

A DISCRETE AND BOUNDED ENVY-FREE CAKE CUTTING PROTOCOL FOR ANY NUMBER OF AGENTS

HARIS AZIZ* AND SIMON MACKENZIE†

Abstract. We consider the well-studied cake cutting problem in which the goal is to find an envy-free allocation based on queries from n agents. The problem has received attention in computer science, mathematics, and economics. It has been a major open problem whether there exists a discrete and bounded envy-free protocol. We resolve the problem by proposing a discrete and bounded envy-free protocol for any number of agents. The maximum number of queries required by the protocol is $n^{n^{n^{n^n}}}$. We additionally show that even if we do not run our protocol to completion, it can find in at most $n^3(n^2)^n$ queries a partial allocation of the cake that achieves proportionality (each agent gets at least $1/n$ of the value of the whole cake) and envy-freeness. Finally we show that an envy-free partial allocation can be computed in at most $n^3(n^2)^n$ queries such that each agent gets a connected piece that gives the agent at least $1/3n$ of the value of the whole cake.

Key words. Fair Division, Elicitation Protocols, Multiagent Resource Allocation, Cake cutting, Communication Complexity.

AMS subject classifications. 91A12, 68Q15

1. Introduction.

“Despite intense efforts over decades, up to this date no one has succeeded in finding a finite bounded cake cutting protocol that guarantees envy-freeness for any number of players.”—Lindner and Rothe [25].

The existence of a discrete and bounded envy-free cake cutting protocol has remained a major open problem for at least two decades. In this paper, we settle the problem by presenting a general discrete envy-free protocol for any number of agents. The protocol is for the cake cutting setting that is a versatile mathematical model for allocation of a heterogeneous divisible good among multiple agents with possibly different preferences over different parts of the cake. The main applications of cake cutting are fair scheduling, resource allocation, and conflict resolution [15]. Cake cutting has been extensively studied within computer science [30] and the social sciences [41]. Since various important divisible resources such as time and land can be captured by cake cutting, the problem of fairly dividing the cake is a fundamental one within the area of fair division and multiagent resource allocation [7, 21, 29, 32, 37, 40, 41].

A cake is represented by an interval $[0, 1]$ and each of the n agents has a valuation function over pieces of the cake that specifies how much that agent values a particular subinterval. The main goal is to divide the cake fairly. Among various fairness concepts proposed by social scientists, a prominent one is *envy-freeness*. An allocation is envy-free if no agent would prefer to take another agent’s allocation instead of his own. Although an envy-free allocation is guaranteed to exist even with $n - 1$ cuts [40]¹, *finding* an envy-free allocation is a challenging problem which has been termed “one of the most important open problems in 20th century mathematics” by Garfunkel [19].

Motivation and Contribution. Since the valuations of agents over the cake can be complex, eliciting each agent’s complete valuations function over the cake is computationally infeasible. A natural approach in cake cutting protocols is to query agents about their valuations of different portions of the cake and based on these queries propose an allocation. A cake cutting protocol is *envy-free* if each agent is guaranteed to be non-*envious* if he reports his real valuations. If a protocol is envy-free, then an honest agent will not be envious even if other agents misreport their valuations. For the case of two agents, the problem has a well known solution in the form of the *Divide and Choose* protocol: one agent is asked to cut the cake into equally preferred pieces and the other agent is asked to choose the preferred piece. For the case of three agents, an elegant and bounded protocol was independently discovered by John L. Selfridge and John H. Conway around 1960 [7]. Since then, an efficient general envy-free protocol for any number of agents has eluded mathematicians, economists, and computer scientists.

In 1995, Brams and Taylor [6] made a breakthrough by presenting an envy-free protocol for *any* number

*Data61, CSIRO and UNSW Australia (haris.aziz@data61.csiro.au).

†Carnegie Mellon University, USA (simonm@andrew.cmu.edu).

¹Su [40] pointed out that the existence of an envy-free cake allocation can be shown via an interesting connection with Sperner’s Lemma.

of agents [22]. Although the protocol is guaranteed to terminate in finite time, there is one drawback of the protocol: the running time or number of queries and even the number of cuts required is unbounded even for four agents. In other words, the number of queries required to identify an envy-free allocation can be arbitrarily large for certain valuations functions. Procaccia [30] terms unboundedness as a “serious flaw”. Brams and Taylor were cognizant of their protocol’s drawback and explicitly mentioned the problem of proposing a bounded envy-free protocol even for $n = 4$. Lindner and Rothe [24] write that “the development of finite bounded envy-free cake cutting protocols still appears to be out of reach, and a big challenge for future research.” The problem has remained open and has been highlighted in several works [3, 6, 7, 11, 16, 34, 23, 29, 30, 25, 32, 33]. Saberi and Wang [33] term the problem as “one of the most important open problems in the field”. Procaccia [29] goes one step further and calls it an interesting challenge within theoretical computer science: “Since the 1940s, the computation of envy-free cake divisions has baffled many great minds across multiple disciplines. Settling this problem once and for all is an important challenge for theoretical computer science.” In recent work, Aziz and Mackenzie [1] proposed a bounded and discrete envy-free protocol for $n = 4$. Despite this progress, the general problem for any number of agents had remained unaddressed.

In this paper, *we present a discrete envy-free protocol for any number of agents that requires a bounded number of queries and hence a bounded numbers of cuts of the cake* thereby closing the central open problem in the field of cake cutting. The bound for the number of queries is $n^{n^{n^{n^n}}}$. Previously, no discrete protocol was known that even uses a bounded number of cuts. Apart from using some classic ideas in cake-cutting, our protocol requires some radically different techniques. Since only three general finite and discrete cake cutting protocols have been presented in the literature [6, 31, 27], our protocol is another addition to the list of finite envy-free protocols and is of interest even if boundedness is not a critical property. More generally, the protocol provides a new constructive argument for existence of envy-free allocations. Our result is surprising because leading experts in the area have conjectured that a bounded and discrete envy-free protocol does not exist: “It is natural to ask whether the envy-free cake cutting problem is inherently difficult: Is it provably impossible to design a bounded envy-free cake cutting algorithm?” [30]. Previously, mathematician Ian Stewart had commented that “However, no discrete procedure with a bounded number cuts (however large) is known for four players, and such schemes probably don’t exist” [38].

We additionally show that even if we do not run our protocol to completion, it can find in at most $n^3(n^2)^n$ queries a partial allocation of the cake that achieves proportionality (with respect to the whole cake) and envy-freeness. Finally, we show that even if we do not run our protocol to completion, it can be tweaked to find in at most $n^3(n^2)^n$ queries an envy-free partial allocation of the cake in which each agent gets a *connected piece* that gives the agent $1/3n$ of the value of the whole cake. For the additional two results, we only need one of our subprotocols called the SubCore Protocol that we define later.

Apart from introducing new techniques, our protocol relies on some ideas from previous protocols. At the heart of our protocol is the idea of domination or ‘irrevocable advantage’ [6]. An agent i dominates another agent j if he is not envious of j even if the unallocated cake is given to j . Note that if agents in subset $S \subset N$ dominate all the agents in $N \setminus S$, then agents in S can simply let the agents in $N \setminus S$ worry about the remaining cake since they will not be envious if a single agent in $N \setminus S$ got all the remaining unallocated cake. Our protocol uses some high level ideas from our previous protocol for 4 agents: (i) We use a ‘Core Protocol’ over the unallocated cake (aka residue) with one agent as the specified cutter that results in an envy-free partial allocation and a smaller residue. The Core Protocol is called repeatedly on the latest residue so as to make the residue even smaller. As the new calls to the Core Protocol are made, the partial allocation remains envy-free. (ii) Every time agents make trims or evaluate pieces of cake, we do not lose such information and in fact keep track of it. (iii) We make agents exchange some parts of their already allocated pieces to obtain new dominations. The exchange was referred to as a *permutation* in [1]. In some steps, our protocol also uses a simple fact: envy-freeness implies proportionality with respect to the allocated cake (i.e., each agent gets at least $1/n$ value of the allocated cake).

Throughout the overall protocol, we have cake that has been allocated and some cake that is the unallocated residue. During the protocol, we ensure that the cake that has been allocated to the agents has been done so in an envy-free manner. Overall, our envy-free protocol (*Main Protocol*) takes as input the cake and a set of agents and returns an envy-free allocation of the whole cake. In order to obtain the envy-free allocation, the Main Protocol calls three other protocols. It repeatedly calls the *Core Protocol* a bounded

number of times. The Core Protocol is the work-horse of the overall protocol that is called repeatedly to further allocate unallocated cake in any envy-free manner. The other protocols—GoLeft and Discrepancy—are used to make one set of agents dominate the others so that the problem reduces to finding an envy-free allocation for fewer agents and the Main Protocol can be called again for this sub-problem.

After calling the Core Protocol many times, if there is still some unallocated cake and there is sufficient difference in the estimation of the value of a cake piece between some agents, then the Discrepancy Protocol is called by the Main Protocol. The goal of the Discrepancy Protocol is to either exploit the differences in valuations or else to ensure that on a very rough scale, agents have similar valuations. If after calling the Discrepancy Protocol, the cake still has not been allocated in an envy-free way, the GoLeft Protocol is called by the Main Protocol. The GoLeft Protocol is used to reallocate pieces of the cake among the agents after additionally adding some crumbs from the residue. The reallocation is useful to obtain new dominations until one set of agents dominates the remaining agents which means that we have reduced the problem to allocating the remainder of the cake among the dominated agents. The different protocols that we present are summarized in Figure 1.

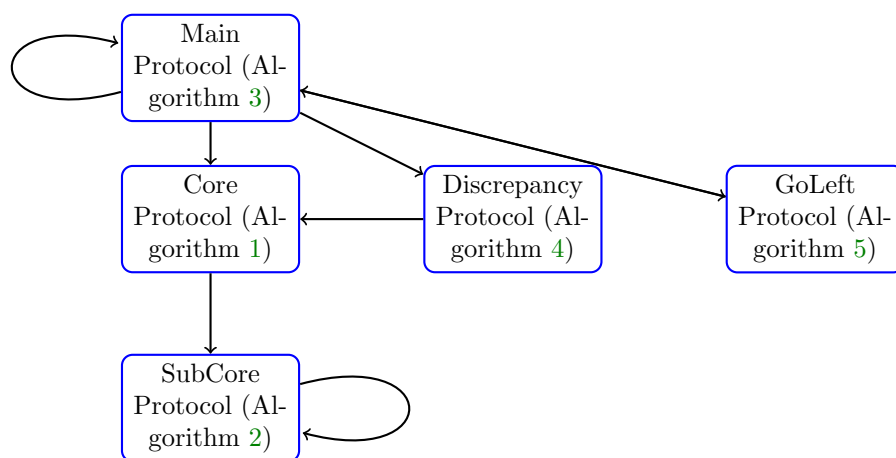


Fig. 1: The envy-free protocol for n agents relies on various protocols. A protocol points to another protocol if it calls the other one. A protocol has a self-loop if it calls itself recursively. The Main Protocol is called to get an envy-free allocation.

Next, we informally summarize what input each protocol takes as a blackbox and then what it outputs.

- **Main Protocol**
 - Input: Cake R and a set of agents N .
 - Output: An envy-free allocation that completely allocates R among agents in N .
- **Core Protocol**
 - Input: Specified cutter agent, agent set N and unallocated cake R .
 - Output: An envy-free allocation of cake $R' \subset R$ for agents in N and updated unallocated cake $R \setminus R'$.
- **SubCore Protocol**
 - Input: Cake R cut into n pieces to be allocated among agents in set $N' \subset N$ with $n' = |N'|$. Additionally, each agent $i \in N'$ has a benchmark value b_i .
 - Output: An envy-free partial allocation of R for agents in N' in which each agent $j \in N'$ gets a piece of value at least b_j .
- **Discrepancy Protocol**
 - Input: Residue R , a specific piece in R , a specific value given by agent u , a set of pieces of cake other than R , and a set of agents N .
 - Output: Possibly modified residue R , a Boolean value called DISCREPANCY, and a partition of N into sets D' and D'' .

- **GoLeft Protocol**

- Input: A set N , a set of ‘snapshots’ (disjoint envy-free allocation for the agents), a set of corresponding ‘extracted pieces’ for each piece in the disjoint allocations, and residue R .
- Output: A set of agents $A \subset N$ such that all agents in $N \setminus A$ dominate all agents in A .

New Techniques and Ideas. Despite some high level similarities with the approach of the recent protocol for four agents [1], our general protocol requires a number of new concepts and techniques. These concepts and techniques will be formalized and explained in the later sections but we mention them briefly and informally.

We define the notion of a *significant advantage or bonus* that informally means that an agent i feels his absolute advantage over another agent is of a high enough fraction of his value of the unallocated cake. We then make use of our Core Protocol that not only reduces the size of the unallocated residue but also guarantees that a given agent has a significant advantage over someone else. The Core Protocol can be used to further allocate the residue so that an agent’s significant advantage over another agent translates into dominance in a bounded number steps. This approach is one of the key ideas for obtaining a bounded protocol rather than simply a finite protocol.

We rely on a notion of discrepancy that is based on the idea of how differently agents view a piece of cake. It is well known in fair division that the more different the agents’ valuations are, the easier it is to achieve fairness. We try to use the structure of agents that have similar valuations at a very coarse level. In case they have high enough discrepancy in valuations over two parts of the unallocated cake, then we simply break the problem into two subproblems where we divide each part in any envy-free manner among those agents who value it highly compared to the other part.

Another key idea to achieve a bounded protocol is that we run the Core Protocol enough times to ensure there are a sufficient number of Core Protocol allocation snapshots (allocation outcomes of various calls of the Core Protocol) that are *isomorphic*.² We will keep track of these snapshots to later make more agents dominate each other.

Further operations such as exchanges of allocated pieces are then implemented on such isomorphic Core snapshots by exploiting their structure. In order to achieve boundedness of the protocol, we also define and work with four carefully chosen bounds that are functions of the number of agents. The biggest of the four bounds is the maximal number of query operations required to complete the protocol.

We use a new technique called *extraction* whereby part of the residue is systematically trimmed off so that it can potentially be combined with some agent’s allocation in the Core Protocol. When a particular extracted piece of cake is given to an agent in addition to his allocated piece in a single snapshot of the Core Protocol, we refer to this as *attachment*. Attachment is helpful to systematically enable more agents to exchange each others’ pieces in snapshots of the Core Protocol without causing any agent to be envious. The exchange of pieces consequently allows for ‘permutations’ and helps achieve agents dominate each other.

In order to methodically track the possible permutations/exchanges, we keep track of a *permutation graph* in which nodes correspond to agents and an agent i points to agent j if i is willing to replace his allocation in a subset of isomorphic snapshots of the Core Protocol with j ’s allocated piece plus a tiny bit more cake without causing envy. The tiny bit more cake to attract agent i can in principle cause agents to be envious so we have to compensate the other agents elsewhere so that envy is not introduced. The attachments and exchanges take place in the GoLeft Protocol. The interplay between the permutation graph and the working set of isomorphic Core snapshots is technically one of the most interesting parts of the overall protocol and is the central part of the GoLeft Protocol.

Related Work. Cake cutting problems originated in the 1940’s when famous mathematicians such as Banach, Knaster, and Steinhaus initiated serious mathematical work on the topic of fair division.³ Since then, the theory of cake cutting algorithms has become a full-fledged field with at least three books written on the topic [4, 7, 32]. The central problem within cake cutting is finding an envy-free allocation [18, 37].

Since the earliest works, mathematicians have been interested in the complexity of cake cutting. Steinhaus [36] wrote that “Interesting mathematical problems arise if we are to determine the minimal number of cuts necessary for fair division.” When formulating efficient cake cutting protocols, a typical goal is to minimize the number of cuts while ignoring the number of valuations queried from the agents. In principle,

²The notion of isomorphic snapshots will be defined later.

³Hugo Steinhaus presented the cake cutting problems to the mathematical and social science communities on Sep. 17, 1947, at a meeting of the Econometric Society in Washington, D.C. [32, 36].

the actual complexity of a problem or a protocol depends on the number of queries. When considering how efficient a protocol is, it is useful to have a formal query model for cake cutting protocols. Robertson and Webb [32] formalized a simple query model in which there are two kinds of queries: Evaluation and Cut. In an Evaluation query, an agent is asked how much he values a subinterval. In a Cut query, an agent is asked to identify an interval, with a fixed left endpoint, of a particular value. Although, the query model of Robertson and Webb [32] is very simple, it is general enough to capture all known protocols in the literature. Note that if the number of queries is bounded, it implies that the number of cuts is bounded in the Robertson and Webb model. The protocol that we present in this paper uses a bounded number of queries in the Robertson and Webb model.

There are various notions of fairness in the cake-cutting literature. Two of the most prominent ones are envy-freeness and proportionality. Proportionality requires that each agent gets at least $1/n$ value of the whole cake. An envy-free allocation of the whole cake also satisfies proportionality. Finding a proportional allocation is an easier task than finding an envy-free allocation. There exists a general discrete protocol for finding a proportional allocation with connected pieces that takes $O(n \log n)$ queries [17].

There is not too much known about the existence of a bounded envy-free protocol for general n except that any envy-free cake cutting algorithm requires $\Omega(n^2)$ queries in the Robertson-Webb model [28, 30]. Also, for $n \geq 3$, there exists no finite envy-free cake cutting algorithm that outputs *contiguous* allocations in which each agent gets a connected piece with no gaps [39]. Brams et al. [8] and Barbanel and Brams [5] presented envy-free protocols for four agents that require 13 and 5 cuts respectively. However, the protocols are not only unbounded but also not finite since they are *continuous* protocols that require the notion of a *moving knife*. An alternative approach is to consider known bounded protocols and see how well they perform in terms of envy-freeness [24]. Apart from the unbounded Brams and Taylor envy-free protocol for n agents, there are other general envy-free protocols by Robertson and Webb [31] and Pikhurko [27] that are also unbounded. Gasarch [20] compared the complexity of the three general unbounded envy-free protocols in the literature. Aziz and Mackenzie [1] proposed a bounded and discrete envy-free protocol for $n = 4$. We use similar ideas but generalizing to the case of any number of agents requires novel subroutines and significantly more abstract arguments.

There are positive algorithmic results concerning envy-free cake cutting when agents have restricted valuations functions [9, 13, 15] or when some part of the cake is left unallocated [34, 33]. There has also been work on *strategyproof* cake cutting protocols for restricted valuation functions [2, 12, 26] as well as strategic aspects of protocols [10].

