

# An Algorithmic Framework for Strategic Fair Division

Simina Brânzei  
UC Berkeley

Ioannis Caragiannis  
U. Patras

David Kurokawa  
CMU

Ariel D. Procaccia  
CMU

## 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 between multiple agents (see, e.g., 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; Cohler et al. 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; Li, Zhang, and Zhang 2015), which is tightly 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 the case of two agents. The first agent cuts the cake into 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). However, taking the 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 that 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). We then explore two ways of circumventing this issue, which give rise to our two main results.

1. We prove that every GCC protocol has at least one  $\epsilon$ -NE for every  $\epsilon > 0$ , in which agents cannot gain more than  $\epsilon$  by deviating, and  $\epsilon$  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  $\epsilon$ -equilibrium must be an  $\epsilon$ -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 complexity model that specifies how a cake cutting protocol may interact with the agents. Their 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  $\epsilon$ -equilibria that are  $\epsilon$ -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 without loss of generality.

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 the 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  $x, y \in [0, 1]$  such that  $x < y$ .

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 that is 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.

## 2.1 Generalized Cut and Choose Protocols

The standard communication model in cake cutting was proposed by Robertson and Webb (1998). The model restricts the interaction between the protocol and the 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  $\alpha$  such that  $V_i([x, y]) = \alpha$ .

However, the communication model does not give much information about the actual implementation of the protocol and what allocations it produces. For example, the protocol could allocate pieces depending on whether a particular cut was made at a specific point (see Algorithm 2).

For this reason, we define a generic class of protocols that are implementable with natural operations, which capture all bounded<sup>1</sup> and discrete cake cutting algorithms, such as cut and choose, Dubins-Spanier, Even-Paz, Successive-Pairs, and Selfridge-Conway (see, e.g., (Procaccia 2013)). 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

