a point $x_i^t > x^{t-1}$ such that

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

Given some $\epsilon > 0$, let θ_{-i} be such that the largest cut at period t , denoted x^t , is to the right of the cut of agent i , and only slightly to the left of 1, such that:

$$x_i^t < x^t$$


$$u_i([x^t, 1]; \theta_i) \geq -\epsilon$$

With such a cut, the rightmost cutter at period t takes almost all of the remaining cake, and the piece that remains is so small that its value for agent i is at least $-\epsilon$. Regardless of what happens in the following rounds, the utility of agent i is at least $-\epsilon$. \square

Since 0 is the largest possible utility, inequality (2) is satisfied too, so rightmost-leaves is NOM.

7.2 Mixed cakes

In the most general cake model, the utility of each agent to every piece of cake may be either positive or negative. So far, we could not adapt the leftmost-leaves mechanism to handle this general setting while keeping its NOM property. We provide an Example in Figure 4 to show that adapting the leftmost-leaves mechanism to this setup is not trivial. In this example, the agent whose preferences we depict gets a utility of 1 if she gets the whole cake, a utility of 1 for each blue piece of length 0.1, and a utility of -1 for each red piece of length 0.1.



(a) There are three truthful cutpoints. In the best-case scenario, x gives a utility of 1, y of 0.5, and z of 2. (b) Cutting the cake at m gives a best-case scenario utility of 3, and thus is an obvious manipulation.

Fig. 4. A mixed cake in which leftmost leaves is OM. Good parts of the cake appear in blue, bad parts in red.

In a mixed cake, unlike in an all-good or all-bad cake, a $\text{CUT}(i; x, \alpha)$ query can be answered by several disconnected values, since the function $u_i([x, y])$ is not a monotonic function of y . For example, in Figure 4(a), the $\text{CUT}(i; 0, 0.5)$ query has three disconnected truthful answers, marked by x, y, z . Each of these answers guarantees Alice a value of 0.5 if she wins in the first round. Otherwise, the other agent wins some piece adjacent to the left end of the cake, and Alice gets the remaining piece, which is adjacent to the right end of the cake. The largest possible value that

Alice can get from a piece adjacent to the right end of the cake is 2 (for the piece $[0.8, 1]$). Therefore the best-case utility of a truthful agent is 2.

However, if Alice answers 0.4 (denoted by m in Figure 4(b)), then in the best case she wins the piece $[0, 0.4]$ in the first round and her utility is 3.

We conjecture that, in the Robertson–Webb model, there is no NOM mechanism that produces a connected and proportional allocation of a mixed cake. We leave further investigation of this interesting question for future research.

8 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 smaller than 0.0001). 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 smaller than 0.0001) and 3 opponents (difference of 16 percentage points, p-value smaller than 0.0001). The Even–Paz modification of leftmost leaves (which is NOM too) was also significantly less manipulated than Banach–Knaster (difference of 35 percentage points, p-value smaller than 0.0001).

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

- Reza Alijani, Majid Farhadi, Mohammad Ghodsi, Masoud Seddighin, and Ahmad S Tajik. 2017. Envy-free mechanisms with minimum number of cuts. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*. 312–318.
- Sergey Avvakumov and Roman Karasev. 2019. Envy-free division using mapping degree. *arXiv preprint* (2019).
- Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2019. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 53–59.
- Haris Aziz and Chun Ye. 2014. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *International Conference on Web and Internet Economics*. Springer, 1–14.
- Salvador Barberá and Bhaskar Dutta. 1982. Implementability via protective equilibria. *Journal of Mathematical Economics* 10, 1 (1982), 49–65.
- Xiaohui Bei, Ning Chen, Guangda Huzhang, Biaoshuai Tao, and Jiajun Wu. 2017. Cake cutting: envy and truth. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 3625–3631.
- Xiaohui Bei, Guangda Huzhang, and Warut Suksompong. 2018. Truthful fair division without free disposal. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 63–69.
- Anna Bogomolnaia, Herve Moulin, and Richard Stong. 2019. Guarantees in Fair Division: beyond Divide & Choose and Moving Knives. *arXiv preprint* (2019).
- Steven Brams, Michael Jones, and Christian Klamler. 2006. Better ways to cut a cake. *Notices of the AMS* 53, 11 (2006), 1314–1321.
- Steven Brams and Alan Taylor. 1996. *Fair division: From cake-cutting to dispute resolution*. Cambridge University Press.
- Steven J Brams, Michael A Jones, and Christian Klamler. 2008. Proportional pie-cutting. *International Journal of Game Theory* 36, 3-4 (2008), 353–367.
- Simina Brânzei and Peter Bro Miltersen. 2015. A dictatorship theorem for cake cutting. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. AAAI Press.