2. Preliminaries.

Model. We consider a cake which is represented by the interval $[0, 1]$. A *piece of cake* is a finite union of disjoint subintervals of $[0, 1]$. We will assume the standard assumptions in cake cutting. Each agent in the set of agents $N = \{1, \dots, n\}$ has his own valuation over subintervals of the interval $[0, 1]$. We denote by $V_i(x, y)$ the value agent i has for interval $[x, y]$. The valuations are (i) *non-negative*: $V_i(Y) \geq 0$ for any subinterval Y ; (ii) *additive*: for all disjoint subintervals Y, Y' , $V_i(Y \cup Y') = V_i(Y) + V_i(Y')$; and (iii) *divisible*, i.e., for every subinterval Y and $0 \leq \lambda \leq 1$, there exists $Y' \subseteq Y$ with $V_i(Y') = \lambda V_i(Y)$. It follows from the divisibility property that the valuations are non-atomic i.e., $V_i(x, x) = 0$.

We will typically denote an allocation by $X = (X_1, \dots, X_n)$ where X_i is the cake allocated to agent i . Two main criteria of fairness are envy-freeness and proportionality.

DEFINITION 1 (Envy-free allocation). *An allocation $X = (X_1, \dots, X_n)$ is envy-free if $V_i(X_i) \geq V_i(X_j)$ for each $i, j \in N$.*

DEFINITION 2 (Proportional allocation). *An allocation $X = (X_1, \dots, X_n)$ is proportional if $V_i(X_i) \geq \frac{1}{n} V_i([0, 1])$ for each $i \in N$.*

In order to ascertain the complexity of a protocol, Robertson and Webb presented a computational framework in which two kinds of queries can be made to the agents: (1) for given $x \in [0, 1]$ and $r \in \mathbb{R}^+$, a CUT query asks an agent to return a point $y \in [0, 1]$ such that $V_i([x, y]) = r$ (2) for given $x, y \in [0, 1]$, EVALUATE query asks an agent to return a value $r \in \mathbb{R}^+$ such that $V_i([x, y]) = r$. A cake cutting protocol specifies how agents interact with queries and cuts. All well known cake cutting protocols can be analyzed in terms of the number of queries required to return a fair allocation. A cake cutting protocol is *bounded* if the number of queries required to return a solution is bounded by a function of n irrespective of the valuations of the agents.

Terms and Conventions. We now define some terms and conventions that we will use in the paper. An allocation is *partial* if it does not necessarily allocate the whole cake. Given an envy-free partial allocation X of the cake and an unallocated residue R , we say that agent i *dominates* agent j if i does not become envious of j even if all of R were to be allocated to j :

$$V_i(X_i) \geq V_i(X_j) + V_i(R).$$

We say that an agent gets a piece *partially*, if he does not necessarily get it completely. In the cake cutting protocols that we will describe, an agent may be asked to trim a piece of cake so that its value equals the value of a less valuable piece. Agents will be asked to trim various pieces of the cake to make a given piece equal to the value of another piece. In Figure 2, we outline the idea of trimming a piece to equal the value of another piece. When an agent trims a piece of cake, he will trim it from the left side: the main piece (albeit trimmed) will be on the right side. The piece minus the trim will be called the partial main piece. When we say that an agent is *guaranteed* to get his second/third/etc. most favoured piece, this guarantee is based on the *ordinal* preferences of the agents over the pieces. By ordinal, we mean that agents simply give a weak ordering over the pieces but do not tell the exact cardinal utility difference between two pieces. If an agent is indifferent between the top three pieces, then we will still say that the agent is guaranteed to get his third most valued piece.

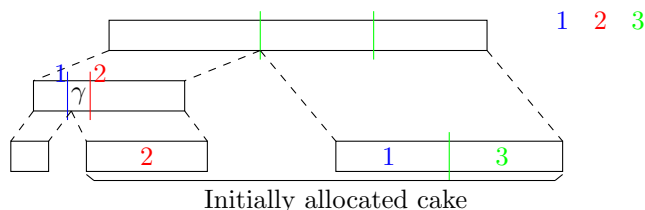


Fig. 2: Example of a trim. Agents 1 and 2 trim their most preferred piece (the left most piece) to the value equal to that of their second most preferred piece. In this instance, agent 2 trims more than agent 1, hence his trim is to the right of agent 1’s trim. Let us assume agent 1 and 3 each get a complete piece with 1 getting his second most preferred piece. Agent 2 is not envious of other agents if he gets the right side of the trimmed piece up till his trim. If agent 2 gets the part to the right of agent 1’s trim instead of his own trim, then he is even happier. Agent 1 is still not envious of agent 2 if 2 gets γ . The reason is that agent 1 thinks that his advantage over agent 2 is of value that is equal to his value for γ .

3. The Protocol. The overall protocol we present is essentially the Main Protocol that calls other protocols. Some protocols such as the SubCore Protocol and the Main Protocol recursively call themselves. The overall idea of the protocol is as follows. We make an agent (the cutter) divide the cake into n equally preferred pieces. The cake is then allocated to the agents with each agent getting a part of one of the pieces. The subroutine will be referred to as the *Core Protocol*. The Core Protocol itself relies on recursively applying the *SubCore Protocol*. When running the Core Protocol, if all the cake is already allocated, we are done. Otherwise, the Core Protocol is run repeatedly on the left-over cake. If some cake is still unallocated, we repeat it with another agent as cutter. After the Core Protocol has been run repeatedly, it is checked whether agents have discrepancy on a part of the cake. Discrepancy implies that agents have radically different valuations over two segments of the cake. This works to our advantage because we can run the *Discrepancy Protocol* in which we let the differing agents concentrate on the segments they prefer much more which reduces our problem to envy-free cake cutting for a smaller number of agents. Otherwise, we exchange some parts of the cake that certain agents hold using the GoLeft Protocol in a systematic way to obtain further dominations between agents. After the GoLeft Protocol has been implemented, some agents dominate all other agents. At this point, we have decomposed the problem into a smaller problem with fewer number of agents. Hence the envy-free algorithm can be called recursively on the unallocated cake to divide the cake among the specified subset of agents. We first present the Core Protocol that is the work-horse of the overall Main Protocol.

3.1. Core Protocol. We first present the Core Protocol that calls the SubCore Protocol. The Core Protocol is helpful in allocating additional residue in an envy-free manner. The main challenge is that just by running the Core Protocol repeatedly on the residue, there is no guarantee that the cake will be allocated completely in a bounded or even finite number of steps. Hence, we will introduce other protocols in addition to the Core protocol that help us allocate the whole cake in an envy-free manner.

Algorithm 1 Core Protocol

Input: Specified cutter (say agent $i \in N$), agent set N such that $i \in N$, and unallocated cake R .

Output: An envy-free allocation of cake $R' \subset R$ for agents in N and updated unallocated cake $R \setminus R'$.

- 1: Generate the tie-breaking values that will be used to break all ties as described in Section 3.1.1.
 - 2: Ask agent i to cut the cake R into n equally preferred pieces.
 - 3: Run SubCore Protocol on the n pieces with agent set $N \setminus \{i\}$ with each agent having a benchmark value as zero. The call gives an allocation to the agents in $N \setminus \{i\}$ such that one of the n pieces is untrimmed and unallocated.
 - 4: Give i one of the unallocated untrimmed pieces from the previous step.
 - 5: **return** envy-free partial allocation (in which each agent gets a connected piece) as well as the unallocated cake.
-

Algorithm 2 SubCore Protocol

Input: Cake cut into n'' pieces (with $n'' \leq n$) to be allocated among agents in set $\{1, \dots, n'\} = N' \subseteq N$ with $n' = |N'|$ and a benchmark value b_j for each $j \in N'$. {We only call SubCore if the benchmark values are such that there exists an envy-free allocation of agents in N' where each agent gets at most one of the pieces giving him at least the specified benchmark value.}

Output: A neat envy-free partial allocation for agents in N' in which each agent $j \in N'$ gets a connected piece of value at least b_j .

- 1: **for** $m = 1$ to n' **do**
- 2: **if** the piece agent m preferred at the launch of the protocol is still unallocated **then**
- 3: we tentatively give agent m that piece and go to the next iteration of the ‘for’ loop.
- 4: **else**
 - the first m agents are contesting for the same $m - 1$ tentatively allocated pieces. We call them the *contested* pieces.
 - For each agent j , set b'_j to be the maximum of b_j and agent j ’s value of the most preferred uncontested piece.
 - Then each agent in $[m]$ is asked to place a trim on all contested pieces of high enough value so that the (contiguous) contested piece on the right hand side of his trim is of the same value as b'_j
- 5: Set W to be the set of agents who trimmed most (had the rightmost trim) in some piece.
- 6: **while** $|W| < m - 1$ **do**
- 7: Ignore the previous trims of agents in W from now on and start over.
- 8: Run SubCore Protocol on the contested pieces with W as the target set of agents (with b'_j as their benchmark value input) and for each contested piece, the part to the left side of the right-most trim by an agent in $[m] \setminus W$ is ignored. {The result of the recursive call of SubCore is an allocation that gives a (partial) contested piece to each of the agents in W .}
- 9: Take any unallocated contested piece a . The current left margin (beyond which the piece is ignored) is by agent $i \in [m] \setminus W$.

$$W \leftarrow W \cup \{i\}.$$

At this point the current allocation of agents in W is tentative and not permanently made. {An agent from $m \setminus W$ has been added to W . For the updated W , each agent in W gets a partial contested piece and no agent envies an unallocated piece. Recall that for each piece, the left side of the right-most trim by an agent in $[m] \setminus W$ is ignored.}

- 10: Update the value b'_j of each agent j in W to equal the value of the piece that they have been tentatively allocated.
 - 11: **end while**
 - 12: Run SubCore on all agents in W and the set of contested pieces, where we ignore the part to the left of the trim made by the agent in $[m] \setminus W$. The benchmark of each $j \in W$ is b'_j . {At this point $|W| = m - 1$ and each agent in W has a tentatively allocated contested piece}
 - 13: The only agent j remaining in $[m] \setminus W$ is tentatively given his most preferred uncontested piece.
 - 14: **end if**
 - 15: **end for**
 - 16: **return** envy-free partial cake for agents in N' (such that each agent gets a connected piece that is on the right hand side of the original piece he trimmed most) as well as the unallocated cake.
-

The Core Protocol asks a specified agent termed as the cutter to cut the unallocated cake into n equally preferred pieces. It then calls the SubCore Protocol that allocates to each of the agents one of the pieces (possibly partially) in an envy-free manner. No agent is given cake from any other piece.

The name of the protocol is the same as the Core Protocol used by Aziz and Mackenzie [1] for the case of 4 agents which returns an envy-free partial allocation in which one agent cuts the cake into four equally preferred pieces and the cutter as well as at least one other agent gets one of these four pieces. The new Core Protocol can be considered as a useful generalization of Core Protocol for the four-agent algorithm of Aziz and Mackenzie [1]. Unlike the Core Protocol for the four-agent case, the new Core Protocol requires a more sophisticated recursive SubCore Protocol. The SubCore Protocol results in an allocation that satisfies some ‘neat’ properties.

DEFINITION 3 (Neat Allocation). *Consider a cake divided into n pieces and $m < n$ agents. An allocation of the cake into m agents is neat if (i) each agent’s allocation is a (not necessarily whole) part of one of the*

pieces; (ii) at least one piece is unallocated; (iii) no agent prefers an unallocated piece over his allocation; and (iv) no agent prefers another agent's allocation over his allocation.

We can prove that the Core Protocol results in an envy-free partial allocation in which the cutter cuts the cake into n pieces, each agent gets a part of exactly one of the pieces, at least one non-cutter agent gets a complete piece and at least one piece is unallocated which no non-cutter agent envies over his allocation (Lemma 14). Since two agents get full pieces, from the cutter's perspective, at least $2/n$ of the cake is allocated after an iteration of the Core Protocol.

Most of the technical work in the Core Protocol is done when it calls the SubCore Protocol. The SubCore Protocol takes as input agents and pieces of cake where the number of pieces is at least as much as the number of agents. Each agent is given a part of exactly one piece. The SubCore Protocol also takes as input the benchmark value for each agent. When we call SubCore for the first time during the Core Protocol, the benchmark value of each agent is zero. The benchmark value of an agent indicates the minimum value an agent expects in an envy-free allocation returned by the SubCore Protocol.

In the SubCore Protocol, the main idea is that we start from a single agent and gradually grow the number of agents in a specified order while making sure that a neat allocation exists for the growing set of agents $\{1, \dots, m\}$ in which at least one agent gets a full piece. We denote $\{1, \dots, m\}$ by $[m]$. All the allocations encountered during the course of the SubCore Protocol are considered tentative until the final step. If the next agent m in the specified order most prefers a piece that is not currently (tentatively) allocated, we can easily handle the new agent as he can get the unallocated piece without causing envy for anyone. Otherwise, we have to do more work to ensure that the previously handled agents as well as the new agent can simultaneously get a neat allocation in which at least one agent gets a full piece. This includes calling the SubCore Protocol recursively for a smaller number of agents.

In the SubCore Protocol, if the m -th agent is also interested in one of the $m - 1$ tentatively assigned pieces, we refer to the $m - 1$ allocated pieces as the contested pieces. Each agent in $[m]$ is asked to trim pieces among the $m - 1$ contested pieces that are of higher value than his most preferred piece outside the $m - 1$ allocated pieces so that the value of the former is the same as the value of the latter. The value is referred to as the benchmark value of the agent. The rationale for asking all the m agents to trim up to their benchmark value is as follows. If $m - 1$ agents each get a part of a contested piece, then one of the m agents who does not get a contested piece can get a most preferred uncontested piece. By asking all agents to place trims according to their benchmark values, the agent j who is forced to get an uncontested piece will not be envious of other agents because each other agent gets a piece to the right of j 's trim on that piece. Note that in order to set the benchmark value of an agent we not only consider the uncontested pieces in that call of the SubCore but also implicitly consider other pieces that were part of the input of the root call to the SubCore. This ensures that agents are not envious or interested in pieces that are unallocated. The set of agents who have the rightmost trim in some piece is referred to as W . Note that W is the set of agents who are guaranteed to have an envy-free allocation where each agent gets a piece that he trimmed most. If $|W| = m - 1$, then if agents are given the piece they win up till their trim, then the corresponding allocation is envy-free for agents in W and gives each agent in W their benchmark value. Nonetheless, we call SubCore recursively on W and the contested pieces with the left aside of trim the agent in $[m] \setminus W$ ignored. Doing this ensures that agents in W get as much of the contested pieces as possible without causing envy. The remaining agent in $[m] \setminus W$ is considered to be 'kicked out' of the contested pieces and is given a piece outside the contested pieces. He is given the most preferred uncontested piece.

We may not be lucky and the number of agents who win some piece may not be $m - 1$, i.e., some agent may have the rightmost trim in multiple contested pieces. The protocol then increases the size of W one by one. W is expanded as follows. The previous trims of agents in W are ignored. The SubCore Protocol is called recursively with W as the target set of agents and for each piece, the left side of the right-most trim by an agent in $[m] \setminus W$ is ignored. Note that by ignoring the trims of agents in W , we have more cake that could potentially be allocated than the allocation where the trims of agents in W are not ignored but each agent in W gets the right hand side of the piece where he had the rightmost trim. In case $|W| < m - 1$, we can continue increasing the size of W while ensuring that agents in W can get an envy-free allocation by getting a partial piece each from among the contested pieces. Take any unallocated contested piece a for which the current left margin (beyond which the piece is ignored) is by agent $i \in [m] \setminus W$. We can add such an i to W . As we increase the size of W , we are able to clear more space for some piece since the non-winner because of which we were ignoring the left part of some piece is now a winner, so we can ignore his previous

trim. When $|W| = m - 1$, we have ensured that all the contested pieces have been partially allocated to one agent each from W . In that case the remaining agent in $[m] \setminus W$ can then be given the most preferred uncontested piece.

Note that for $n = 2$, the Core Protocol coincides with the well-known Divide and Choose protocol. Next, we further explain the SubCore Protocol with the help of an example.

EXAMPLE 1 (SubCore Protocol). *In order to illustrate how the SubCore Protocol works, let us explain the SubCore Protocol when $n' = 3$ and there are $n = 4$ pieces a, b, c, d .*

In the for loop, when $m = 2$, if agent 2 has the same most preferred piece as agent 1's tentative piece a , then piece a is the contested piece. Both agents place trim on that piece to equal the value of the second most preferred piece. Let us assume that 2's trim is to the right of 1's trim. In that case the cake to the left of 1's trim is ignored temporarily as if the cake to the left did not exist. Agent 1 can take a full unallocated piece b that is at least as preferred by him as 2's tentative piece. Agent 1 is considered being kicked out of his piece and gets a piece of exactly his benchmark value. Since there is exactly one contested piece and the piece is 'won' by agent 2, we call SubCore on piece a and agent 2 with the cake to the left of a 's trim ignored. Hence agent 2 gets a up till 1's trim.

We now increment m to handle agent 3. If agent 3's most preferred piece among the four pieces is not one currently held (even partially) by 1 and 2, then we can simply give 3 that piece. Otherwise, all agents 1, 2, 3 are asked to place a trim over the contested pieces among $\{a, b\}$ to make the value equal to the most preferred piece among c and d . Let us assume that 3 places the rightmost trim over b and over a . In that case, $W = \{3\}$. The fact that 3 is in W means that there exists a neat allocation for the agent in $W = \{3\}$ where he can get one of the pieces which are contested. We now ignore the trims of agent 3 on pieces a and b and ignore the left side of the rightmost trim by agents in $\{1, 2\}$. We call SubCore Protocol on the pieces a, b with the left side of 1 and 2's trims ignored and with agent $W = \{3\}$ as input. It returns a neat allocation (with respect to the current left margins on pieces beyond which they are ignored) for agent 3 where 3 will be assigned one of the pieces a or b up to the non-winner (not in W) agent's rightmost trim. Let us say 3 gets piece a . Since the allocation is neat (with respect to the current left margins), 3 does not envy any of the other pieces up to the non-winner's trims, meaning a new agent can be allocated a contested piece. Let us say this agent is agent 1, and that 1 is now the rightmost trimmer of piece b and his trim coincided with the current left margin of b beyond which b is ignored. 1 is added to W and 1 is allocated b up till the current left margin. Since now $|W| = m - 1$, we can make one last recursive call for agents 1 and 3 with the left part of the trims by agent 2 is ignored. This allocates piece a to 3 and piece b to 1. The uncontested pieces c and d are untrimmed. Agent 2 can be seen as having been 'kicked out' of the contest, and therefore is left with his most preferred uncontested 'benchmark' piece, let us say c . Piece d has remained unallocated and untrimmed. At this point the allocation of the SubCore Protocol is returned.

3.1.1. Dealing with ties. To avoid having to mention the special case of ties in each argument concerning the Core/SubCore protocol, we will deal with them in a modular way.