<sup>1</sup>In 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 = 1/3) then
  agent 1 Chooses from {[0, x], [x, 1]}
end if
```

**Algorithm 2:** A non-GCC protocol.

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:  $<$ ,  $\leq$ ,  $=$ ,  $\geq$ ,  $>$ .

More formally, a *generalized cut and choose (GCC)* protocol is implemented exclusively with the following types of 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 the 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.
- **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 discrete version of Dubins-Spanier belongs to the class of GCC protocols — but first we must describe the original 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.

## 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, we note that 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 General-

ized 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 becomes pervasive, is that agents can obtain a significant amount 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 agent  $i$  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  $\epsilon$ -NE (resp.,  $\epsilon$ -SPNE) is a relaxed solution concept where an agent cannot gain more than  $\epsilon$  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  $\epsilon$ .

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  $\epsilon$ -SPNE.*

The proof of Theorem 1 is relegated to Appendix A. 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  $\epsilon$ -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 Appendix B. 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 equilibrium* (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

- 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.

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  $\epsilon$ -NE must be  $\epsilon$ -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.

- Agent 3 values the untrimmed pieces at  $1/7$  and 0, respectively, 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 therefore 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 Appendix C.

**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, that is, 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.

## 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. *The American Mathematical Monthly* 102(1):9–18.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University 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 the American Mathematical Society* 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.
- Cohler, Y. J.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2011. Optimal envy-free cake cutting. In *Proc. of 25th AAAI*, 626–631.
- 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.
- Li, M.; Zhang, J.; and Zhang, Q. 2015. Truthful cake cutting mechanisms with externalities: Do not make them care for others too much! In *Proc. of 24th IJCAI*, 589–595.
- 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.
- Shoham, Y., and Leyton-Brown, K. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- 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.

## A Proof of Theorem 1

Let  $\epsilon > 0$ , and let  $f(n)$  be an upper bound on the number of operations (i.e., on the height of the game tree) of the protocol. Define a grid,  $\mathcal{G}_1$ , such that every cell on the grid is worth at most  $\frac{\epsilon}{2f(n)^2}$  to each agent. For every  $n$ , let  $K$  denote the maximum number of cut operations, where  $0 \leq K \leq f(n)$ . For each  $i \in \{1, \dots, K\}$ , we define the grid  $\mathcal{G}_i$  so that the following properties are satisfied:

- The grids are nested, i.e.,  $\{0, 1\} \subset \mathcal{G}_1 \subset \mathcal{G}_2 \subset \dots \subset \mathcal{G}_K$ .
- There exists a unique point  $z \in \mathcal{G}_{i+1}$  between any two consecutive points  $x, y \in \mathcal{G}_i$ , such that  $x < z < y$  and  $z \notin \mathcal{G}_i$ , for every  $i \in \{1, \dots, K-1\}$ .
- Each cell on  $\mathcal{G}_i$  is worth at most  $\frac{\epsilon}{2f(n)^2}$  to any agent, for all  $i \in \{1, \dots, K\}$ .

Having defined the grids, we compute the backward induction outcome on the discretized cake, where the  $i$ -th *Cut* operation can only be made on the grid  $\mathcal{G}_i$ . We will show that this outcome is an  $\epsilon$ -SPNE, even though agents could deviate by cutting anywhere on the cake. On the continuous cake, the agents play a perturbed version of the idealized game from the grid  $\mathcal{G}$ , but maintain a mapping between the perturbed game and the idealized version throughout the execution of the protocol, such that each cut point from the continuous cake is mapped to a grid point that approximates it within a very small (additive) error. Thus when determining the next action, the agents use the idealized grid as a reference. The order of the cuts is the same in the ideal and perturbed game, however the values of the pieces may differ by at most  $\epsilon/f(n)$ .

We start with the following useful lemma. (For ease of exposition, in the following we refer to  $[x, y]$  as the segment between points  $x$  and  $y$ , regardless of whether  $x < y$  or  $y \leq x$ .)

**Lemma 1.** *Given a sequence of cut points  $x_1, \dots, x_k$  and nested grids  $\mathcal{G}_1 \subset \dots \subset \mathcal{G}_k$  with cells worth at most  $\frac{\epsilon}{4f(n)^2}$  to each agent, there exists a map  $\mathcal{M} : \{x_1, \dots, x_k\} \rightarrow \mathcal{G}_k$  such that:*

- (1) For each  $i \in \{1, \dots, k\}$ ,  $\mathcal{M}(x_i) \in \mathcal{G}_i$ .
- (2) The map  $\mathcal{M}$  is order-preserving. Formally, for all  $i, j \in \{1, \dots, k\}$ ,  $x_i < x_j \iff \mathcal{M}(x_i) < \mathcal{M}(x_j)$  and  $x_i = x_j \iff \mathcal{M}(x_i) = \mathcal{M}(x_j)$ .
- (3) The piece  $[x_i, \mathcal{M}(x_i)]$  is “small”, that is:  $V_l([x_i, \mathcal{M}(x_i)]) \leq \frac{k\epsilon}{2f(n)^2}$ , for each agent  $l \in N$ .

*Proof.* We prove the statement by induction on the number of cut points  $k$ .

