Computational Fair Division

Simina Brânzei

Ph.D. Dissertation

Department of Computer Science
Aarhus Universitet
Denmark

Computational Fair Division

A Dissertation
Presented to the Faculty of Science and Technology
of
Aarhus University
in Partial Fulfillment of the Requirements
for the Ph.D. Degree

by
Simina Brânzei

Main Supervisor: Peter Bro Miltersen
Dissertation Committee:
Moshe Babaioff, Gerth Stølting Brodal, Edith Elkind

February, 2015

Abstract

Fair division is a fundamental problem in economic theory and one of the oldest questions faced through the history of human society. The high level scenario is that of several participants having to divide a collection of resources such that everyone is satisfied with their allocation – e.g. two heirs dividing a car, house, and piece of land inherited. The literature on fair division was developed in the 20th century in mathematics and economics, but computational work on fair division is still sparse.

This thesis can be seen as an excursion in computational fair division divided in two parts. The first part tackles the cake cutting problem, where the cake is a metaphor for a heterogeneous divisible resource such as land, time, mineral deposits, and computer memory. We study the equilibria of classical protocols and design an algorithmic framework for reasoning about their game theoretic properties. In our framework, the protocols are built from simple instructions that can be executed on a computer. Moreover, we prove an impossibility theorem for truthful mechanisms in the classical query model, which is similar in spirit to the Gibbard-Sattherthwaite theorem of social choice theory. We also study alternative and richer models, such as externalities in cake cutting, simultaneous cake cutting, and envy-free cake cutting.

The second part of the thesis tackles the fair allocation of multiple goods, divisible and indivisible. In the realm of divisible goods, we investigate the well known Adjusted Winner procedure, obtaining several novel characterizations of the protocol and giving a complete picture of its pure Nash equilibria and their efficiency. For indivisible goods, we study the competitive equilibrium from equal incomes solution concept for valuations with perfect complements and valuations with perfect substitutes. We obtain characterizations of when competitive equilibria exist, as well as polynomial time algorithms and hardness results for computing the equilibria.

Resumé

Retfærdig deling af resurser er et fundamentalt problem i økonomisk teori og i øvrigt et af de ældste samfundsmæssige problemer som er håndteret i menneskehedens historie. Overordnet set betragtes en situation hvor adskillige deltagere skal dele en samling resurser så de hver især er tilfredse med deres allokering. Et eksempel er to arvinger der skal dele en bil, et hus, og et stykke jord. Den videnskabelige literatur om retfærdig deling blev etableret i det 20. århundrede indenfor de matematiske og økonomiske videnskaber, men de beregningsmæssige aspekter af emnet er endnu ikke grundigt studeret. Denne afhandling omhandler sådanne beregningsmæssige aspekter af retfærdig deling og består af to dele.

Den første del håndterer kagedelingsproblemet. Kagen i kagedelingsproblemet er en metafor for en heterogen og delelig resurse såsom land, tid, mineralforekomster eller computerhukommelse. We studerer ligevægte af klassiske protokoller og for de protokoller der består af simple instruktioner der kan eksekveres på en computer udvikler vi en algoritmisk ramme til at argumentere om deres spilteoretiske egenskaber. Desuden viser vi et negativt resultat om eksistensen af manipulationsresistente mekanismer i en klassisk model for diskrete kagedelingsprotokoller; dette resultat har lighedspunkter med Gibbards og Satterthwaites sætning fra teorien om sociale valg. Vi studerer også alternative og rigere modeller, såsom kagedeling med eksternaliteter, simultan kagedeling, og misundelsesfri kagedeling.

Den anden del håndterer retfærdig deling af adskillige goder, delelige såvel som udelelige. For delelige goder studerer vi den velkendte "Adjusted winner" protokol, opnår adskillige nye karakterisationer af protokollen, og giver et fuldstændigt billede af protokollens rene Nash ligevægte og deres effektivitet. For udelelige goder studerer vi løsningsbegrebet "competitive equilibrium from equal incomes" for præferencer med perfekte komplementer og præferencer med perfekte substitutter. Vi karakteriserer de situationer hvor en ligevægt eksisterer og giver polynomieltids algoritmer og hårdhedsresultater for at beregne en ligevægt.

Acknowledgments

There are many people to whom I would like to say thank you. First, I would like to express my deepest gratitude to my advisor, Peter Bro Miltersen, for his support during the past three and a half years and for giving an example of how a world class researcher should be. I would like to thank Peter for his inspiring enthusiasm and brilliant insight into the topic, for so often finding time to discuss research and give sharp technical advice, and for offering me the freedom and encouragement to develop my own ideas while contributing his expertise to both the high level directions and technical analysis. Last, but not least, I am grateful to Peter for his friendship during the past years and for the many interesting discussions ranging from movies to the universe and life in general.

I would like to thank Ariel Procaccia for hosting my exchange at Carnegie Mellon University and being an amazing mentor throughout my time there (taking care of me as if I were his own student), and for his insightful guidance, contagious optimism, and deep involvement in all aspects of the projects that we worked on together. I've had the chance to work with outstanding researchers and I'm very grateful for all the collaborations that followed: Margareta Ackerman, Haris Aziz, Eric Balkanski, Ioannis Caragiannis, Yiling Chen, Agatha Chrobak, Xiaotie Deng, Clara Forero, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Hadi Hosseini, Nicholas Jennings, David Kurokawa, Kate Larson, Tomasz Michalak, Peter Bro Miltersen, Jamie Morgenstern, Ariel Procaccia, Talal Rahwan, Piotr Szczepanski, Michael Wooldridge, *Jie Zhang. A special thanks goes to Kate Larson, who was my advisor during my undergraduate and master's at the University of Waterloo, for being a very inspiring role model and for her early encouragement when I was a clueless undergrad trying to figure out what research meant.*

I am grateful to Moshe Babaioff, Gerth Stølting Brodal, and Edith Elkind for reading this thesis, Amir Ronen for supporting my IBM grant, Shai Ben-David for introducing me to theoretical machine learning, Yoram Bachrach, Arthur Carvalho, Lachlan Dufton, Sune Jakobsen, Kevin Leyton-Brown, Yuezhou Lv, Samuel Ranellucci, Pingzhong Tang, and Bo Zheng for useful discussions about my work, Sofia Ceppi, Bart de Keijzer, Mona Rahn for making conferences more fun, the members of the COST action on Computational Social Choice for inviting me to several great workshops, Hanna Ewa Makaruk and Robert Michal Owczarek for many invited talks and remaining unique mentors and deeply caring friends.

I would also like to thank the wonderful staff at Aarhus University, including Mads Berggreen, Jan Frederiksen, Ellen Kjemtrup, Dorthe Haagen Nielsen, Søren Poulsen, Julie Sara Terndrup (who went out of her way to help me find housing in Aarhus), the staff at Tsinghua University for their hospitality during my exchange there, and AnnMarie Zanger (at Carnegie Mellon University). In addition, I am grateful to Anders Møller for being an attentive chair of the Ph.D. committee, Ole Lehrmann Madsen for helping me write a grant application (together with Mads), Susanne Bødker and Olivier Danvy for their wise advice as my formal support group, and Peyman Afshani and Ivan Damgård for teaching great courses.

This work was generously supported from the Danish National Research Foundation and The National Science Foundation of China (under grants 61061130540 and 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation and from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research

Council. I am also very grateful to Google, IBM, IFAAMAS, Microsoft Research and ACM-W for additional financial support.

I would like to thank all my friends and colleagues in the Mathematical Computer Science group at Aarhus University (past and present) for a welcoming and intellectually rich environment: Aris, Branislav, Dominik, Elad, Elias, Gudmund, Joshua, Kevin, Kristoffer, Navid, Peter, Rasmus, Sven, Søren, Thomas. I am also grateful to the (past) members of the Computer Science Club at the University of Waterloo (including Alex, Ali, Alina, Caelyn, David $\times$ *2, Holden, Mark* $\times$ *2, Michael, Mr. Biggs, Tor) for sharing their passion for computer science and mathematics.*

Going back in time, thank you Arie for an amazing year at the IBM Toronto Lab: you believed in me at a time when I did not believe myself and had a profound influence on my later development. And Dina, for your kindness – I hope to meet again.

I am grateful to all my other wonderful friends (many of which are far away) for the joy that they brought in my life. To name some of them: Aizhan, Alessandra, Alexandra $\times$ *2, Antigoni, Barbora, Carina, Clara, Faiza, Georgia, Ira, Irina, John, Laura, Ligia, Olga, Sergiu, Simona, Stephanie, Tereza. Thank you Nikola for understanding and caring in what turned out to be a similar journey, Margareta for your close friendship of so many years and for the persistence that you had with me a long time ago, without which I would not be here today, and for which I am forever grateful.*

Last, but not least, I would like to thank my family for their love and unwavering support, for the courage and hope that they inspire in me, for making me feel special, and simply being. This thesis is dedicated to them with all my love.

Simina Brânzei,
Aarhus, February 27, 2015.

Contents

Chapter 1

Preface

My main research area during the past three and a half years has been computational fair division, with an emphasis on models such as cake cutting and the allocation of multiple divisible and indivisible goods.

Cake cutting is the oldest formal model of fair division; it encapsulates the essence of the fair division problem, having been studied in an extensive body of literature starting with World War II. In more recent works, problems such as the allocation of multiple goods have also been developed and studied extensively; the latter can be seen as variants of the cake cutting problem with stronger informational assumptions and sometimes non-additive valuations.

The primary focus of this work has been the study of *protocols* for fair division, with an emphasis on the following questions:

(i) Given that the participants in a fair division scenario are strategic, can one ensure that the protocols have "good" outcomes in the equilibrium?

(ii) Can such protocols be implemented efficiently on a computer?

(iii) How do richer and alternative representations change the landscape of what can be computed?

This thesis initiated the direction of studying equilibrium outcomes of classical fair division protocols as an answer to the first question ([28, 32]). More specifically, Chapter 3 studies the equilibria of the Dubins-Spanier procedure—a well-known protocol for computing *proportional* allocations—for a basic class of strategies. Chapter 4 generalizes these results, developing an algorithmic framework for fair division—GCC, from generalized cut and choose protocols—in which the procedures are built from basic instructions that can be executed on a computer. We show the GCC framework captures all the known discrete cake cutting protocols and that it is well suited for rational agents. In particular, every GCC protocol has approximate subgame-perfect equilibria (and in fact exact equilibria when the tie-breaking rule is set correctly). Moreover, we design a specific GCC protocol for any number of agents which has Nash equilibria that are always envy-free. Chapter 4 also offers an answer to question (ii), since the GCC framework is designed to be implemented as a computer program, while standard cake cutting procedures are specified only through a communication model, without any specifications on their internal implementation.

Additionally, we show that a framework such as GCC is in fact necessary, as the standard model of communication for cake cutting is ill-suited for rational agents. This is formalized in Chapter 5, where we find a strong impossibility result similar to the classical *dictatorship* results of social choice theory, in particular the Gibbard-Satterthwaite theorem [66, 114], which is a cornerstone of social choice theory and mechanism design. Our main theorem is that every strategyproof protocol for two "hungry" agents[1] in the standard query model is a dictatorship. For more than two hungry agents, a similar impossibility holds: there always exists an agent that does not get anything.

For the third question, I investigated the model of "simultaneous cake cutting" (Chapter 6), where the communication between the center and the participants occurs in parallel, as is often the case in resource allocation instances on the internet, while Chapter 7 explores the model of cake cutting with externalities, where the agents are affected by the allocations of others.

Part II of the thesis tackles the allocation of multiple goods and answers questions (i) and (ii) in two well known scenarios. That is, Chapter 10 focuses on analyzing the strategic outcomes of a well known protocol: the Adjusted Winner procedure, which has been patented and used in real world scenarios for the fair allocation of resources among two parties (e.g. divorce settlements, inheritances). We show that Adjusted Winner has pure Nash equilibria, with surprisingly good properties with respect to the ground truth (i.e. the true but possibly hidden valuations of the parties), and that it admits several succinct characterizations that shed further light on the properties of the protocol.

Finally, Chapter 11 studies the existence and computation of competitive equilibria from equal incomes (CEEI) for the allocation of indivisible goods for valuations with perfect complements and perfect substitutes. An approximate variant of the CEEI solution has been studied intensively recently as a desirable solution to allocate goods in real settings (such as allocating courses to students at the Wharton Business School at the University of Pennsylvania). In addition to answering the relevant computational questions, for valuations with perfect complements we find a very succinct characterization of the instances that admit a CEEI solution for indivisible goods.

The necessary background on fair division is given in Chapter 1 and Chapter 9. For coherence I included only papers strictly about fair division and the authors are in alphabetical order as standard in theoretical areas (except for 12):

[1] An agent is said to be hungry if it values of all of the resource

1. “Equilibrium Analysis in Cake Cutting” [28], Simina Brânzei and Peter Bro Miltersen. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)* (Chapter 3).

2. “An Algorithmic Framework for Strategic Fair Division” [32], Simina Brânzei, Ioannis Caragiannis, David Kurokawa, and Ariel D. Procaccia. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI 2016)* (Chapter 4).

3. “A Dictatorship Theorem for Cake Cutting” [29], Simina Brânzei and Peter Bro Miltersen. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)* (Chapter 5).

4. “Simultaneous Cake Cutting” [13], Eric Balkanski, Simina Brânzei, David Kurokawa, and Ariel D. Procaccia. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI 2014)* (Chapter 6).

5. “Externalities in Cake Cutting” [31], Simina Brânzei, Ariel D. Procaccia, and Jie Zhang. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI 2013)* (Chapter 7).

6. “A Note on Envy-Free Cake Cutting with Polynomial Valuations” [33], Simina Brânzei. In *Information Processing Letters, 2015* (Chapter 8).