To this end we introduce what we will call imaginary values. Each agent has a list of such values. Each agent's values are on a different scale (for example agent 2's values are much smaller than agent 1's). Our aim is to define those values so that no agent is indifferent between 2 pieces, and so that agents never place trims on the exact same location (where locations with different imaginary values are distinct). We do so whilst ensuring that when agents equalise pieces of cake, they still prefer the pieces that were trimmed to the one they used as a benchmark to equalise.

To do so generate for each agent j $E = n^3(n^2)^n$ values $\epsilon_1^j, \epsilon_2^j \dots \epsilon_E^j$ with the properties that $\epsilon_i^j \ll E \times \epsilon_{i+1}^j$ and $\epsilon_E^j \ll E \times \epsilon_1^{j+1}$. Let L_j be j 's list of these epsilon values ordered from largest to smallest.

The protocol now internally represents each piece of cake as a number with a physical part and imaginary part. All the cuts that are actually made on the cake are in reality only dependent on the physical part of the value. The imaginary part is used internally by the protocol to break ties in a consistent way. Agents only look at the imaginary value of a piece of cake when required to break a tie on the physical part. Agents do not disagree on the imaginary value of a piece. This means that if we query agent l and j on a piece, they might disagree on the physical value of the piece, but will agree on the imaginary value.

Let us now explain how those imaginary values are associated with the pieces generated by the protocol. When the Core protocol asks the cutter to cut n equal pieces, the cutter will add the ϵ from the top of his list L every-time he creates a new piece and then discard that ϵ from L . Every other agent also adds

ϵ values to the pieces. In the SubCore protocol, agents generate new pieces of cake when placing trims on larger pieces to make them equal to a smaller piece. When an agent is asked to make larger pieces equal to a smaller piece b , he will add the epsilon value at the top of his list L to the piece he least preferred before the trim (excluding b), then discard that value from L . He then proceeds to add the second value to the second least preferred piece, and so on until he reaches the most preferred piece that he trimmed. This ensures that while the physical value of the pieces becomes the same from the agent's perspective, the actual order of the pieces remains unchanged.

CLAIM 1. *Agents have a strict order on the pieces that were generated by the cutter. The order between two pieces $a >_i b$ can only change if i 's valuation of the physical part of a becomes strictly less than his valuation of the physical part of b .*

Proof. New pieces are generated in the SubCore protocol when an agent j is asked to equalise pieces $c_1 \dots c_l$ to a piece c_z . The pieces will now have equal physical value from the perspective of agent j but since the former larger pieces are allocated bigger imaginary values, the order is preserved. \square

CLAIM 2. *When taking into account both the physical and imaginary part of the valuation of pieces an agent cannot have the same value for 2 pieces. Moreover, 2 agents cannot place a trim on the same point (they may place on the same physical point but the imaginary value breaks the tie). In other words, with imaginary values taken into account, ties do not happen in the protocol.*

Proof. Each ϵ value is unique and cannot be obtained by recombining the others through addition and subtraction, which are the only operations available to the protocol. \square

EXAMPLE 2. *Imagine that agent j is asked to equalise piece a to equal b . Let us have $v_j(a) = 0.2$ and $v_j(b) = 0.1 + 0.001i$ where i just represents the fact that it is an imaginary value. j will trim a so that it is equal to $0.1 + 0.001i$, then add the next ϵ value ϵ_l on top of its list to piece a so that $v_j(a) = 0.1 + (0.001 + \epsilon_l)i$.*

DEFINITION 4 (Snapshots). *When we run the Core Protocol we end up with an envy-free partial allocation of the cake and a residue. We call the partial allocation a snapshot or a core snapshot.*

The algorithm will keep track of certain snapshots which we will label p_j with $j \in \{1, \dots, C'\}$. The pieces allocated in each snapshot p_j are labelled c_{jk} where $k \in \{1, \dots, n\}$ indicates which agent got the piece. Since the allocation in each snapshot is envy-free, each agent thinks he got at least as much value for his piece as any other allocated piece. Therefore each agent thinks he has some (possible zero or more) bonus value over another agent in the Core snapshot. Our protocol will make use of these bonuses. In our protocol, we repeatedly call the Core Protocol over the resultant residue thereby making the residue smaller until the set of snapshots have enough structure that we will exploit later.

3.2. Groundwork for the Main Protocol.

Parameters. We will use four parameters to represent the bounds we work with. These are

$$C \ll C' \ll B \ll B'.$$

These are constant once n is fixed, however they are dependent on n . Since the Main Protocol calls itself on a strict subset of the agents, we will use the notation B'_{n-1} to label the bound on the Main Protocol run on $n - 1$ agents. The bounds correspond to the following concepts.

- C' is the number of snapshots generated by the main algorithm that we label and keep track of. All subsequent partial allocations generated by runs of the Core Protocol simply serve the purpose of making the residue smaller from the cutter's perspective. Once the main algorithm has extracted all the pieces it needs from the residue and if a discrepancy has not been successfully exploited, the GoLeft Protocol will be run.
- When we run GoLeft we look at a subset of C isomorphic snapshots from the C' total snapshots.
- The value B' corresponds to the total number of queries required to run the whole protocol.
- B is used to define what we call significant pieces or values. Running the Core Protocol guarantees that a piece of cake is made smaller from the perspective of the cutter. The value B is a bound on the number of times the algorithm allows us to run the Core Protocol to make the residue smaller.

For the sake of achieving our boundedness results, the following values of the parameters work. (i) $C = n^{n^n}$ (ii) $C' = n^{n^{n^n}}$ (iii) $B = n^{n^{n^n}}$ (iv) $B' = n^{n^{n^{n^n}}}$. We have not optimized the values of the bounds so we expect that our general algorithmic approach works for better bounds.

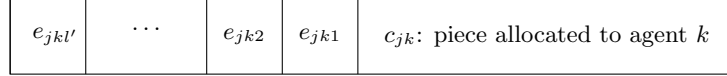


Fig. 3: For any allocated piece in some Core snapshot, we consider extracted pieces to be attached to the piece. In the figure, piece c_{jk} is allocated to agent k in snapshot p_j . The pieces $e_{jk1}, e_{jk2}, \dots, e_{jkl'}$ are extracted from the residue and are in consideration for attachment to the piece c_{jk} in the order $e_{jk1}, e_{jk2}, \dots, e_{jkl'}$.

In the protocol we often want to make a piece of cake (the residue) smaller. To do so we run the Core Protocol on it. To make descriptions more succinct we define a function which converts the number of times we run the Core Protocol on a piece of cake to how much smaller it is from the cutter's perspective. Note that when the Core Protocol is run, at least $2/n$ value of the cake is allocated to the agents because the cutter cuts the cake into n equally preferred pieces and at least two pieces are fully allocated. By running the Core Protocol B times with the same cutter, the cutter thinks only $f(B) = (\frac{n-2}{n})^B$ value of the cake is unallocated.

DEFINITION 5 (Bound function f). *Let $f(B) = (\frac{n-2}{n})^B$.*

For any piece of cake c in a given snapshot, each agent has a bonus value which corresponds to how much more cake he thinks he got in that snapshot than he would have got had he been allocated c instead.

DEFINITION 6 (Bonus value). *An agent i 's bonus value on piece of cake c_{jk} in snapshot p_j is the value $V_i(c_{ji}) - V_i(c_{jk})$, where c_{ji} is the piece that was allocated to agent i in p_j by the Core Protocol.*

In a Core snapshot, a bonus value is *significant* for an agent if we can make the residue smaller than that value from the agent's perspective in a bounded number of steps. Note that the significance of a piece of cake does not depend on the absolute value that an agent ascribes to it but it is relative to the value of the unallocated cake (residue).

DEFINITION 7 (Significant value). *An agent i thinks a value is significant if the value is more than or equal to $V_i(R)f(B)$ where R is the unallocated residue. A piece is significant for an agent if it has significant value for him.*

Note that if an agent i finds a piece of cake significant, he will still find it significant if the residue becomes smaller than before. We also observe the following about the Core Protocol and the cutter getting a significant advantage over at least one other agent.

REMARK 1. *Note that when the Core Protocol is run once, the cutter has a significant bonus over at least one agent (the agent who gets the smallest valued piece from the cutter's perspective). This bonus results in a domination of the cutter over the agent in $k = (\log n)(\frac{n-2}{n}) + 1$ iterations of the Core Protocol on the residue with same cutter. The reason is that the cutter i has advantage at least $V_i(R)/(n-2)$ over the agent who gets the smallest valued piece from i 's perspective. This is the worst case when equal value of residue comes from each of the maximum of $n-2$ trimmed off pieces. In k iterations of the Core Protocol with i as cutter, agent i 's value of the residue is $V_i(R)(\frac{n-2}{n})^k$ which is less than $V_i(R)/(n-2)$.*

If we run the Core Protocol repeatedly with the same agent as cutter and we are lucky that the cutter has a significant advantage over each agent in some snapshot, then in $(\log n)(\frac{n-2}{n}) + 1$ extra iterations of the Core Protocol with the same cutter, we can make the cutter dominate each agent. This simplifies our problem, because we now only need to allocate the remaining cake among the dominated agents in an envy-free manner. In general, we may not be so lucky that the cutter dominates all other agent due to which we have to do more work to ensure that a set of agents dominates the other agents. In this additional work, the process of extraction is crucial:

Extraction. For each piece allocated in each of the Core snapshots and for each agent, the Main Protocol tries to associate a piece of the cake extracted from the residue. For the C' snapshots, there are $C'n$ pieces allocated to the agents. For each of the pieces, other agents who do not get the piece *extract* corresponding pieces from the residue that is the unallocated cake. Take for example, a piece c allocated to some agent (say agent 1) in a snapshot. In the snapshot in which c is allocated to 1, each agent i gets a piece that is of value (according to i) at least as much as c . In the residue, each agent i is asked to put a trim so that the cake from

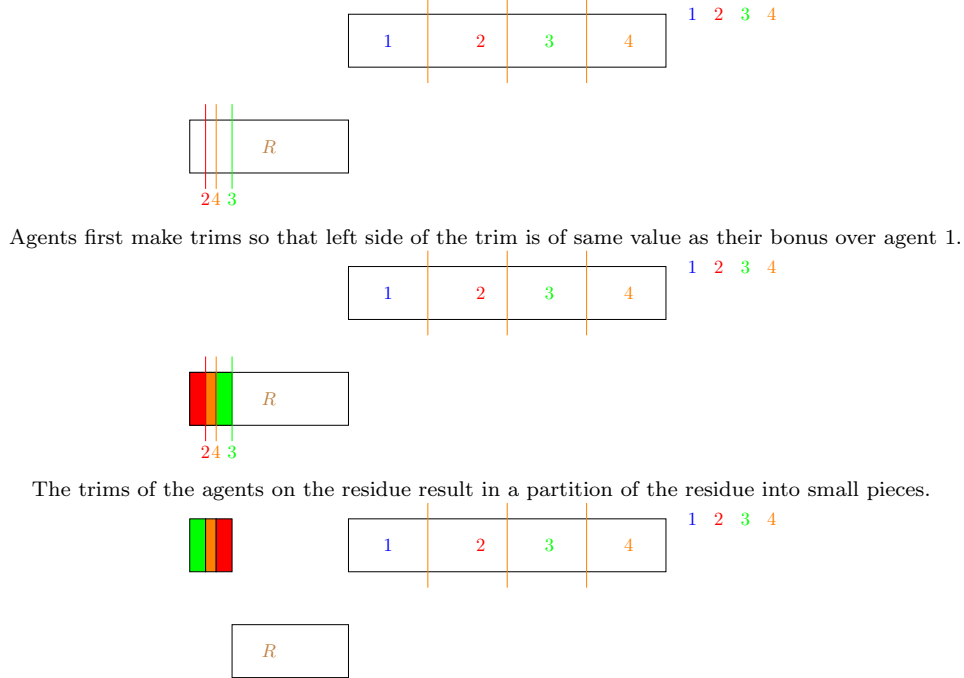


Fig. 4: Process of extraction for piece allocated to agent 1 in the Core snapshot

the left extreme of the residue to the trim is of value equal to i 's value of his piece minus his value for c . The trims of the agents on the residue give rise to pieces of cake corresponding to the agents' bonus value over piece c . These pieces are *extracted* (cut away from the residue) and *associated with c* (kept in consideration with piece c). When agents make trim marks on the residue in accordance with their bonus (over a piece in a snapshot), each pair of successive trim marks gives rise to a separate piece that can be extracted. For piece c_{jk} in snapshot p_j , we will denote by $e_{jk1}, e_{jk2}, \dots, e_{jkl'}$ the set of pieces that are extracted in the same order with e_{jk1} extracted first (see Figure 3). We say that a piece e_{jkl} is extracted by agent i if the right hand extreme of the piece coincided with the trim of agent i on the residue. Note that $l' \leq n - 1$ because for each allocated piece in a snapshot, at most $n - 1$ other agents can put trim marks on the residue so as to obtain at most $n - 1$ extracted pieces. Each extracted piece has a clear corresponding allocated piece in a Core snapshot with which it is associated. Note that it can be the case that not all $n - 1$ agents place a trim we allow extraction only for agents whose bonus values are not significant.

In our protocol, pieces are extracted from the residue only if each agent finds the pieces *not* significant. The reason is that we want to keep sufficient residue to do further extractions on the residue as well as maintain structure.

EXAMPLE 3. *Figure 4 illustrates the process of extraction and association. We focus on a single Core snapshot in which each of the four agents are allocated a piece. Since the Core allocation is envy-free, each of agents 2, 3, 4 think they got at least as much value as the piece that 1 got. Now agents 2, 3, and 4 are asked to place a trim mark each on the residue to indicate that the residue to the left of their trim mark is equal to their advantage over agent 1. The trim marks result in three slices of cake that are extracted. Among the extractions, since 2 has the leftmost trim, his piece is extracted first, then 4 and then 3. The extracted pieces are associated with 1's allocated piece in the snapshot. In case the pieces will be added to 1's piece, the piece extracted by agent 2 will be attached first, then the piece by agent 4 and then 3.*

Later on, the extracted pieces may be attached to c so that piece c is now attractive to other agents because of the additional extracted pieces combined with c . We clarify that when pieces are extracted and associated with a given piece in a Core snapshot, such extracted pieces have not yet been allocated to any

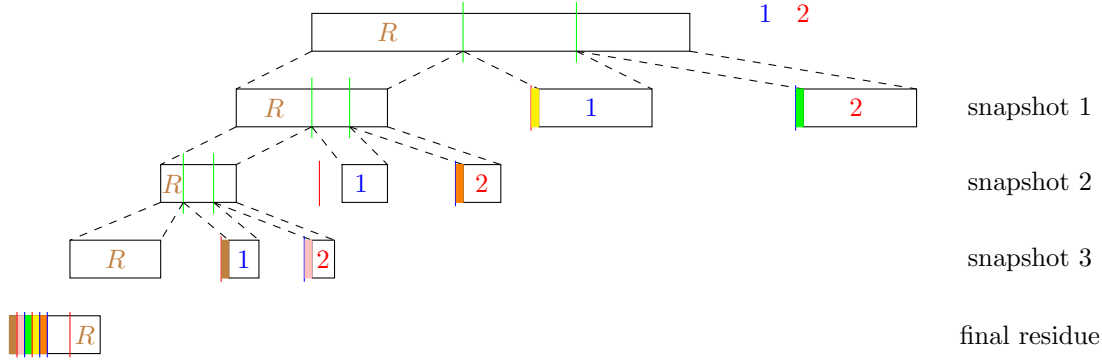


Fig. 5: Snapshots with extracted pieces. The Core Protocol is called thrice resulting in three snapshots. The Core Protocol is always called on the residue—cake that has not yet been allocated. Since the example involves only two agents, it is not the order of trims in extraction but only whether a trim induces a significant piece or not which determines isomorphism. Snapshot 1 and 3 are isomorphic since for each allocated piece in the snapshots, the other agent was able to extract a piece from the final residue to reflect his bonus/advantage. For the second snapshot a piece was not extracted because agent 2’s bonus value was significant. The extracted pieces present and the order in which they were extracted is the same for snapshot 1 and 3, making them isomorphic. Note that this figure is just to demonstrate isomorphism. In reality, for two agents, the Core Protocol coincides with Divide and Choose which means that there is no residue.

particular agent. Extracted pieces can only be allocated after they are officially attached to their associated piece.

For the protocol we need to restrict our focus on a subset of the snapshots where the same set of agents extracted pieces in the same order:

DEFINITION 8 (Isomorphic snapshots/snapshot pieces/ extracted pieces).

We call two snapshots p_j and $p_{j'}$ isomorphic to each other if for each agent $i \in N$ and for any two pieces c_{ji} and $c_{j'i}$ allocated to agent i in the corresponding snapshots p_j and $p_{j'}$, the set of agents who extracted cake from the residue and associated to c_{ji} and $c_{j'i}$ is the same and the agents in the set extracted pieces in the same order.

We extend this notion to allocated pieces in snapshots. Two allocated pieces of cake belonging to two isomorphic snapshots are isomorphic if they were allocated to the same agent.

We also extend the notion of isomorphism to extracted pieces. We say that for two isomorphic snapshots p_j and $p_{j'}$, two extracted pieces e_{jkl} and $e_{j'kl}$ are isomorphic if they are associated respectively to isomorphic pieces c_{jk} and $c_{j'k}$ and that e_{jkl} and $e_{j'kl}$ were extracted by the same agent.

In Figure 5, we illustrate isomorphic snapshots.

Our protocol will keep track of a set of isomorphic snapshots, progressively discarding (not changing) some so that we can make manipulations on the ones we keep in an envy-free way. The pieces of cake we are working with are labelled c_{jk} for allocated pieces or e_{jkl} for extracted pieces. Suppose that the given set of snapshots is S . We use the notation c_k, S and e_{kl}, S to denote the set of pieces of cake c_{jk} or e_{jkl} in snapshots in S . Abusing the notation, we will simply say c_k or e_{kl} to mean the set of pieces which are in the snapshots we are working with. Note that we will generally use index j for the snapshot number, k for the piece number in a given snapshot, and l for the l -th extracted piece.

3.3. The Main Protocol. The Main Protocol is the engine which runs our overall envy-free protocol. It is responsible for allocating the whole cake among all the agents in an envy-free manner. The Main Protocol works recursively. If the number of agents is four or less, a previously known bounded envy-free algorithm can directly be called. Otherwise, the Main Protocol divides part of the cake in an envy-free manner and identifies a set of agents $N \setminus A$ that all dominate agents in A with respect to the cake that is unallocated. The Main Protocol then recursively calls itself to divide the remaining cake among agents in A .

