

An Algorithmic Framework for Strategic Fair Division

Simina Brânzei

University of California Berkeley
simina.branzei@gmail.com

Ioannis Caragiannis

University of Patras
caragian@ceid.upatras.gr

David Kurokawa

Carnegie Mellon University
dkurokaw@cs.cmu.edu

Ariel D. Procaccia

Carnegie Mellon University
arielpro@cs.cmu.edu

Abstract

We study the paradigmatic fair division problem of fairly allocating a divisible good among agents with heterogeneous preferences, commonly known as *cake cutting*. Classic cake cutting protocols are susceptible to manipulation. Do their strategic outcomes still guarantee fairness? To address this question we adopt a novel algorithmic approach, proposing a concrete computational model and reasoning about the game-theoretic properties of algorithms that operate in this model. Specifically, we show that each protocol in the class of *generalized cut and choose (GCC) protocols* — which includes the most important discrete cake cutting protocols — is guaranteed to have approximate subgame perfect Nash equilibria, or even exact equilibria if the protocol’s tie-breaking rule is flexible. We further observe that the (approximate) equilibria of proportional protocols — which guarantee each of the n agents a $1/n$ -fraction of the cake — must be (approximately) proportional, thereby answering the above question in the positive (at least for one common notion of fairness).

1 Introduction

A large body of literature deals with the so-called *cake cutting* problem — a misleadingly childish metaphor for the challenging and important task of *fairly* dividing a heterogeneous divisible good among multiple agents (see the recent survey by Procaccia (2013) and the books by Brams and Taylor (1996) and Robertson and Webb (1998)). In particular, there is a significant amount of AI work on cake cutting (Procaccia 2009; Caragiannis, Lai, and Procaccia 2011; Brams et al. 2012; Bei et al. 2012; Aumann, Dombb, and Hassidim 2013; Kurokawa, Lai, and Procaccia 2013; Brânzei, Procaccia, and Zhang 2013; Brânzei and Miltersen 2013; Chen et al. 2013; Balkanski et al. 2014; Brânzei and Miltersen 2015; Segal-Halevi, Hassidim, and Aumann 2015), which is closely intertwined with emerging real-world applications of fair division more broadly (Goldman and Procaccia 2014; Kurokawa, Procaccia, and Shah 2015).

Going back to the word “fairly”, two formal notions of fairness have emerged as the most appealing and well-studied in the context of cake cutting: *proportionality*, in which each of the n agents receives at least a $1/n$ -fraction of the entire cake according to its valuation; and *envy-freeness*,

which stipulates that no agent would wish to swap its own piece with that of another agent. At the heart of the cake cutting endeavor is the design of cake cutting *protocols*, which specify an interaction between agents — typically via iterative steps of manipulating the cake — such that the final allocation is guaranteed to be proportional or envy-free.

The simplest cake cutting protocol is known as *cut and choose*, and is designed for two agents. The first agent cuts the cake in two pieces that it values equally; the second agent then chooses the piece that it prefers, leaving the first agent with the remaining piece. It is easy to see that this protocol yields a proportional and envy-free allocation (in fact these two notions coincide when there are only two agents and the entire cake is allocated). However, taking a game-theoretic point of view, it is immediately apparent that the agents can often do better by disobeying the protocol when they know each other’s valuations. For example, in the cut and choose protocol, assume the first agent only desires a specific small piece of cake, whereas the second agent uniformly values the cake. The first agent can obtain its entire desired piece, instead of just half of it, by carving that piece out.

So how would strategic agents behave when faced with the cut and choose protocol? A standard way of answering this question employs the notion of *Nash equilibrium*: each agent would use a strategy that is a best response to the other agent’s strategy. To set up a Nash equilibrium, suppose that the first agent cuts two pieces that the second agent values equally; the second agent selects its more preferred piece, and the one less preferred by the first agent in case of a tie. Clearly, the second agent cannot gain from deviating, as it is selecting a piece that is at least as preferred as the other. As for the first agent, if it makes its preferred piece even bigger, the second agent would choose that piece, making the first agent worse off. Interestingly enough, in this equilibrium the tables are turned; now it is the second agent who is getting exactly half of its value for the whole cake, while the first agent generally gets more. Crucially, the equilibrium outcome is also proportional and envy-free. In other words, even though the agents are strategizing rather than following the protocol, the outcome in equilibrium has the same fairness properties as the “honest” outcome!

With this motivating example in mind, we would like to make general statements regarding the equilibria of cake cutting protocols. We wish to identify a general family of

cake cutting protocols — which captures the classic cake cutting protocols — so that each protocol in the family is guaranteed to possess (approximate) equilibria. Moreover, we wish to argue that these equilibrium outcomes are fair. Ultimately, our goal is to be able to reason about the fairness of cake divisions that are obtained as outcomes when agents are presented with a standard cake cutting protocol and behave strategically.