7. “The Adjusted Winner Procedure: Characterizations and Equilibria” [10], Haris Aziz, Simina Brânzei, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen. *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)* (Chapter 10).

8. “Characterization and Computation of Equilibria for Indivisible Goods” [37], Simina Brânzei, Hadi Hosseini, and Peter Bro Miltersen. *Proceedings of the Eighth International Symposium on Algorithmic Game Theory (SAGT 2015)* (Chapter 11).

In addition to fair division, I also worked on other topics, which resulted in the following completed papers during my Ph.D.:

9. “Verifiably Truthful Mechanisms” [34], Simina Brânzei and Ariel D. Procaccia. In *Proceedings of the Sixth Conference on Innovations in Theoretical Computer Science (ITCS 2015).*

10. “The Fisher Market Game: Equilibrium and Welfare” [36], Simina Brânzei, Yiling Chen, Xiaotie Deng, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. In *Proceedings of the Twenty-Third Conference on Artificial Intelligence (AAAI 2014).*

11. "The Authorship Dilemma: Alphabetical or Contribution?" [1], Margareta Ackerman and Simina Brânzei. In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)* (extended abstract). Full version under journal submission.

12. "Implementation and Computation of a Value for Generalized Characteristic Function Games" [97], Tomasz Michalak, Piotr Szczepanski, Agatha Chrobak, Simina Brânzei, Talal Rahwan, Michael Wooldridge, and Nicholas R. Jennings. In *ACM Transactions on Economics and Computation, 2014*.

13. "How Bad is Selfish Voting?" [30], Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI 2013)*.

14. "Equilibria of Chinese Auctions", Simina Brânzei, Clara Forero, Kate Larson, and Peter Bro Miltersen. *Manuscript* (arXiv:1208.0296), 2012.

Part I

Cake Cutting

Chapter 2

Background

How should one fairly allocate resources among multiple economic players? The question of fair division is as old as civil society itself (Moulin [100]), with written instances of the problem dating back to thousands of years ago. For example, records of a fair division protocol can be found in books such as the Bible, which has references to the division of land and estate, and Hesiod's "Theogony" (cca. 750 B.C.), where a protocol known today as *Cut-and-Choose* is described through a dispute over a pile of meat between Prometheus and Zeus. Fair division has been studied in an extensive body of literature in economics, mathematics, political science [21, 100, 112, 127], and more recently, in computer science, as the fair allocation of resources is arguably relevant to the design of multiagent systems [44, 109]. Examples include manufacturing and scheduling, airport traffic, and industrial procurement [45, 109]. More recently, the problem of fair division is also motivated by the allocation of computational resources (such as CPU, memory, bandwidth) among users of shared computing systems[1] ([69, 79]), and has emerged as an important topic in artificial intelligence [6, 15, 27, 31, 41, 43, 46, 84, 108].

Cake cutting is a fundamental model in fair division. The *cake* is a metaphor for a heterogeneous divisible resource, such as land, time, memory in shared computing systems, clean water, greenhouse gass emissions, fossil fuels and other natural deposits. The problem is to fairly divide the resource among multiple participants, such that everyone is satisfied with their allocation.

The cake cutting literature typically represents the cake as the interval $[0, 1]$. There is a set of players $N = \{1, \ldots, n\}$, and each player $i \in N$ is endowed with a private *valuation function* V_i that assigns a value to every subinterval of $[0, 1]$. These values are induced by a non-negative continuous *value density function* v_i, so that for an interval I, $V_i(I) = \int_{x \in I} v_i(x)\, dx$. By

[1] In shared computing environments, resources such as CPU and memory get multiplexed such that each user can use their computing unit at their own pace and without concern for the activity of others accessing the system (this problem was stated by Fernando Corbato (1962) in the context of developing time-sharing operating systems).

definition, V_i satisfies the first two properties below; the third is an assumption that is made without loss of generality.

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

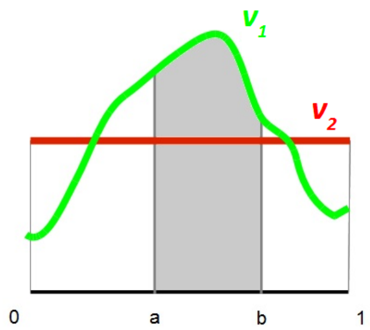

Figure 2.1: Cake cutting instance with two value density functions. The valuation of player 1 for the interval $[a, b]$ is given by the shaded area.

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

A *piece of cake* is a finite union of disjoint intervals. A piece is *contiguous* if it consists of a single interval. An *allocation* $A = (A_1, \ldots, A_n)$ is a partition of the cake among the players, that is, each player i receives the piece A_i, the pieces are disjoint, and $\bigcup_{i \in N} A_i = [0,1]$.

There are many fairness notions in the literature, such as proportionality and envy-freeness; we define a few of them next. Note that a protocol is said to have a property such as envy-freeness if each player i is *guaranteed* not to be envious by behaving *truthfully* in the protocol (i.e., not misrepresenting its private valuation function), regardless of what the other players do.

2.1 Fairness and Efficiency Properties

Definition 1 (Proportionality)**.** *An allocation $A = (A_1, \ldots, A_n)$ is said to be* proportional *if $V_i(A_i) \geq 1/n$ for all $i \in N$.*

Proportional allocations are guaranteed to exist on every instance with $n-1$ cuts (i.e. the allocation is contiguous), but no protocol can compute such allocations on general inputs with as few as $n-1$ cuts. The best possible proportional protocol (due to Even and Paz [109]) uses $O(n \log n)$ cuts. The best value that can be guaranteed with $n-1$ cuts is $\frac{1}{2n-2}$ and there exist such protocols (see Chapter 9.3, Robertson and Webb [112]).

Definition 2 (Envy-Freeness)**.** *An allocation $A = (A_1, \ldots, A_n)$ is said to be* envy-free *if $V_i(A_i) \geq V_i(A_j)$ for all $i, j \in N$.*

Note that envy-freeness implies proportionality when the entire cake is allocated. Without this requirement, one can obtain an envy-free allocation by simply throwing away the whole resource, which is clearly not proportional. On the other hand, proportionality is a weaker notion that does not necessarily guarantee envy-freeness.

Contiguous envy-free allocations are guaranteed to exist for any cake cutting instance (Stromquist [117]). This result is used several times in the thesis and we show next a proof given by Simmons [116] that works for hungry players and is based on Sperner's Lemma. The result is presented as described in a paper by Su [119].

The first useful notion is that of an n-simplex. A 0-simplex is a point, 1-simplex a line, 2-simplex a triangle, 3-simplex a tetrahedron, 4-simplex a pentatope. More generally, an n simplex is defined as the convex hull of $n+1$ affinely independent points in $\mathbb{R}^m$, where $m \geq n$. These points form the vertices of the simplex. A k-face of the n-simplex is defined as the k-simplex defined by a subset of $k+1$ vertices, and a facet is defined as any face spanned by n vertices.

A triangulation of an n-simplex $\mathcal{S}$ is a set of distinct smaller n-simplices whose union is $\mathcal{S}$, such that any two of the smaller simplices either intersect in a face common to both, or do not intersect at all. Let there be a numbering of the facets of $\mathcal{S}$ by $1, \ldots, n+1$. Given a triangulation of the simplex $\mathcal{S}$, consider a labelling as described next.

Definition 3 (Sperner labelling of an n-simplex)**.** *Given a triangulation of an n-simplex $\mathcal{S}$, label each vertex by one of the numbers in $\{1, \ldots, n+1\}$ such that the interior vertices can be labelled by any number and on the boundary of $\mathcal{S}$, none of the vertices on facet i are labelled i.*

Note that in a Sperner labelling the main vertices have distinct labels.

We say that an elementary simplex is *fully labelled* if all its vertices are labelled by different numbers. The following theorem holds.

Theorem 1 (Sperner's Lemma)**.** *Any Sperner-labelled triangulation of an n-simplex must contain an odd number of fully labelled elementary n-simplices. In particular, there is at least one.*

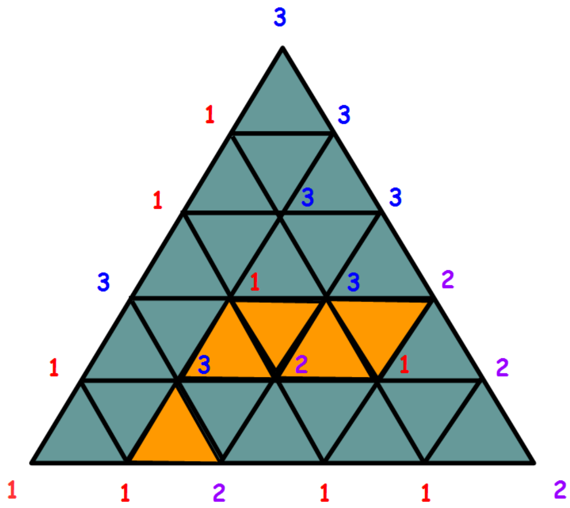

Figure 2.2: Sperner labelling for $n = 2$. There are five fully labelled triangles (highlighted).

An example of a Sperner labelling is given in Figure 2.2.

The existence of contiguous envy-free allocations can be proved by applying Sperner's Lemma as follows. Note that by definition of the cake cutting model, the preference sets are closed (in particular, no individual points of cake are valued strictly positively).

Theorem 2. *A contiguous envy-free division of the cake is guaranteed to exist for hungry players.*

Proof. (sketch) Let $N = \{1, \ldots, n\}$ be the set of players. Consider any division of the cake using $n-1$ cuts and let $x_1, \ldots, x_n$ denote the sizes (i.e. length) of the pieces generated this way, from left to right. Thus the first cut is at x_1, the second cut at $x_1 + x_2$, etc. Then $x_i \geq 0$ for all $i \in N$ and $x_1 + \ldots x_n = 1$. Let $\mathcal{S}$ denote the space of all possible partitions of the cake in this way (using $n-1$ cuts) and observe that $\mathcal{S}$ constitutes an $n-1$ simplex in $\mathbb{R}^n$. Given a partition of the cake $\mathbf{x} = (x_1, \ldots, x_n)$, we say that player i prefers a piece (demarcated by two adjacent cuts) if there is no other such piece that the player likes better.

Triangulate the simplex $\mathcal{S}$ by barycentric subdivision and assign to each vertex an owner chosen from one of the players, such that each elementary simplex is fully labelled by the set of players $1, \ldots, n$. An example for $n = 3$ is given in Figure 2.3 (In this example, the names of the players are $\{A, B, C\}$, and so each elementary triangle is labelled ABC).

Next, obtain a secondary labelling, in which each vertex is labelled by the index of the piece preferred by its owner. For example, in Figure 2.3,

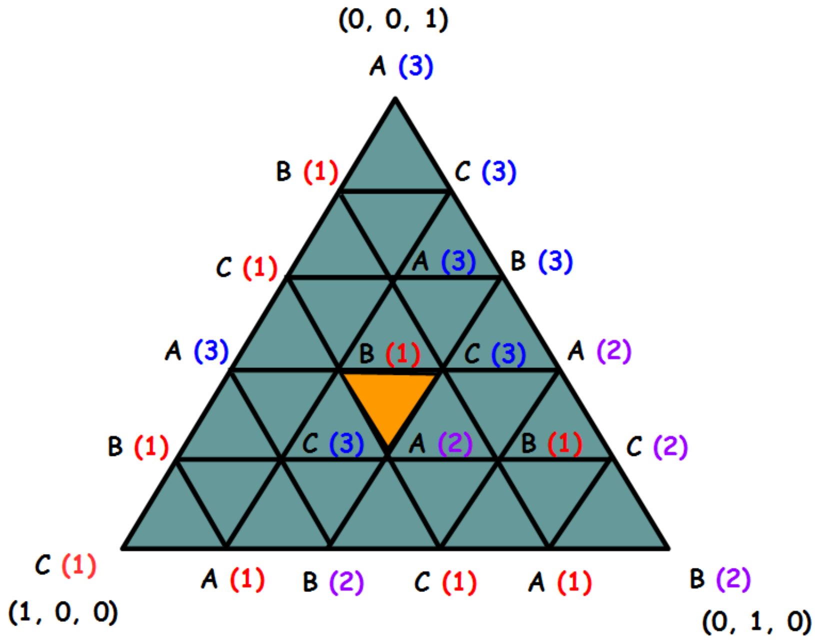

Figure 2.3: Sperner coloring. Each vertex is assigned one of the players as owner (A, B, or C). Afterwards, each vertex is labelled with the index of the piece preferred by its owner. By applying Sperner's Lemma, there exists a triangle (highlighted) for which every vertex is colored with a different number

the leftmost point (with coordinates $(1, 0, 0)$) is owned by player C and the favorite piece of player C is indexed 1, since it contains the entire cake. On the other hand, the top vertex (with coordinates $(0, 0, 1)$) is owned by player A and the favorite piece of this player is indexed by 3 (since again it consists of the full cake). This second labelling is a Sperner labelling; by applying Sperner's Lemma, we obtain that there exists an elementary simplex that is fully labelled, which is equivalent to each owner preferring a different piece. By taking a sequence of arbitrarily small triangulations and using the fact that the space of preferences is closed, we obtain that there exists a limit point $\mathbf{x}^* = (x_1^*, \ldots, x_n^*)$ at which the players prefer different pieces. This gives a contiguous envy-free allocation. □

The next fairness notion requires that the players are equally satisfied with their pieces.

Definition 4 (Equitability)**.** *An allocation $A = (A_1, \ldots, A_n)$ is said to be* equitable *if there exists a constant c such that $V_i(A_i) = c$, for all $i \in N$.*

One of the most demanding fairness notions, which combines the requirements of envy-freeness, proportionality, and equitability, is that of a perfect allocation.