- Simina Brânzei and Noam Nisan. 2017. The query complexity of cake cutting. *arXiv preprint arXiv:1705.02946* (2017).
- Yiling Chen, John Lai, David Parkes, and Ariel Procaccia. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77, 1 (2013), 284–297.
- Richard Cole, Vasilis Gkatzelis, and Gagan Goel. 2013. Mechanism design for fair division: Allocating divisible items without payments. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. ACM, 251–268.
- Vincent P. Crawford. 1977. A game of fair division. *Review of Economic Studies* 44, 2 (1977), 235–247.
- Lester Dubins and Edwin Spanier. 1961. How to cut a cake fairly. *Amer. Math. Monthly* 68, 1 (1961), 1–17.
- Shimon Even and Azaria Paz. 1984. A note on cake cutting. *Discrete Applied Mathematics* 7, 3 (1984), 285 – 296.
- Egor Ianovski. 2012. Cake Cutting Mechanisms. *arXiv preprint* (2012).
- David Kurokawa, John K Lai, and Ariel D Procaccia. 2013. How to cut a cake before the party ends. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. 555–561.
- Maria Kyropoulou, Josué Ortega, and Erel Segal-Halevi. 2019. Fair cake-cutting in practice. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 547–548.
- Minming Li, Jialin Zhang, and Qiang Zhang. 2015. Truthful cake cutting mechanisms with externalities: Do not make them care for others too much!. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*. AAAI Press, 589–595.
- Shengwu Li. 2017. Obviously strategy-proof mechanisms. *American Economic Review* 107, 11 (2017), 3257–87.
- Vijay Menon and Kate Larson. 2017. Deterministic, strategyproof, and fair cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 352–358.
- Frédéric Meunier and Shira Zerbib. 2019. Envy-free cake division without assuming the players prefer nonempty pieces. *Israel Journal of Mathematics* 234, 2 (2019), 907–925.
- John Moore and Rafael Repullo. 1988. Subgame perfect implementation. *Econometrica* (1988), 1191–1220.
- Elchanan Mossel and Omer Tamuz. 2010. Truthful fair division. In *Algorithmic Game Theory*, Spyros Kontogiannis, Elias Koutsoupias, and Paul Spirakis (Eds.). Springer, 288–299.
- Hervé Moulin. 2004. *Fair division and collective welfare*. MIT press.
- Antonio Nicolò and Yan Yu. 2008. Strategic divide and choose. *Games and Economic Behavior* 64, 1 (2008), 268 – 289.
- Ariel Procaccia. 2016. Cake cutting algorithms. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jerome Lang, and Ariel Procaccia (Eds.). Cambridge University Press, 311–330.
- Jack Robertson and William Webb. 1998. *Cake-cutting algorithms: Be fair if you can*. AK Peters/CRC Press.
- Erel Segal-Halevi. 2018. Fairly dividing a cake after some parts were burnt in the oven. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems*. 1276–1284.
- Erel Segal-Halevi and Balázs R Sziklai. 2018. Resource-monotonicity and population-monotonicity in connected cake-cutting. *Mathematical Social Sciences* 95 (2018), 19–30.
- Erel Segal-Halevi and Balázs R Sziklai. 2019. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory* 68, 2 (2019), 363–401.
- Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16, 1 (1948), 101–104.
- Peter Troyan and Thayer Morrill. 2019. Obvious manipulations. *Journal of Economic Theory* (2019), 104970.