1.1 Model and Results

To set the stage for a result that encompasses classic cake cutting protocols, we introduce (in Section 2) the class of *generalized cut and choose (GCC)* protocols. A GCC protocol is represented by a tree, where each node is associated with the action of an agent. There are two types of nodes: a *cut node*, which instructs the agent to make a cut between two existing cuts; and a *choose node*, which offers the agent a choice between a collection of pieces that are induced by existing cuts. Moreover, we assume that the progression from a node to one of its children depends only on the relative positions of the cuts (in a sense to be explained formally below). We argue that classic protocols — such as Dubins-Spanier (1961), Selfridge-Conway (see (Robertson and Webb 1998)), Even-Paz (1984), as well as the original cut and choose protocol — are all GCC protocols. We view the definition of the class of GCC protocols as one of our main contributions.

In Section 3, we observe that GCC protocols may not have exact Nash equilibria (NE), then explore ways of circumventing this issue, which give rise to our two main results.

1. We prove that every GCC protocol has at least one ϵ -NE for every $\epsilon > 0$, in which agents cannot gain more than ϵ by deviating, and ϵ can be chosen to be arbitrarily small. In fact, we establish this result for a stronger equilibrium notion, (approximate) *subgame perfect Nash equilibrium (SPNE)*, which is, intuitively, a strategy profile where the strategies are in NE even if the game starts from an arbitrary point.
2. We slightly augment the class of GCC protocols by giving them the ability to make *informed tie-breaking* decisions that depend on the entire history of play, in cases where multiple cuts are made at the exact same point. While, for some valuation functions of the agents, a GCC protocol may not possess any exact SPNE, we prove that it is always possible to modify the protocol’s tie-breaking scheme to obtain SPNE.

In Section 4, we observe that for any proportional protocol, the outcome in any ϵ -equilibrium must be an ϵ -proportional division. We conclude that under the classic cake cutting protocols listed above — which are all proportional — strategic behavior preserves the proportionality of the outcome, either approximately, or exactly under informed tie-breaking.

One may wonder, though, whether an analogous result is true with respect to envy-freeness. We give a negative answer, by constructing an envy-inducing SPNE under the Selfridge-Conway protocol, a well-known envy-free protocol for three agents. However, we are able to design a curious

GCC protocol in which every NE outcome is a contiguous envy-free allocation and vice versa, that is, the set of NE outcomes coincides with the set of contiguous envy-free allocations. It remains open whether a similar result can be obtained for SPNE instead of NE.

1.2 Related Work

The notion of GCC protocols is inspired by the Robertson-Webb (1998) model of cake cutting — a concrete query model that specifies how a cake cutting protocol may interact with the agents. This model underpins a significant body of work in theoretical computer science and AI, which focuses on the complexity of achieving different fairness or efficiency notions in cake cutting (Edmonds and Pruhs 2006a; 2006b; Woeginger and Sgall 2007; Deng, Qi, and Saberi 2012; Aumann, Dombb, and Hassidim 2013; Procaccia 2009; Kurokawa, Lai, and Procaccia 2013). In Section 2, we describe the Robertson-Webb model in detail, and explain why it is inappropriate for reasoning about equilibria.

In the context of the strategic aspects of cake cutting, Nicolò and Yu (2008) were the first to suggest equilibrium analysis for cake cutting protocols. Focusing exclusively on the case of two agents, they design a specific cake cutting protocol whose unique SPNE outcome is envy-free. And while the original cut and choose protocol also provides this guarantee, it is not “procedural envy free” because the cutter would like to exchange roles with the chooser; the two-agent protocol of Nicolò and Yu aims to solve this difficulty. Brânzei and Miltersen (2013) also investigate equilibria in cake cutting, but in contrast to our work they focus on one cake cutting protocol — the Dubins-Spanier protocol — and restrict the space of possible strategies to *threshold strategies*. Under this assumption, they characterize NE outcomes, and in particular they show that in NE the allocation is envy-free. Brânzei and Miltersen also prove the existence of ϵ -equilibria that are ϵ -envy-free; again, this result relies on their strong restriction of the strategy space, and applies to one specific protocol.

Several papers by computer scientists (Chen et al. 2013; Mossel and Tamuz 2010; Maya and Nisan 2012) take a mechanism design approach to cake cutting; their goal is to design cake cutting protocols that are *strategyproof*, in the sense that agents can never benefit from manipulating the protocol. This turns out to be an almost impossible task (Zhou 1991; Brânzei and Miltersen 2015); positive results are obtained by either making extremely strong assumptions (agents’ valuations are highly structured), or by employing randomization and significantly weakening the desired properties. In contrast, our main results, given in Section 3, deal with strategic outcomes under a large class of cake cutting protocols, and aim to capture well-known protocols; our result of Section 4 is a positive result that achieves fairness “only” in equilibrium, but without imposing any restrictions on the agents’ valuations.