*Base case:* We consider a few cases. If  $x_1 \in \mathcal{G}_1$ , then define  $\mathcal{M}(x_1) := x_1$ . Otherwise, let  $R(x_1) \in \mathcal{G}_1$  be the leftmost point on the grid  $\mathcal{G}_1$  to the right of  $x_1$ . If  $R(x_1) \neq 1$ , define  $\mathcal{M}(x_1) := R(x_1)$ ; else, let  $L(x_1)$  denote the rightmost point on  $\mathcal{G}_1$  strictly to the left of 1 and define  $\mathcal{M}(x_1) := L(x_1)$ . To verify the properties of the lemma, note that:

- (1)  $\mathcal{M}(x_1) \in \mathcal{G}_1$ .
- (2) The map  $\mathcal{M}$  is order-preserving since there is only one point.

- (3)  $V_l([x_1, \mathcal{M}(x_1)]) \leq \frac{\epsilon}{2f(n)^2}$  for each agent  $l \in N$  since the grid  $\mathcal{G}_1$  has (by construction) the property that each cell is worth at most  $\frac{\epsilon}{2f(n)^2}$  to each agent, and the interval  $[x_1, \mathcal{M}(x_1)]$  is contained in a cell.

*Induction hypothesis:* Assume that a map  $\mathcal{M}$  with the required properties exists for any sequence of  $k-1$  cut points.

*Induction step:* Consider any sequence of  $k$  cut points  $x_1, \dots, x_k$ . By the induction hypothesis, we can map each cut point  $x_i$  to a grid representative  $\mathcal{M}(x_i) \in \mathcal{G}_i$ , for all  $i \in \{1, \dots, k-1\}$ , in a way that preserves properties 1–3. We claim that the map  $\mathcal{M}$  on the points  $x_1, \dots, x_{k-1}$  can be extended to the  $k$ -th point,  $x_k$ , such that the entire sequence  $\mathcal{M}(x_1), \dots, \mathcal{M}(x_k)$  satisfies the requirements of the lemma. We consider four exhaustive cases.

- (a) There exists  $i \in \{1, \dots, k-1\}$  such that  $x_k = x_i$ . Then define  $\mathcal{M}(x_k) := \mathcal{M}(x_i)$ .
- (b) There exists  $i \in \{1, \dots, k-1\}$  such that  $x_i < x_k$ , but  $\mathcal{M}(x_i) \geq x_k$ . Let  $x_j$  be the rightmost cut such that  $x_j < x_k$ ; because  $\mathcal{M}$  is order-preserving, it holds that  $\mathcal{M}(x_j) \geq x_k$ . Let  $R(\mathcal{M}(x_j))$  be the leftmost point on  $\mathcal{G}_k$  strictly to the right of  $\mathcal{M}(x_j)$ , and set  $\mathcal{M}(x_k) := R(\mathcal{M}(x_j))$ .

Now let us check the conditions. Condition (1) holds by definition. Condition (2) holds because  $\mathcal{M}(x_k) > \mathcal{M}(x_j)$ , and for every  $t$  such that  $x_t > x_k$ ,  $\mathcal{M}(x_t) > \mathcal{M}(x_j)$  and  $\mathcal{M}(x_t) \in \mathcal{G}_{k-1}$ , whereas  $\mathcal{M}(x_k)$  uses a “new” point of  $\mathcal{G}_k \setminus \mathcal{G}_{k-1}$  that is closer to  $\mathcal{M}(x_j)$ . For condition (3), we have that for every  $l \in N$ ,

$$\begin{aligned} V_l([x_k, \mathcal{M}(x_k)]) &\leq V_l([x_j, \mathcal{M}(x_k)]) \\ &= V_l([x_j, \mathcal{M}(x_j)]) + V_l([\mathcal{M}(x_j), \mathcal{M}(x_k)]) \\ &\leq \frac{(k-1)\epsilon}{2f(n)^2} + \frac{\epsilon}{2f(n)^2} \leq \frac{k\epsilon}{2f(n)^2}, \end{aligned}$$

where the third transition follows from the induction assumption.

- (c) There exists  $i \in \{1, \dots, k-1\}$  such that  $x_i > x_k$ , but  $\mathcal{M}(x_i) \leq x_k$ . This case is symmetric to case (b).
- (d) For every  $x_i$  such that  $x_i < x_k$ ,  $\mathcal{M}(x_i) \leq x_k$ , and for every  $x_j$  such that  $x_j > x_k$ ,  $\mathcal{M}(x_j) \geq x_k$ . Let  $x_i$  and  $x_j$  be the rightmost and leftmost such cuts, respectively; without loss of generality they exist, otherwise our task is even easier.

Let  $R(x_k)$  be the leftmost point in  $\mathcal{G}_k$  such that  $R(x_k) \geq x_k$ , and let  $L(x_k)$  be the rightmost point in  $\mathcal{G}_k$  such that  $L(x_k) \leq x_k$ . Assume first that  $\mathcal{M}(x_j) > R(x_k)$ ; then set  $\mathcal{M}(x_k) := R(x_k)$ . This choice obviously satisfies the three conditions, similarly to the base of the induction.

Otherwise,  $R(x_k) = \mathcal{M}(x_j)$  (notice that it cannot be the case that  $R(x_k) > \mathcal{M}(x_k)$ ); then set  $\mathcal{M}(x_k) := L(x_k)$ . Let us check that this choice is order-preserving (as the other two conditions are trivially satisfied). Note that  $\mathcal{M}(x_j) \in \mathcal{G}_{k-1}$ , so  $R(x_k) \in \mathcal{G}_{k-1}$ . Therefore, it must hold that  $L(x_k) \in \mathcal{G}_k \setminus \mathcal{G}_{k-1}$  — it is the new point that we have added between  $R(x_k)$ , and the rightmost

point the left of it on  $\mathcal{G}_{k-1}$ . Since it is also the case that  $\mathcal{M}(x_i) \in \mathcal{G}_{k-1}$ , we have that  $\mathcal{M}(x_i) < \mathcal{M}(x_k) < \mathcal{M}(x_j)$ .

By induction, we can compute a mapping with the required properties for  $k$  points. This completes the proof of the lemma.  $\square$

Now we can define the equilibrium strategies. Let  $x_1, \dots, x_k$  be the history of cuts made at some point during the execution of the protocol. By Lemma 1, there exists an order-preserving map  $\mathcal{M}$  such that each point  $x_i$  has a representative point  $\mathcal{M}(x_i) \in \mathcal{G}_i$  and the piece  $[x_i, \mathcal{M}(x_i)]$  is “small”, i.e.

$$V_l([x_i, \mathcal{M}(x_i)]) \leq \frac{k\epsilon}{2f(n)^2} \leq \frac{\epsilon}{2f(n)}$$

for each agent  $l \in N$  — using  $k \leq f(n)$ .

Consider any history of cuts  $(x_1, \dots, x_k)$ . Let  $i$  be the agent that moves next. Agent  $i$  computes the mapping  $(\mathcal{M}(x_1), \dots, \mathcal{M}(x_k))$ . If the next operation is:

- *Choose*: agent  $i$  chooses the available piece (identified by the symbolic names of the cut points it contains and their order) which is optimal in the idealized game, given the current state and the existing set of ordered ideal cuts,  $\mathcal{M}(x_1), \dots, \mathcal{M}(x_k)$ . Ties are broken according to a fixed deterministic scheme which is known to all the agents.
- *Cut*: agent  $i$  computes the optimal cut on  $\mathcal{G}_{k+1}$ , say at  $x_{k+1}^*$ . Then  $i$  maps  $x_{k+1}^*$  back to a point  $x_{k+1}$  on the continuous game, such that  $\mathcal{M}(x_{k+1}) = x_{k+1}^*$ . That is, the cut  $x_{k+1}$  (made in step  $k+1$ ) is always mapped by the other agents to  $x_{k+1}^* \in \mathcal{G}_{k+1}$ . Agent  $i$  cuts at  $x_{k+1}$ .

We claim that these strategies give an  $\epsilon$ -SPNE. The proof follows from the following lemma, which we show by induction on  $t$  (the maximum number of remaining steps of the protocol):