Algorithm 3 Main Protocol**Input:** Cake R and a set of agents N with $n = |N|$.**Output:** An envy-free allocation that completely allocates R among agents in N .

Base Case

1: **if** $|N| \leq 4$ **then**
2: allocate the residue among the agents in N using one of the known constant-time envy-free protocols for $|N| \leq 4$.
3: **else**

Generate Core Snapshots

4: **for** all $j = 1$ to C' **do**
5: For some $i \in N$ who has acted as cutter in the Core Protocol the least number of times, run Core protocol(i, N, R); Update R to the cake that is unallocated. The Core Protocol allocation gives us snapshot p_j and n pieces of cake c_{jk} with $k : 1, \dots, n$.
6: **end for** {We have now generated C' Core snapshots. After generating the first n snapshots, we already have an envy-free allocation in which each agent gets $1/n$ value of the original cake.}
7: **while** for agent $i \in N$ and some piece c_{jk} , $V_i(R)f(B)^2 < V_i(c_{ji}) - V_i(c_{jk}) < V_i(R)f(B)^{\frac{1}{2}}$ **do** {We ensure that the bonus values are either larger than significant or smaller than significant by a large factor dependent on our bound B .}
8: Run Core Protocol(i, N, R).
9: **end while**
10: **if** some set of agents $N \setminus A$ dominates all agents in set A **then** Call Main Protocol(R, A) to divide the unallocated cake R .
11: **return** allocation of the cake to the agents.
12: **end if**

Extraction

13: Define Boolean a and set it to 0 {This Boolean is used to reset the whole trimming process when running the Discrepancy Protocol causes a bonus value to become significant}
14: **while** $a = 0$ **do**
15: {We now attempt to extract pieces corresponding to the non-significant bonus values from the residue, resetting the whole process when a discrepancy in agents' valuations has forced us to shrink the residue}
16: Set a to 1.
17: **for** all pieces of cake c_{jk} **do**
18: **for** all agents i **do**
19: **if** i 's bonus value is not significant on piece c_{jk} **then**
20: Ask i to place a trim on the residue such that the piece stretching from the left of the residue to the trim is equal to the agent's bonus value for that piece, which we will label $b_{c_{jk}}^i$.
21: **end if**
22: **end for**
23: Label from left to right the pieces delimited by those trims with labels e_{jkl} with $l : 1, \dots, m$ where m is the number of agents who placed a trim on the residue (did not have significant bonus) and j and k are used as previously {if two or more agents' trims coincide we break ties lexicographically and ascribe empty cake to the latter agents' extractions.}
24: {For a given piece of cake c_{jk} the m agents with non-significant bonus value have now delimited a piece of cake corresponding to that value from the residue, and those pieces have been labelled.}
25: **for** all $l : 1, \dots, m$ **do**
26: **if** no agent thinks the piece e_{jkl} is significant **then** {Everyone agrees piece is insignificant}
27: Extract piece e_{jkl} from the residue, associate it to piece c_{jk} and update R to $R - e_{jkl}$ {See Figure 4}.
28: **end if**

Discrepancy

29: **if** some agents think e_{jkl} is significant but others do not (agents think the piece is discrepant) **then**
30: {There is a large discrepancy between agents' valuation, this causes a problem because some agents might think we are taking too much away from the residue. However discrepancy can be exploited or eliminated.}
31: Run Discrepancy($e_{jkl}, b_{c_{jk}}^u, \{e_{jkl}\}, R$) with the discrepant piece e_{jkl} , the bonus value $b_{c_{jk}}^u$ of the agent u who made the rightmost trim delimiting the piece, the set of previously extracted pieces of cake $\{e_{jkl}\}$ and the residue R as input. This will return the Boolean value DISCREPANCY and when that value is 1 the protocol also returns two sets $D \subset N$ and $D' \subset N$ such that D and D' partition N .
32: **if** DISCREPANCY = 1 **then** {The discrepancy can be exploited}
33: Main Protocol(e_{jkl}, D).
34: Main Protocol(R, D').
35: **return** allocation of the cake to the agents.
36: **else** {The discrepancy cannot be exploited but $b_{c_{jk}}^u$ is now a significant bonus value.}
37: Add e_{jkl} to R and Set a to 0.
38: **end if**
39: **end if**

GoLeft

40: Run GoLeft($N, \{p_j\}, \{c_{jk}\}, \{e_{jkl}\}, R$) Protocol with the set of labelled snapshots $\{p_j\}$, the set of allocated pieces of cake $\{c_{jk}\}$, the set of extracted pieces $\{e_{jkl}\}$ and the residue R as input. The output of GoLeft is a set $A \subset N$.
41: **for** all agents i **do**
42: Run Core Protocol(i, N, R) $2B$ times each time on the updated smaller residue R .
43: **end for**
44: {The previous for loop converts significant bonus values into dominance with respect to the residue R .}

Recursion of the Main Protocol

45: Call Main Protocol($R, A \subset N$), that is the Main Protocol on a subset of agents $A \subset N$ output by the GoLeft Protocol and with R as the input cake.
46: **end if**
47: **return** allocation of the cake to the agents.

If the number of agents is more than four, the Main Protocol (Algorithm 3) calls the Core Protocol sequentially on the updated residue so that the residue becomes smaller than before after each call of the Core Protocol. The repeated calls of the Core Protocol help generate a number of snapshots each containing n pieces of cake. In each snapshot, each agent has been allocated a piece of cake. Envy-freeness is maintained throughout for the allocated cake, and in the Core Protocol each agent thinks he got the highest value piece.

REMARK 2. *When we call the Core Protocol n times each time with a different cutter, we argue that we get an envy-free allocation in which agents get at least $1/n$ value of the original cake.*

Suppose an agent i is the cutter in the $j \leq n$ -th call of the Core Protocol. Then in the first $j - 1$ calls of the Core Protocol, i gets $1/n$ value of the allocated cake. This follows from the envy-freeness of the allocated cake. For the remaining cake, i again gets $1/n$ value of the unallocated cake in the j -th call because i is the cutter. Hence i gets $1/n$ value of the original cake. Since the cake is always allocated via Core, it is allocated in an envy-free manner. Hence after the Main Protocol has made the first n calls of the Core Protocol, we have an envy-free allocation in which agents get at least $1/n$ value of the original cake.

After the first n calls to the Core Protocol, the Main Protocol does not stop making calls to the Core Protocol if there is still some unallocated cake. After calling the Core Protocol C' times, C' Core snapshots are obtained. The Core Protocol may be further called (in the while loop in step 7) to make the residue even smaller. This ensures that each agent considers each piece in the first C' snapshots significant or smaller than significant by a large factor dependent on our bound B .

For each piece of cake c in the C' snapshots, agents can ascribe what we refer to as a *bonus value*, which corresponds to how much more value they got in that snapshot than their value for piece c . In the Main Protocol, for each piece of cake c in the snapshots, we ask all agents with a non-significant bonus value to make a cut on the residue equal to their bonus value. The pieces obtained from these cuts are then taken from the residue and associated (but not yet attached) to piece of cake c . We refer to this process as extraction. A piece is extracted only if all agents find it insignificant. The extracted pieces will potentially be attached to their associated piece of cake c in the GoLeft Protocol. Most of them however will be sent back to the residue or shared amongst a subset of agents in an envy-free way. The process of attaching the extracted pieces to their associated piece in the GoLeft Protocol is designed to make that piece desirable to the agent whose bonus value was used to extract the piece from the residue. The Main Protocol also calls the Discrepancy Protocol in case there is some piece in consideration for extraction that some agents consider significant and others do not. The goal of the Discrepancy Protocol is to exploit any such discrepancy and to ensure that when the GoLeft Protocol is called by the Main Protocol, then there is no discrepancy in how the extracted pieces are viewed, i.e., no actually extracted piece is considered significant by some agent.

3.4. Discrepancy Protocol. When pieces are being extracted from the residue during the Main Protocol, it may be the case that one of the pieces e_{jkl} in consideration for extraction is significant for some agent. In that case, the piece is not extracted and the Discrepancy Protocol (Algorithm 4) is called in line 28 that either exploits or ‘eliminates’ this discrepancy. The discrepant piece e_{jkl} is kept aside from the residue. On the other hand, all previously extracted pieces are added back to the residue. Since the difference between a piece that is just above significant or just below significant can be arbitrarily small, the Core Protocol is used to create a *gap* so that the discrepant piece either has value at least $V_i(R)n$ or value at most $V_i(R)/n$. This gap is highly useful because the goal of the Discrepancy Protocol is to either ensure that (1) everyone thinks that the discrepant piece is significant or (2) the problem of finding an envy-free allocation can be broken into two sub-problems where some agents D are allocated the discrepant piece and the rest D' are allocated the residue. In this, we use the fact mentioned earlier that envy-freeness implies proportionality.

In case of (1), the process of extraction in the Main Protocol is reset and the discrepant piece as well as all the pieces that had been extracted are sent back to the residue. This may appear to be a waste of work but when we do this, we have ensured that at least one agent has significant advantage over another’s piece for a given snapshot. This ‘setback’ can only happen $C'n^2$ times before each agent dominates each other agent in which case the remaining cake can be allocated arbitrarily without causing envy.

We note that Discrepancy makes the residue smaller as it calls the Core Protocol. When the residue becomes smaller, it may be that a piece of cake that was not significant for an agent becomes significant because significance is defined with respect to the residue. Hence any agent i who thinks that $\frac{V_i(R)}{n} \leq V_i(e_{jkl}) \leq V_i(R)n$ will eventually think that $V_i(e_{jkl}) \geq V_i(R)n$ when the Core Protocol has been run to reduce the residue.

Algorithm 4 Discrepancy Protocol

Input: Residue R , *discrepant* piece e_{jkl} , bonus value $b_{c_{jk}}^u$ on c_{jk} of the agent u who wanted to extract the discrepant piece (but could not extract), set of extracted pieces of cake $\{e_{jkl} : \text{extracted pieces}\}$, agent set N .

Output: Possibly modified residue R , Boolean value called DISCREPANCY, set of agents D , set of agents D' .

1: Take out the discrepant piece e_{jkl} from R ; Reinsert all the *extracted* pieces back into the residue, relabel the aggregate piece R .

Run Core

2: Run Core Protocol(i, N, R) for every agent i B times iteratively on the updated residue R

3: **while** for some agent i , it is the case that $\frac{V_i(R)}{n} \leq V_i(e_{jkl}) \leq V_i(R)n$ **do**
 { *The piece might be discrepant but agents who do not classify the piece the same might have arbitrarily close valuation of the piece if they are lying next to the bound. This while loop ensures that there is a gap between agents who think the piece is significant and others* }

4: Run Core Protocol(i, N, R) B times iteratively on the updated residue R .

5: **end while**

Exploit Discrepancy

6: **if** e_{jkl} still has some agents consider it significant and others not **then** { *In this case we can exploit the discrepancy and ‘separate’ the agents* }

7: Set D to $\{i \in N : V_i(e_{jkl}) \geq V_i(R)n\}$.

8: Set D' to $\{i \in N : V_i(e_{jkl}) \leq \frac{V_i(R)}{n}\}$.

{ *Because of the while loop line 3, we have no agents satisfying $\frac{V_i(R)}{n} \leq V_i(e_{jkl}) \leq V_i(R)n$; these inequalities come from proportionality. We want to ensure that agents sharing the residue think that even if someone else gets all of the discrepant piece (or residue), they will still be better off by getting their proportional share of the residue (or discrepant piece). Note that $N = D \cup D'$.* }

9: Set DISCREPANCY to 1.

10: **return** R , DISCREPANCY, D and D' .

11: **else**

Cannot Exploit Discrepancy — but all agents now on the same page

{ *This means that we cannot exploit a discrepancy but now all agents think that the bonus value of whoever made the trim delimiting the discrepant piece is significant, including agent u who made the trim* }

12: Set DISCREPANCY to 0.

{ *$b_{c_{jk}}^u$ is now considered significant by agent u* }

13: **return** R and DISCREPANCY.

14: **end if**

3.5. Groundwork for the GoLeft Protocol. The GoLeft Protocol is the heart of our overall protocol and is crucial to allocate the cake that is still not allocated. All the other steps in the Main Protocol can be viewed as preparing the ground for the GoLeft Protocol to work. By calling the Core Protocol a sufficient number of times in the Main Protocol, enough C' Core snapshots are obtained that are helpful to identify C isomorphic snapshots in the GoLeft Protocol (these isomorphic snapshots constitute the working set of snapshots over which the GoLeft Protocol operates). Moreover, by calling Discrepancy before GoLeft, it is ensured that all agents are on the same page: all agents consider all the extracted pieces as insignificant.

The GoLeft Protocol (Algorithm 5) is called by the Main Protocol. The goal of the GoLeft Protocol is to identify a set of agents $N \setminus A$ that dominate agents in A . This means that the remaining residue can be allocated among agents in A in an envy-free manner without worrying about agents in $N \setminus A$ envying them. The goal of the GoLeft protocol is achieved by attaching extracted pieces to the pieces in the working set of snapshots in a methodical manner while maintaining envy-freeness of the allocated cake. In order to maintain envy-freeness, the working set of snapshots is modified in various ways. When all extracted pieces for the working set of snapshots have been attached, we will show that we obtain a set of agents that all dominate the other agents. For example, if we end up with Core snapshots in which for one isomorphic allocated piece, there are a total of less than $n - 1$ extracted pieces that have all been attached and are held by a certain agent i , then agents who did not manage to extract pieces corresponding to the main piece have a significant advantage over i as well as all other agents who extracted pieces before i for that main piece. This significant advantage translates into dominance.

The GoLeft protocol operates on a working set of isomorphic snapshots. The main operations of the GoLeft Protocol are to attach extracted pieces to the Core snapshot allocation and to implement exchanges in which there is a sequence of agents a_0, a_1, \dots, a_{k-1} where each agent a_i in the sequence gets the pieces of agent $a_{i+1 \bmod k}$. When operations such as attachments and exchanges happen, the isomorphic snapshots change but remain isomorphic nonetheless.

Before we give an overview of the steps of the GoLeft Protocol, we establish some concepts and mathematical structures we will work with. The two key structures are the *working set S of isomorphic snapshots* and the *permutation graph* that is defined with respect to the working set of isomorphic snapshots. During

Algorithm 5 GoLeft Protocol

Input: Set N , Set C' of snapshots p_j , set of corresponding pieces $\{c_{jk}\}$, the set of extracted pieces $\{e_{jkl}\}$, and residue R .

Output: A set of agents $A \subset N$ such that all agents in $N \setminus A$ dominate all agents in A .

Isomorphic Snapshots and build Permutation Graph

- 1: Select from the C' snapshots a set of size C of *isomorphic snapshots* (this follows by Claim 3 which uses a pigeon-hole argument).
- 2: Relabel the C snapshots p_j with j now ranging from 1 to C .
- 3: Declare set S in which we add all C snapshots $\{p_j : j \in \{1, \dots, C\}\}$. *{Out of all the snapshots that the Main Protocol generated, we will focus on isomorphic snapshots. The isomorphic snapshots that are selected are put in a set S which will shrink in size as the algorithm proceeds.}*
- 4: Build the permutation graph with nodes corresponding to the agents. Throughout the algorithm, we maintain two sets: (1) T (set of nodes/agents such that the isomorphic pieces held by them in S have not had $n - 1$ attachments) and (2) T' (set of nodes/agents such that the isomorphic pieces held by them in S have had $n - 1$ attachments). T and T' partition the nodes. Initially, (1) all nodes are placed in T so that T' is empty; (2) each node points to itself in the permutation graph (3) each agent owns his originally allocated set of pieces allocated in the Core snapshots in the working set S . Thus each node has the following associated information: agent and his allocation in snapshots in S . *{The permutation graph will keep track of which permutation we can do next in order to make the protocol progress. If the allocated pieces associated with a node have had $n - 1$ attachments, the node is moved from T to T' .}*
- 5: **while** there is a node in T **do**

Cycle and Exchange

- 6: Find a cycle in the permutation graph which involves a node from T . *{Since nodes in T all have in-degree exactly 1 and each node points to each node in T' , such a cycle exists as proved in Claim 4.}*
- 7: For the agents in the cycle, exchange the allocation of the agents in the snapshots in S as follows: if i points to j , give i the pieces of cake that j was currently allocated in snapshots S .
- 8: If i points to j in the selected cycle, do as follows. Transfer all in-edges of j to i . Replace j in set T or T' by i . *{ This update reflects the permutation of the allocation. At this point, each node that was in the cycle has a self-loop. }*
- 9: Take a node/agent i in the cycle that is from T .

Separation - the only way the while loop exits

- 10: **if** there is a node in the located cycle that is from T (has had less than $n - 1$ attachments) but has no extracted pieces to be attached **then**
- 11: {Found a separation of agents where one set dominates the other. The reason no extracted piece is present is because all agents who have *not* had their corresponding extracted piece attached have a significant advantage over agents who already have had an extracted piece attached}. Focus on the set of isomorphic pieces C associated with the node. All agents who have been given an element of C are placed in A .
- 12: **return** A .
- 13: **else** we now know that there is an agent i that was in the cycle from T that holds pieces in S that still have extracted pieces to be attached. We focus on this agent i in the attachment phase.
- 14: **end if**

Attachment — (to be done in a subset of the snapshots)

{The next part of the protocol will attach in a subset of the snapshots the set of isomorphic extracted pieces in $e_{k(l+1)}$ to the set of pieces c_k , thus making pieces in set c_k desirable to the agent who extracted pieces in $e_{k(l+1)}$. The set of pieces c_k and its current attachments are currently owned by agent i . At this point for each isomorphic piece c_k , associated pieces up till i 's trim have already been attached.}

- 15: **Index the agents in the order which they extracted the pieces associated to c_k** so that 1 originally got c_k , 2 made the next extraction and so on. In doing so, we index agent i as l .
- 16: Declare $S' \leftarrow \emptyset$.