2 The Model

The cake cutting literature typically represents the cake as the interval $[0, 1]$. There is a set of agents $N = \{1, \dots, n\}$,

and each agent $i \in N$ is endowed with a *valuation function* V_i that assigns a value to every subinterval of $[0, 1]$. These values are induced by a non-negative continuous *value density function* v_i , so that for an interval I , $V_i(I) = \int_{x \in I} v_i(x) dx$. By definition, V_i satisfies the first two properties below; the third is an assumption that is made w.l.o.g.

1. **Additivity:** For every two disjoint intervals I_1 and I_2 , $V_i(I_1 \cup I_2) = V_i(I_1) + V_i(I_2)$.
2. **Divisibility:** For every interval $I \subseteq [0, 1]$ and $0 \leq \lambda \leq 1$ there is a subinterval $I' \subseteq I$ such that $V_i(I') = \lambda V_i(I)$.
3. **Normalization:** $V_i([0, 1]) = 1$.

Note that valuation functions are non-atomic, i.e., they assign zero value to points. This allows us to disregard the boundaries of intervals, and in particular we treat intervals that overlap at their boundary as disjoint. We sometimes explicitly assume that the value density functions are *strictly positive*, i.e., $v_i(x) > 0$ for all $x \in [0, 1]$ and for all $i \in N$; this implies that $V_i([x, y]) > 0$ for all $x < y, x, y \in [0, 1]$.

A *piece of cake* is a finite union of disjoint intervals. We are interested in allocations of disjoint pieces of cake X_1, \dots, X_n , where X_i is the piece allocated to agent $i \in N$. A piece is *contiguous* if it consists of a single interval.

We study two fairness notions. An allocation X is *proportional* if for all $i \in N$, $V_i(X_i) \geq 1/n$; and *envy-free* if for all $i, j \in N$, $V_i(X_i) \geq V_i(X_j)$. Note that envy-freeness implies proportionality when the entire cake is allocated.

2.1 Generalized Cut and Choose Protocols

The standard communication model in cake cutting was proposed by Robertson and Webb (1998); it restricts the interaction between the protocol and agents to two types of queries:

- **Cut query:** $Cut_i(x, \alpha)$ asks agent i to return a point y such that $V_i([x, y]) = \alpha$.
- **Evaluate query:** $Evaluate_i(x, y)$ asks agent i to return a value α such that $V_i([x, y]) = \alpha$.

Note that in the RW model, a protocol could allocate pieces depending on whether a particular cut was made at a specific point (see Algorithm 2). More generally, a protocol in the RW model has a property such as envy-freeness if, roughly speaking, it gathers enough information so that there *exists* an allocation such that for *any* valuation function consistent with the answers to the queries, the allocation is envy-free. Since the RW model does not specify how the allocation is computed, there need not exist a succinct representation of the allocation that arises as the outcome of a protocol, which makes it difficult to analyze the strategic properties of protocols in the RW model.

For this reason, we define a generic class of protocols that are implementable with natural operations, which capture all bounded¹ and discrete cake cutting algorithms, such as cut and choose, Dubins-Spanier, Even-Paz, Successive-Pairs, and Selfridge-Conway (see, e.g., (Procaccia 2013)).

¹In the sense that the number of operations is upper-bounded by a function that takes the number of agents n as input.

```
agent 1 Cuts in {[0, 1]} // @x
agent 1 Cuts in {[0, 1]} // @y
agent 1 Cuts in {[0, 1]} // @z
if ( $x < y < z$ ) then
  agent 1 Chooses from {[x, y], [y, z]}
end if
```

Algorithm 1: A GCC protocol. The notation “// @ x ” assigns the symbolic name x to the cut point made by agent 1.

```
agent 1 Cuts in {[0, 1]} // @x
if ( $x = \frac{1}{3}$ ) then
  agent 1 Chooses from {[0, x], [x, 1]}
end if
```

Algorithm 2: A non-GCC protocol.

At a high level, the standard protocols are implemented using a sequence of natural instructions, each of which is either a *Cut* operation, in which some agent is asked to make a cut in a specified region of the cake; or a *Choose* operation, in which some agent is asked to take a piece from a set of already demarcated pieces indicated by the protocol. In addition, every node in the decision tree of the protocol is based exclusively on the execution history and absolute ordering of the cut points, which can be verified with any of the following operators: $<$, \leq , $=$, \geq , $>$.

Formally, a *generalized cut and choose (GCC)* protocol is implemented exclusively with the following instructions:

- **Cut:** The syntax is “ i Cuts in S ”, where $S = \{[x_1, y_1], \dots, [x_m, y_m]\}$ is a set of contiguous pieces (intervals), such that the endpoints of every piece $[x_j, y_j]$ are 0, 1, or cuts made in previous steps of the protocol. Agent i can make a cut at any point $z \in [x_j, y_j]$, for some $j \in \{1, \dots, m\}$.
- **Choose:** The syntax is “ i Chooses from S ”, where $S = \{[x_1, y_1], \dots, [x_m, y_m]\}$ is a set of contiguous pieces, such that the endpoints of every piece $[x_j, y_j] \in S$ are 0, 1, or cuts made in the previous steps of the protocol. Agent i can choose any *single* piece $[x_j, y_j]$ from S to keep from that point on.
- **If-Else Statements:** The conditions depend on the result of choose queries and the absolute order of all the cut points made in the previous steps.