Definition 5 (Perfect Allocation)**.** *An allocation $A = (A_1, \ldots, A_n)$ is said to be* perfect *if $V_i(A_j) = 1/n$ for all $i, j \in N$.*

Perfect allocations are guaranteed to exist and were shown to require at most $n(n-1)^2$ cuts by Alon [2]. However, there exists no protocol that can compute perfect partitions on general inputs and approximations have been devised instead [112].

Additional properties of allocations, often studied in addition to fairness, are *social welfare maximization* and *Pareto efficiency* (also known as *Pareto optimality*), which are introduced next.

Definition 6 (Social Welfare)**.** *The* social welfare *of an allocation $A = (A_1,$ $\ldots, A_n)$ is defined as the sum of the utilities of the players:*

$$SW(A) = \sum_{i=1}^{n} V_i(A_i)$$

Allocation A maximizes social welfare if there is no allocation B with a social welfare higher than that of A.

Pareto efficiency is a weaker notion compared to social welfare maximization and is often viewed as a minimal requirement for an allocation of resources to be deemed acceptable.

Definition 7 (Pareto efficiency)**.** *An allocation $A = (A_1, \ldots, A_n)$ is Pareto efficient if there exists no other allocation $A' = (A'_1, \ldots, A'_n)$, which is strictly better for one player without degrading the other ones. That is, it is not the case that $V_i(A'_i) > V_i(A_i)$ and $V_j(A'_j) \geq V_j(A_j)$ for all $j \neq i$.*

Cake cutting protocols are typically divided in two classes, of *discrete* and *continuous* (or moving-knife) procedures. Recently direct revelation protocols were also studied [43, 95, 109] in the context of mechanism design. Next we give examples of several fair division protocols, some of which are referenced in later chapters. All the classical protocols consist of a sequence of steps (queries) between a center and the players, at the end of which an allocation is output. Note that the valuations are private to the players and the fairness properties are guaranteed for a player i (regardless of the other players' strategies) provided that i reveals his valuation truthfully during the interaction.

2.2 Discrete Protocols

Discrete protocols are classified in *bounded*, *unbounded*, and *discrete infinite*. A discrete protocol is said to be *bounded* if on any cake cutting instance with

n players, it terminates in at most $F(n)$ steps (i.e. queries), for some function $F : \mathbb{N} \to \mathbb{N}$. A protocol is said to be *unbounded* if it always terminates after a finite number of steps, but the runtime depends on the valuations, i.e. the number of steps cannot be bounded by any function $F(n)$. Finally, a protocol is said to be *discrete infinite* if it requires an infinite (countable) number of steps; such a protocol guarantees a property in the limit.

Next we present several discrete bounded protocols: Cut-and-Choose, Selfridge-Conway, and Even-Paz.

One of the simplest cake cutting protocols is known as *Cut-and-Choose* [109]; it works for two players as follows:

> **Cut-and-Choose**: *Player 1 starts by dividing the cake in two contiguous pieces that he values equally. Player 2 chooses the piece that he prefers, then player 1 takes the remainder.*

It can be observed that *Cut-and-Choose* produces allocations that are both proportional and envy-free (in fact, the two notions coincide in the case of two players when the entire cake is allocated). Player 2 is not envious because he chooses his favorite piece, while player 1 is not envious because the two pieces generated are worth exactly 1/2 according to his valuation.

For $n = 3$ players, the simplest envy-free protocol is due to Selfridge and Conway [112]:

Selfridge-Conway:

1. *Player* 1 *cuts the cake into three equal parts according to his value.*
2. *Player* 2 *trims the most valuable of the three pieces such that there is a tie with the two most valuable pieces.*
3. *Set aside the trimmings.*
4. *Player* 3 *chooses one of the three pieces to keep.*
5. *Player* 2 *chooses one of the remaining two pieces to keep — with the stipulation that if the trimmed piece is not taken by player* 3*, player* 2 *must take it.*
6. *Player* 1 *takes the remaining piece.*
7. *Denote by* $T \in \{2, 3\}$ *the player that received the trimmed piece, and* $NT = \{2, 3\} \setminus \{T\}$.
8. *Players* NT *now cuts the trimmings into three equal parts in the player's value.*
9. *Players* T, 1*, and* NT *choose one of the three pieces to keep in that order.*

To see why the protocol is envy free, note that the division of three pieces in steps 4, 5, and 6 is trivially envy free. For the division of the trimmings in step 9, player T is not envious because it chooses first, and player NT is not envious because it was the one that cut the pieces (presumably, equally according to its value). In contrast, player 1 may prefer the piece of trimmings that player T received in step 9, but overall player 1 cannot envy T, because at best T was able to "reconstruct" one of the three original pieces that was trimmed at step 2, which player 1 values as much as the untrimmed piece it received in step 6.

The question of whether there exists a general envy-free protocol for $n \geq 4$ players had been open for over 50 years. In 1992, Brams and Taylor announced a positive answer to this question, by giving the first envy-free protocol for any number of players [20]. However, from a computational point of view, the Brams-Taylor protocol is problematic, because the runtime (i.e. the number of queries that the center needs to ask before being able to output an envy-free allocation) depends on the valuations. The question of whether there exists a bounded envy-free protocol for cake cutting is still open. The description of the Brams-Taylor protocol is not included here as it is very involved and outside the scope of this work; the interested reader can find descriptions in Robertson-Webb [112] and Brams-Taylor [21].

Next we present a well-known protocol for computing a proportional allocation for any number of players due to Evan and Paz [61].

Even-Paz: *For ease of exposition, assume that n is a power of two.*

1. *Each player marks the midpoint of the cake (according to their own valuations).*
2. *Divide the players in two subsets of equal size, such that all the cuts made by the players in the first subset are to the left of the cuts made by the players in the second subset.*
3. *Recurse (going back to step 1) with the players in the first subset with the piece of cake to the left of their rightmost cut, and the second subset with the piece to the right of their leftmost cut. Whenever a piece is claimed by a single player, allocate it to them.*

It can be verified that the procedure is proportional; whenever the set of players is halved, the cake remaining for any of the two subsets of players is worth at least half of the total value before the halving.

Finally, an example of an unbounded discrete protocols is the Brams-Taylor procedure for computing an envy-free allocation for $n \geq 4$ players (see Robertson-Webb [112], Chapter 10.3 and Brams-Taylor [21], Chapter 7.4).

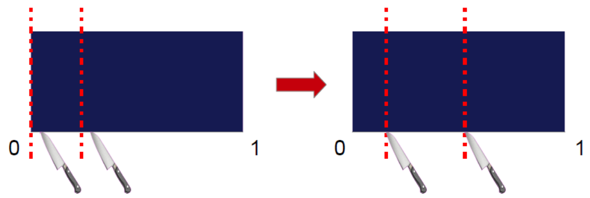

Figure 2.4: Austin's procedure

2.3 Continuous (Moving-Knife) Protocols

Continuous procedures involve one or more moving-knives that slide across the cake; the knives can be held either by the players or by the center.

We begin by illustrating continuous protocols with the Austin moving knife procedure, which computes a *perfect* allocation for two players.

> **Austin's procedure**: *A referee slowly moves a knife from left to right across the cake. At any point, a player can call stop. When a player called, a second knife is placed at the left edge of the cake. The player that shouted stop – say 1 – then moves both knives parallel to each other. While the two knives are moving, player 2 can call stop at any time. After 2 called stop, a randomly selected player gets the portion between player 1's knives, while the other one gets the two outside pieces*

To see that Austin's procedure results in a perfect allocation, note that player 1 can guarantee 1/2 if the first knife cuts the cake $50\% - 50\%$ and the division remains $50\% - 50\%$ throughout. On the other hand, player 2 can guarantee 1/2 if it calls stop when the middle piece is worth exactly 1/2.

Player 2 can find such a point since:

- The player didn't call stop exactly when the knife started moving, i.e. the left piece was worth less than 50% then.
- When the knife reaches 1, the piece between the knives is the complement of the initial piece.

By applying the intermediate value theorem, there exists a location of the knives where player 2's value is exactly 1/2 for both pieces, and the conclusion follows.

More generally, in the case of two players, exact allocations in any given ratio $k_1 : k_2$ can be computed using a moving knife procedure as outlined next (see Robertson and Webb, Chapter 5.4). Perfect partitions represent the special case of $k_1 = k_2$.

Moving Knife Exact Division for 2 players in ratio $k_1 : k_2$:

- ***Step 1****: One player moves two parallel knives across the cake, from left to right, always maintaining a value of* $\frac{k_1}{k_1+k_2}$ *between the knives. If at any point the other player agrees on the values of the pieces, the necessary exact shares are produced.*
- ***Step 2****: If no agreement is found in* Step 1*, repeat, this time keeping a value of* $\frac{k_2}{k_1+k_2}$ *between the knives. At some point the players will agree and the exact shares are produced.*

It can be seen that the procedure works by considering a few cases and applying again the intermediate value theorem. In general, one can devise protocols that compute approximations to some given ratios $k_1 : k_2 : \ldots : k_n$ for $n \geq 3$ players.

Next, we present the Dubins-Spanier moving knife protocol for computing a proportional allocation.

Dubins-Spanier: *A referee holds a knife and moves it slowly across the cake, from the left to the right endpoint. When the knife reaches a point such that one of the players has valuation exactly* $1/n$ *for the piece to the left of the knife, that player shouts* Cut!*. The first player to do so receives the left piece and exits. The remaining* $n-1$ *players repeat the procedure on the leftover cake, except that now they call cut when the perceived value of the piece to the left of the knife is* $1/(n-1)$ *(of the remaining cake).*

To argue that proportionality is guaranteed, note that the first player to be allocated gets a piece worth exactly $1/n$ according to his valuation. Moreover, the remaining cake is worth at least $(n-1)/n$ of the total value, and by an inductive argument, each player receives at least $1/n$ of the total. Unlike Austin's procedure, which cannot be discretized (no discrete protocol can compute a perfect partition even for two players), there exists a discrete analogue of the Dubins-Spanier protocol.

Finally, we present the Brams-Taylor-Zwicker protocol for computing an envy-free allocation when $n = 4$; it is bounded (in the sense that the number of steps has the same upper bound regardless of the valuations), but requires moving knives.

Brams-Taylor-Zwicker:

- ***Step 1****: Players* 1 *and* 2 *generate four pieces (not necessarily contiguous), each worth* $1/4$ *to both of them. This is accomplished by running Austin's procedure twice.*

- ***Step 2:*** *Player* 3 *trims one of the four pieces so that the two best pieces are equal in value (according to his measure). The trimmings are set aside and denoted by* Cake 2*, while the remainder represents* Cake 1*.*
- ***Step 3:*** (Allocation of Cake 1)*. The players choose from the four pieces in the order* 4*,* 3*,* 2*,* 1*, subject to the constraint that player* 3 *must take the trimmed piece if available.*
- ***Step 4:*** (Allocation of Cake 2)*. Let* $T \in \{3, 4\}$ *denote the player that took the trimmed piece and* NT *the other player from* $\{3, 4\}$*. Players* T *and* NT *run Austin's procedure twice on* Cake 2 *to generate four pieces equal from their point of view. The four pieces are allocated in the order* T*,* 1*,* NT*,* 2*.*

Saberi and Wang [113] gave a bounded protocol for $n = 5$ players, which also requires moving knives. Their algorithm can be discretized to obtain an ϵ-envy-free allocation in $O(\text{polylog}(1/\epsilon))$ steps.

2.4 The Robertson-Webb Model

All the discrete cake cutting protocols operate in a query model known as the Robertson-Webb model (Figure 2.5). The model was formalized mathematically by Woeginger and Sgall [126] and allows the following two types of queries between the protocol and the players:

- ***Cut***$(i; \alpha)$: Player i cuts the cake at a point y where $V_i([0, y]) = \alpha$. The point y becomes a *cut point*.
- ***Eval***$(i; y)$: Player i returns $V_i([0, y])$ where y is a previously made cut point.

The queries made by a protocol in the Robertson-Webb model may depend on the outputs of previous queries. At termination, the cut points define a partition of the cake into a finite set of intervals that the protocol allocates to the players in some specified way.

An alternative formalization was given by Procaccia [109]:

- ***Cut***$_i(x, \alpha)$:
  Player i outputs y such that $V_i([x, y]) = \alpha$.
- ***Eval***$_i(x, y)$:
  Player i outputs α such that $V_i([x, y]) = \alpha$.

The formalization by Procaccia is slightly more permissive (thus not equivalent to the one by Woeginger and Sgall), because it allows protocols in which the center can define cut points on the cake, such as:

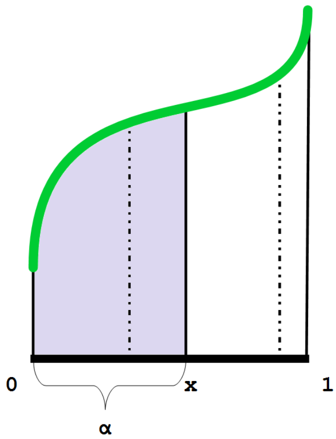

Figure 2.5: Cut query in the Robertson-Webb model. A designated player must report the point x such that its valuation for the interval $[0, x]$ is exactly α

"*Allocate* $[0, 0.5)$ *to player 1 and* $[0.5, 1]$ *to player 2*"

The Robertson-Webb model was employed in a body of work studying the complexity of cake cutting [58, 59, 84, 108, 126], where it was shown that the complexity of proportionality is $\Theta(n \log n)$ (the Even-Paz protocol gives the upper bound). The complexity of envy-freeness is still open; a lower bound of $\Omega(n^2)$ was given by Procaccia, which strictly separated envy-freeness from proportionality.