Attachment Subphase A — making it agreeable for agents from $l + 1$ to n

- 17: **for** agent i ranging from $l + 1$ to n **do** *{Ask agent i who does not (so far) find pieces in c_k desirable to discard enough snapshots where being allocated c_k instead of whichever piece he is holding would be most disadvantageous. In other words, he reserves advantage over agents in $\{1, \dots, l\}$.}*
- 18: Ask i to choose $\frac{|S|}{n-l+1}$ snapshots from S for which i values the difference between his bonus value for c_k and the extracted pieces currently attached to c_k the most.
- 19: Remove the chosen snapshots from S and add them to S' .
- 20: **end for**
- 21: All extracted pieces that have not been attached from snapshots in S' are put back into the residue R . The new aggregate piece is labelled R . *{The pieces that were associated to a piece in a discarded snapshot but not attached will never be attached, we therefore send them back to the residue. The fact that they were not significant means that the size of the residue is hardly affected by this}*

Attachment Subphase B — making it agreeable for agents from 1 to l

- 22: **for** all agents i ranging from 1 to l **do** *{In this part we ensure that agents $i \in \{1, \dots, l\}$ will not be envious of agent $l + 1$ even if $l + 1$ is allocated c_k with the additional $e_{k(l+1)}$ as well as the intermediate attachments.}*
- 23: Ask i to choose the $\frac{|S|n}{ln+1}$ snapshots from S for which he values the piece $e_{k(l+1)}$ the most.
- 24: Remove the chosen snapshots from S and add them to S'' .
- 25: **end for**
- 26: All extracted pieces that have not been attached from snapshots in S'' are put back into the residue R . The new aggregate piece is labelled R . *{Again extracted pieces that will never be attached are sent back to the residue}*
- 27: All pieces $e_{k(l+1)}$ in snapshots in S'' are aggregated into a piece a .
- 28: Run Main Protocol($a, \{1, \dots, l\}$).
- 29: $S'' \leftarrow \emptyset$.

Attachment — now happening

- 30: We now attach the piece of cake $e_{k(l+1)}$ to c_k in the snapshots still in S , meaning that if agent $l + 1$ were to move to piece c_k , he would also get $e_{k(l+1)}$. In other words c_k is now desirable to $l + 1$ because of the attachments. However in snapshots where we keep agent l as the agent allocated piece c_k , the piece $e_{k(l+1)}$ is sent back to the residue, as it could not be given to l in an envy-free way *{This update reflects that we made a set of allocated pieces desirable to a new agent}*.
 - 31: Remove the self-loop of i . Replace it with an edge going from the agent who extracted the piece that the protocol just attached to c_{jk} .
 - 32: If the pieces held by i have had $n - 1$ attachments, delete i from T and place it in T' and make every node point to i .
 - 33: **end while**
-

the GoLeft Protocol, each structure gets updated based on the information on the other structure. The algorithm makes progress when agents exchange their pieces with each other.

Working Set of Isomorphic Snapshots. During the course of the GoLeft Protocol, the isomorphic snapshots in S get changed in the following way: (1) some subset of S is removed from S (2) agents exchange their pieces and (3) extracted pieces are attached to pieces in the snapshots in S .

When we update S , we maintain isomorphism and other invariant properties as follows. If an agent holds a piece c_{jk} , he holds the whole set c_k of isomorphic pieces in set of snapshots S . Note that we are abusing notation here as c_k refers to a set of pieces allocated to agent k in the working set of snapshots S . If an agent holds an extracted piece e_{jkl} , he holds the whole set e_{kl} of isomorphic extracted pieces associated with S as well. When we implement an exchange, we are essentially making $|S|$ exchanges – one in each of the snapshots. These are exchanges that not only involve the pieces in the snapshots but also involve those extracted pieces that have been *attached* to the pieces. When we attach an extraction to a piece in a snapshot in S , we simultaneously attach isomorphic extractions to the corresponding isomorphic allocated piece in each of the snapshots.

We also ensure that the extracted pieces are attached in the appropriate order so that each associated piece gets attached after the previous associated pieces have been attached. For example, if we were focussing on snapshot in Figure 4, the extracted associated piece due to agent 3 will get attached after the extracted associated piece due to agent 4. Note that if we attach an extracted piece e_{jkl} to an allocated piece c_{jk} (along with its previously attached extracted pieces) in a snapshot in the working set S , we perform a similar attachment for all such isomorphic extracted pieces in the set e_{kl} to their corresponding pieces in set c_k . Also, if an agent i currently holds an extracted piece, he also holds all earlier extracted pieces as well.

We make a crucial point about a consistency condition that we enforce.

REMARK 3. *We enforce a consistency condition whereby an agent cannot hold extracted pieces beyond the extraction he himself made. Therefore, when we attach an extracted piece to agent j 's piece in a snapshot to attract agent i to it, agent j does not actually hold the latest attachment because it is beyond j 's extraction. However in an exchange, if i were to get j 's piece along with its attachments, then i will also get the latest attachments that were originally extracted by i himself.*

In the GoLeft Protocol, we operate on an increasingly small set of isomorphic snapshots. The goal is to reach a set of isomorphic snapshots that can be used to find a set of agents who all dominate other agents. As the GoLeft Protocol proceeds, we discard snapshots to allow us to focus on others where extracted pieces from E have been ‘attached’ in an envy-free way to the allocated pieces in C to which they were associated. The set of isomorphic snapshots S becomes smaller when a set of isomorphic extractions $e_{k(l+1)}$ are attached to set of isomorphic pieces c_k in the snapshots S with each particular extracted piece in $e_{k(l+1)}$ getting attached to its corresponding piece in c_k . By *discarding* a set of Core snapshots, we mean that these snapshots and their allocations are not further worked upon and their associated unattached extracted pieces are sent back to the residue. Intuitively, the purpose of discarding snapshots will be to preserve the remaining advantages of agents over other agents.

The GoLeft procedure gradually attaches the extracted pieces to the allocated pieces in the working set of snapshots. The reason we call the protocol ‘go left’ is because we can visualize that for each piece in the set of snapshots we work with, we want to add the next extracted pieces to the isomorphic allocated pieces that are kept to the left of the isomorphic allocated pieces (see Figure 6). By attaching the next extracted pieces, we are ‘going left’.

How do we know which set of isomorphic pieces in S can have their next extraction? Which set of agents can exchange their currently held isomorphic pieces along with their current attachments? For this we work in tandem with the permutation graph.

Permutation Graph. The permutation graph keeps track of which agent is willing to move to which piece in the snapshots. The high level idea is that nodes of the permutation graph correspond to the agents and an agent i points to another agent j if he will be as happy taking j 's allocated pieces along with the attachments on those pieces.

DEFINITION 9 (Permutation graph). *The permutation graph is a directed graph where the set of nodes correspond to the set of agents. Hence when we refer to the graph, we will use agents and nodes interchangeably. The arcs of the permutation graph depend on the current state of the working set of isomorphic snapshots S . In particular, they depend on which extracted pieces have been attached to the originally allo-*

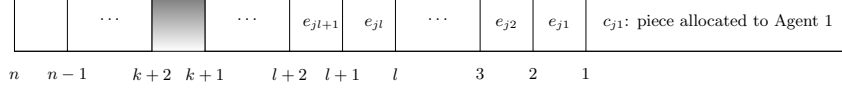


Fig. 6: Illustration of the GoLeft Protocol on a particular piece of cake that is originally allocated to agent 1 and in which the pieces were extracted in this particular case by agents in set $\{2, \dots, k + 1\}$ and in order $2, \dots, k + 1$. The piece c_{j1} was originally allocated to agent 1 in snapshot p_j . When we refer to ‘extractions up till agent l ’s extracted piece we will mean extracted pieces $e_{j1}, \dots, e_{j(l-1)}$. During the course of the GoLeft protocol, the pieces extracted by the agents $2, 3, \dots, k + 1$ will be attached to c_{j1} . This attachment can be viewed as going left. Agents $k + 2$ to n will not go left (have corresponding pieces extracted/attached) and are the prospective dominators because they find the shaded space between the trims of $k + 2$ and $k + 1$ significant. If during the course of the GoLeft Protocol, all the extracted pieces are successfully attached, then agents in $\{k + 2, \dots, n\}$ will dominate agents in set $A = \{1, \dots, k + 1\}$.

cated pieces in the isomorphic snapshots in S . Agent i points to agent j if j holds isomorphic pieces in S that have had all attachments up till i ’s extracted pieces. We build the initial permutation graph with each node pointing only to itself. Throughout the protocol, we ensure that each node in the permutation graph has in-degree at least one.

The permutation graph itself gets updated when isomorphic extractions are attached to a isomorphic pieces in S . The process of extracted pieces being attached to an allocated piece results in the aggregated piece becoming attractive to a new agent who then wants to point to the agent holding that piece (see Figures 7 and 8).

The permutation graph also suggests a natural way to exchange pieces. It has a similar idea as the trading graph used in top trading cycles algorithm for housing markets. If there is a cycle in the graph, we have the possibility of exchanging the allocations of agents in the cycle by giving an agent the piece of the agent he points to [35]. The permutation graph is more intricate because updates on the permutation graph reflect simultaneous updates on the working set of isomorphic snapshots S . Also by Remark 3, an agent in an exchange does not offer what he holds but can offer more because of the additional attachments beyond his own extraction.

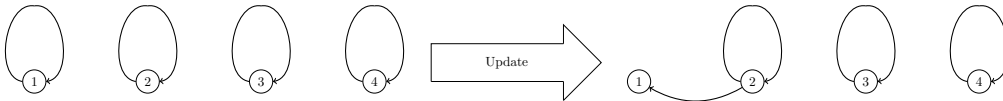


Fig. 7: Update on the permutation graph after the leftover isomorphic pieces of cake associated with node 1 and held by agent 1 have had their first associated piece of cake attached, making the pieces desirable to agent 2.



Fig. 8: Update on the permutation graph after the leftover isomorphic pieces of cake associated with node 2 and held by agent 2 have had their first associated piece of cake attached, making the piece desirable to agent 1. This results in a non-trivial cycle that allows for agents 1 and 2 to exchange their currently held pieces.

3.6. The GoLeft Protocol. Now that we have established the important ideas in the GoLeft Protocol, let us tie up the ideas and give an overview of how the protocol works. When the GoLeft Protocol starts, it

first identifies a working set S of C Core snapshots from out of the C' Core snapshot that we focus on. The protocol then constructs a permutation graph corresponding to the working set of isomorphic snapshots.

In the permutation graph, each node i corresponds to an agent i who holds a set of isomorphic pieces along with their attached extracted pieces in the working set of isomorphic snapshots S . We divide the nodes of the permutation graph into sets T and T' . Set T is the set of nodes/agents such that the isomorphic pieces held by them in S have not had $n - 1$ attachments). T' is the set of nodes/agents such that the isomorphic pieces held by them in S have had $n - 1$ attachments.

The protocol identifies a cycle in the permutation graph that includes at least one node i from T . Such a cycle always exists. In each of the working set S of isomorphic snapshots, we implement an exchange of pieces held by agents in the cycle: each agent in the cycle is given the piece corresponding to the node that the agent points to in the cycle. After implementing the exchange, the permutation graph is updated to reflect the exchange. In the exchange, if an agent gets an inferior piece, he always gets the additional extracted pieces associated with the inferior piece up till the agent's extractions. Hence in each snapshot in S , each agent's value from his allocation in the snapshot is preserved even if he gets a different piece than in the original Core snapshot. For any agent i , as long as no agent gets extracted pieces beyond i 's extraction, i will not be envious. In the GoLeft protocol, it can be the case that some agent j gets extracted pieces beyond i 's extracted pieces but before any such attachments in the last part of the GoLeft protocol, we ensure that no envy arises.

After implementing the cycle, we focus on a node $i \in T$ that was in the cycle. For agent/node i we know that for all snapshots in the working set S , agent i has been allocated the original isomorphic pieces c_k as well as all associated pieces up till i 's extracted piece. If the piece of cake agent i is currently allocated in the snapshots S has no more extracted pieces left to attach to it, but it has not had $n - 1$ attachments, this means that all agents who have not had their corresponding piece extracted/attached have a significant advantage over agents who have had an extracted piece attached. In this case, the GoLeft Protocol returns the set of dominated agents to the Main Protocol and we are left with a smaller envy-free allocation problem because it involves fewer number of agents.

In case node i does not lead to an exit from the GoLeft Protocol, we know that there are associated pieces that can still be attached to the isomorphic pieces held by i in the working set of Core snapshots S . We focus on the next set of associated pieces $e_{k(l+1)}$ that we are interested to attach to the pieces c_k that have already had associated pieces $e_{k2}, e_{k3}, \dots, e_{kl}$ attached in their corresponding main pieces c_k . Additionally attaching pieces $e_{k(l+1)}$ to pieces c_k is useful in making the agent who extracted them, interested in the pieces c_k because of the additional $e_{k(l+1)}$ as well as the previous attachments. However, naively attaching the pieces can be problematic and spoil the envy-freeness of the allocation that we maintain. We deal with the issue as follows.

- The agents who did not extract pieces associated with the c_k pieces as well as agents who extracted pieces that have not been attached are asked to 'reserve' a big enough subset $S' \subset S$ of snapshots in which they value the difference between their bonus value for c_k and the extracted pieces currently attached to c_k the most. These snapshots S' are removed from S and their remaining unattached associated pieces sent back to the residue. By maintaining the advantages in the snapshots S' , such agents will not be envious even if some agent in $\{1, \dots, l\}$ additionally gets all other extracted pieces $e_{k(l+1)}$ in the remaining snapshots in S . The advantages reserved due to the discarded snapshots in S' are crucial for the next phase where agents in $\{1, \dots, l\}$ get extracted pieces $e_{k(l+1)}$ from subset S'' of the working set.
- The agents indexed from 1 to l who have all already had their extracted pieces attached to c_k are asked to choose a high enough fraction of the snapshots in S in which they value the $e_{k(l+1)}$ pieces. We call these snapshots S'' . The $e_{k(l+1)}$ pieces from S'' are bunched together and the Main Protocol is called to divide this cake in an envy-free way among the agents indexed from 1 to l where l is strictly less than n . Since envy-freeness implies proportionality, they derive enough value that they will not be envious if the agent indexed $l + 1$ gets all other pieces in set $e_{k(l+1)}$. The corresponding set of snapshots S'' are then discarded.

Hence each time we attach isomorphic extracted pieces $e_{k(l+1)}$ to isomorphic pieces c_k , we discard snapshots $S' \cup S''$ from the working set S and still maintain an envy-free allocation. Note that in the snapshots that remain in S , agents may currently hold a different isomorphic piece than they previously did, but since they also hold the corresponding attachments associated with the isomorphic piece, each

agent's total value in each isomorphic snapshot in S stays the same. In Figure 9, we show the states of the permutation graph and the corresponding representative Core snapshot as well as the corresponding extracted pieces.

When the protocol attaches extracted pieces $e_{k(l+1)}$ to allocated pieces c_k currently held by agent l , it deletes the incoming edge of node/agent l and replaces it by an edge coming from agent $l + 1$ who extracted pieces in $e_{k(l+1)}$. Intuitively, $l + 1$ is now willing to be allocated c and its attached pieces instead of his current pieces in S . We delete previous edges to ensure that until termination, nodes in T have in-degree strictly 1 which guarantees that no matter the cycle involving a node in T found by the protocol, we will make progress towards termination. By attaching enough extracted pieces in the appropriate order, the GoLeft Protocol finally arrives at a point, where there is some isomorphic set of pieces c_k in the set S for which all possible associate pieces have been attached but there is some set of agents $N \setminus A$ who do not have associated pieces. The reason agents in $N \setminus A$ could not extract such pieces is because they had a unanimous significant advantage over the agent indexed 1 who got the pieces c_k . By gradually attaching (unanimously insignificant) associated piece to pieces c_k and ensuring that all agents who did extract corresponding pieces do get some isomorphic piece in c_k (along with the associated insignificant attachments), we make sure that agents in $N \setminus A$ now dominate agents in A . At this point, we can return from the GoLeft Protocol.

EXAMPLE 4 (GoLeft Protocol). *In Figure 9, we demonstrate how the permutation graph along with the working set of isomorphic snapshots changes in the GoLeft Protocol. Note that even when the representative snapshot changes, there still exist snapshots isomorphic to the previous representative snapshots but these snapshots have been discarded from the working set of snapshots. The coloured/shaded pieces represent the pieces given by the Core Protocol to each player. The small pieces on the left of the coloured pieces are extracted pieces, each labelled by the agent who extracted it. At first the extracted pieces are associated with a specific allocated piece. Then they are attached to it (represented by the dotted lines). Finally, when a coloured/shaded piece is reallocated to a new agent, the extracted pieces attached to it are also allocated to the new agent (in the diagram we now aggregate the extracted piece to the main piece). In the second state of the isomorphic snapshot, agent 2 points to agent 1 because the piece extracted by agent 2 has been attached to 1's held piece. In the third state of the isomorphic snapshot, agent 3 points to agent 2 because the piece extracted by agent 3 has been attached to 2's held piece. In the fourth state of the isomorphic snapshot, agent 1 points to agent 3 because the piece extracted by agent 1 has been attached to 3's held piece. In the fifth state, the agents 1, 2, 3 exchange their currently held piece and are allocated cake up to their extracted piece. In the fifth (last) state of the isomorphic snapshot, agent 1 holds a piece up till his extraction but neither agent 2 or 4 extracted pieces for the piece that agent 1 holds. This means that agents 2 and 4 have a significant advantage over agent 1. Initially the piece was held by 3 and still is in discarded isomorphic snapshots. This implies significant advantage of 2 and 4 over 3. Therefore agent 2 and 4 can be made to dominate 1 and 3.*

4. Argument for Boundedness and Envy-freeness. In the previous section, we presented the Main Protocol that relies on a few other protocols. We will now argue why it is bounded and envy-free.