A GCC protocol uniquely identifies every contiguous piece by the symbolic names of all the cut points contained in it. For example, Algorithm 1 is a GCC protocol. Algorithm 2 is not a GCC protocol, because it verifies that the point where agent 1 made a cut is exactly $1/3$, whereas a GCC protocol can only verify the ordering of the cut points relative to each other and the endpoints of the cake. Note that, unlike in the communication model of Robertson and Web (1998), GCC protocols cannot obtain and use information about the valuations of the agents — the allocation is only decided by the agents’ *Choose* operations.

As an illustrative example, we now discuss why the dis-

crete variant of the Dubins-Spanier protocol² belongs to the class of GCC protocols — but first we must describe the original Dubins-Spanier protocol. Dubins-Spanier is a proportional (but not envy-free) protocol for n agents, which operates in n rounds. In round 0, each agent makes a mark x_i^1 such that the piece of cake to the left of the mark is worth $1/n$, i.e., $V_i([0, x_i^1]) = 1/n$. Let i^* be the agent that made the leftmost mark; the protocol allocates the interval $[0, x_{i^*}^1]$ to agent i^* ; the allocated interval and satisfied agent are removed. In round t , the same procedure is repeated with the remaining $n - t$ agents and the remaining cake. When there is only one agent left, it receives the remaining cake. To see why the protocol is proportional, first note that in round t the remaining cake is worth at least $1 - t/n$ to each remaining agent, due to the additivity of the valuation functions and the fact that the pieces allocated in previous rounds are worth at most $1/n$ to these agents. The agent that made the leftmost mark receives a piece that it values at $1/n$. In round $n - 1$, the last agent is left with a piece of cake worth at least $1 - (n - 1)/n = 1/n$.

The protocol admits a GCC implementation as follows. For the first round, each agent i is required to make a cut in $\{[0, 1]\}$, at some point denoted by x_i^1 . The agent i^* with the leftmost cut $x_{i^*}^1$ can be determined using *If-Else* statements whose conditions only depend on the ordering of the cut points x_1^1, \dots, x_n^1 . Then, agent i^* is asked to choose “any” piece in the singleton set $\{[0, x_{i^*}^1]\}$. The subsequent rounds are similar: at the end of every round the agent that was allocated a piece is removed, and the protocol iterates on the remaining agents and remaining cake. Note that agents are not constrained to follow the protocol, i.e., they can make their marks (in response to cut instructions) wherever they want; nevertheless, an agent can guarantee a piece of value at least $1/n$ by following the Dubins-Spanier protocol, regardless of what other agents do.

While GCC protocols are quite general, a few well-known cake cutting protocols are beyond their reach. For example, the Brams-Taylor (1995) protocol is an envy-free protocol for n agents, and although its individual operations are captured by the GCC formalism, the number of operations is not bounded as a function of n (i.e., it may depend on the valuation functions themselves). Its representation as a GCC protocol would therefore be infinitely long. In addition, some cake cutting protocols use *moving knives* (see, e.g., (Brams, Taylor, and Zwicker 1997)); for example, they can keep track of how an agent’s value for a piece changes as the piece smoothly grows larger. These protocols are not discrete, and, in fact, cannot be implemented even in the Robertson-Webb model.

We also note that the GCC model is *incomparable* to the RW model. Indeed, given a protocol in the RW model, it may not be possible to implement it as a GCC protocol because the RW model does not indicate a specific allocation, as discussed above. Conversely, cut queries in the GCC model cannot in general be translated into cut queries in the RW

²In fact, the discrete variant of Dubins-Spanier was invented much earlier by Banach and Knaster and is better known as the “last diminisher” procedure (see Steinhaus 1948).

model, as in the latter model cuts are associated with a specific value.

2.2 The Game

We study GCC protocols when the agents behave strategically. Specifically, we consider a GCC protocol, coupled with the valuation functions of the agents, as an *extensive-form game of perfect information* (see, e.g., (Shoham and Leyton-Brown 2008)). In such a game, agents execute the *Cut* and *Choose* instructions strategically. Each agent is fully aware of the valuation functions of the other agents and aims to optimize its overall utility for the chosen pieces, given the strategies of other agents.

While the perfect information model may seem restrictive, the same assumption is also made in previous work on equilibria in cake cutting (Nicolò and Yu 2008; Brânzei and Miltersen 2013). More importantly, it underpins foundational papers in a variety of areas of microeconomic theory, such as the seminal analysis of the Generalized Second Price (GSP) auction by Edelman et al. (2007). A common justification for the complete information setting, which is becoming increasingly compelling as access to big data gets pervasive, is that agents can obtain significant amounts of information about each other from historical data.