**Lemma 2.** *Given a point in the execution of the protocol from which there are at most  $t$  operations left until termination, it is  $\frac{t\epsilon}{f(n)}$ -optimal to play on the grid.*

*Proof.* Consider any history of play, where the cuts were made at  $x_1, \dots, x_k$ . Without loss of generality, assume it is agent  $i$ 's turn to move.

*Base case:*  $t = 1$ . The protocol has at most one remaining step. If it is a cut operation, then no agent receives any utility in the remainder of the game regardless of where the cut is made. Thus cutting on the grid ( $\mathcal{G}_k$ ) is optimal. If it is a choose operation, then let  $Z = \{Z_1, \dots, Z_s\}$  be the set of pieces that  $i$  can choose from. Agent  $i$ 's strategy is to map each piece  $Z_j$  to its equivalent  $\mathcal{M}(Z_j)$  on the grid  $\mathcal{G}_k$ , and choose the piece that is optimal on  $\mathcal{G}_k$ . Recall that  $V_q([x_j, \mathcal{M}(x_j)]) \leq \frac{\epsilon}{2f(n)}$  for each agent  $q \in N$ . Thus if a piece is optimal on the grid, it is  $\frac{\epsilon}{f(n)}$ -optimal in the continuous game (adding up the difference on both sides). It follows that  $i$  cannot gain more than  $\frac{\epsilon}{f(n)}$  in the last step by deviating from the optimal piece on  $\mathcal{G}_k$ .

*Induction hypothesis:* Assume that playing on the grid is  $\frac{(t-1)\epsilon}{f(n)}$ -optimal whenever there are at most  $t-1$  operations

left on every possible execution path of the protocol, and there exists one path that has exactly  $t-1$  steps.

*Induction step:* If the current operation is *Choose*, then by the induction hypothesis, playing on the grid in the remainder of the protocol is  $\frac{(t-1)\epsilon}{f(n)}$ -optimal for all the agents, regardless of  $i$ 's move in the current step. Moreover, agent  $i$  cannot gain by more than  $\frac{\epsilon}{f(n)}$  by choosing a different piece in the current step, compared to piece which is optimal on  $\mathcal{G}_k$ , since  $V_i([x_l, \mathcal{M}(x_l)]) \leq \frac{\epsilon}{2f(n)}$  for all  $l \in \{1, \dots, k\}$ .

If the current operation is *Cut*, then the following hold:

1. By construction of the grid  $\mathcal{G}_{k+1}$ , agent  $i$  can induce any given branch of the protocol using a cut in the continuous game if and only if the same branch can be induced using a cut on the grid  $\mathcal{G}_{k+1}$ .
2. Given that the other agents will play on the grid for the remainder of the protocol, agent  $i$  can change the size of at most one piece that it receives down the road by at most  $\frac{\epsilon}{f(n)}$  by deviating (compared to the grid outcome), since  $V_j([x_l, \mathcal{M}(x_l)]) \leq \frac{\epsilon}{2f(n)}$  for all  $l \in \{1, \dots, k+1\}$  and for all  $j \in N$ .

Thus by deviating in the current step, agent  $i$  cannot gain more than  $\frac{t\epsilon}{f(n)}$ .  $\square$

Since  $t \leq f(n)$ , the overall loss of any agent is bounded by  $\epsilon$  by Lemma 2. We conclude that playing on the grid is  $\epsilon$ -optimal for all the agents, which completes the proof of the theorem.  $\square$

## B Proof of Theorem 2

Before we begin, we take this moment to formally introduce the auxiliary concept of a *mediated game* in an abstract sense. We will largely distance ourselves from the specificity of GCC games here and work in a more general model. We do this for two purposes. First, it allows for a cleaner view of the techniques; and second, we believe such general games may be of independent interest. We begin with a few definitions.

**Definition** In an extensive-form game, an *action tuple* is a tuple of actions that describe an outcome of the game. For example, the action tuple  $(a_1, \dots, a_r)$  states that  $a_1$  was the first action to be played,  $a_2$  the second, and  $a_r$  the last.

**Definition** Given an action tuple, the  $k^{\text{th}}$  action is said to be SPNE if the subtree of the game tree rooted where the first  $k-1$  actions are played in accordance to the action tuple is induced by some SPNE strategy profile. Furthermore, call such an action tuple *k-SPNE*.

Note that if the  $k^{\text{th}}$  action is SPNE, so too are all actions succeeding it in the action tuple. To clarify Definition B, note that strategies of an extensive-form game are defined on every possible node of the game tree, so a  $k$ -SPNE action tuple can be equivalently defined as being an SPNE of the subgame rooted at the  $k^{\text{th}}$  action.

With these definitions in hand, we can now describe the games of interest.

**Definition** We call an extensive-form game a *mediated game* if the following conditions hold:

1. The set of agents consists of a single special agent, referred to as the *mediator*, and some finite number  $n$  of other regular agents. Intuitively, the mediator is an agent who is overseeing the proper execution of a protocol.
2. The height  $h$  of the game tree is bounded.
3. Every agent's utility is bounded.
4. Starting from the first or second action, the mediator plays every second action (and only these actions).
5. Every action played by the mediator shares the same action space:

$$\{0, \dots, n\} \times \left( [0, 1]^2 \cup 2^{\{1, \dots, h\}} \right).$$

This represents the agent who plays next (0 represents ending the game), and the interval which represents their action space or the allowed pieces they may choose from.

6. The mediator's utility is binary (i.e. it is in  $\{0, 1\}$ ) and is described entirely by the notion of *allowed edges*. This is a set of edges in the game tree such that the mediator's utility is 1 iff it plays edges only in this set. Importantly, this set has the property that for every allowed edge, each grandchild subtree (i.e. subtree that represents the next mediator's action) must have at least one allowed edge from its root. Intuitively, these edges are the ones that follow the protocol the mediator is implementing.
7. A regular agent's utility is continuous<sup>2</sup> in the action tuple.
8. *Allowed-edges-closedness*: given a convergent sequence of action tuples where the mediator plays only allowed edges, the mediator must play only allowed edges in the limit action tuple as well.