We illustrate the Robertson-Webb model with the well known *Cut and Choose* protocol, which computes an envy-free and proportional allocation for two players (Algorithm 1).

$x \leftarrow Cut\left(1; \frac{1}{2}\right)$
$\alpha \leftarrow Eval\,(2; x)$
if $\left(\alpha \geq \frac{1}{2}\right)$ **then**
  allocate $[0, x]$ to Player 1 and $[x, 1]$ to Player 2
else
  allocate $[0, x]$ to Player 2 and $[x, 1]$ to Player 1
end if

Algorithm 1: Cut-and-Choose protocol implemented in the Robertson-Webb model

2.5 Truthful Mechanisms

The classical cake cutting literature assumes that the players are honest and do not misrepresent their private value density function when answering the queries from the center. It is not hard to observe that even the simplest protocols can be manipulated; for instance, consider the example in Figure 2.6, where player 1 values the cake uniformly, while player 2 only values the interval $[\alpha, 1]$, where $1 > \alpha \gg 0$. Then instead of cutting at $1/2$ as instructed, player 1 can shift his cut point to α and receive a much larger piece, namely the entire interval $[0, \alpha]$.

A recent body of literature in computer science has investigated the design of *strategy-proof* mechanisms, i.e. mechanisms in which a player never benefits by misrepresenting their private valuation function, regardless of the strategies of the other players.

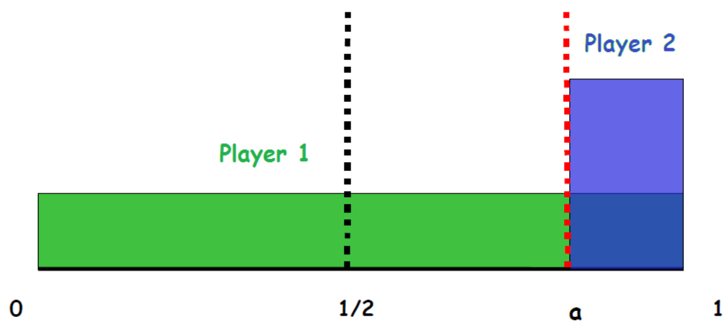

Figure 2.6: Manipulation in *Cut-and-Choose*. Player 1 has incentives to shift its cut point to the left of 1/2, which is the recommended strategy

The earliest such work is by Brams [26], who considers a weak version of strategy proofness. There, the players are risk averse and report their true valuations if there exists a choice of valuations of the other players such that the outcome would be worse by misreporting. Chen *et al.* [43] design strategy-proof mechanisms to compute envy-free and proportional allocations for restricted classes of valuation functions. The main results include a polynomial-time deterministic mechanism which computes an envy-free and proportional allocation for piecewise uniform valuations, and a polynomial-time randomized mechanism which is truthful in expectation, universally proportional, and universally envy-free for piecewise linear valuations. Maya and Nisan [95] study incentive-compatibility and Pareto-efficiency for two players, and provide characterizations and lower bounds on the social welfare attainable by any deterministic or randomized mechanism.

The work on strategy-proofness is focused on *direct revelation* mechanisms, in which the protocol (i.e. mechanism) consists of two steps:

1. Each player submits their entire value density function to the center.

2. The center outputs an allocation based on the reports.

Clearly, direct revelation protocols cannot always be implemented in the Robertson-Webb query model, while every Robertson-Webb protocol can be simulated by a direct revelation mechanism. We describe next one of the positive results in mechanism design for cake cutting, namely the protocol designed by Chen *et al.* [43] for the family of *piecewise uniform* valuations.

We say that a valuation function V_i is *piecewise constant* if the associated value density function v_i is piecewise constant; that is, the interval $[0, 1]$ can be partitioned into a finite number of intervals such that v_i is constant on each interval. We say that a valuation function V_i is *piecewise uniform* if the associated value density function, v_i, is piecewise constant and, moreover, on each interval, v_i is either zero or a fixed constant c_i (player-dependent). An example is given in Figure 2.7.

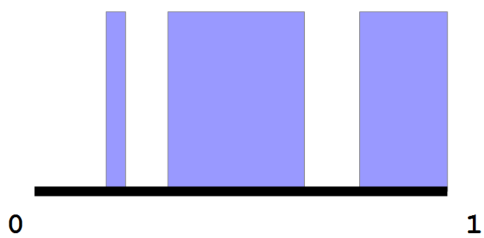

Figure 2.7: Piecewise uniform valuation

Piecewise constant valuations can be used to approximate to an arbitrary degree of accuracy any other value density function. While piecewise uniform valuations are much more restricted, they can still be used to represent problems where intervals (of time for example) are either acceptable or unacceptable [43]. In a recent paper, Kurokawa *et al.* [84] showed that piecewise uniform valuations are also useful for reasoning about the complexity of cake cutting; in particular, if there exists a bounded envy-free protocol for piecewise uniform valuations, then there exists an envy-free protocol for general valuations.

In [43], Chen *et al.* designed a truthful protocol for piecewise uniform valuations that runs in polynomial time (in the number of players and intervals in the representation) and computes a proportional, envy-free, and Pareto efficient allocation. The protocol is stated as Algorithm 2.

Theorem 3 (Chen, Lai, Parkes, Procaccia)**.** *Assume that the players have piecewise uniform valuation functions. Then there exists a truthful, proportional, envy-free Pareto-efficient, and polynomial time mechanism (Algorithm 2).*

Some notation required to understand the algorithm is as follows. Given a set of players $S \subseteq N$ with piecewise uniform valuations and a piece of cake X, let $D(S, X)$ denote the cake valued strictly positively by at least one of the players in S. Denote by:

$$\text{avg}(S, X) = \frac{\text{len}(D(S, X))}{|S|}$$

the average length of intervals in X desired by at least one player in S. We say that an allocation is *exact* with respect to S and X if it gives each player in S a piece of cake of length $\text{avg}(S, X)$ comprised only of desired intervals. Algorithm 2 is a recursive algorithm which computes exact allocations with respect to a selected subset of players, and then recurses on the remaining players and leftover cake. The proof that the implementation can be done in polynomial time is relatively involved and omitted from this introduction.

input: *Set* $N = \{1, \dots, n\}$ *of players with piecewise uniform valuations* $\mathbf{V} = (V_1, \dots, V_n)$
output: *Proportional, envy-free, and Pareto efficient allocation*
SUBROUTINE($N, [0, 1], \mathbf{V}$) // *call to a recursive subroutine*
....
procedure SUBROUTINE($S, X, \mathbf{V}$):
input: Set S of players, piece X of cake, valuations $\mathbf{V}$
if $S = \emptyset$ **then**
    return
end if
$S \leftarrow \text{argmin}_{S' \subseteq S} \text{avg}(S', X)$ // *breaking ties arbitrarily*
Let $E_1, \dots, E_n$ be an exact allocation with respect to S_{min} and X // *breaking ties arbitrarily*
for each $i \in S_{min}$ **do**
    $A_i \leftarrow E_i$
end for
SUBROUTINE($S \setminus S_{min}, X \setminus X_{min}, \mathbf{V}$)

Algorithm 2: Truthful and polynomial time mechanism for piecewise uniform valuations, which computes a proportional, envy-free, and Pareto efficient allocation (due to Chen, Lai, Parkes, Procaccia [43])

It is important to note that Algorithm 2 can discard cake and it is not clear that this behavior can be eliminated. That is, does there exist a truthful, polynomial time mechanism that computes an envy-free and Pareto efficient allocation of the entire cake for piecewise uniform valuations? The question of whether similar guarantees can be obtained for *hungry* players (e.g. with piecewise constant valuations) remains open.

Moving to *randomized mechanisms*, we say that a mechanism $\mathcal{M}$ is *truthful in expectation* if on every instance, the expected utility of a player i (taken over all random coin tosses of $\mathcal{M}$) is the best possible when behaving truthfully, regardless of the strategies of the other players.

Mossel and Tamuz [99] showed a randomized *direct revelation* protocol that is truthful in expectation and computes a *perfect* allocation, that is, an allocation $A = (A_1, \ldots, A_n)$ where $V_i(A_j) = 1/n$, $\forall i, j \in N$:

> **Mossel-Tamuz Mechanism**: *Given input valuations* $V_1, \ldots, V_n$*, find a perfect partition* $A = (A_1, \ldots, A_n)$ *and allocate it using a random permutation* π *over* $\{1, \ldots, n\}$ *(i.e. player* i *receives the piece* A_{π_i}*).*

Perfect partitions are guaranteed to exist with at most $n(n-1)^2$ cuts [2], so the Mossel-Tamuz mechanism is well-defined, but not constructive. For the restricted case of piecewise constant valuations, perfect partitions can be trivially computed: for each interval in the representation, divide it in n pieces of equal length and give a piece to each player.

Other recent work on socially optimal cake cutting includes papers by Caragiannis *et al.* [40], Cohler *et al.* [46], Brams *et al.* [27], Bei *et al.* [15]. Aziz and Ye [8] study Pareto efficiency and give several algorithms for piecewise uniform and piecewise constant representations, subject to additional constraints such as strategy-proofness and fairness. Zivan [128] studies the efficiency of cake cutting algorithms depending on the amount of trust among the players. Richer representations, such as PUML valuations, where the players have a minimum length requirement on the contiguous pieces received [41] blend the cake cutting problem and the allocation of indivisible resources.

Chapter 3

Equilibrium Analysis

A standard assumption in classical cake cutting protocols is that the players do not know each other's preferences. However, in many real-world settings, the participants *do know* each other's preferences. For example, when countries divide land at the end of a war, it is usually common knowledge which areas of land are preferred by which country. Thus, when a general protocol is employed to produce an allocation of the cake, the players may be able to improve their utility by being strategic if they know the others' valuations during the execution of the algorithm. In this chapter, we initiate the study of equilibria of classical cake cutting protocols.

While the classical protocols are not necessarily strategy-proof, they are often very simple, elegant, and designed so that the players can easily implement them by following a sequence of natural steps. One of the most intuitive and best-known procedures for computing a proportional allocation of the cake is the Dubins-Spanier procedure (introduced in Chapter I). Note that to make the protocol completely precise, a tie breaking rule has to be specified for the case of two players calling cut simultaneously. No matter how such tie breaking is defined, it is easily verified that the allocation produced by the Dubins-Spanier moving-knife procedure is proportional. However, it is not necessarily envy-free.

In this chapter, we consider the strategic version of Dubins-Spanier, which we refer to as the *moving knife game*. The players know each other's valuations and compete against each other to maximize their allocations. In the moving knife game, a player would like to delay as much as possible the moment of calling cut, since the longer they wait, the better the piece to the left of the knife becomes. However, if they wait for too long, someone else may call cut before they do and take the piece instead. The moving knife game is related to games of timing [86], such as war of attrition models, in which the decision of each player is when to quit and victory belongs to the player that held on longer, and preemption games, in which the players prefer to stop first.

It seems very challenging to characterize all the equilibria of the moving knife game if it is modeled as a continuous time extensive form game in the obvious way. Instead, we analyze the game when the players are restricted to use *threshold strategies*, defined as follows. The moving knife game proceeds in n rounds, corresponding to each of the time intervals between consecutive cut points. Each player has n thresholds, one for every round. A player calls cut in a given round when the value of the piece to the left of the knife is equal to the player's threshold for that round. Note that threshold strategies is a simple generalization of the prescribed behavior in the orginal Dubins-Spanier protocol – in particular, the classical Dubins-Spanier procedure outlined above can be viewed as playing the moving knife game with all players using the sequence of thresholds $\left(\frac{1}{n}, \frac{1}{n-1}, \frac{1}{n-2}, \ldots, 1\right)$.

Our main result is a direct correspondence between the equilibria of the moving knife game and envy-free allocations of the cake with contiguous pieces, when players are restricted to threshold strategies.

That is, every pure Nash equilibrium of the moving knife game induces an envy-free allocation with contiguous pieces which contains the entire cake. Moreover, every envy-free allocation with contiguous pieces of the entire cake can be mapped to a pure Nash equilibrium of the corresponding moving knife game, when ties are broken in a particular way. This result can be viewed as an affirmative answer to the natural question: "Can fair allocations arise as equilibria of simple and natural protocols?" The question of designing a game such that its equilibria correspond to desirable allocations of the cake was also considered by Ianovski [77].

3.1 Moving Knife Game

We now introduce the moving knife game. Given a cake with the corresponding value density functions, a knife moves continuously from the left to the right endpoint of the cake. The game is divided in n rounds. Each player i has a strategy consisting of n thresholds, $T_i = [t_{i,1}, \ldots, t_{i,n}] \in [0,1]^n$, one threshold for each round. Player i calls cut in round j when the piece to the left of the knife is worth exactly $t_{i,j}$ according to i's valuation. The player to call cut first receives the piece to the left of the knife. When multiple players call cut simultaneously, the piece is given to the player who comes first in a tie-breaking rule $\pi = (\pi_1, \ldots, \pi_n)$, which is a fixed permutation of $(1, \ldots, n)$. Once a player has received a piece, he exits and the game continues from that point on with the remaining players and leftover cake.

Given a tuple (N, v, T, π), where N is a set of players, v are the value density functions, T are the strategies, and π is a tie-breaking rule, the *induced allocation* $X = (X_1, \ldots, X_n)$ results from playing the moving knife game under the tie-breaking rule π, such that each player $\sigma_i \in N$ receives the piece X_i, for some ordering σ of the players.

Finally, we say that a player is *active* at a round if the player has not exited the game in the previous rounds.

We illustrate the game with the following example.