In more detail, the game can be represented by a tree (called a *game tree*) with *Cut* and *Choose* nodes:

- In a *Cut* node defined by “ i cuts in S ”, where $S = \{[x_1, y_1], \dots, [x_m, y_m]\}$, the strategy space of agent i is the set S of points where i can make a cut at this step.
- In a *Choose* node defined by “ i chooses from S ”, where $S = \{[x_1, y_1], \dots, [x_m, y_m]\}$, the strategy space is the set $\{1, \dots, m\}$, i.e., the indices of the pieces that can be chosen by the agent from the set S .

The strategy of an agent defines an action for *each* node of the game tree where it executes a *Cut* or a *Choose* operation. If an agent deviates, the game can follow a completely different branch of the tree, but the outcome will still be well-defined.

The strategies of the agents are in *Nash equilibrium* (NE) if no agent can improve its utility by unilaterally deviating from its current strategy, i.e., by cutting at a different set of points and/or by choosing different pieces. A *subgame perfect Nash equilibrium* (SPNE) is a stronger equilibrium notion, which means that the strategies are in NE in every subtree of the game tree. In other words, even if the game started from an arbitrary node of the game tree, the strategies would still be in NE. An ϵ -NE (resp., ϵ -SPNE) is a relaxed solution concept where an agent cannot gain more than ϵ by deviating (resp., by deviating in any subtree).

3 Existence of Equilibria

It is well-known that finite extensive-form games of perfect information can be solved using *backward induction*: starting from the leaves and progressing towards the root, at each node the relevant agent chooses an action that maximizes its utility, given the actions that were computed for the node’s children. The induced strategies form an SPNE. Unfortunately, although we consider finite GCC protocols, we also

need to deal with *Cut* nodes where the action space is infinite, hence naïve backward induction does not apply.

In fact, it turns out that not every GCC protocol admits an exact NE — not to mention SPNE. For example, consider Algorithm 1, and assume that the value density function of agent 1 is strictly positive. Assume there exists a NE where agent 1 cuts at x^*, y^*, z^* , respectively, and chooses the piece $[x^*, y^*]$. If $x^* > 0$, then the agent can improve its utility by making the first cut at $x' = 0$ and choosing the piece $[x', y^*]$, since $V_1([x', y^*]) > V_1([x^*, y^*])$. Thus, $x^* = 0$. Moreover, it cannot be the case that $y^* = 1$, since the agent only receives an allocation if $y^* < z^* \leq 1$. Thus, $y^* < 1$. Then, by making the second cut at any $y' \in (y^*, z^*)$, agent 1 can obtain the value $V_1([0, y']) > V_1([0, y^*])$. It follows that there is no exact NE where the agent chooses the first piece. Similarly, it can be shown that there is no exact NE where the agent chooses the second piece, $[y^*, z^*]$. This illustrates why backward induction does not apply: the maximum value at some *Cut* nodes may not be well defined.

3.1 Approximate SPNE

One possible way to circumvent the foregoing example is by saying that agent 1 should be happy to make the cut y very close to z . For instance, if the agent’s value is uniformly distributed over the case, cutting at $x = 0, y = 1 - \epsilon, z = 1$ would allow the agent to choose the piece $[x, y]$ with value $1 - \epsilon$; and this is true for any ϵ .

More generally, we have the following theorem.

Theorem 1. *For any n -agent GCC protocol \mathcal{P} with a bounded number of steps, any n valuation functions V_1, \dots, V_n , and any $\epsilon > 0$, the game induced by \mathcal{P} and V_1, \dots, V_n has an ϵ -SPNE.*

The proof of Theorem 1 is relegated to the full version.³ In a nutshell, the high-level idea of our proof relies on discretizing the cake — such that every cell in the resulting grid has a very small value for each agent — and computing the optimal outcome on the discretized cake using backward induction. At every cut step of the protocol, the grid is refined by adding a point between every two consecutive points of the grid from the previous cut step. This ensures that any ordering of the cut points that can be enforced by playing on the continuous cake can also be enforced on the discretized instance. Therefore, for the purpose of computing an approximate SPNE, it is sufficient to work with the discretization. We then show that the backward induction outcome from the discrete game gives an ϵ -SPNE on the continuous cake.

3.2 Informed Tie-Breaking

Another approach for circumventing the example given at the beginning of the section is to change the *tie-breaking* rule of Algorithm 1, by letting agent 1 choose even if $y = z$ (in which case agent 1 would cut in $x = 0, y = 1, z = 1$, and get the entire cake). Tie-breaking matters: even the Dubins-Spanier protocol fails to guarantee SPNE existence due to a curious tie-breaking issue (Brânzei and Miltersen 2013).