Note that appending meaningless actions (that affect no agent's utility) to a branch of the game tree will not affect the game in any impactful way. Thus, for the sake of convenience, we will assume for any game we consider all leaves of the game occur at the same depth (often denoted by  $r$ ).

We now give a series of definitions and lemmas that culminate in the main tool used in the proof of Theorem 2: all mediated games have an SPNE.

**Definition A** A sequence of action tuples  $(a_1^i, \dots, a_r^i) \mid_i$  is said to be *consistent* if for every  $j$  the agent who plays action  $a_j^i$  is constant throughout the sequence and, moreover, its action spaces are always subsets of  $[0, 1]$  or always the same subset of  $\{1, \dots, h\}$  throughout the sequence.

**Lemma 3.** *Let  $(a_1^i, \dots, a_r^i) \mid_i$  be a sequence of action tuples in a mediated game. Then there is a convergent subsequence.*

*Proof.* Due to the finite number of agents and bounded height of the game, we can find an infinite consistent subsequence  $\vec{b}^i \mid_i = (b_1^i, \dots, b_r^i) \mid_i$ . It suffices to show this subsequence has a convergent subsequence of its own. It is fairly

<sup>2</sup>The notions of convergence, compactness and continuity, which we will utilize often, necessarily assumes our action spaces are defined as metric spaces. Applicable metrics for the action spaces are not difficult to find, but are cumbersome to describe fully. We therefore will not belabour this point much further.

clear that we can find a convergent subsequence via compactness arguments, but there is a slight caveat: we must show that the limit action tuple is legal. That is, if the limit action tuple is  $(a_1, \dots, a_r)$  we must show that for every  $i < r$  such that the mediator plays action  $i$ , action  $i+1$  is played by the agent prescribed by  $a_i$ , and within the bounds prescribed by it. We will prove this by induction.

*Base hypothesis:* First 0 actions have a convergent subsequence — this is vacuously true.

*Induction hypothesis:* Assume there exists a subsequence such that the first  $k$  actions converge legally.

*Induction step:* We wish to show that there exists a subsequence such that the first  $k+1$  actions converge. By the inductive assumption, there exists a subsequence  $\vec{c}^i \mid_i$  such that the first  $k$  actions converge. Now suppose  $p$  plays the  $k+1^{\text{th}}$  action. If  $p$  is the mediator, then the action space is indifferent to actions played previously and is compact. Thus, the  $\vec{c}^i \mid_i$  must have a convergent subsequence such that the  $k+1^{\text{th}}$  element of the action tuple converges and so we are done.

Alternatively, if  $p$  is a regular agent, the action space is not necessarily indifferent to previous actions. If the action spaces are always the same subset of  $\{1, \dots, h\}$ , then we are clearly done. We therefore need only consider the case where the action spaces will be contained in  $[0, 1]$ . Due to the compactness of this interval, there will be a convergent subsequence of  $\vec{c}^i \mid_i$  such that the  $k+1^{\text{th}}$  action converges to some  $\gamma \in [0, 1]$ . Call this subsequence  $\vec{d}^i \mid_i$ .

We argue that  $\gamma$  is in the limit action space of the  $k+1^{\text{th}}$  action. For purposes of contradiction, assume this is false. Let  $\delta$  be the length from  $\gamma$  to the closest point in the limit action space (i.e. the action space in the limit given by the  $k^{\text{th}}$  action played by the mediator). Then there exists some  $M$  such that after the  $M^{\text{th}}$  element in  $\vec{d}^i \mid_i$ , the closest point in the  $k+1^{\text{th}}$  action space to  $\gamma$  is at least  $\delta/2$  away. Moreover, there exists some  $N$  such that after the  $N^{\text{th}}$  element in  $\vec{d}^i \mid_i$  the  $k+1^{\text{th}}$  action is no further than  $\delta/3$  to  $\gamma$ . Elements of  $\vec{d}^i \mid_i$  after element  $\max(M, N)$  then simultaneously must have the  $k+1^{\text{th}}$  action space be at least  $\delta/2$  away from  $\gamma$  and have a point at most  $\delta/3$  away from  $\gamma$ . This is a clear contradiction.  $\square$

**Lemma 4.** *For every  $k$ , if we have a convergent sequence of action tuples where the  $k^{\text{th}}$  action from the end is SPNE, then the  $k^{\text{th}}$  action from the end for the limit action tuple is also SPNE. That is, for every  $k$ , convergent sequences of  $(r-k+1)$ -SPNE action tuples are  $(r-k+1)$ -SPNE.*

*Proof.* We prove the result by induction on  $k$ .

*Base Case ( $k=0$ ):* This is vacuously true.

*Induction hypothesis ( $k=m$ ):* Assume convergent sequences of  $(r-m+1)$ -SPNE action tuples are  $(r-m+1)$ -SPNE.

*Induction step ( $k=m+1$ ):* Let  $\vec{a}^i \mid_i = (a_1^i, \dots, a_r^i) \mid_i$  be a convergent sequence of  $(r-m)$ -SPNE action tuples with the limit action tuple  $(a_1, \dots, a_r)$ . We wish to show that if all actions before the last  $m+1$  actions play their limit actions, then the remaining  $m+1$  actions are SPNE — note that by

Lemma 3 we know that the limit sequence is a valid action tuple.

Let  $p$  be the agent that commits the  $m + 1^{th}$  action from the end. If  $p$  is the mediator, then by the definition of mediated games the desired statement is true (specifically via the allowed-edges-closedness condition). Now suppose instead that  $p$  is not the mediator, and simply a regular agent. We show if the  $m + 1^{th}$  action from the end took on some other valid value  $\alpha \neq a_{r-m}$ , there exists SPNE strategies for the remaining  $m$  actions such that  $p$  achieves a utility no higher than had it stuck with the limit action of  $a_{r-m}$ .

So suppose the  $m + 1^{th}$  action from the end in the  $i^{th}$  element of the sequence is  $\alpha^i$  such that  $\lim_{i \rightarrow \infty} \alpha^i = \alpha$ . Since  $\vec{a}^i |_i$  is a sequence of  $(r - m)$ -SPNE action tuples, we can construct the sequence:

$$\vec{b}^i |_i = (a_1^i, \dots, a_{r-m-1}^i, \alpha^i, \tilde{a}_1^i, \dots, \tilde{a}_m^i) |_i$$