4.1. Argument for Boundedness. We first show that the Main Protocol is bounded.

LEMMA 10. *The Core Protocol is bounded by $(n^2)^{n+1}$.*

Proof. The protocol can execute each of the steps without encountering a problem. In the Core Protocol, the cutter makes $n - 1$ trims, and the other agents are asked to make at most one trim on the pieces. This takes n^2 queries. The most costly procedure is having to call at most $2n$ times the SubCore Protocol for less than n number of agents. The recursion tree has at most n^2 branches and a depth of n . Hence there are at most $n(n^2)^n$ nodes in the tree with each node requiring n^2 steps. Hence the overall time of the Core Protocol is bounded by $n^3(n^2)^n$. \square

LEMMA 11. *The Discrepancy Protocol is bounded by $Bn \times n^3(n^2)^n$.*

Proof. The Discrepancy Protocol can execute each of its steps without encountering a problem. Since at least one agent finds the discrepant piece significant and since we call the Core Protocol B times before the while loop with each agent as cutter, at least one agent i thinks that the discrepant piece is of at least as much value as R . Hence after the while loop, such an agent i thinks that the discrepant piece is of value at least n times the value of R . The while loop from line 3 to line 5 will run at most once for each agent. If an agent values the piece within the gap delimited by the inequality, running the Core Protocol with the agent

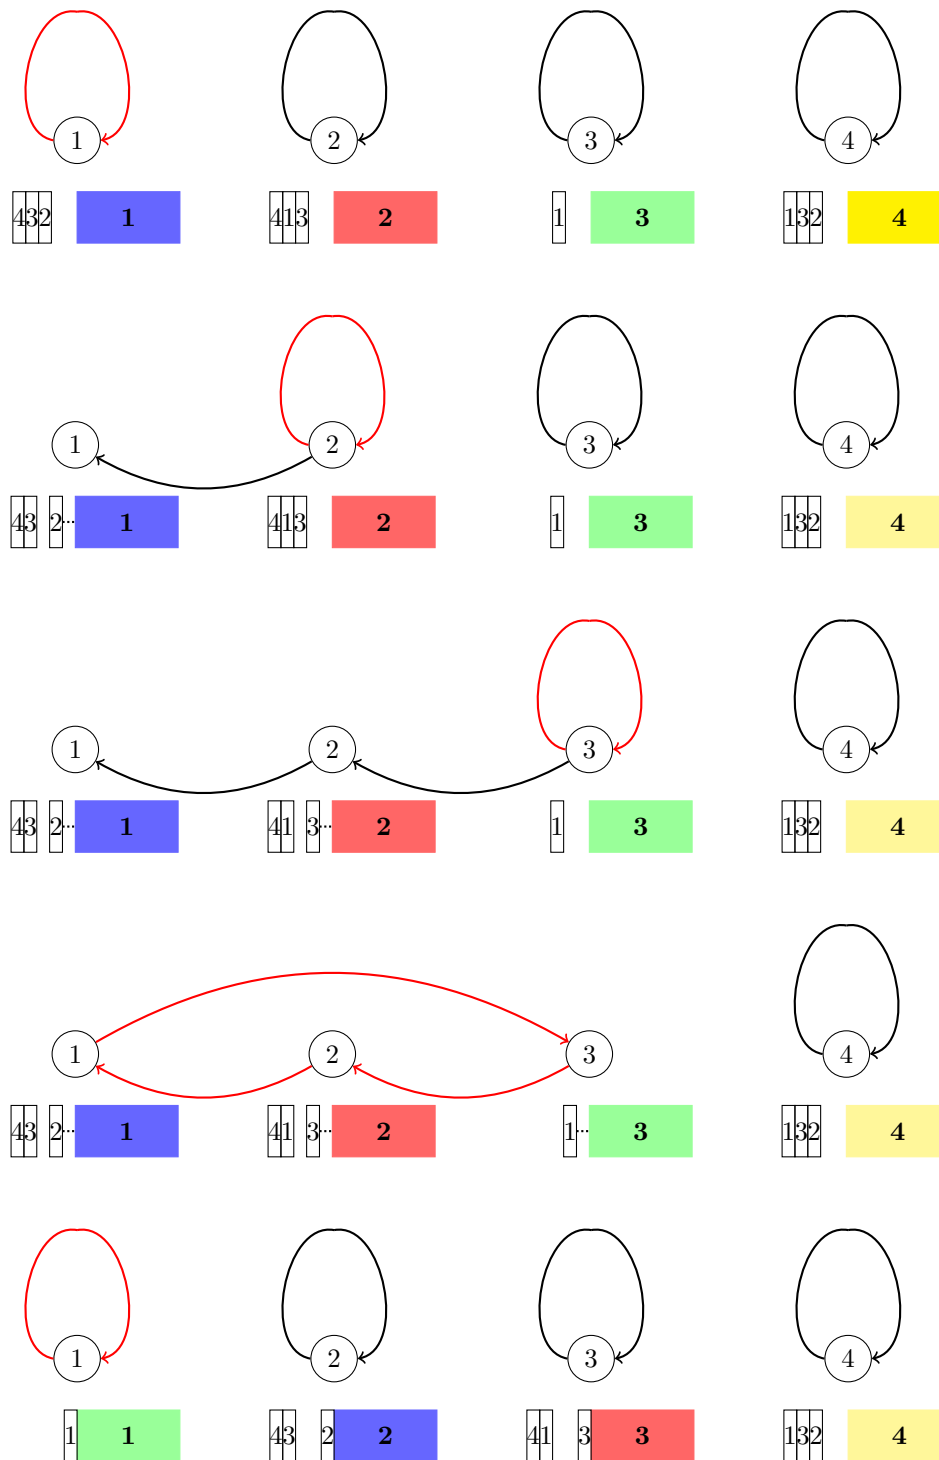


Fig. 9: Permutation graph along with corresponding state of an allocation snapshot representative of the working set of isomorphic snapshots.

as cutter B times will ensure that the residue is small enough for the piece to be on the right hand side of the inequality. Therefore the while loop is bounded by $Bn \times Core$ where $Core$ is the number of operations required to run the Core Protocol. The rest of the Discrepancy Protocol does not make any queries. The Discrepancy Protocol is therefore bounded by $Bn \times n^3(n^2)^n$ where $n^3(n^2)^n$ is the bound on the Core protocol as proved by Lemma 10. \square

LEMMA 12. *The GoLeft Protocol is bounded by $B'_{n-1} + 2Cn^3 + 2n^2 + C'$.*

Proof. First let us show that the GoLeft Protocol can execute all steps. The two steps which might be problematic are finding isomorphic snapshots in line 1 and finding a cycle in line 6. We handle the two cases by a couple of claims.

CLAIM 3. *Out of the C' snapshots generated, the GoLeft procedure can find C isomorphic snapshots in line 1.*

Proof. For this we use a simple counting argument. The isomorphism is entirely dependent on the order and existence of the extracted pieces. For each allocated piece of cake, there could be up to $n - 1$ extracted pieces. Each extracted piece might be extracted by any agent or not exist. This gives us $(n + 1)^{n-1}$ possibilities for each allocated piece. Since a snapshot has n allocated pieces, this gives $(n + 1)^{n^2 - n}$ different possible configurations. We are given n^{n^n} snapshots and desire to find n^{n^n} isomorphic snapshots. For $n > 4$ we have $n^{n^n} > (n + 1)^{n^2 - n} n^{n^n}$, which implies we can find a set of isomorphic snapshots of size C . \square

Therefore isomorphic snapshots can always be found.

We now argue that a cycle in line 6 can be found.

CLAIM 4. *In the permutation graph, at each step of the GoLeft Protocol, there exists a cycle containing at least one node from T .*

Proof. We observe that the permutation graph always contains a cycle since each node has in-degree at least 1. For each node s we can travel to the node t pointing to it. This process of backtracking can occur at most as many times as the number of nodes until we cycle.

We now show that it always contains a cycle with at least one node from T . Recall that each node in the graph points to nodes in T' if T' is non-empty. Each node in T has in-degree exactly 1 whereas each node in T' has in-degree $n - 1$. We start from any node in T and backtrack according to the incoming arc. In case we only encounter nodes from T , we will eventually cycle so that we have found a cycle consisting of nodes only from T . Otherwise, we may reach a node t' from T' . Since the node $t \in T$ from which we started backtracking points to each node in T' and hence to t' , we have again found a cycle consisting of node $t \in T$. \square

Now that we handled the two possibly problematic parts of the GoLeft Protocol, let us go over the other steps. The isomorphic snapshots can be found in C' steps. The while loop from line 5 to line 33 runs until a node in T is found such that it has had less than $n - 1$ attachments but has no extracted pieces to be attached (the if condition in step 10 is true). To make sure this happens, we require that there always be a node in T so that either we make progress in terms of attaching new sets of extractions or we already encounter the separation condition (Step 10). We prove this as follows:

CLAIM 5. *T' cannot include all nodes.*

Proof. A node being in T' means the agent associated with the node holds a piece of cake with no significant bonus. At least one piece of cake must have a significant bonus from the cutter's perspective (see e.g., Remark 1). \square

Each time the loop runs, either we get a separation of the agents given that the if condition in step 10 is true or a new set of extracted pieces is attached to a set of allocated pieces. There are only Cn^2 extracted pieces in the isomorphic snapshots that we are considering.

The for loop ranging from line 17 to line 20 just goes through a subset of the agents and therefore cannot run more than n times.

The for loop ranging from line 22 to line 25 also goes through a subset of the agents and therefore cannot run more than n times.

The while loop contains the two for loops and therefore its overall complexity is $(Cn^2 + n) \times 2n$.

Finally the call to the Main Protocol in line 28 calls the Main Protocol on a strict subset of the agents and is therefore bounded by B'_{n-1} .

This gives an overall complexity of $(Cn^2 + n) \times 2n + B'_{n-1} + C'$.

THEOREM 13. *The Main Protocol is bounded by $B' = n^{n^{n^{n^n}}}$ Robertson and Webb operations.*

Proof. We first establish the following claim.

CLAIM 6. *When the agents are asked to place a trim on the residue in line 19, they are always asked to delimit a value which is (much) smaller than their valuation of the residue.*

Proof. For the first extractions this is straightforward, seeing as we only ask an agent to place a trim if he thinks his bonus value is not significant. However as we extract pieces, the residue could become much smaller. This does not happen due to the bounds set in place. A non-significant piece of cake is valued to be less than $(\frac{n-2}{n})^B$ of the residue. We extract for C' snapshots and there are only $C'n^2$ pieces that could potentially be extracted. $C'n^2(\frac{n-2}{n})^B \ll 1$, hence extracting the pieces barely affects the residue, and we can therefore always extract a non-significant value from it. \square

The protocol can therefore always execute line 19.

The initial statement in line 7 runs the Core Protocol C' times. It is therefore upper bounded by $C' \times n^3(n^2)^n$ where $n^3(n^2)^n$ is the complexity of the Core Protocol, as proved by Lemma 10.

The while loop from line 7 to 9 checks that each agent values a piece outside a certain gap. The gap is based on the residue and if an agent values a piece within that gap, by running the Core Protocol B times with the agent as cutter we can ensure that the piece will never be in that gap again. The while loop will therefore run a maximum of $BnC'n$ times. Its complexity is therefore bounded by $BnC'n \times n^3(n^2)^n$.

The while loop from line 14 to 39 runs at most once for each bonus value piece since every time it repeats it labels a new bonus value as significant and a value can never be made to no longer be significant. Therefore the number of times it runs is bounded by $C'n^2$.

The for loop going from line 16 to 38 simply goes through all the pieces of cake and therefore runs $C'n$ times. The for loop from line 17 to 21 goes through all the agents and therefore runs n times. The for loop from line 23 to 37 goes through all the agents and therefore runs n times. The statement of line 28 is bounded by the complexity of the Discrepancy Protocol which according to Lemma 11 is bounded by $Bn \times n^3(n^2)^n$.

The Main Protocol calls in line 30 and in line 31 each are on a strict subset of the agents and by assumption are bounded by B'_{n-1} .

We therefore have a while loop running $C'n^2$ with each run of the loop taking $C'n \times (n + n \times (Bn \times n^3(n^2)^n + B'_{n-1}))$ steps, upper bounding the complexity of that while loop by $C'n^2 \times (C'n \times (n + n \times (Bn \times n^3(n^2)^n + B'_{n-1})))$

The call to GoLeft in line 40 is bounded by the complexity of the GoLeft Protocol which by Lemma 12 is bounded by $B'_{n-1} + 2Cn^3 + 2n^2$. The call to the Main Protocol in line 44 is on a strict subset of the agents and therefore is bounded by B'_{n-1} . This gives us an overall complexity of $C' \times n^3(n^2)^n + BnC'n \times n^3(n^2)^n + C'n^2 \times (C'n \times (n + n \times (Bn \times n^3(n^2)^n + B'_{n-1}))) + B'_{n-1} + 2Cn^3 + 2n^2 + B'_{n-1}$ which for $n > 4$ can be verified to be less than B' .

4.2. Argument for Envy-freeness.

We now argue for envy-freeness.

LEMMA 14. *The Core Protocol results in an envy-free partial allocation in which the cutter cuts the cake into n pieces, each agent gets a part of one of the pieces, and the cutter as well as at least one non-cutter agent gets a complete piece.*

Proof. It is sufficient to prove that the SubCore for $n - 1$ agents results in a neat allocation in which one agent gets a full piece. The cutter can take the last unallocated piece. No non-cutter agent i is envious of a cutter agent because i 's piece is at least as preferred as any unallocated piece. The cutter is indifferent among all pieces so is not envious of any non-cutter agent.

In order to prove the statement, it is sufficient to prove by induction the following statement. Suppose there are m agents and at least $m + 1$ pieces and there exists an envy-free allocation of the m agents where

each agent partially gets one of the pieces giving him at least the specified benchmark value, then the SubCore will find a neat allocation for m agents where at least one of the m agents gets a full piece. When we call the SubCore for the first time, the benchmark values are zero, so there does exist an envy-free allocation in which each agent i gets a piece of value at least b_i .

Base Case. If there is exactly one agent, he can be given the most preferred piece among the pieces. If there are multiple most preferred pieces, he is given the most preferred piece with the lowest index. The value achieved by the agent is at least b_i .

Induction. Suppose we get a neat allocation for $m - 1$ agents where one agent gets a full piece. If the m -th agent most prefers an unallocated piece, then we are already done. No agent is envious of another. Since the first $m - 1$ agents already achieved their benchmark values, they still achieve their benchmark value. As for agent m , he gets a most preferred piece so he trivially achieves his benchmark value because the benchmark value of an agent is either zero or a value of the one of the pieces.

If agent m is not interested in any of the contested pieces, the induction follows easily. This is because m 's most preferred piece is outside of the allocated pieces so far, and by giving m his most preferred piece we preserve neatness and the condition that one agent gets a full piece.

Now suppose that agent m is also interested in one of the $m - 1$ contested pieces. In that case, each agent in $[m]$ is asked to trim the highly valued contested pieces to equal the value of the most preferred uncontested piece (this value is kept track of by b'_j). Thus for each agent in $[m]$ the most preferred uncontested piece provides a benchmark value that they expect to get in any envy-free allocation of the agents in $[m]$.

We first claim that each of the $m - 1$ contested pieces has at least one trim even if the trim is on the extreme left margin. Suppose a contested piece a has no trim not even on the left margin. This means that each agent in $[m]$ considers some uncontested piece more preferred than piece a . But this contradicts our hypothesis that the allocation for agents $[m - 1]$ agents is neat.

When each agent in $[m]$ is asked to trim the contested pieces to the value of his most preferred uncontested piece, we distinguish between two cases: (1) we do not get multiple winners in W i.e., each contested piece has one designated winner and the the number of designated winners for the contested pieces is $n - 1$ and (2) some agent in W is a multiple winner, i.e., he is the designated winner with the rightmost trim on more than one contested piece. We can ensure thanks to claim 2 that each piece has exactly one designated winner.

We first deal with (1) which is the easier case. Suppose we do not get multiple winners so that $|W| = m - 1$. Each agent in W gets the benchmark value if he gets the piece where he had the rightmost trim. The allocation is also envy-free for agents in W . Nonetheless we call SubCore recursively on agents in W and the contested pieces. By the induction hypothesis, we get an envy-free allocation in which each agent in W gets a piece of at least benchmark value. The only remaining agent in $[m] \setminus W$ can be given a most preferred uncontested piece which is a full piece because it was never trimmed. All the agents get their benchmark value so the value achieved is at least b_i .

No agent i who is a winner of a contested piece envies another winner j of a contested piece because i 's trim on j 's piece is not on the right of j 's trim on j 's piece. No agent i who is a winner of a contested piece envies any uncontested piece because i 's benchmark value was exactly the value of the most preferred uncontested piece. Finally agent j who is given an uncontested piece is not envious of any winner i of a contested piece because j 's trim on i 's piece is not on the right of i 's trim on i 's piece. Also j is not envious of an unallocated uncontested piece because j got a most preferred uncontested piece.

We now deal with (2). Suppose there is at least one agent who wins multiple pieces. In that case $|W| < m - 1$. Note that if an agent in W gets the right side of the piece where he trimmed the most, he gets his benchmark value — the value of his most preferred uncontested piece. Since the agents in $[m]$ were asked to trim the contested pieces, agents in W only have trims on the contested pieces.

We now show that when $|W| < m - 1$, we can always increase W by one by adding an agent i from $[m] \setminus W$ and still ensure that $W \cup \{i\}$ admits a neat allocation (with respect to the current left margins) in

which each agent gets a partial contested piece. When $|W|$ becomes $m - 1$, we give the remaining agent in $[m] \setminus W$ the most preferred uncontested piece.

Suppose that $|W| < m - 1$ which means that there are multi-winners from W in the contested pieces.

CLAIM 7. *The SubCore protocol run on the contested pieces with the set of contested pieces, agents in W and values b'_j as input will return a neat allocation where each agent j in W gets at least value b'_j*

Proof. The SubCore protocol would fail to run only if an agent cannot get a piece of value equal to his benchmark value. However we know from the previous tentative allocation that as long as each agent gets at least his benchmark value, they all have a piece they are guaranteed to get in an envy-free manner. \square

Since $|W| \leq m - 1$, by the induction hypothesis, the SubCore will result in a neat allocation (with respect to the current left margins beyond which the left part of the piece is ignored) of the contested pieces. Since the allocation is neat (with respect to the current left margins), this means the current left margin coincides with a trim of some non-winner.

Each of the agents in W got a contested piece because SubCore was only called on W and the contested pieces. We now argue the following claim

CLAIM 8. *After we run SubCore on line 12, an additional agent from $[m] \setminus W$ can be added to W .*

Proof. We focus on an arbitrary unallocated contested piece a . We only consider that part of a that can be allocated and not the left part which is supposed to be ignored left of the rightmost non-winner's trim mark.

Case 1 Contested unallocated piece a has a marginal trim (on the current left margin) by a non-winner $i \in [m] \setminus W$. In that case i can be given a . We add i to W and W will still admit a neat (with respect to the current left margins) allocation in which each agent gets a contested piece.

Case 2 Contested unallocated piece a has no marginal trim (on the current left margin) by a non-winner. Let us refer to the current call of the SubCore protocol as SubC. We want to show that Case 2 cannot happen when we call a new SubCore protocol referred to as Sub1 on a set of contested pieces C_1 and agents W_1 . Let a be the contested piece satisfying the conditions of Case 2. Since a is in C_1 , either a is one of the agents in W_1 's most preferred piece or at some previous call of the SubCore protocol i was 'kicked out' from contested pieces and given piece a , thus making a a contested piece. Let us call this previous call Sub2. Let us say that Sub2 was called on agents W_2 and contested pieces C_2 such that $a \notin C_2 \subset C_1$.

We will now argue that in Sub1 i must get piece a , showing that Case 2 cannot happen. First observe that for any piece b in C_1 but not in C_2 , agent i prefers a to b . This follows from him being kicked out to a rather than b . Therefore Sub1 will not allocate b to i .