To accommodate more powerful tie-breaking rules, we slightly augment GCC protocols, by extending their ability to compare cuts in case of a tie. Specifically, we can assume without loss of generality that the *If-Else* statements of a GCC protocol are specified only with weak inequalities (as an equality can be specified with two inequalities and a strong inequality via an equality and weak inequality), which involve only pairs of cuts. We consider *informed GCC protocols*, which are capable of using *If-Else* statements of the form “if $[x < y$ or $(x = y$ and history of events $\in \mathcal{H})$ then”. That is, when cuts are made in the same location and cause a tie in an *If-Else*, the protocol can invoke the power to check the entire history of events that have occurred so far. We can recover the $x < y$ and $x \leq y$ comparisons of “uninformed” GCC protocols by setting \mathcal{H} to be empty or all possible histories, respectively. Importantly, the history can include where cuts were made exactly, and not simply where in relation to each other.

We say that an informed GCC protocol \mathcal{P}' is *equivalent up to tie-breaking* to a GCC protocol \mathcal{P} if they are identical, except that some inequalities in the *If-Else* statements of \mathcal{P} are replaced with informed inequalities in the corresponding *If-Else* statements of \mathcal{P}' . That is, the two protocols are possibly different only in cases where two cuts are made at the exact same point.

For example, in Algorithm 1, the statement “if $x < y < z$ then” can be specified as “if $x < y$ then if $y < z$ then”. We can obtain an informed GCC protocol that is equivalent up to tie-breaking by replacing this statement with “if $x < y$ then if $y \leq z$ then” (here we are not actually using augmented tie-breaking). In this case, the modified protocol may feel significantly different from the original — but this is an artifact of the extreme simplicity of Algorithm 1. Common cake cutting protocols are more complex, and changing the tie-breaking rule preserves the essence of the protocol.

We are now ready to present our second main result.

Theorem 2. *For any n -agent GCC protocol \mathcal{P} with a bounded number of steps and any n valuation functions V_1, \dots, V_n , there exists an informed GCC protocol \mathcal{P}' that is equivalent to \mathcal{P} up to tie-breaking, such that the game induced by \mathcal{P}' and V_1, \dots, V_n has an SPNE.*

Intuitively, we can view \mathcal{P}' as being “undecided” whenever two cuts are made at the same point, that is, $x = y$: it can adopt either the $x < y$ branch or the $x > y$ branch — there *exists* an appropriate decision. The theorem tells us that for any given valuation functions, we can set these tie-breaking points in a way that guarantees the existence of an SPNE. In this sense, the tie-breaking of the protocol is *informed* by the given valuation functions. Indeed, this interpretation is plausible as we are dealing with a game of perfect information.

The proof of Theorem 2 is somewhat long, and has been relegated to the full version. This proof is completely different from the proof of Theorem 1; in particular, it relies on real analysis instead of backward induction on a discretized space. The crux of the proof is the development of an auxiliary notion of *mediated games* (not to be confused with Monderer and Tennenholtz’s *mediated equilib-*

³Available from: <http://procaccia.info/research>.

rium (Monderer and Tennenholtz 2009)) that may be of independent interest. We show that mediated games always have an SPNE. The actions of the mediator in this SPNE are then reinterpreted as a tie-breaking rule under an informed GCC protocol. In the context of the proof it is worth noting that some papers prove the existence of SPNE in games with infinite action spaces (see, e.g., (Harris 1985; Hellwig and Leininger 1987)), but our game does not satisfy the assumptions required therein.

4 Fair Equilibria

The existence of equilibria (Theorems 1 and 2) gives us a tool for predicting the strategic outcomes of cake cutting protocols. In particular, classic protocols provide fairness guarantees when agents act honestly; but do they provide any fairness guarantees in equilibrium?

We first make a simple yet crucial observation. In a proportional protocol, every agent is guaranteed a value of at least $1/n$ regardless of what the others are doing. Therefore, in every NE (if any) of the protocol, the agent still receives a piece worth at least $1/n$; otherwise it can deviate to the strategy that guarantees it a utility of $1/n$ and do better. Similarly, an ϵ -NE must be ϵ -proportional, i.e., each agent must receive a piece worth at least $1/n - \epsilon$. Hence, classic protocols such as Dubins-Spanier, Even-Paz, and Selfridge-Conway guarantee (approximately) proportional outcomes in any (approximate) NE (and of course this observation carries over to the stronger notion of SPNE).

One may wonder, though, whether the analogous statement for envy-freeness holds; the answer is negative. We demonstrate this via the Selfridge-Conway protocol — a 3-agent envy-free protocol, which is given in its truthful, non-GCC form as Algorithm 3. To see why the protocol is envy free, note that the division of three pieces in steps 4, 5, and 6 is trivially envy free. For the division of the trimmings in step 9, agent i is not envious because it chooses first, and agent j is not envious because it was the one that cut the pieces (presumably, equally according to its value). In contrast, agent 1 may prefer the piece of trimmings that agent i received in step 9, but overall agent 1 cannot envy i , because at best i was able to “reconstruct” one of the three original pieces that was trimmed at step 2, which agent 1 values as much as the untrimmed piece it received in step 6.