Example 1. Let $N = \{1, 2\}$. Consider the following value density functions:

- $v_1(x) = 1, \forall x \in [0, 1]$
- $v_2(x) = \frac{1}{4}, \forall x \in \left[0, \frac{1}{3}\right]$ and $v_2(x) = \frac{11}{8}, \forall x \in \left[\frac{1}{3}, 1\right]$

Let $T = [T_1, T_2]$, where $T_1 = \left[\frac{1}{2}, \frac{2}{3}\right]$ and $T_2 = \left[\frac{1}{12}, \frac{2}{3}\right]$. Then in:

- *Round* 1: Player 2 calls cut first at $\frac{1}{3}$, since $V_2\left(\left[0, \frac{1}{3}\right]\right) = t_{2,1} = \frac{1}{12}$. Player 1 does not get to call cut in this round, since:
$$V_1\left(\left[0, \frac{1}{3}\right]\right) < t_{1,1} = \frac{1}{2}.$$
- *Round* 2: Player 1 is the only one left, and the leftover cake is $\left[\frac{1}{3}, 1\right]$. Player 1 calls cut at 1, since:
$$V_1\left(\left[\frac{1}{3}, 1\right]\right) = t_{2,1} = \frac{2}{3}.$$

The induced allocation is $X = (X_1, X_2)$, where player 2 receives $X_1 = \left[0, \frac{1}{3}\right]$ and player 1 receives $X_2 = \left[\frac{1}{3}, 1\right]$.

A strategy profile $T = [T_1, \ldots, T_n] \in [0, 1]^{n \times n}$ is a *pure Nash equilibrium* under a tie-breaking rule π if no player $i \in N$ can receive a better allocation by deviating to $T_i^{'} \neq T_i$. That is, $u_i(T) \geq u_i(T_i^{'}, T_{-i}), \forall T_i^{'} \in [0, 1]^n$.

In the following example, we illustrate how the players can be strategic during the execution of the moving knife game, i.e., we illustrate that it is not necessarily a Nash equilibrium that all players play the Dubins-Spanier strategy $[1/n, \ldots, 1/2, 1]$. Consider the scenario where player 1 has a uniform valuation over the cake and just wants as much of it as possible, while player 2 only likes a very thin slice at the right end. Then player 1 can delay the moment of calling cut, since he knows that player 2 is following the Dubins-Spanier recommendation and will only call cut close to the right endpoint. A precise version of this example follows:

Example 2. Let $N = \{1, 2\}$. Consider the following value density functions:

- $v_1(x) = 1, \forall x \in [0, 1]$
- $v_2(x) = 0, \forall x \in \left[0, \frac{3}{4}\right]$ and $v_2(x) = 4, \forall x \in \left[\frac{3}{4}, 1\right]$.

Under the Dubins-Spanier protocol, player 1 calls cut first at $\frac{1}{2}$. The resulting allocation is $X = (X_1, X_2)$, with $X_1 = \left[0, \frac{1}{2}\right]$ and $X_2 = \left[\frac{1}{2}, 1\right]$, with utilities: $V_1(X_1) = \frac{1}{2}$ and $V_2(X_2) = 1$.

However, player 1 can improve his utility by waiting and calling cut at $\frac{3}{4}$ instead. Then the allocation is $X' = (X_1', X_2')$, with $X_1' = \left[0, \frac{3}{4}\right]$ and $X_2' = \left[\frac{3}{4}, 1\right]$. The new utilities are $V_1(X_1') = \frac{3}{4}$ and $V_2(X_2') = 1$.

3.2 Exact Equilibria

In this section, we analyze the pure Nash equilibria of the moving knife game, for any fixed hungry valuations (i.e. $v_i(x) > 0$, $\forall x \in [0, 1]$, $\forall i \in N$).

First, the original result of Dubins and Spanier immediately yields the following proposition.

Proposition 1. *In any pure Nash equilibrium of the moving knife game, each player's utility is at least* $1/n$ *and the entire cake is allocated to the players.*

Proof. Suppose a player gets a smaller utility in Nash equilibrium. Then he can deviate by playing the strategy prescribed in the original Dubins-Spanier protocol, i.e., $[1/n, \ldots, 1/2, 1]$, improving his utility to at least $1/n$, and contradicting that a Nash equilibrium is played. Also, if the entire cake is not allocated, the last player's last threshold is strictly smaller than 1. He can therefore deviate to threshold 1 and receive a larger utility, contradicting that a Nash equilibrium is played. □

Now we show that the existence of Nash equilibrium crucially depends on the tie breaking rule used. That is, there exist tie-breaking rules where the moving knife game does not have a pure Nash equilibrium:

Proposition 2. *There exist a tie breaking rules and value density functions so that the corresponding moving knife game does not have a pure Nash equilibrium.*

Proof. Let $N = \{1, 2\}$, with tie-breaking order $(1, 2)$, and value density functions:

- $v_1(x) = \frac{1}{4}$, $\forall x \in \left[0, \frac{4}{5}\right]$ and $v_1(x) = 37.5x - 29.75$, $\forall x \in \left[\frac{4}{5}, 1\right]$
- $v_2(x) = 1$, $\forall x \in [0, 1]$.

Assume there exists a profile of threshold strategies in equilibrium, T, such that the first cut is made at $x \in (0, 1]$. We analyze the case where the cut at x is made in round 1; the case where the cut is made in round 2 is similar.

First, T must be such that both players call cut at x simultaneously. Otherwise, if $t_{1,1} = V_1([0, x])$, while $t_{2,1} < V_2([0, x])$, then player 1 can increase his threshold to $t'_{1,1} = t_{1,1} + \epsilon$, for small enough $\epsilon > 0$, and receive a strictly better piece, $[0, x']$, where $x' > x$.

Similarly, if $t_{2,1} = V_2([0, x])$, while $t_{1,1} > V_1([0, x])$, then player 2 can increase his threshold and get a strictly better piece $[0, x']$. Thus $t_{1,1} = V_1([0, x])$ and $t_{2,1} = V_2([0, x])$. Since the tie-breaking rule is $(1, 2)$, players 1 and 2 receive pieces $[0, x]$ and $[x, 1]$, respectively.

In addition, we have that:

$$V_1([0, x]) \geq V_1([x, 1]), \tag{3.1}$$

since otherwise player 1 can deviate by setting $t_{1,1} = 1$ – the deviation would ensure that player 1 receives a better piece in round 2. Similarly, it can be shown that:

$$V_2([x, 1]) \geq V_2([0, x]). \tag{3.2}$$

However, inequalities (3.1) and (3.2) cannot be met simultaneously for the given valuations. Thus the pure Nash equilibrium T cannot exist. □

We show that in a pure Nash equilibrium, then in each of the first $n - 1$ rounds, the player who is allocated a piece has a competitor that calls cut simultaneously in that round.

Proposition 3. *Let a moving knife game with hungry valuations be given. Let T be a profile of threshold strategies in equilibrium under a deterministic tie-breaking rule. Then, in every round except the last one, the player who is allocated the piece has an (active) competitor that calls cut simultaneously.*

Proof. Let $X = (X_1, \ldots, X_n)$ be the allocation induced by T, such that player σ_i receives the piece $X_i = [x_{i-1}, x_i]$. It follows by Proposition 1 that X contains the entire cake and $X_i \neq \emptyset$, $\forall i \in N$. Assume by contradiction that there exists a round $i < n$ in which only player σ_i calls cut at x_i. Then it must be the case that

$$t_{\sigma_j,i} > V_{\sigma_j}([x_{i-1}, x_i]), \forall \sigma_j \in N \setminus \{\sigma_i\}.$$

By the continuity of the valuation functions, there exists $\epsilon > 0$ such that by deviating to threshold:

$$t'_{\sigma_i,i} = t_{\sigma_i,i} + \epsilon$$

in round i, player σ_i is guaranteed a strictly better piece, $[x_{i-1}, x'_i]$, where $x'_i > x_i$. This is a contradiction with T being in equilibrium. Thus every player who receives a piece in the first $n - 1$ rounds has a competitor that calls cut simultaneously in that round. □

Finally, every profile of threshold strategies in equilibrium induces an envy-free allocation. We first show the following proposition.

Proposition 4. *Let a moving knife game with hungry valuations be given. Let T be a profile of threshold strategies under a deterministic tie-breaking rule, such that each player σ_i receives a piece in round i and in every round except the last, there exist two active players who call cut simultaneously. Then if some player σ_i deviates to $T^{'}_{\sigma_i} \neq T_{\sigma_i}$ and receives a new piece in some round k under $T^{'} = (T^{'}_{\sigma_i}, T_{-\sigma_i})$, then the set of cuts made in the first $k-1$ rounds are the same under T and $T^{'}$.*

Proof. Let $X = (X_1, \ldots, X_n)$ be the allocation induced by T, where the piece $X_i = [x_{i-1}, x_i]$ is given to player σ_i. Let $T^{'}_{\sigma_i}$ be the new sequence of thresholds used by player σ_i, where

$$t^{'}_{\sigma_i,k} = 1, \forall k \in \{1, \ldots, j\}$$

Since player σ_i did not receive a piece in the first $i-1$ rounds under T, and does not call cut before other players under $T^{'} = (T^{'}_{\sigma_i}, T_{-i})$, it follows that the allocation $X^{'}$ (induced by $T^{'}$) is identical to X for the first $i-1$ pieces. If $j < i$, then the statement of the proposition follows immediately.

Otherwise, $j \geq i$. By condition 3 of the proposition, there exists a player $\sigma_{r_1} \neq \sigma_i$ who also calls cut at x_i in round i, and is second after σ_i in the tie-breaking rule among the players that call cut at x_i. That is,

$$t_{\sigma_{r_1},i} = V_{\sigma_{r_1}}([x_{i-1}, x_i]).$$

Then player σ_{r_1} receives the piece $[x_{i-1}, x_i]$ under $T^{'}$.

The allocations made in rounds $i+1, \ldots, r_1 - 1$ are identical under $T^{'}$ and T, since the same players that received the pieces

$$X_{i+1}, \ldots, X_{r_1-1}$$

under T continue to call cut at the points:

$$x_{i+1}, \ldots, x_{r_1-1},$$

respectively, and to win the ties (if any) under $T^{'}$. The piece X_{r_1} is taken by some player σ_{r_2}, which was second in the tie for receiving the piece X_{r_1} under T.

Iteratively, it can be shown that in the first j rounds, the same cuts are made under T and $T^{'}$, and this set is $\{x_1, \ldots, x_j\}$. □

Theorem 4. *Consider a moving knife game with hungry valuations and deterministic tie-breaking. Then every pure Nash equilibrium of the game induces an envy-free allocation.*

Proof. Let T be a profile of threshold strategies in equilibrium under tie-breaking rule $\pi = (\pi_1, \ldots, \pi_n)$. Let $X = (X_1, \ldots, X_n)$ be the induced allocation, such that piece $X_i = [x_{i-1}, x_i]$ is given to player σ_i, $\forall i \in N$.

Assume by contradiction that X is not envy-free. Since the empty allocation is envy-free, it follows by Proposition 1 that X contains the entire cake. Then there exists a player σ_i such that

$$V_{\sigma_i}([x_{j-1}, x_j]) > V_{\sigma_i}([x_{i-1}, x_i]),$$

for some $j \in N \setminus \{i\}$. By continuity of the valuation functions, there exists $\epsilon > 0$ such that

$$V_{\sigma_i}([x_{j-1}, x_j - \epsilon]) > V_{\sigma_i}([x_{i-1}, x_i]).$$

We consider two cases:

1. $(j < i)$: Then player σ_i can deviate to strategy profile $T^{'}_{\sigma_i}$, where

$$t^{'}_{\sigma_i,k} = \begin{cases} V_{\sigma_i}([x_{j-1}, x_j - \epsilon]) & \text{if k = j} \\ t_{\sigma_i,k} & \text{otherwise} \end{cases}$$

   Under $T^{'} = (T^{'}_{\sigma_i}, T_{-\sigma_i})$, player σ_i is guaranteed to receive the piece $[x_{j-1}, x_j - \epsilon]$, since no other player calls cut before x_j in round j. This deviation improves σ_i's utility, contradiction with T being in equilibrium.

2. $(j > i)$: Then player σ_i can deviate to strategy profile $T^{'}_{\sigma_i}$, where

$$t^{'}_{\sigma_i,k} = \begin{cases} V_{\sigma_i}([x_{j-1}, x_j - \epsilon]) & \text{if k = j} \\ 1 & \text{otherwise} \end{cases}$$

   By Proposition 4, the same cuts are made under T and $T^{'}$ in the first $j-1$ rounds, and this set is $\{x_1, \ldots, x_{j-1}\}$.

   Then player σ_i receives the piece $[x_{j-1}, x_j - \epsilon]$ in round j, which strictly improves σ_i's utility, since:

$$V_{\sigma_i}([x_{j-1}, x_j - \epsilon]) > V_{\sigma_i}([x_{i-1}, x_i]).$$

   This is a contradiction with T being in equilibrium.

From Case 1 and 2, it follows that the assumption must have been false, and so the induced allocation is envy-free. □

We can now characterize the set of pure Nash equilibria as follows.

Theorem 5. *Consider a moving knife game with hungry valuations. A strategy profile T is in Nash equilibrium under a deterministic tie-breaking rule if and only if the induced allocation contains the entire cake and is envy-free and in every round except the last one, the player who is allocated the piece has an active competitor that calls cut simultaneously.*

Proof. Let T be a profile of thresholds strategies.

($\Rightarrow$): If T is a pure Nash equilibrium under some deterministic tie-breaking rule, then by Proposition 1 and Theorem 4, it follows that the induced allocation contains the entire cake and is envy-free. Also, by Proposition 3, in every round except the last, the player who is allocated the piece has an active competitor that calls cut simultaneously.