For any remaining piece c both in C_1 and C_2 , we will now show that i cannot get c . To show this, look at the agent j_0 who got c in Sub2. When we call Sub1, either j_0 is in W_1 or is not. If he is not, then he is a non-winner, and has placed a trim on c which must be to the right of or on the part of c he was allocated in Sub2. This follows from the fact that if j_0 is a non-winner he is asked to equalise c to his most preferred uncontested piece, which will be worse than the part of c j_0 got in Sub2 (else neatness breaks). Agent i was not envious of the lesser trimmed c when allocated a , and if indifferent ties were broken in favour of a through the placement of the imaginary values. This means that if he is allocated c in Sub1, he would rather get a .

Finally, let us imagine j_0 is in W_1 . For i to get c , j_0 must get a different piece. If j_0 gets a piece outside of C_2 , then this is a contradiction since previously his most preferred piece outside of C_2 caused j_0 to place a trim on c that was to the right of i 's trim corresponding to i 's valuation of a . If j_0 is getting a piece in C_2 , he must be taking the piece that was allocated to another agent j_1 in Sub2. This is only possible if j_1 is in W_1 . This means that in Sub1, j_1 must be getting another piece. We can continue this argument until we reach an agent j_k who gets a piece outside of C_2 , and therefore must have been kicked out by agent j_{k-1} placing their trim the same place or to the left of where it was in Sub2. This is a contradiction since in Sub2, j_k got the aforementioned piece trimmed to the left of his most preferred piece outside of C_2 . This concludes the argument that i cannot get a piece c in C_2 .

Since i also cannot get a piece b different from a outside of C_2 , i must get allocated a by Sub1.

This proves the claim. \square

Hence, we can always grow the set W until its size is $m - 1$. Previously we were maintaining a set W that had a neat envy-free allocation with respect to the current left margins and with respect to the contested pieces. The allocation was not necessarily a globally neat allocation because we were ignoring the left-side of the rightmost non-winner’s trims and also uncontested pieces. However, when $|W|$ becomes $m - 1$, the contested pieces have all been allocated so there is no unallocated piece for which some left part is ignored. The agents who get contested pieces get at least their benchmark value which is the value of the most preferred uncontested piece. Hence the agents who get contested pieces are not envious of any uncontested piece. The only remaining agent in $[m] \setminus W$ can be given the highest value uncontested piece that also happens to be a full piece because uncontested pieces remain untrimmed.

This completes the proof. \square

EXAMPLE 5. We will now walk through the critical part of the SubCore protocol to convey the structure of the proof and make the intuition clearer. Imagine that we have 5 agents. A cutter agent has cut the cake into 5 pieces. The SubCore protocol iteratively introduces each agent with the for loop line 1, until all agents have been allocated a piece and the allocation is neat. We will show through this example how introducing a new agent works. Let us say that so far 3 agents have been introduced and a neat allocation computed for them. We are now adding a fourth agent. The situation can be visualised in Figure 10.

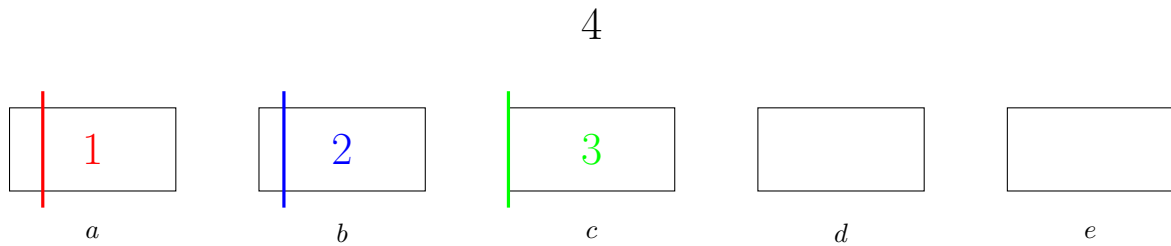


Fig. 10

Each agent trims most the piece that they are tentatively holding. In the diagram the trims are of the colour of the agent who placed them. We now query agent 4 as to which piece he prefers. The easy case is if he prefers piece d or e since then we could give him that piece and neatness would still be satisfied. We therefore dwell on the case where he prefers either a , b or c . In such a scenario, a , b and c will be referred to as the contested pieces. As we will see each, one agent will get ‘kicked out’ of these pieces and be forced to pick either d or e . However, since only one is kicked out, all agents are guaranteed to get their full preferred piece out of $\{d, e\}$. This gives us a lower bound (as long as the assumption made above that only one agent gets kicked out holds) for the value each agent is guaranteed to get. The value an agent j is guaranteed to get is kept track of by the variable b'_j in the protocol. In line 4, we ask each agent to place a trim on the contested pieces so that they are equal to this lower bound. This may look as so as in Figure 11.

In this diagram all the trims have been placed, and the agents who placed the rightmost trims are said to ‘win’ the piece. Their trims are highlighted in the diagram. These agents are added to the set W in line 5. If each trim had been won by a separate agent, we end up in an easy case of the protocol since each agent except one wins a separate piece, and the one who does not gets kicked out. However, in the situation depicted by the diagram, agent 1 has won both pieces a and b and agent 4 won c . We therefore do not know who gets kicked out between 2 and 3. However we do know that if we were to continuously make agent 1 and 4’s trims move left in an envy-free way, we could eventually ensure that a piece is won by someone new. Of course the protocol cannot move the trims continuously, but it is nonetheless possible to ensure that a new agent gets a contested piece. This is done in line 8 when we call the protocol recursively. This is what the recursive call would look like in the following situation (Figure 12):

The greyed out pieces of cake are not called as part of the input. Only the winner agents are called as part of the input (agent 1 and 4). We keep track of the value they are guaranteed to get through the inputs

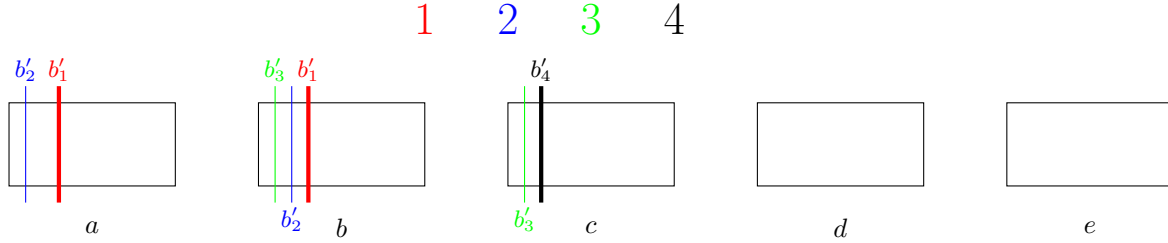


Fig. 11

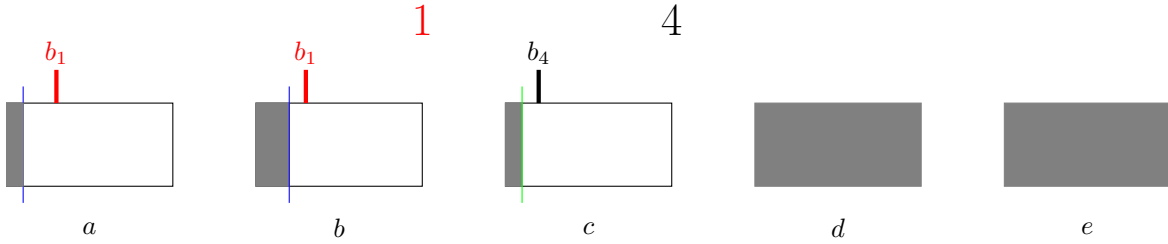


Fig. 12

b_1 and b_4 . This is important since pieces d and e are no longer part of the input. Now there are two concerns which we may have from this recursive call: (1) Does it actually manage to return a neat allocation? (2) Does the allocation returned allow us to expand the set W of winners? Concern (1) is addressed by Claim 7. Concern (2) is addressed by Claim 8. Once the protocol returns, the situation may look like (Figure 13).

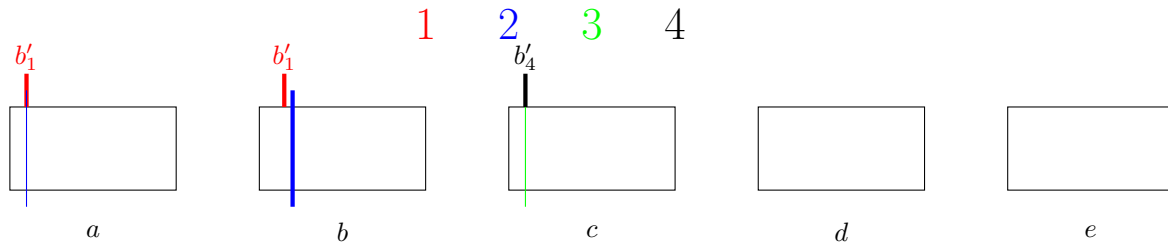


Fig. 13

The recursive call allocated a to 1 and c to 4, leaving b unallocated and with agent 2's trim to the right of agent 1's updated b'_i value. This means that we can now add agent 2 to the set of winners.

LEMMA 15. *Assuming envy-freeness is preserved the Core Protocol, envy-freeness is preserved in the Discrepancy Protocol.*

Proof. Discrepancy does not make allocations but simply distinguishes between 2 cases. Therefore as long as envy-freeness is preserved in the Core Protocol it is also preserved in the Discrepancy Protocol. \square

LEMMA 16. *Assuming envy-freeness is preserved in the Core Protocol and the Main Protocol with $n - 1$ or less agents, envy-freeness is preserved in the GoLeft Protocol.*

Proof. GoLeft performs a series of operations allowing us to attach a set of extracted piece e_{kl} to a set of allocated piece c_k in a subset S (which gets smaller and smaller as the algorithm progresses) of the snapshots and then permutes the allocation on these snapshots. The first part of the protocol uses the permutation graph to find what permutation is permissible in an envy-free way, and what extracted piece to attach next to allow further permutations. It updates the permutation graph to reflect what permutations

will be allowed and desirable after the operation of attaching e_{kl} to c_k has been executed. Let us first show that the operations to attach the extracted piece to the allocated piece do not break envy-freeness, and second that the permutation graph is correct, that is that on the subset of snapshots where the piece was attached, we can permute agents in a cycle of the permutation graph without breaking envy-freeness.

1. Assume we are trying to attach pieces $e_{k(l+1)}$ to the corresponding pieces in c_k which implies that pieces e_{km} with $m \leq l$ have already been attached to c_k . To do this we take the following two steps.

- (a) Make sure that for each agent $o \in \{l+1, \dots, n\}$, even if all remaining extraction up till o 's extractions $e_{k(l+1)}, \dots, e_{ko}$ are given to some agent $z \in \{1, \dots, l+1\}$, o will not be envious of z . We do so in phase ‘Attachment — making it agreeable for agents from $l+1$ to n ’ of the GoLeft Protocol.

Note that if all $e_{l(l+1)}$ pieces are given to agent $l+1$, agents in $\{l+2, \dots, n\}$ will not be envious of $l+1$ in any case because the pieces reflect a part of their advantage over $l+1$. Essentially, we want to ensure something weaker: that agent $o \in \{l+1, \dots, n\}$ will not be envious if the remaining $e_{l(l+1)}$ pieces are given to some agent $z \in \{1, \dots, l\}$.

We achieve the above goal by asking all agents o to choose $\frac{|S|}{n-l+1}$ snapshots where they value the difference between their bonus value and the extracted pieces which have so far been attached to c_k (let us call this value b) most. This leaves $\frac{|S|}{n-l+1}$ in $S \setminus S'$ (which is relabelled S). In any snapshot j and piece of cake c_{jk} , for any agent o , b is greater than the sum of all $V_o(e_{jkp})$ with $p \in \{l+1, \dots, o\}$. Since we ask agent o to discard as many snapshots with highest value b as are snapshots left in $S \setminus S'$, the advantage accumulated in the discarded snapshots amounts to more than the sum of all extracted pieces e_{kp} in the remaining snapshots. Hence, even if some agent $z \in \{1, \dots, l\}$ takes all the $e_{k(l+1)}$ pieces, the discarding agent will not be envious.

- (b) Ensure that for each agent $r \in \{1, \dots, l\}$, even if the remaining pieces in the set of pieces $e_{k(l+1)}$ are given to some agent $z \in \{l+1, \dots, n\}$, r will not be envious of z . We do so in phase ‘Attachment — making it agreeable for agents from 1 to l ’ of the GoLeft Protocol. To do this we ask r to choose $\frac{|S|n}{ln+1}$ snapshots j for which $V_j(e_{k(l+1)})$ is greatest and add them to a piece a which will then be shared amongst all agents in $\{1, \dots, l\}$. This is possible in an envy-free way because agents in $\{l+1, \dots, n\}$ are willing to give all pieces $e_{k(l+1)}$ to agent r due to the above argument. These snapshots are discarded and once all agents in $\{1, \dots, l\}$ have chosen we are left with $\frac{|S|}{ln+1}$ snapshots in $S \setminus S''$. Since r gets at least $1/n$ value of piece a , and that the $e_{jk(l+1)}$ pieces he values most were added in a , even if agent $l+1$ gets all pieces $e_{k(l+1)}$ in $S \setminus S''$, r will not be envious of him as these do not sum up to what he got from a . This means we can attach $e_{k(l+1)}$ to c_k and that agent $l+1$ can be allocated c_k in the snapshots in $S \setminus S''$ (which will be relabelled S) without envy-freeness being broken. When $e_{k(l+1)}$ pieces are attached to c_k , then c_k is equally desirable to $l+1$ as $l+1$'s current allocation in S .

2. Each edge in the permutation graph corresponds to moving some agent i from c_k to c'_k without changing his value in any of the isomorphic snapshots in S . Hence, when an exchange happens, each agent in the cycle gets exactly as good an allocation as before. Some agent j may possibly get envious because some other agent got extra attached pieces beyond j 's extraction. However, we have already argued above why agents are not envious because of additional extractions that are attached.

These arguments ensure that we end up with snapshots where some extracted pieces have been attached or shared by a subset of agents and for which agents' allocations have been shifted around, but for which envy-freeness is always preserved. \square

LEMMA 17. *The subset of agents A returned by GoLeft can be made to be dominated with respect to R by agents in $N \setminus A$ in $2B$ steps.*

Proof. In order to prove the lemma, we prove a couple of claims.

CLAIM 9. *For any agent i , once the residue has reached a value R , the protocol will never add cake to the residue in such a way that it becomes of value bigger than $(1 + \frac{1}{n})R$.*

Proof. All extracted pieces which are attached to a piece of cake in one of the snapshots are seen as non-significant by all agents. This means that for all agents i and each extracted piece e_{jkl} , we have $v_i(e_{jkl}) \leq f(B)v_i(R)$. In the protocol, once the residue is made smaller (either through discrepancy or

in the main protocol after running GoLeft), all pieces that had previously been extracted from it can no longer be reattached to it (in the case of discrepancy they are reattached before making the residue smaller, and in the second case the protocol no longer reattaches anything extracted from the residue). This means that when pieces that were extracted from the residue are reattached to it, each agent considers the pieces to satisfy the above inequality. Since there are at most Cn^3 such pieces, then for each agent i , $\sum_{e_{jkl}} v_i(e_{jkl}) \leq f(B)v_i(R)Cn^3 \leq \frac{v_i(R)}{n}$. \square

CLAIM 10. *Once an agent i has been allocated a piece c_{jk} by the GoLeft procedure, agents with a significant bonus value on c_{jk} can be made to dominate i with respect to R in $2B$ calls to the Core protocol.*

In other words the significant advantage obtained is maintained.

Proof. Before the GoLeft procedure works on the working set of isomorphic snapshots they are all envy-free. When the GoLeft procedure gives some of the extracted cake to an agent (either by attaching it to an allocated piece of cake (line 30) or by sharing it amongst a subset of agents (line 28), the extra amount of cake given to the agent is much smaller than a significant value from the perspective of all agents (line 7 in the Main Protocol). The values are such that when an agent has a significant bonus value over another, no matter what happens in the GoLeft procedure, the significant bonus cannot be diminished enough for the protocol to not be capable of making the residue smaller than that value in $2B$ calls to the Core protocol. To be more precise, the gap enforced by line 7 ensures that the following inequality holds: $n(\text{extracted pieces}) \leq nV_i(R)f(B)^2 < V_i(R)f(B)^{\frac{1}{2}} \leq \text{significant bonus}$.

The GoLeft routine terminates when the permutation graph is deleted. This means that there exists a set of isomorphic pieces of cake c_k such that all agents who extracted a piece from the residue and attached it to c_k have been allocated an element of c_k at least once, whilst at least one other agent had a significant bonus on c_k . According to Lemma 10, the agents with a significant bonus on elements of c_k can be made to dominate any agent having been allocated a member of c_k in $2B$ steps.

LEMMA 18. *Assuming the Main Protocol with $n-1$ or less agents is envy-free, that the Core protocol, the Discrepancy Protocol and the GoLeft Protocol are envy-free, envy-freeness is preserved in the Main Protocol.*

Proof. The Main Protocol does not make allocations without calling a sub-protocol. Assuming that all the sub-protocols it calls are envy-free, the fact that it calls them on all agents preserves overall envy-freeness. The Main Protocol also calls itself on a subset of the agents sharing either the residue or some discrepant piece. The two possible instances where this may occur are

1. The GoLeft Protocol has terminated, the Core Protocol has been run enough times to ensure that significant bonus values have been converted into domination on the residue R , and a subset of agents $N \setminus A$ dominate agents in subset A with respect to the residue R . We then run the Main Protocol with the residue and A as input. Since according to Lemma 17 agents in $N \setminus A$ dominate agents in A , they are willing to give all the residue to either agents in A , therefore no matter how agents in A share R , agents in $N \setminus A$ will not be envious.
2. The Discrepancy Protocol has been called and a discrepant piece can be exploited. In that case the Main Protocol will call 2 instances of itself. First on the discrepant pieces and agents who think it is significant (set D). Second on the residue and agents who think the discrepant piece is small (set D'). Each agent $i \in D$ thinks that the discrepant piece e_{jkl} is at least n times bigger than R . This means that if they think an agent s gets all of R but none of e_{jkl} , they will not be envious of him since i is guaranteed $\frac{1}{n}$ value of e_{jkl} . Therefore the agents sharing e_{jkl} are not envious of those sharing R . The same argument applies the other way around. \square

Based on the lemmata above, we have given the overall argument for envy-freeness. In each recursive call of the Main Protocol, the number of agents among which the cake is to be allocated decreases by at least one. Hence when the number of agents is four or less, we can use a known bounded envy-free protocol to allocate the residue in an envy-free manner.

THEOREM 19. *The Main Protocol allocates all the cake in an envy-free manner.*

5. Simplification of the SubCore and Core Protocol. Based on the insights of the SubCore Protocol, one can design a very simple non-recursive version of the SubCore protocol (Algorithm 6). The simplified protocol is simply a brute force protocol that iterates of permutations of agents and permutations of pieces. The simplification idea is based on the structural properties of the outcome of the SubCore Protocol

as well as some recent elegant insights about constructing a Core type protocol for the dual problem of envy-free chore division [14].