where the  $\tilde{a}_j^i$  are SPNE actions such that  $p$  achieves at most the utility achieved by instead playing  $a_{r-m}^i$ . Via Lemma 3,  $\vec{b}^i |_i$  must have a convergent subsequence — call  $\vec{c}^i |_i$  and indexed by increasing function  $\sigma$ . That is,  $\vec{c}^i = \vec{b}^{\sigma(i)}$ .  $\vec{c}^i |_i$  is then a convergent sequence of  $(r - m + 1)$ -SPNE action tuples and thus, by the inductive assumption, its limit action tuple is also an  $(r - m + 1)$ -SPNE.

Now consider the limit action tuple  $(a_1, \dots, a_r)$  (of  $\vec{a}^i |_i$ ) and the limit action tuple of  $\vec{c}^i |_i$  denoted by  $(c_1, \dots, c_r)$ . Note that:

1.  $\forall i < r - m: a_i = c_i$ .
2. By the continuity requirement of mediated games (where  $V_p$  is the utility function of  $p$ ):

$$\begin{aligned} & V_p(a_1, \dots, a_r) \\ &= \lim_{i \rightarrow \infty} V_p(a_1^i, \dots, a_r^i) \\ &= \lim_{i \rightarrow \infty} V_p(a_1^{\sigma(i)}, \dots, a_r^{\sigma(i)}) \\ &\geq \lim_{i \rightarrow \infty} V_p(a_1^{\sigma(i)}, \dots, a_{r-m-1}^{\sigma(i)}, \alpha^{\sigma(i)}, \tilde{a}_{r-m+1}^{\sigma(i)}, \dots, \tilde{a}_r^{\sigma(i)}) \\ &= \lim_{i \rightarrow \infty} V_p(c_1^i, \dots, c_r^i) \\ &= V_p(c_1, \dots, c_r). \end{aligned}$$

These two points imply that we can set SPNE strategies for the remaining  $m$  actions such that the utility of  $p$  playing  $\alpha$  is less than or equal to if it plays  $a_{r-m}$  for the  $m + 1^{th}$  action from the end (when the actions preceding the  $m + 1^{th}$  action from the end are those given in the limit action tuple  $(a_1, \dots, a_r)$ ). As the  $\alpha$  was arbitrary, the  $m + 1^{th}$  action from the end of  $(a_1, \dots, a_r)$  can be made an SPNE action, which completes the proof.  $\square$

**Lemma 5.** *All mediated games have an SPNE.*

*Proof.* We prove the lemma via induction on the height of the game tree. Note that this is possible as mediated games (like extensive-form games) are recursive: the children of a node of a mediated game are mediated games.

*Base case* (at most 0 actions): This is vacuously true.

*Induction hypothesis* (at most  $k$  actions): Assume we have shown that any mediated game with a game tree of height at most  $k$  has an SPNE.

*Induction step* (at most  $k + 1$  actions): Let  $p$  be the agent that commits the first action. If  $p$  is the mediator, any action that is an allowed edge will be SPNE; and if no such action exists, any action will be SPNE (as the mediator is doomed to a utility of 0). Now suppose  $p$  is not the mediator.

Assume by the inductive assumption, once  $p$  makes its move, all remaining (at most)  $k$  actions are SPNE actions. By the definition of a mediated game,  $p$ 's utility is bounded. Then the least upper bound property of  $\mathbb{R}$  implies that  $p$ 's utility as a function of the first action must have a supremum  $S$ . Via the axiom of choice, we construct a sequence of possible actions for the first action that approaches  $S$  in  $p$ 's utility. That is, we have some sequence  $x^i |_i$  such that if  $p$  plays  $x^i$  for the first action, it achieves some utility  $f(x^i)$  — where  $\lim_{i \rightarrow \infty} f(x^i) = S$ . Moreover, let  $g(x^i)$  map the action  $x^i$  to a tuple of the remaining actions — which are SPNE. By Lemma 3  $(x^i, g(x^i)) |_i$  must have a convergent subsequence  $(y^i, g(y^i)) |_i$  that converges to  $(y, g(y))$  — where  $y$  is a legal first action and  $g(y)$  are legal subsequent actions.

Notice that  $(y_i, g(y_i)) |_i$  is a convergent sequence of 2-SPNE action tuples and thus by Lemma 4,  $(y, g(y))$  is a 2-SPNE action tuple as well. Furthermore, note that by the continuity requirement of mediated games,  $y$  must give  $p$  a utility of  $S$ . Therefore, this must be an SPNE action and so we are done.  $\square$

With this machinery in hand, we are now ready to complete the proof of Theorem 2. Our main task is to make a formal connection between mediated games and (informed) GCC protocols.

*Proof of Theorem 2.* Suppose we have a  $n$ -agent GCC protocol  $\mathcal{P}$  with a bounded number of steps and set valuations of the agents  $V_1, \dots, V_n$ . Then we wish to prove that 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.

Outfit  $\mathcal{P}$  as a game  $M$ , such that all but the final condition of mediated games are satisfied — that is, the mediator enforces the rules of  $\mathcal{P}$  and achieves utility 1 if it follows the rules of  $\mathcal{P}$  and 0 otherwise. More explicitly, the mediator plays every second action and upon examination of the history of events (i.e. the ordering of the cuts made thus far, and results of choose queries), decides the next agent to play and their action space based on the prescription of  $\mathcal{P}$ . To see how all but the last condition is satisfied, we go through them in order.

1. This is by definition.
2. The height of the tree is twice the height of the GCC protocol.
3. The mediator's utility is bounded by 1 by definition, and all other agent's utilities are bounded by 1 as that is their value of the entire cake.
4. This is by definition.

5. When the mediator wishes to ask a *Cut* query to agent  $i$  in the interval  $[a, b]$ , it plays the action  $(i, (a, b))$ , whereas when it wishes to ask a *Choose* query to agent  $i$  giving them the choice between the  $x_1^{th}, \dots, x_k^{th}$  pieces from the left, it plays the action  $(i, \{x_1, \dots, x_k\})$ . This method of giving choose queries deviates slightly from the definition given in Section 2.1, but the two representations are clearly equivalent.
6. The allowed edges are ones that follow the rules of  $\mathcal{P}$ .