($\Leftarrow$): If T verifies the conditions of the theorem, then we claim it is a pure Nash equilibrium. Let $X = (X_1, \ldots, X_n)$ be the induced allocation, where piece $X_i = [x_{i-1}, x_i]$ is given to player σ_i, $\forall i \in N$.

Assume by contradiction that there exists a player σ_i who can improve by deviating to $T^{'}_{\sigma_i} \neq T_{\sigma_i}$. Let k be the round in which σ_i receives a piece when playing $T^{'}_{\sigma_i}$. Since σ_i does not receive a piece in the first $k-1$ rounds under $T^{'}_{\sigma_i}$, we can assume without loss of generality that:

$$t^{'}_{\sigma_i,l} = 1, \forall l \in \{1, \ldots, k-1\}.$$

By Proposition 4, the cut made in round $k-1$ was at x_{k-1}, and one of the following conditions holds:

- $x_k = 1$, or
- there exists a player $\sigma_j \neq \sigma_i$ who calls cut at x_k (in round k) when σ_i plays $T^{'}_{\sigma_i}$.

Thus the highest value that σ_i can receive in round k is $V_{\sigma_i}(X_k)$. By envy-freeness of X, we have that:

$$V_{\sigma_i}(X_i) \geq V_{\sigma_i}(X_k),$$

thus the deviation does not improve σ_i's utility. Thus, T is an equilibrium, which concludes the proof of the theorem. $\square$

Next, we show that for every moving knife game with strictly positive value density functions, a pure Nash equilibrium is guaranteed to exist for *some* deterministic tie-breaking rule. In fact, we show that for any envy-free allocation of the cake, there exists is a pure Nash equilibrium that induces this allocation. This implies existence of a pure Nash equilbrium, as an envy-free allocation of the cake with $n-1$ cuts is guaranteed to exist (see Stromquist [118]).

Theorem 6. *Consider a moving knife game with hungry valuations. Given any envy free allocation of the cake with $n-1$ cuts, there exists a deterministic tie-breaking rule π such that the game has a pure Nash equilibrium inducing this allocation.*

Proof. In an envy free allocation of the cake with $n-1$ cuts, each player gets a contiguous pice. That is, there exists a permutation $\pi = (\pi_1, \ldots, \pi_n)$ of N and numbers x_i such that player π_i receives the piece $X_i = [x_{i-1}, x_i]$. Now use π as the tie-breaking rule for the moving knife game and consider the strategy sets:

$$T_i = [t_{i,1}, \ldots, t_{i,n}],$$

where

$$t_{i,k} = V_i([x_{k-1}, x_k]), \forall i, j \in N.$$

It can be verified that the strategies in T verify the conditions of Theorem 5. That is, the induced allocation is envy-free, contains the entire cake, and in every round except the last, the player winning the piece has a competitor who calls cut simultaneously. Thus the set of strategies T are in equilibrium under the tie-breaking rule π. □

This completes the proof of our main result: Any pure Nash equilibrium of the moving nice game induces an envy-free allocation and any envy-free allocation is induced by some pure Nash equilibrium.

3.3 Achieving tie breaking rule independence

The dependence of the existence of Nash equilibrium on the tie breaking rule is an annoying (but unavoidable) flaw of our main result: The tie-breaking rule requires information about the valuation functions of the players in order for a non-trivial pure Nash equilibrium to exist. Clearly, in many natural settings, the tie-breaking rule is given exogenously. For example, when countries divide land at the end of a war, some countries may have higher priority than others due to prior bilateral agreements that have been signed.

It is interesting to understand the special cases where a pure Nash equilibrium is guaranteed to exist, no matter which tie breaking rule is used. We have first the following simple observation.

Proposition 5. *Consider a moving knife game with players that have* identical *hungry valuations. Then the game has a pure Nash equilibrium under* every *deterministic tie-breaking rule.*

Proof. Consider an envy-free division of the cake with $n-1$ cuts, $[x_0, \ldots, x_n]$. The players have identical value density functions, and so

$$V_i([x_{j-1}, x_j]) = \frac{1}{n}, \forall i, j \in N.$$

For any tie-breaking rule π, construct an allocation $X = (X_1, \ldots, X_n)$, such that player π_i receives the piece $X_i = [x_{i-1}, x_i]$. By applying Theorem 8 to the envy-free allocation X, it follows that the game has a pure Nash equilibrium under π. □

Next, we show that for arbitrary strictly positive value density functions and every possible tie-breaking rule, including, for example, randomized or round-dependent rules, there exists an *approximate* equilibrium in pure strategies such that the induced allocation is *approximately* envy-free and contains the entire cake.

We say that a set of strategies $T = [T_1, \ldots, T_n] \in [0,1]^{n \times n}$ is an ϵ-equilibrium if for every $i \in N$, player i cannot improve his utility by more than ϵ by deviating to $T_i^{'} \neq T_i$. That is, $u_i(T_i^{'}, T_{-i}) \leq u_i(T) + \epsilon$.

Theorem 7. *Consider a moving knife game with hungry valuations. Then for every tie-breaking rule, the game has an ϵ-equilibrium in pure strategies such that the induced allocation is ϵ-envy-free and contains the entire cake.*

Proof. Let $\epsilon > 0$ and $X = (X_1, \ldots, X_n)$ an envy-free allocation of the entire cake, where player π_i receives the piece $X_i = [x_{i-1}, x_i]$, $\forall i \in N$.

Starting from X, we construct an allocation $Z = (Z_1, \ldots, Z_n)$, where player π_i receives the piece $Z_i = [z_{i-1}, z_i]$, $\forall i \in N$, such that Z is induced by an ϵ-equilibrium T, contains the entire cake, and is ϵ-envy-free. The idea of the proof is similar to that of Theorem 8. To avoid tie-breaking, we construct the thresholds such that for every round, the active players would call cut immediately *after* the player who is supposed to receive an allocation in that round. Thus, if a player π_i deviates to a new sequence of thresholds $T_{\pi_i}^{'}$ and receives a new piece in round $k \neq i$, then the set of cuts made in rounds $\{1, \ldots, k-1\}$ are *approximately the same* under T and $T^{'} = (T_{\pi_i}^{'}, T_{-\pi_i})$. That is, the following hold for the allocation induced by $T^{'}$:

- If π_i still receives a piece in round i, then the improvement cannot be larger than ϵ, since another active player will call cut immediately after π_i's expected cut point in T.

- If π_i receives a piece in round $k < i$, then π_i's new piece is a subset of Z_k, and so the improvement cannot be greater than ϵ by ϵ-envy-freeness of Z.

- If π_i receives a piece in round $k > i$, then the set of cuts made in rounds $\{1, \ldots, k-1\}$ are approximately the same under T and $T^{'}$, and so π_i's new piece is approximately a subset of Z_k. Again the improvement cannot be better than ϵ by ϵ-envy-freeness of Z.

Formally, the profile of threshold strategies T is defined as follows. Let $z_n = x_n$. The valuation functions are continuous and bounded, thus there exists $z_{n-1} \in (x_{n-2}, x_{n-1})$ such that:

$$V_j([z_{n-1}, x_{n-1}]) < \frac{\epsilon}{2}, \forall j \in N.$$

We construct a set of points:

$$y_{1,k}, \ldots, y_{n,k}, z_k$$

for all rounds k, such that the threshold of each player j is set to call cut at $y_{j,k} \in [z_{k-1}, x_{k-1})$ in round k. Define $y_{j,n} = z_n$, $\forall j \in N$.

Consider round $n-1$, and let

$$y_{j,n-1} = \begin{cases} z_{n-1} & \text{if } j = \pi_{n-1} \\ \frac{z_{n-1}+x_{n-1}}{2} & \text{otherwise} \end{cases}$$

We now construct z_{n-2}. For each $j \in N$, there exists

$$z_{j,n-2} \in (x_{n-3}, x_{n-2})$$

such that

$$V_j\left([z_{j,n-2}, x_{n-2}]\right) < V_j([y_{j,n-1}, x_{n-1}])$$

Define $z_{n-2} = \max_{j \in N} z_{j,n-2}$. For all $j \in N$, we have:

$$\begin{aligned} V_j([z_{n-2}, x_{n-2}]) &< V_j([y_{j,n-1}, x_{n-1}]) \\ &< V_j([z_{n-1}, x_{n-1}]) < \frac{\epsilon}{2} \end{aligned}$$

Iteratively, for all rounds k from $n-2$ to 1, we construct points

$$y_{1,k}, \ldots, y_{n,k}, z_{k-1}$$

in a manner similar to the construction for round $n-1$, such that the following conditions are met:

- $z_{k-1} < x_{k-1}$, if $k \in \{2, \ldots, n-1\}$, and $z_{k-1} = x_{k-1}$, if $k = 1$
- $x_{k-1} < z_k \leq y_{1,k}, \ldots, y_{n,k} < x_k$
- $V_j([z_{k-1}, x_{k-1}]) < V_j([y_{j,k}, x_k])$, $\forall j \in N$.

Consider the profile of threshold strategies T, given by:

$$t_{j,k} = V_j([z_{k-1}, y_{j,k}]), \forall j, k \in N$$

Let Z be the allocation induced by T, where player π_i receives the piece $Z_i = [z_{i-1}, z_i]$, $\forall i \in N$. We claim that T is an ϵ-equilibrium and Z is ϵ-envy-free.

First, we show that T is an ϵ-equilibrium. Assume by contradiction that there exists player π_i who can improve his utility more than ϵ by deviating to $T^{'}_{\pi_i}$. Let k be the round in which player π_i is allocated a piece under $T^{'} = (T^{'}_{\pi_i}, T_{-\pi_i})$. We show by induction that in each previous round $l < k$, a cut is made in the interval $[z_l, x_l)$. For $l = 1$ the statement trivially holds, since:

$$0 < t_{j,1} = V_j([0, y_{j,1}]) < V_j([0, x_1]), \forall j \in N$$

Assume the property holds for all rounds $1, \ldots, l-1$. By the induction hypothesis, a cut was made in round $l-1$ in the interval $[z_{l-1}, x_{l-1})$. For each player j, the threshold in round l is such that:

$$\begin{aligned} t_{j,l} &= V_j([z_{l-1}, y_{j,l}]) \\ &= V_j([z_{l-1}, x_{l-1}]) + V_j([x_{l-1}, x_l]) - V_j([y_{j,l}, x_l]) \\ &< V_j([x_{l-1}, x_l]) \end{aligned}$$

Note that the inequality:

$$V_j([z_{l-1}, x_{l-1}]) < V_j([y_{j,l}, x_l])$$

holds by Condition 3. Thus, in round l, every remaining player j will call cut

- no earlier than $y_{j,l}$ if in the previous round the cut was made at z_{l-1}
- strictly before x_l if in the previous round the cut was made at x_{l-1}

Thus the statement also holds for round l. It follows that the cut in round $k-1$ was made in the interval $[z_{k-1}, x_{k-1})$. Moreover, all the remaining players will call cut before x_k in round k. Then, using the envy-freeness of allocation X, we can bound the utility of π_i as follows:

$$\begin{aligned} u_i(T') &< V_{\pi_i}([z_{k-1}, x_k]) \\ &= V_{\pi_i}([z_{k-1}, x_{k-1}]) + V_{\pi_i}([x_{k-1}, x_k]) \\ &\leq \frac{\epsilon}{2} + V_{\pi_i}([x_{i-1}, x_i]) \\ &\leq \frac{\epsilon}{2} + V_{\pi_i}([x_{i-1}, x_i]) + \left(V_{\pi_i}([z_{i-1}, x_{i-1}]) - V_{\pi_i}([z_i, x_i]) + \frac{\epsilon}{2}\right) \\ &= V_{\pi_i}([z_{i-1}, z_i]) + \epsilon \\ &= u_i(T) + \epsilon \end{aligned}$$

Thus player i cannot improve by more than ϵ by deviating.

Finally, we show that the induced allocation is ϵ-envy-free. For every two players π_i and π_j the following hold:

$$\begin{aligned} V_{\pi_i}(Z_i) &= V_{\pi_i}([z_{i-1}, z_i]) \\ &= V_{\pi_i}([x_{i-1}, x_i]) + V_{\pi_i}([z_{i-1}, x_{i-1}]) - V_{\pi_i}([z_i, x_i]) \\ &\geq V_{\pi_i}([x_{j-1}, x_j]) - \frac{\epsilon}{2} \\ &\geq V_{\pi_i}([z_{j-1}, z_j]) + V_{\pi_i}([z_j, x_j]) - V_{\pi_i}([z_{j-1}, x_{j-1}]) - \frac{\epsilon}{2} \\ &\geq V_{\pi_i}([z_{j-1}, z_j]) - \epsilon \\ &= V_{\pi_i}(Z_j) - \epsilon \end{aligned}$$

where the inequality:

$$V_{\pi_i}([x_{i-1}, x_i]) \geq V_{\pi_i}([x_{j-1}, x_j])$$

holds by envy-freeness of X. Thus the induced allocation, Z, is ϵ-envy-free and contains the entire cake. □

3.4 The Generalized Game

In this section we introduce and briefly discuss a natural generalization of the moving knife game, in which the players can receive multiple pieces of cake. This generalization is motivated by several other moving knife procedures [112] in which the players can receive more than one piece of cake (see, e.g., the moving knife scheme of Brams *et al.* [25], which can use as many as eleven cuts to produce an envy-free allocation for four players).