Algorithm 6 (Simplified) SubCore Protocol

Input: Cake cut into n'' pieces (with $n'' \leq n$) to be allocated among agents in set $\{1, \dots, n'\} = N' \subseteq N$ with $n' = |N'|$ and a benchmark value b_j for each $j \in N'$. {We only call SubCore if the benchmark values are such that there exists an envy-free allocation of agents in N' where each agent gets at most one of the pieces giving him at least the specified benchmark value.}

Output: A neat envy-free partial allocation for agents in N' in which each agent $j \in N'$ gets a connected piece of value at least b_j .

```

1: for each permutation  $(a_1, \dots, a_{n'})$  of  $N'$  do
2:   for each permutation  $(q_1, \dots, q_{(n''-n')}, p_1, \dots, p_{n'})$  of the  $n''$  pieces  $(c_1, \dots, c_{n''})$  do
3:     for  $j = 1$  to  $n'$  do
4:       Give  $p_j$  to  $a_j$ . Ask  $a_j$  to trim any of the pieces  $p_{j+1}, \dots, p_{n'}$  if needed so that their value for  $a_j$  does not exceed the value  $a_j$  has for his allocation  $p_j$ . {Pieces  $q_1, \dots, q_{(n''-n')}$  are untrimmed and unallocated.}
5:     end for
6:     if the allocation corresponding to the permutation of agents and pieces is envy-free and satisfies the benchmark requirements
7:       then
8:         return the allocation.
9:       else
10:        Undo the allocation.
11:      end if
12:    end for

```

THEOREM 20. *The (Simplified) SubCore Protocol returns an envy-free allocation in which the cutter and one other agent gets a full piece. It takes number of queries $O(n^{2n})$ where n is the number of pieces.*

Proof. Consider an allocation returned by the SubCore Protocol. In the allocation $n'' - n'$ pieces are unallocated and untrimmed whereas one agent gets a full piece. Suppose the $n'' - n'$ unallocated pieces are $q_1, \dots, q_{(n''-n')}$ and a_1 gets a full piece that is p_1 . Since the original SubCore allocation is envy-free, a_1 does not envy other agents and hence can trim off pieces among $p_2, \dots, p_{n'}$ of higher values to his value of p_1 . We know that when p_1 is allocated to a_1 , there exists a neat allocation of the possibly trimmed off pieces $p_2, \dots, p_{n'}$ for agents in $a_2, \dots, a_{n'}$ where each agent gets his benchmark value. By the correctness of the SubCore Protocol, we know that we can call SubCore on the trimmed pieces $p_2, \dots, p_{n'}$ and one agent gets one of these pieces completely. Let this agent be a_2 who gets p_2 . Since a_2 is not envious of the other agents in a SubCore allocation, he trims the pieces among $p_3, \dots, p_{n'}$ of higher value to equalize their value to his value for p_2 . We can repeat this argument so that each a_j gets p_j and is asked to trim any of the pieces $p_{j+1}, \dots, p_{n'}$ if needed so that their value for a_j does not exceed the value a_j has for his allocation p_j . Eventually for one permutation of agents and pieces, the allocation achieved is the same as the allocation returned by the original SubCore Protocol. \square

Note that the Core Protocol can call the simplified version of SubCore instead of the original SubCore. In what follows, we also present a self-contained version of simplified Core (Algorithm 7). which is based on exactly the same idea as the simplified SubCore.

6. Partial allocations. In Remark 2, we gave an argument that the first n calls of the Core Protocol in the Main Protocol ensure that we obtain a partial allocation that is envy-free and gives each agent $1/n$ value of the whole cake. As a result we answer an open problem posed by Segal-Halevi et al. [34] whether there exists a bounded algorithm that returns a proportional and envy-free partial allocation. We also note the following.

LEMMA 21. *If the SubCore Protocol runs in time $f(n)$, then there exists an algorithm that returns a proportional and envy-free partial allocation, and takes $O(nf(n))$ time.*

Note that the partial allocation may give disconnected pieces to the agents. We now give an additional result concerning partial allocation of the cake in which agents get connected pieces.

It is possible to use the SubCore Protocol to obtain envy-free partial allocations guaranteeing that each agent receives a *connected* piece of value $1/3n$ of the original cake. We specify Algorithm 8 that achieves this objective. Previously, Segal-Halevi et al. [34] showed that in bounded time, one can obtain an envy-free partial allocation in which each agent gets a connected piece of value $1/2^{n-1}$ of the whole cake. Note that our guarantee of $O(1/n)$ is optimal up to a constant factor.

In each iteration of the for loop in the algorithm, at least one agent who was not holding any piece now holds a piece. In particular agent i who is asked to make divisions of value $1/3n$ can get one of the pieces.

Algorithm 7 (Simplified) Core Protocol**Input** Agent set N , a cutter $i \in N$ and an unallocated cake R .**Output** An envy-free allocation of cake $R' \subseteq R$ for agents in N and an updated unallocated cake $R \setminus R'$.

```

1: Ask agent  $i$  to cut the cake  $R$  into  $n$  equally preferred pieces  $(p_1, \dots, p_n)$ .
2: for each permutation  $(a'_2, \dots, a'_n)$  of  $N \setminus \{i\}$  do
3:   for each permutation  $(p'_1, \dots, p'_n)$  of the  $n$  pieces  $(p_1, \dots, p_n)$  do
4:     Give  $p'_1$  to  $i$ 
5:     for  $j = 2$  to  $n$  do
6:       Give  $p'_j$  to  $a'_j$ . Ask  $a'_j$  to trim any of the pieces  $p'_{j+1}, \dots, p'_n$  if needed so the value of the pieces
       does not exceed the value  $a'_j$  has for her allocation  $p'_j$ .
7:     end for
8:     if the allocation  $p$  corresponding to the permutation of agents and pieces is envy-free then
9:       return the allocation  $p$  (which is called a Core allocation)
10:    else
11:      Reattach the trimmed parts to regain the original pieces.
12:    end if
13:  end for
14: end for

```

Algorithm 8 ConnectedPieces Protocol**Input:** Agent set N , unallocated cake R .**Output:** An envy-free partial allocation of cake for agents in N guaranteeing each agent $1/3n$ of the cake and a connected piece.

```

1: for each agent  $i \in N$  do
2:   if  $i$  is currently not holding a piece then
3:     for each connected piece  $c_j$  in the residue  $R$  do
4:       Ask  $i$  to decompose  $c_j$  into pieces of value  $1/3n$  and some leftover piece  $\alpha_j$  of value less than  $1/3n$ . {This does not change the
       cake except that it has trim marks on it. It is still the same cake with possibly some gaps in it because of agents in the previous
       iterations holding a piece.}
5:     end for
6:     Pick  $n$  of the pieces of value  $1/3n$  and call the set of those  $n$  pieces  $R'$ 
7:     Run SubCore( $R'$ ,  $N \setminus \{i\}$ ) and give  $i$  one of the unallocated pieces.
8:     Agent  $i$  believes the allocation he got from the previous step is of value  $1/3n$ , if another agent believes that the previous step
     gave him a bigger piece than what he was previously holding and bigger than  $1/3n$  then he keeps that piece, returning any previous
     piece he was previously holding to the residue. If an agent does not keep the piece given by the SubCore protocol, the piece is
     returned to the residue  $R$ .
9:   end if
   {At any point of the algorithm, there are agents who do not hold any cake and agents who hold exactly one connected piece.
   At the start no one holds a piece. Agents who hold a connected piece will keep holding a connected piece. They may exchange
   their current connected piece with a more preferred connected piece later on. }
10: end for
11: return envy-free partial allocation (in which each agent gets a connected piece)

```

At any point in the algorithm we should view the cake as having at most $k < n$ gaps or discontinuities as a result of k agents holding connected pieces from the cake. When a new agent i makes trims to create pieces of value $1/3n$, this does not change the cake or the gaps in it. The algorithm has n calls of the SubCore Protocol.

We observe that for any agent holding a piece of cake, he already gets value $1/3n$ of the original cake. For any agent not holding a piece of cake, he thinks that no agent who is holding a piece has a piece of value $1/3n$. Hence, he thinks that at least $2/3$ value of the cake is unallocated.

THEOREM 22. *The ConnectedPieces Protocol allocates the cake partially in an envy-free way such that each agent has a connected piece of value at least $1/3n$ of the original cake.*

Proof. We show that the algorithm proceeds correctly and terminates with each agent holding one connected piece. The algorithm iteratively calls SubCore with n pieces created by each agent. The only reason why the algorithm may not proceed is if an agent i cannot produce n divisions of value $1/3n$ of the original cake. When a new agent i is asked to make n divisions of value exactly $1/3n$, we want to show he can make these divisions by appropriate trim marks. Assume that at this point k agents are holding a piece of cake which means that there are k gaps (discontinuities) in the cake. So for an agent i who wants to create divisions of value $1/3n$, there can be two problems: (1) too much value of the cake is allocated and (2) the

pieces that are allocated lead to discontinuities which prevents the agent from creating n divisions of value $1/3n$. For the first case, less than $k/3n$ value of the cake is allocated from i 's perspective. Hence at least $2/3$ value of the cake is still unallocated for i . This means that there is enough cake to generate n pieces of value $1/3n$. For the second case, the discontinuities result in at most $k+1/3n$ of the cake not being used. This is because the k gaps divide R into at most $k+1$ pieces and less than $1/3n$ of cake can be lost per piece due to non-divisibility by $1/3n$ of the value of each connected piece. This means that there is still more than $1/3$ of the cake that can be used for the divisions. Since we require n divisions each of value $1/3n$, there is enough cake to obtain these divisions.

We now argue that the final allocation is envy-free. At some point an agent i gets a piece to hold. He thinks this piece is strictly more preferred than any piece that was held in previous iterations by any agent. In the same iteration i thinks he got at least as valuable a piece as any piece allocated in the iteration because of envy-freeness of SubCore. So i thinks he has a best piece. In case at a latter iteration, he thinks someone got a better piece than his own currently held piece, then due to envy-freeness of SubCore, i can get at least preferred a piece as well. \square

We also note the following.

LEMMA 23. *If the SubCore Protocol runs in time $f(n)$, then there exists an algorithm that takes $O(nf(n))$ time and allocates the cake partially in an envy-free way such that each agent has a connected piece of value at least $1/3n$ of the original cake.*

7. Discussion. In this paper we presented the first general discrete and bounded envy-free protocol. Now that boundedness has been established, the paper opens the door to new work on finding the optimal bound. Getting a clearer understanding of the complexity of envy-freeness in terms of the number of queries is an interesting direction for future work.

The Core Protocol we defined can be used as an oracle in other envy-free protocols that allocate a large enough part of the cake. If the Core Protocol is called n times on the updated residue each time with a different agent as cutter, then we obtain an envy-free allocation in which each agent gets $1/n$ of the total value of the cake. Hence, after the first n calls of the Core Protocol, our Main Protocol already identifies an allocation that satisfies proportionality (with respect to the whole cake) and envy-freeness. If the Main Protocol is taking too much time, it can be timed out and still give an envy-free allocation with the proportionality guarantee.

It will also be interesting to use the new techniques presented in the paper for finding fair allocations for other notions of fairness.

Acknowledgments. The authors thank Sajid Aziz, Simina Brânzei, Omer Lev, Abdallah Saffidine, Erel Segal-Halevi, Xin Huang, Jade Tan-Holmes, Tommaso Urli, and the reviewers of FOCS 2016 for comments. They also thank Steven Brams and Ariel Procaccia for pointers to the literature. Haris Aziz is supported by a Julius Career Award. He thanks Ulle Endriss for introducing the subject to him at the COST-ADT Doctoral School on Computational Social Choice, Estoril, 2010.

REFERENCES

- [1] H. AZIZ AND S. MACKENZIE, *A discrete and bounded envy-free cake cutting protocol for four agents*, in Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC), ACM Press, 2016, pp. 454–464.
- [2] H. AZIZ AND C. YE, *Cake cutting algorithms for piecewise constant and piecewise uniform valuations*, in Proceedings of the 10th International Workshop on Internet and Network Economics (WINE), 2014, pp. 1–14.
- [3] B. BARBANEL AND A. D. TAYLOR, *Preference relations and measures in the context of fair division*, American Mathematical Monthly, 123 (1995), pp. 2061–2070.
- [4] J. B. BARBANEL, *The Geometry of Efficient Fair Division*, Cambridge University Press, 2005.
- [5] J. B. BARBANEL AND S. J. BRAMS, *Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond*, Mathematical Social Sciences, 48 (2004), pp. 251–269.
- [6] S. J. BRAMS AND A. D. TAYLOR, *An envy-free cake division protocol*, The American Mathematical Monthly, 102 (1995), pp. 9–18.
- [7] S. J. BRAMS AND A. D. TAYLOR, *Fair Division: From Cake-Cutting to Dispute Resolution*, Cambridge University Press, 1996.
- [8] S. J. BRAMS, A. D. TAYLOR, AND W. S. ZWICKER, *A moving-knife solution to the four-person envy-free cake division*, Proceedings of the American Mathematical Society, 125 (1997), pp. 547–554.

- [9] S. BRANZEI, *A note on envy-free cake cutting with polynomial valuations*, Information Processing Letters, 115 (2015), pp. 93–95.
- [10] S. BRANZEI AND P. B. MILTERSEN, *A dictatorship theorem for cake cutting*, in Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI), AAAI Press, 2015, pp. 482–488.
- [11] C. BUSCH, M. S. KRISHNAMOORTHY, AND M. MAGDON-ISMAIL, *Hardness results for cake cutting*, Bulletin of the EATCS, 86 (2005), pp. 85–106.
- [12] Y. CHEN, J. K. LAI, D. C. PARKES, AND A. D. PROCACCIA, *Truth, justice, and cake cutting*, Games and Economic Behavior, 77 (2013), pp. 284–297.
- [13] Y. J. COHLER, J. K. LAI, D. C. PARKES, AND A. D. PROCACCIA, *Optimal envy-free cake cutting*, in Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI), AAAI Press, 2011, pp. 626–631.
- [14] S. DEGHANI, A. F. M. HAJIAGHAYI, AND Y. YAMI, *Envy-free chore division for an arbitrary number of agents*, in Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2018.
- [15] X. DENG, Q. QI, AND A. SABERI, *Algorithmic solutions for envy-free cake cutting*, Mathematics of Operations Research, 60 (2012), pp. 1461–1476.
- [16] J. EDMONDS AND K. PRUHS, *Cake cutting really is not a piece of cake*, in Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2006, pp. 271–278.
- [17] S. EVEN AND A. PAZ, *A note on cake-cutting*, Discrete Applied Mathematics, 7 (1984), pp. 285–296.
- [18] G. GAMOW AND M. STERN., *Puzzle-math*, Macmillan, 1958.
- [19] S. GARFUNKEL, *For all practical purposes social choice*, COMAP, (1988).
- [20] W. GASARCH, *Which unbounded protocol for envy free cake cutting is better?*, Tech. Report arXiv:1507.08497, arXiv.org, 2015.
- [21] R. GUO, *Cross-Border Management: Theory, Method and Application*, Springer, 2015.
- [22] W. HIVELY, *Dividing the spoils*, Discover Magazine, March (1995).
- [23] D. KUROKAWA, J. LAI, AND A. D. PROCACCIA, *How to cut a cake before the party ends*, in Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI), AAAI Press, 2013, pp. 555–561.
- [24] C. LINDNER AND J. ROTHE, *Degrees of guaranteed envy-freeness in finite bounded cake-cutting protocols*, in Proceedings of the 5th International Workshop on Internet and Network Economics (WINE), vol. 5929 of Lecture Notes in Computer Science (LNCS), Springer-Verlag, 2009, pp. 149–159.
- [25] C. LINDNER AND J. ROTHE, *Cake-cutting: Fair division of divisible goods*, in Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division, J. Rothe, ed., Springer-Verlag, 2015, ch. 7.
- [26] A. MAYA AND N. NISAN, *Incentive compatible two player cake cutting*, in Proceedings of the 8th International Workshop on Internet and Network Economics (WINE), vol. 7695 of Lecture Notes in Computer Science (LNCS), Springer-Verlag, 2012, pp. 170–183.
- [27] O. PIKHURKO, *On envy-free cake division*, The American Mathematical Monthly, 107 (2000), pp. 736–738.
- [28] A. D. PROCACCIA, *Thou shalt covet thy neighbor’s cake*, in Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI), AAAI Press, 2009, pp. 239–244.
- [29] A. D. PROCACCIA, *Cake cutting: Not just child’s play*, Communications of the ACM, 56 (2013), pp. 78–87.
- [30] A. D. PROCACCIA, *Cake cutting algorithms*, in Handbook of Computational Social Choice, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, eds., Cambridge University Press, 2016, ch. 13.
- [31] J. M. ROBERTSON AND W. WEBB, *Near exact and envy-free cake division*, Ars Combinatorica, 45 (1997), pp. 97–108.
- [32] J. M. ROBERTSON AND W. A. WEBB, *Cake Cutting Algorithms: Be Fair If You Can*, A. K. Peters, 1998.
- [33] A. SABERI AND Y. WANG, *Cutting a cake for five people*, in Proceedings of the 5th International Conference on Algorithmic Aspects in Information and Management (AAIM), Lecture Notes in Computer Science (LNCS), Springer-Verlag, 2009, pp. 292–300.
- [34] E. SEGAL-HALEVI, A. HASSIDIM, AND Y. AUMANN, *Waste makes haste: Bounded time protocols for envy-free cake cutting with free disposal*, in Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), IFAAMAS, 2015, pp. 901–908.
- [35] T. SÖNMEZ AND M. U. ÜNVER, *Matching, allocation, and exchange of discrete resources*, in Handbook of Social Economics, J. Benhabib, M. O. Jackson, and A. Bisin, eds., vol. 1, Elsevier, 2011, ch. 17, pp. 781–852.
- [36] H. STEINHAUS, *The problem of fair division*, Econometrica, 16 (1948), pp. 101–104.
- [37] I. STEWART, *Division without envy*, Scientific American, (1999).
- [38] I. STEWART, *How to cut a cake : and other mathematical conundrums*, Oxford University Press, 2006.
- [39] W. STROMQUIST, *Envy-free cake divisions cannot be found by finite protocols*, The Electronic Journal of Combinatorics, 15 (2008).
- [40] F. E. SU, *Rental harmony: Sperner’s lemma in fair division*, American Mathematical Monthly, 10 (1999), pp. 930–942.
- [41] W. THOMSON, *Children crying at birthday parties. Why?*, Economic Theory, 31 (2007), pp. 501–521.