We construct an example by specifying the valuation functions of the agents and their strategies, and arguing that the strategies are in SPNE. The example will have the property that the first two agents receive utilities of 1 (i.e. the maximum value). Therefore, we can safely assume their play is in equilibrium; this will allow us to define the strategies only on a small part of the game tree. In contrast, agent 3 will deviate from its truthful strategy to gain utility, but in doing so will become envious of agent 1.

In more detail, suppose after agent 2 trims the three pieces we have the following.

- Agent 1 values the first untrimmed piece at 1, and all other pieces and the trimmings at 0.
- Agent 2 values the second untrimmed piece at 1, and all other pieces and the trimmings at 0.

- 1: Agent 1 cuts the cake into three equal parts in the agent’s value.
- 2: Agent 2 trims the most valuable of the three pieces such that there is a tie with the two most valuable pieces.
- 3: Set aside the trimmings.
- 4: Agent 3 chooses one of the three pieces to keep.
- 5: Agent 2 chooses one of the remaining two pieces to keep — with the stipulation that if the trimmed piece is not taken by agent 3, agent 2 must take it.
- 6: Agent 1 takes the remaining piece.
- 7: Denote by $i \in \{2, 3\}$ the agent which received the trimmed piece, and $j = \{2, 3\} \setminus \{i\}$.
- 8: Agent j now cuts the trimmings into three equal parts in the agent’s value.
- 9: Agents i , 1, and j choose one of the three pieces to keep in that order.

Algorithm 3: Selfridge-Conway: an envy-free protocol for three agents.

- Agent 3 values the untrimmed pieces at $1/7$ and 0, the trimmed piece at $1/14$, and the trimmings at $11/14$.

Now further suppose that if agent 3 is to cut the trimmings (i.e. take on the role of j in the protocol), then the first two agents always take the pieces most valuable to agent 3. Thus, if agent 3 does not take the trimmed piece it will achieve a utility of at most $1/7 + (11/14)(1/3) = 119/294$ by taking the first untrimmed piece, and then cutting the trimmings into three equal parts. On the other hand, if agent 3 takes the trimmed piece of worth $1/14$, agent 2 cuts the trimmings into three parts such that one of the pieces is worth 0 to agent 3, and the other two are equivalent in value (i.e. they have values $(11/14)(1/2) = 11/28$). Agents 1 and 3 take these two pieces. Thus, in this scenario, agent 3 receives a utility of $1/14 + 11/28 = 13/28$ which is strictly better than the utility of $119/294$. Agent 3 will therefore choose to take the trimmed piece. However, in this outcome agent 1, from the point of view of agent 3, receives a piece worth $1/7 + 11/28 = 15/28$ and so agent 3 will indeed be envious.

The foregoing example shows that envy-freeness is not guaranteed when agents strategize, and so it is difficult to produce envy-free allocations when agents play to maximize their utility. A natural question to ask, therefore, is whether there are any GCC protocols such that all SPNE are envy-free, and existence of SPNE is guaranteed. This remains an open question, but we do give an affirmative answer for the weaker solution concept of NE in the following theorem, whose proof appears in the full version of the paper.

Theorem 3. *There exists a GCC protocol \mathcal{P} such that on every cake cutting instance with strictly positive valuation functions V_1, \dots, V_n , an allocation X is the outcome of a NE of the game induced by \mathcal{P} and V_1, \dots, V_n if and only if X is an envy-free contiguous allocation that contains the entire cake.*

Crucially, an envy-free contiguous allocation is guaranteed to exist (Stromquist 1980), hence the set of NE of protocol \mathcal{P} is nonempty.

Theorem 3 is a positive result *à la* implementation theory

(see, e.g., (Maskin 1999)), which aims to construct games where the NE outcomes coincide with a given specification of acceptable outcomes for each constellation of agents' preferences (known as a *social choice correspondence*). Our construction guarantees that the NE outcomes coincide with (contiguous) envy-free allocations, i.e. in this case the envy-freeness criterion specifies which outcomes are acceptable.

That said, the protocol \mathcal{P} constructed in the proof of Theorem 3 is impractical: its Nash equilibria are unlikely to arise in practice. This further motivates efforts to find an analogous result for SPNE. If such a result is indeed feasible, a broader, challenging open question would be to characterize GCC protocols that give rise to envy-free SPNE, or at least provide a sufficient condition (on the protocol) for the existence of such equilibria.

Acknowledgments

Brânzei was supported by the Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation, and the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Brânzei was also partially supported by a fellowship from the Simons Institute for the Theory of Computing. Caragiannis was partially supported by Caratheodory research grant E.114 from the University of Patras. Kurokawa and Procaccia were partially supported by the National Science Foundation under grants CCF-1215883, CCF-1525932, and IIS-1350598, and by a Sloan Research Fellowship.