Informally, a *generalized moving knife game* is a moving knife game where each player $i \in N$ can receive up to $m_i \in \mathbb{N}^*$ pieces, and the game has $M \in \mathbb{N}^*$ rounds. In the generalized game, a strategy of player i consists of a sequence of M thresholds:

$$T_i = [t_{i,1}, \ldots, t_{i,M}] \in [0, 1]^M,$$

such that player i calls cut in round k when the piece to the left of the knife is worth $t_{i,k}$ according to i's valuation. The moving knife game introduction in Section 2 is an instance of the generalized game where the budget of each player is one and the number of rounds is n. A particularly relevant instance of the generalized moving knife game is the one-round moving knife game (with $M = 1$ and $m_i = 1, \forall i \in N$), which is related to war of attrition models (see, e.g., the war of attrition in continuous time analyzed by Hendricks *et al* [**?**]).

In the case of one-round moving knife games with hungry valuations, this is the unique pure Nash equilibrium.

Proposition 6. *In a one-round moving knife game with hungry valuations, every pure Nash equilibrium of the game induces the empty allocation.*

Proof. Assume by contradiction that there exists a one-round game with hungry valuations, continuous valuations, and deterministic tie-breaking such that the game has a non-trivial pure Nash equilibrium. Without loss of generality, let us assume that the tie-breaking rule is $(1, \ldots, n)$.

Let T be a profile of threshold strategies in equilibrium. Then there exists $x \in [0, 1]$ such that $V_i([0, x]) = t_i$ for some player $i \in N$, and the following hold:

- $V_j([0, x]) < t_j, \forall j \in \{1, \ldots, i - 1\}$
- $V_j([0, x]) \leq t_j, \forall j \in \{i + 1, \ldots, n\}$.

The utilities under T are:

$$u_i(T) = t_i$$

and

$$u_j(T) = 0, \forall j \in N \setminus \{i\}.$$

Then any player $j \in N \setminus \{i\}$, can strictly improve their utility by deviating to threshold:

$$t_j^{'} = \frac{V_j([0, x/2])}{2},$$

since

$$u_j(T^{'}, T_{-j}) = t_j^{'} > 0,$$

where $T^{'} = (T_j^{'}, T_{-j})$. □

More generally, the result holds for all moving knife games with strictly positive value density functions where the number of rounds is small enough (i.e. $M < \sum_{i=1}^{n} m_i$).

Finally, when the players have symmetric value density functions, i.e. $v_i(x) = v_j(x), \forall i, j \in N$, and the number of rounds is large enough to allow all the players to receive the number of pieces they are entitled to, then the generalized moving knife game has a non-trivial pure Nash equilibrium for every deterministic tie-breaking rule.

Proposition 7. *Consider a generalized moving knife game with symmetric and hungry valuations, where the number of rounds is equal to the total number of pieces that the players are entitled to. Then the game has a pure Nash equilibrium for every deterministic tie-breaking rule.*

Proof. Let M be the number of rounds and m_i the maximum number of pieces that player i is entitled to receive. Then we have that $M = \sum_{i=1}^{n} m_i$.

Let $\pi = (\pi_1, \ldots, \pi_n)$ be the tie-breaking rule. Since the players have identical value density functions, there exists a partition of the cake in M contiguous pieces, $X = (X_1, \ldots, X_M)$, such that

$$V_i(X_j) = \frac{1}{M}, \forall i \in N.$$

Define the following thresholds:

$$t_{i,k} = \frac{1}{M}, \forall i \in N, k \in \{1, \ldots, M\}.$$

It can be easily verified that the strategies are in equilibrium, and the utility of each player under T is:

$$u_i(T) = \frac{m_i}{M}, \forall i \in N.$$

Note that the equilibrium allocation of each player is directly influenced by their budget, i.e. players with higher budget receive proportionally larger pieces. □

3.5 Discussion and Future Work

We studied the strategic version of the Dubins-Spanier protocol when the players have simple threshold strategies. Our main technical result is the existence of a direct correspondence between the non-trivial pure Nash equilibria of the moving knife game and the envy-free allocations of the cake with contiguous pieces. A characterization of the equilibria in the generalized moving knife game is left open. If one requires that the induced allocations have desirable properties, related to proportionality and envy-freeness, then the existence of such equilibria depends on whether envy-free allocations with a given number of cuts and ordering of the players exist. In particular, we are interested in the existence of mixed-strategy equilibria with uncountably infinite support, such that the entire cake is allocated with positive probability.

It would also be interesting to understand the outcomes of the game under richer strategy spaces. We note that generalizations in which each player has $n!$ thresholds (to account not only for the round number, but also for the players that have been allocated in the previous rounds) do not necessarily have envy-free equilibria. However, this does not preclude the existence of envy-free equilibria in the corresponding continuous time extensive form game.

In addition, this work initiates the direction of understanding the consequences of strategic behaviour in classical cake cutting protocols. For example, it would be interesting to understand whether protocols that compute fair allocations in the classical model (such as Brams-Talor) have fair equilibria under complete information.

Chapter 4

An Algorithmic Framework for Strategic Fair Division

In the previous chapter we analyzed the outcomes obtained in the Dubins-Spanier protocol when the players are strategic and have threshold strategies. We would like to make general statements regarding the equilibria of cake cutting protocols and without restrictions on the strategies of the players. We wish to identify a general family of cake cutting protocols — which captures the classic cake cutting protocols — so that each protocol in the family is guaranteed to possess (approximate) equilibria. Moreover, we wish to argue that these equilibrium outcomes are fair. Ultimately, our goal is to be able to reason about the fairness of cake divisions that are obtained as outcomes when players are presented with a standard cake cutting protocol and behave strategically.

We begin with a motivating example using the simplest cake cutting protocol, *Cut-and-Choose*. Recall that in this protocol, the first player cuts the cake into two pieces that it values equally; the second player then chooses the piece that it prefers, leaving the first player with the remaining piece. So how would strategic players behave when faced with the cut and choose protocol? A standard way of answering this question employs the notion of *Nash equilibrium*: each player would use a strategy that is a best response to the other player's strategy. To set up a Nash equilibrium, suppose that the first player cuts two pieces that the second player values equally; the second player selects its more preferred piece, and the one less preferred by the first player in case of a tie. Clearly, the second player cannot gain from deviating, as it is selecting a piece that is at least as preferred as the other. As for the first player, if it makes its preferred piece even bigger, the second player would choose that piece, making the first player worse off. Interestingly enough, in this equilibrium the tables are turned; now it is the second player who is getting exactly half of its value for the whole cake, while the first player generally gets more. Crucially, the equilibrium outcome is also proportional and envy-free.

In other words, even though the players are strategizing rather than following the protocol, the outcome in equilibrium has the same fairness properties as the "honest" outcome!

4.1 Model, Results, and a New Algorithmic Paradigm

To set the stage for a result that encompasses classic cake cutting protocols, we introduce (in Section 7.2) the class of *generalized cut and choose (GCC)* protocols. A GCC protocol is represented by a tree, where each node is associated with the action of a player. There are two types of nodes: a *cut node*, which instructs the player to make a cut between two existing cuts; and a *choose* node, which offers the player a choice between a collection of pieces that are induced by existing cuts. Moreover, we assume that the progression from a node to one of its children depends only on the relative positions of the cuts (in a sense to be explained formally below). We argue that classic protocols — such as Dubins-Spanier [55], Selfridge-Conway (see [112]), Even-Paz [61], as well as the original cut and choose protocol — are all GCC protocols. We view the definition of the class of GCC protocols as one of our main contributions.

In Section 4.4, we observe that GCC protocols may not have exact Nash equilibria (NE). We then explore two ways of circumventing this issue, which give rise to our two main results.

1. We prove that every GCC protocol has at least one ϵ-NE for every $\epsilon > 0$, in which players cannot gain more than ϵ by deviating, and ϵ can be chosen to be arbitrarily small. In fact, we establish this result for a stronger equilibrium notion, (approximate) *subgame perfect Nash equilibrium (SPNE)*, which is, intuitively, a strategy profile where the strategies are in NE even if the game starts from an arbitrary point.

2. We slightly augment the class of GCC protocols by giving them the ability to make *informed tie-breaking* decisions that depend on the entire history of play, in cases where multiple cuts are made at the exact same point. While, for some valuation functions of the players, a GCC protocol may not possess any exact SPNE, we prove that it is always possible to modify the protocol's tie-breaking scheme to obtain SPNE.

In Section 4.7, we observe that for any proportional protocol, the outcome in any ϵ-equilibrium must be an ϵ-proportional division. We conclude that under the classic cake cutting protocols listed above — which are all proportional — strategic behavior preserves the proportionality of the outcome, either approximately, or exactly under informed tie-breaking.

One may wonder, though, whether an analogous result is true with respect to envy-freeness. We give a negative answer, by constructing an envy-inducing

SPNE under the Selfridge-Conway protocol, a well-known envy-free protocol for three players. However, we are able to design a curious GCC protocol in which every NE outcome is a contiguous envy-free allocation and vice versa, that is, the set of NE outcomes coincides with the set of contiguous envy-free allocations. It remains open whether a similar result can be obtained for SPNE instead of NE.

Taking a broader perspective, our approach involves introducing a concrete computational model that captures well-known algorithms, and reasoning about the game-theoretic guarantees of *all* algorithms operating in this model. This approach appears distinct from related ones, where concrete query models are defined in order to evaluate the computational complexity of economic methods [17, 73], or restrictions *on the output of the algorithm* — such as the well-known *maximal-in-range* restriction [53] — give rise to desirable game-theoretic properties. Perhaps the most closely related approach was taken by Tennenholtz [121] in his work on program equilibrium (later extended by Fortnow [64]), but there the strategies are the programs themselves, whereas in our work a common algorithm (the GCC protocol) induces the players' strategies. We therefore believe that our conceptual contributions may be of independent interest to researchers working in other areas of algorithmic game theory, such as auction design.

The notion of GCC protocols is inspired by the Robertson-Webb [112] model of cake cutting — a concrete complexity model that specifies how a cake cutting protocol may interact with the players. Their model underpins a significant body of work in theoretical computer science and AI, which focuses on the complexity of achieving different fairness or efficiency notions in cake cutting [6, 50, 58, 59, 84, 108, 126]. In Section 7.2, we describe the Roberston-Webb model in detail, and explain why it is inappropriate for reasoning about equilibria.

In the context of the strategic aspects of cake cutting, Nicolò and Yu [104] were the first to suggest equilibrium analysis for cake cutting protocols. Focusing exclusively on the case of two players, they design a specific cake cutting protocol whose unique SPNE outcome is envy-free. And while the original cut and choose protocol also provides this guarantee, it is not "procedural envy free" because the cutter would like to exchange roles with the chooser; the two-player protocol of Nicoló and Yu aims to solve this difficulty.

4.2 Generalized Cut and Choose Protocols

Recall that the standard communication model in cake cutting was proposed by Robertson and Webb [112] (but formalized explicitely in a paper by Woeginger and Sgall [126]. We focus on the slightly augmented version proposed by Procaccia [109], where the interaction between the protocol and the players to the following two types of queries:

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

However, the communication model does not give much information about the actual implementation of the protocol and what allocations it produces. For example, the protocol could allocate pieces depending on whether a particular cut was made at an irrational point (see Algorithm 4).

For this reason, we define a generic class of protocols that are implementable with natural operations, which capture all bounded[1] and discrete cake cutting algorithms, such as cut and choose, Dubins-Spanier, Even-Paz, Successive-Pairs, and Selfridge-Conway (see, e.g., [109]). At a high level, the standard protocols are implemented using a sequence of natural instructions, each of which is either a *Cut* operation, in which some player is asked to make a cut in a specified region of the cake; or a *Choose* operation, in which some player is asked to take a piece from a set of already demarcated pieces indicated by the protocol. In addition, every node in the decision tree of the protocol is based exclusively on the execution history and absolute ordering of the cut points, which can be verified with any of the following operators: $<, \leq, =, \geq, >$.

More formally, a *generalized cut and choose (GCC)* protocol is implemented exclusively with the following types of instructions:

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

A GCC protocol uniquely identifies every contiguous piece by the symbolic names of all the cut points contained in it. For example, Algorithm 1 is a GCC protocol. Algorithm 4 is not a GCC protocol, because it verifies that the point where player 1 made a cut is exactly 1/3, whereas a GCC protocol can only verify the ordering of the cut points relative to each other and the endpoints

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

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

Algorithm 3: A GCC protocol. The notation "// @x" assigns the symbolic name x to the cut point made by player 1.

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

Algorithm 4: A non-GCC protocol.

of the cake. Note that, unlike in the communication model of Robertson and Web [112], GCC protocols cannot obtain and use information about the valuations of the players — the allocation is only decided by the players' *Choose* operations.

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

While GCC protocols are quite general, a few well-known cake cutting protocols are beyond their reach. For example, the Brams-Taylor [20] protocol is an envy-free protocol for n players, and although its individual operations are captured by the GCC formalism, the number of operations is not bounded as a function of n (i.e., it may depend on the valuation functions themselves). Its representation as a GCC protocol would therefore be infinitely long. In addition, some cake cutting protocols use *moving knives* (see, e.g., [25]); for example, they can keep track of how a player's value for a piece changes as the piece smoothly grows larger. These protocols are not discrete, and, in fact, cannot be implemented even in the Robertson-Webb model.

4.3 The Game

We study GCC protocols when the players behave strategically. Specifically, we consider a GCC protocol, coupled with the valuation functions of the players, as an *extensive-form game of perfect information* (see, e.g., [115]). In such a game, players execute the *Cut* and *Choose* instructions strategically. Each player is fully aware of the valuation functions of the other players and aims to optimize its overall utility for the chosen pieces, given the strategies of other players.

While the perfect information model may seem restrictive, we note that the same assumption is also made in previous work on equilibria in cake cutting [28, 104]. More importantly, it underpins foundational papers in a variety of areas of microeconomic theory, such as the seminal analysis of the Generalized Second Price (GSP) auction by Edelman et al. [57]. A common justification for the complete information setting, which is becoming increasingly compelling as access to big data becomes pervasive, is that players can obtain a significant amount of information about each other from historical data.