7. This property is only relevant when considering *Cut* nodes. To establish it, first consider the action in a single *Cut* node, and fix all the other actions. We claim that for every  $\epsilon > 0$  there exists  $\delta = \delta(\epsilon) > 0$  that is independent of the choice of actions in other nodes such that moving the cut by at most  $\delta$  changes the values by at most  $\epsilon$ . Indeed, let us examine how pieces change as the cut point moves. As long as the cut point moves without passing any other cut point, one piece shrinks as another grows. As the cut point approaches another cut point, the induced piece — say  $k$ 'th from the left — shrinks. When the cut point passes another cut point  $x$ , the  $k$ 'th piece from the left grows larger, or it remains a singleton and another piece grows if there are multiple cut points at  $x$ . In any case, it is easy to verify that the sizes of various pieces received in *Choose* nodes change by at most  $\delta$  if the cut point is moved by  $\delta$ . Furthermore, note that the number of steps is bounded by  $r$  and — since the value density functions are continuous — there is an upper bound  $M$  on the value density functions such that if  $y - x \leq \delta'$  then  $V_i([x, y]) \leq M\delta'$  for all  $i \in N$ . Therefore, choosing  $\delta \leq \epsilon/(Mr)$  is sufficient. Finally,  $V_1, \dots, V_n$  are continuous even in the actions taken in multiple *Cut* nodes, because we could move the cut points sequentially.

We now alter  $M$  such that at every branch induced by a comparison of cuts via an *If-Else*, we allow in the case of a tie to follow either branch. Formally, suppose at a branch induced by the statement “if  $x \leq y$  then  $A$  else  $B$ ” we now set in the case of  $x = y$  the edges for both  $A$  and  $B$  as allowed. Then we claim the property of allowed-edges-closedness is satisfied.

To see this, let us consider action tuples. An action tuple where the mediator in  $M$  only plays on allowed edges can be viewed as a trace of an execution of  $\mathcal{P}$  which records the branch taken on every *If-Else* statement — though when there is a tie the trace may follow the “incorrect” branch. A convergent sequence of such action tuples at some point in the sequence must then keep the branches it chooses in the execution of  $\mathcal{P}$  constant — unless in the limit, the cuts compared in a branch that is not constant coincide. Thus, we have that in the limit, if a branch is constant, the mediator always takes an allowed edge trivially, and otherwise due to our modification of  $M$  the mediator still takes an allowed edge. Furthermore, for all actions of the mediator that are not induced by *If-Else* statements, the mediator clearly still plays on allowed edges and so we have proved the claim.

Now as  $M$  is a mediated game, it has an SPNE  $S$  by Lemma 5. Let  $\mathcal{P}'$  be the informed GCC protocol equivalent to  $\mathcal{P}$  up to tie-breaking such that for every point in the game

tree of  $M$  that represents the mediator branching on an “if  $x \leq y$  then  $A$  else  $B$ ” statement in the original protocol  $\mathcal{P}$ ,  $\mathcal{P}'$  chooses the  $A$  or  $B$  that  $S$  takes in the event of a tie. Then the informedness of the tie-breaking is built into  $\mathcal{P}'$  and we immediately see that the SPNE actions of the regular agents in  $M$  correspond to SPNE actions in  $\mathcal{P}'$ .  $\square$

### C Proof of Theorem 3

The proof of the theorem uses the Thieves Protocol given by Algorithm 4. In this protocol, agent 1 first demarcates a contiguous allocation  $X = \{X_1, \dots, X_n\}$  of the entire cake, where  $X_i$  is a contiguous piece that corresponds to agent  $i$ . This can be implemented as follows. First, agent 1 makes  $n$  cuts such that the  $i$ -th cut is interpreted as the left endpoint of  $X_i$ . The left endpoint of the leftmost piece is reset to 0 by the protocol. Then, the rightmost endpoint of  $X_i$  is naturally the leftmost cut point to its right or 1 if no such point exists. Ties among overlapping cut points are resolved in favor of the agent with the smallest index; the corresponding cut point is assumed to be the leftmost one. Notice that every allocation that assigns nonempty contiguous pieces to all agents can be demarcated in this way.

After the execution of the demarcation step,  $X$  is only a tentative allocation. Then, the protocol enters a verification round, where each agent  $i$  is allowed to *steal* some non-empty strict subset of a piece (say,  $X_j$ ) demarcated for another agent. If this happens (i.e., the if-condition is true) then agent  $i$  takes the stolen piece and the remaining agents get nothing. This indicates the failure of the verification and the protocol terminates. Otherwise, the pieces of  $X$  are eventually allocated to the agents, i.e., agent  $i$  takes  $X_i$ .

We will require two important characteristics of the protocol. First, it guarantees that no state in which some agent steals can be a NE; this agent can always steal an even more valuable piece. Second, stealing is beneficial for an envious agent.

*Proof of Theorem 3.* Let  $\mathcal{P}$  be the Thieves protocol given by Algorithm 3 and  $\mathcal{E}$  be any NE of  $\mathcal{P}$ . Denote by  $X$  the contiguous allocation of the entire cake obtained during the demarcation step, where  $X_i = [x_i, y_i]$  for all  $i \in N$ , and let  $w_i$  and  $z_i$  be the cut points of agent  $i$  during its verification round. Assume for the sake of contradiction that  $X$  is not envy-free. Let  $k^*$  be an envious agent, where  $V_{k^*}(X_{j^*}) > V_{k^*}(X_{k^*})$ , for some  $j^* \in N$ . There are two cases to consider:

*Case 1:* Each agent  $i$  receives the piece  $X_i$  in  $\mathcal{E}$ . This means that, during its verification round, each agent  $i$  selects its cut points from the set  $\bigcup_{j=1}^n \{x_j, y_j\}$ . By the non-envy-freeness condition for  $X$  above (and by the fact that the valuation function  $V_{k^*}$  is strictly positive), there exist  $w'_{k^*}, z'_{k^*}$  such that  $x_{j^*} < w'_{k^*} < z'_{k^*} < y_{j^*}$  and  $V_{k^*}([w'_{k^*}, z'_{k^*}]) > V_{k^*}([x_{k^*}, y_{k^*}])$ . Thus, agent  $k^*$  could have been better off by cutting at points  $w'_{k^*}$  and  $z'_{k^*}$  in its verification round, contradicting the assumption that  $\mathcal{E}$  is a NE.

*Case 2:* There exists an agent  $i$  that did not receive the piece  $X_i$ . Then, it must be the case that some agent  $k$  stole a non-empty strict subset  $[w'_k, z'_k] = [w_k, z_k] \cap Z_j$  of another piece  $X_j$ . However, agent  $k$  could have been better off at the