References

- Aumann, Y.; Dombb, Y.; and Hassidim, A. 2013. Computing socially-efficient cake divisions. In *Proc. of 12th AAMAS*, 343–350.
- Balkanski, E.; Brânzei, S.; Kurokawa, D.; and Procaccia, A. D. 2014. Simultaneous cake cutting. In *Proc. of 28th AAAI*, 566–572.
- Bei, X.; Chen, N.; Hua, X.; Tao, B.; and Yang, E. 2012. Optimal proportional cake cutting with connected pieces. In *Proc. of 26th AAAI*, 1263–1269.
- Brams, S. J., and Taylor, A. D. 1995. An envy-free cake division protocol. *Am Math Monthly* 102(1):9–18.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge Univ Press.
- Brams, S. J.; Feldman, M.; Morgenstern, J.; Lai, J. K.; and Procaccia, A. D. 2012. On maxsum fair cake divisions. In *Proc. of 26th AAAI*, 1285–1291.
- Brams, S. J.; Taylor, A. D.; and Zwicker, W. S. 1997. A moving-knife solution to the four-person envy-free cake-division problem. *Proceedings of AMS* 125(2):547–554.
- Brânzei, S., and Miltersen, P. B. 2013. Equilibrium analysis in cake cutting. In *Proc. of 12th AAMAS*, 327–334.
- Brânzei, S., and Miltersen, P. B. 2015. A dictatorship theorem for cake cutting. In *Proc. of 24th IJCAI*, 482–488.
- Brânzei, S.; Procaccia, A. D.; and Zhang, J. 2013. Externalities in cake cutting. In *Proc. of 23rd IJCAI*, 55–61.
- Caragiannis, I.; Lai, J. K.; and Procaccia, A. D. 2011. Towards more expressive cake cutting. In *Proc. of 22nd IJCAI*, 127–132.
- Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77:284–297. Preliminary version in AAAI'10.
- Deng, X.; Qi, Q.; and Saberi, A. 2012. Algorithmic solutions for envy-free cake cutting. *Operations Research* 60(6):1461–1476.
- Dubins, L. E., and Spanier, E. H. 1961. How to cut a cake fairly. *American Mathematical Monthly* 68(1):1–17.
- Edelman, B.; Ostrovsky, M.; and Schwarz, M. 2007. Internet advertising and the Generalized Second-Price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1):242–259.
- Edmonds, J., and Pruhs, K. 2006a. Balanced allocations of cake. In *Proc. of 47th FOCS*, 623–634.
- Edmonds, J., and Pruhs, K. 2006b. Cake cutting really is not a piece of cake. In *Proc. of 17th SODA*, 271–278.
- Even, S., and Paz, A. 1984. A note on cake-cutting. *Discrete Applied Mathematics* 7:285–296.
- Goldman, J., and Procaccia, A. D. 2014. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges* 13(2):41–46.
- Harris, C. 1985. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica* 53(3):613–628.
- Hellwig, H., and Leininger, W. 1987. On the existence of subgame-perfect equilibrium in infinite-action games of perfect information. *Journal of Economic Theory* 45:55–75.
- Kurokawa, D.; Lai, J. K.; and Procaccia, A. D. 2013. How to cut a cake before the party ends. In *Proc. of 27th AAAI*, 555–561.
- Kurokawa, D.; Procaccia, A. D.; and Shah, N. 2015. Leximin allocations in the real world. In *Proc. of 16th EC*, 345–362.
- Maskin, E. 1999. Nash equilibrium and welfare optimality. *Review of Economic Studies* 66:23–38. First circulated in 1977.
- Maya, A., and Nisan, N. 2012. Incentive compatible two player cake cutting. In *Proc. of 8th WINE*, 170–183.
- Monderer, D., and Tennenholtz, M. 2009. Strong mediated equilibrium. *Artificial Intelligence* 173(1):180–195.
- Mossel, E., and Tamuz, O. 2010. Truthful fair division. In *Proc. of 3rd SAGT*, 288–299.
- Nicolò, A., and Yu, Y. 2008. Strategic divide and choose. *Games and Economic Behavior* 64(1):268–289.
- Procaccia, A. D. 2009. Thou shalt covet thy neighbor's cake. In *Proc. of 21st IJCAI*, 239–244.
- Procaccia, A. D. 2013. Cake cutting: Not just child's play. *Communications of the ACM* 56(7):78–87.
- Robertson, J. M., and Webb, W. A. 1998. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters.
- Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2015. Envy-free cake-cutting in two dimensions. In *Proc. of 29th AAAI*, 1021–1028.
- Shoham, Y., and Leyton-Brown, K. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16:101–104.
- Stromquist, W. 1980. How to cut a cake fairly. *American Mathematical Monthly* 87(8):640–644.
- Woeginger, G. J., and Sgall, J. 2007. On the complexity of cake cutting. *Discrete Optimization* 4:213–220.
- Zhou, L. 1991. Inefficiency of strategy-proof allocation mechanisms in pure exchange economies. *Social Choice and Welfare* 8:247–254.