```

Agent 1 demarcates a contiguous allocation  $X$  of the cake
for  $i = 2, \dots, n, 1$  do
  // Verification of envy-freeness for agent  $i$ 
  Agent  $i$  Cuts in  $\{[0, 1]\}$  // @  $w_i$ 
  Agent  $i$  Cuts in  $\{[w_i, 1]\}$  // @  $z_i$ 
  for  $j = 1$  to  $n$  do
    if  $\emptyset \neq ([w_i, z_i] \cap X_j) \subsetneq X_j$  then
      // Agent  $i$  steals a non-empty strict subset of  $X_j$ 
      Agent  $i$  Chooses from  $\{[w_i, z_i] \cap X_j\}$ 
      exit // Verification failed: protocol terminates
    end if
  end for
  // Verification successful for agent  $i$ 
end for
for  $i = 1$  to  $n$  do
  Agent  $i$  Chooses from  $\{X_i\}$ 
end for

```

**Algorithm 4:** Thieves Protocol: Every NE induces a contiguous envy-free allocation that contains the entire cake and vice versa.

node in the game tree reached in its verification round by making the following marks:  $w'_k = \frac{x_j + w''_k}{2}$  and  $z'_k = \frac{z''_k + y_j}{2}$ . Since either  $x_j \leq w''_k < z''_k < y_j$  or  $x_j < w''_k < z''_k \leq y_j$  (recall that  $[w''_k, z''_k]$  is a non-empty strict subset of  $X_j$  and the valuation function  $V_k$  is strictly positive), it is also true that  $V_k([w'_k, z'_k]) > V_k([w''_k, z''_k])$ , again contradicting the assumption that  $\mathcal{E}$  is a NE.

So, the allocation computed by agent 1 under every NE  $\mathcal{E}$  is indeed envy-free; this completes the proof of the first part of the theorem.

We next show that every contiguous envy-free allocation of the entire cake is the outcome of a NE. Let  $Z$  be such an allocation, with  $Z_i = [x_i, y_i]$  for all  $i \in N$ . We define the following set of strategies  $\mathcal{E}$  for the agents:

- At every node of the game tree (i.e., for every possible allocation that could be demarcated by agent 1), agent  $i \geq 2$  cuts at points  $w_i = x_i$  and  $z_i = y_i$  during its verification round.
- Agent 1 specifically demarcates the allocation  $Z$  and cuts at points  $w_1 = x_1$  and  $z_1 = y_1$  during its verification round.

Observe that  $[w_i, z_i] \cap Z_j$  is either empty or equal to  $Z_j$  for every pair of  $i, j \in N$ . Hence, the verification phase is successful for every agent and agent  $i$  receives the piece  $Z_i$ .

We claim that this is a NE. Indeed, consider a deviation of agent 1 to a strategy that consists of the demarcated allocation  $Z'$  (and the cut points  $w'_1$  and  $z'_1$ ). First, assume that the set of pieces in  $Z'$  is different from the set of pieces in  $Z$ . Then, there is some agent  $k \neq 1$  and some piece  $Z'_j$  such that the if-condition  $\emptyset \subset [x_k, y_k] \cap Z'_j \subsetneq Z'_j$  is true. Hence, the verification round would fail for some agent  $i \in \{2, \dots, k\}$  and agent 1 would receive nothing. So, both  $Z'$  and  $Z$  contain the same pieces, and may differ only in the way these pieces are tentatively allocated to the agents. But in this case the maximum utility agent 1 can get is  $\max_j V_1(Z'_j)$ , either

by keeping the piece  $Z'_1$  or by stealing a strict subset of some other piece  $Z'_j$ . Due to the envy-freeness of  $Z$ , we have:

$$\max_j V_1(Z'_j) = \max_j V_1(Z_j) = V_1(Z_1),$$

hence, the deviation is not profitable in this case either.

Now, consider a deviation of agent  $i \geq 2$  to a strategy that consists of the cut points  $w'_i$  and  $z'_i$ . If both  $w'_i$  and  $z'_i$  belong to  $\bigcup_{j=1}^n [x_j, y_j]$ , then  $[w'_i, z'_i] \cap Z_j$  is either empty or equal to  $Z_j$  for some  $j \in N$ . Hence, the deviation will leave the allocation unaffected and the utility of agent  $i$  will not increase. If instead one of the cut points  $w'_i$  and  $z'_i$  does not belong to  $\bigcup_{j=1}^n [x_j, y_j]$ , this implies that the condition

$$\emptyset \subset [w'_i, z'_i] \cap Z_j \subsetneq Z_j$$

is true for some  $j \in N$ , i.e., agent  $i$  will steal the piece  $[w'_i, z'_i] \cap Z_j$ . However, the utility  $V_i([w'_i, z'_i] \cap Z_j)$  of agent  $i$  cannot be greater than  $V_i(Z_j)$ , which is at most  $V_i(Z_i)$  due to the envy-freeness of  $Z$ . Hence, again, this deviation is not profitable for agent  $i$ .

We conclude that  $\mathcal{E}$  is a NE; this completes the proof of the theorem.  $\square$