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

- In a *Cut* node defined by "*i cuts in S*", where $S = \{[x_1, y_1], \ldots, [x_m, y_m]\}$, the strategy space of player i is the set S of points where player i can make a cut at this step.

- In a *Choose* node defined by "*i chooses from S*", where $S = \{[x_1, y_1], \ldots, [x_m, y_m]\}$, the strategy space is the set $\{1, \ldots, m\}$, i.e., the indices of the pieces that can be chosen by the player from the set S.

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

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

4.4 Existence of Equilibria

It is well-known that finite extensive-form games of perfect information can be solved using *backward induction*: starting from the leaves and progressing towards the root, at each node the relevant player chooses an action that maximizes its utility, given the actions that were computed for the node's children. The induced strategies form an SPNE. Unfortunately, although we consider finite GCC protocols, we also need to deal with *Cut* nodes where the action space is infinite, hence naïve backward induction does not apply.

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

4.5 Approximate SPNE

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

More generally, we have the following theorem.

Theorem 8. *For any n-player GCC protocol $\mathcal{P}$ with a bounded number of steps, any n valuation functions $V_1, \ldots, V_n$, and any $\epsilon > 0$, the game induced by $\mathcal{P}$ and $V_1, \ldots, V_n$ has an ϵ-SPNE.*

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

of theorem 8. Let $\epsilon > 0$, and let $f(n)$ be an upper bound on the number of operations (i.e., on the height of the game tree) of the protocol. Define a grid, $\mathcal{G}_1$, such that every cell on the grid is worth at most $\frac{\epsilon}{2f(n)^2}$ to each player. For every n, let K denote the maximum number of cut operations, where $0 \leq K \leq f(n)$. For each $i \in \{1, \ldots, K\}$, we define the grid $\mathcal{G}_i$ so that the following properties are satisfied:

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

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

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

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

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

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

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

1. $\mathcal{M}(x_1) \in \mathcal{G}_1$.

2. The map $\mathcal{M}$ is order-preserving since there is only one point.

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

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

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

1. There exists $i \in \{1, \ldots, k-1\}$ such that $x_k = x_i$. Then define $\mathcal{M}(x_k) := \mathcal{M}(x_i)$.

2. There exists $i \in \{1, \ldots, k-1\}$ such that $x_i < x_k$, but $\mathcal{M}(x_i) \geq x_k$. Let x_j be the rightmost cut such that $x_j < x_k$; because $\mathcal{M}$ is order-preserving, it holds that $\mathcal{M}(x_j) \geq x_k$. Let $R(\mathcal{M}(x_j))$ be the leftmost point on $\mathcal{G}_k$ strictly to the right of $\mathcal{M}(x_j)$, and set $\mathcal{M}(x_k) := R(\mathcal{M}(x_j))$. Now let us check the conditions. Condition (1) holds by definition. Condition (2) holds because $\mathcal{M}(x_k) > \mathcal{M}(x_j)$, and for every t such that $x_t > x_k$, $\mathcal{M}(x_t) > \mathcal{M}(x_j)$ and $\mathcal{M}(x_t) \in \mathcal{G}_{k-1}$, whereas $\mathcal{M}(x_k)$ uses a "new" point of $\mathcal{G}_k \setminus G_{k-1}$ that is closer to $\mathcal{M}(x_j)$. For condition (3), we have that for every $l \in N$,

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

   where the third transition follows from the induction assumption.

3. There exists $i \in \{1, \ldots, k-1\}$ such that $x_i > x_k$, but $\mathcal{M}(x_i) \leq x_k$. This case is symmetric to case (b).

4. For every x_i such that $x_i < x_k$, $\mathcal{M}(x_i) \leq x_k$, and for every x_j such that $x_j > x_k$, $\mathcal{M}(x_j) \geq x_k$. Let x_i and x_j be the rightmost and leftmost such cuts, respectively; without loss of generality they exist, otherwise our task is even easier. Let $R(x_k)$ be the leftmost point in $\mathcal{G}_k$ such that $R(x_k) \geq x_k$, and let $L(x_k)$ be the rightmost point in $\mathcal{G}_k$ such that $L(x_k) \leq x_k$. Assume first that $\mathcal{M}(x_j) > R(x_k)$; then set $\mathcal{M}(x_k) := R(x_k)$. This choice obviously satisfies the three conditions, similarly to the base of the induction. Otherwise, $R(x_k) = \mathcal{M}(x_j)$ (notice that it cannot be the case that $R(x_k) > \mathcal{M}(x_k)$); then set $\mathcal{M}(x_k) := L(x_k)$. Let us check that this choice is order-preserving (as the other two conditions are trivially satisfied). Note that $\mathcal{M}(x_j) \in \mathcal{G}_{k-1}$, so $R(x_k) \in \mathcal{G}_{k-1}$. Therefore, it must hold that $L(x_k) \in G_k \setminus G_{k-1}$ — it is the new point that we have added between $R(x_k)$, and the rightmost point the left of it on $\mathcal{G}_{k-1}$. Since it is also the case that $\mathcal{M}(x_i) \in \mathcal{G}_{k-1}$, we have that $\mathcal{M}(x_i) < \mathcal{M}(x_k) < \mathcal{M}(x_j)$.

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

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

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

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

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

- *Choose*: player i chooses the available piece (identified by the symbolic names of the cut points it contains and their order) which is optimal in the idealized game, given the current state and the existing set of ordered ideal cuts, $\mathcal{M}(x_1), \ldots, \mathcal{M}(x_k)$. Ties are broken according to a fixed deterministic scheme which is known to all the players.

- *Cut*: player i computes the optimal cut on $\mathcal{G}_{k+1}$, say at x^*_{k+1}. Then i maps x^*_{k+1} back to a point x_{k+1} on the continuous game, such that $\mathcal{M}(x_{k+1}) = x^*_{k+1}$. That is, the cut x_{k+1} (made in step $k+1$) is always mapped by the other players to $x^*_{k+1} \in \mathcal{G}_{k+1}$. Player i cuts at x_{k+1}.

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

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

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

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

Induction hypothesis: Assume that playing on the grid is $\frac{(t-1)\epsilon}{f(n)}$-optimal whenever there are at most $t - 1$ operations left on every possible execution path of the protocol, and there exists one path that has exactly $t - 1$ steps.

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

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

1. By construction of the grid $\mathcal{G}_{k+1}$, player i can induce any given branch of the protocol using a cut in the continuous game if and only if the same branch can be induced using a cut on the grid $\mathcal{G}_{k+1}$.

2. Given that the other players will play on the grid for the remainder of the protocol, player i can change the size of at most one piece that it receives down the road by at most $\frac{\epsilon}{f(n)}$ by deviating (compared to the grid outcome), since $V_j([x_l, \mathcal{M}(x_l)]) \leq \frac{\epsilon}{2f(n)}$ for all $l \in \{1, \ldots, k+1\}$ and for all $j \in N$.

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

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

4.6 Informed Tie-Breaking

Another approach for circumventing the example given at the beginning of the section is to change the *tie-breaking* rule of Algorithm 1, by letting player 1 choose even if $y = z$ (in which case player 1 would cut in $x = 0, y = 1, z = 1$, and get the entire cake). Tie-breaking matters: Even the Dubins-Spanier protocol fails to guarantee SPNE existence due to a curious tie-breaking issue [28].

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

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

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

We are now ready to present our second main result.

Theorem 9. *For any n-player GCC protocol $\mathcal{P}$ with a bounded number of steps and any n valuation functions $V_1, \ldots, V_n$, there exists an informed GCC protocol $\mathcal{P}'$ that is equivalent to $\mathcal{P}$ up to tie-breaking, such that the game induced by $\mathcal{P}'$ and $V_1, \ldots, V_n$ has an SPNE.*

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

The proof of Theorem 9 is somewhat long, and has been relegated to Appendix 4.8. This proof is completely different from the proof of Theorem 8; in particular, it relies on real analysis instead of backward induction on a discretized space. The crux of the proof is the development of an auxiliary notion of *mediated games* (not to be confused with Monderer and Tennenholtz's *mediated equilibrium* [98]) that may be of independent interest. We show that mediated games always have an SPNE. The actions of the mediator in this SPNE are then reinterpreted as a tie-breaking rule under an informed GCC protocol. In the context of the proof it is worth noting that some papers prove the existence of SPNE in games with infinite action spaces (see, e.g., [72, 74]), but our game does not satisfy the assumptions required therein.

4.7 Fair Equilibria

The existence of equilibria (Theorems 8 and 9) gives us a tool for predicting the strategic outcomes of cake cutting protocols. In particular, classic protocols provide fairness guarantees when players act honestly; but do they provide any fairness guarantees in equilibrium?

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

One may wonder, though, whether the analogous statement for envy-freeness holds; the answer is negative. We demonstrate this via the Selfridge-Conway protocol — the 3-player envy-free protocol, which was given in its truthful, non-GCC form in Section 2.2. We construct an example by specifying the valuation functions of the players and their strategies, and arguing that the strategies are in SPNE. The example will have the property that the first two players receive utilities of 1 (i.e. the maximum value). Therefore, we can safely assume their play is in equilibrium; this will allow us to define the strategies only on a small part of the game tree. In contrast, player 3 will deviate from its truthful strategy to gain utility, but in doing so will become envious of player 1.

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

- Player 1 values the first untrimmed piece at 1, and all other pieces and the trimmings at 0.
- Player 2 values the second untrimmed piece at 1, and all other pieces and the trimmings at 0.
- Player 3 values the untrimmed pieces at 1/7 and 0, respectively, the trimmed piece at 1/14, and the trimmings at 11/14.

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

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

Theorem 10. *There exists a GCC protocol $\mathcal{P}$ such that on every cake cutting instance with strictly positive valuation functions $V_1, \ldots, V_n$, an allocation X is the outcome of a NE of the game induced by $\mathcal{P}$ and $V_1, \ldots, V_n$ if and only if X is an envy-free contiguous allocation that contains the entire cake.*

Crucially, an envy-free contiguous allocation is guaranteed to exist [117], hence the set of NE of protocol $\mathcal{P}$ is nonempty.

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

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

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

Proof of Theorem 10. Let $\mathcal{P}$ be the Thieves protocol given by Algorithm 3 and $\mathcal{E}$ be any NE of $\mathcal{P}$. Denote by X the contiguous allocation of the entire cake obtained during the demarcation step, where $X_i = [x_i, y_i]$ for all $i \in N$, and let w_i and z_i be the cut points of player i during its verification round. Assume for the sake of contradiction that X is not envy-free. Let k^* be an envious player, where $V_{k^*}(X_{j^*}) > V_{k^*}(X_{k^*})$, for some $j^* \in N$. There are two cases to consider:

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

Case 2: There exists a player i that did not receive the piece X_i. Then, it must be the case that some player k stole a non-empty strict subset $[w''_k, z''_k] = [w_k, z_k] \cap Z_j$ of another piece X_j. However, player k could have been better off at the node in the game tree reached in its verification round by making the following marks: $w'_k = \frac{x_j + w''_k}{2}$ and $z'_k = \frac{z''_k + y_j}{2}$. Since either $x_j \le w''_k < z''_k < y_j$ or $x_j < w''_k < z''_k \le y_j$ (recall that $[w''_k, z''_k]$ is a non-empty strict subset of X_j and the valuation function V_k is strictly positive), it is also true that $V_k([w'_k, z'_k]) > V_k([w''_k, z''_k])$, again contradicting the assumption that $\mathcal{E}$ is a NE.

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

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

Algorithm 5: Thieves Protocol: Every NE induces a contiguous envy-free allocation that contains the entire cake and vice versa.

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

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

- Player 1 specifically demarcates the allocation Z and cuts at points $w_1 = x_1$ and $z_1 = y_1$ during its verification round.

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

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

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

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

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

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

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

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

□

The theorem is a positive result *à la* implementation theory (see, e.g., [93]), which aims to construct games where the NE outcomes coincide with a given specification of acceptable outcomes for each constellation of players' preferences (known as a *social choice correspondence*). Our construction guarantees that the NE outcomes coincide with (contiguous) envy-free allocations, that is, in this case the envy-freeness criterion specifies which outcomes are acceptable.

That said, the protocol $\mathcal{P}$ constructed in the proof of Theorem 10 is impractical: its Nash equilibria are unlikely to arise in practice. This further motivates efforts to find an analogous result for SPNE. If such a result is indeed feasible, a broader, challenging open question would be to characterize GCC protocols that give rise to envy-free SPNE, or at least provide a sufficient condition (on the protocol) for the existence of such equilibria.

4.8 Proof of Theorem 7

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

Definition 8. *In an extensive-form game, an* action tuple *is a tuple of actions that describe an outcome of the game. For example, the action tuple* $(a_1, ..., a_r)$ *states that* a_1 *was the first action to be played,* a_2 *the second, and* a_r *the last.*

Definition 9. *Given an action tuple, the k^{th} action is said to be SPNE if the subtree of the game tree rooted where the first $k-1$ actions are played in accordance to the action tuple is induced by some SPNE strategy profile. Furthermore, call such an action tuple k-*SPNE*.*

Note that if the k^{th} action is SPNE, so too are all actions succeeding it in the action tuple. To clarify Definition 9, note that strategies of an extensive-form game are defined on every possible node of the game tree, so a k-SPNE action tuple can be equivalently defined as being an SPNE of the subgame rooted at the k^{th} action.

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

Definition 10. *We call an extensive-form game a* mediated game *if the following conditions hold:*

1. *The set of players consists of a single special player, referred to as the* mediator*, and some finite number n of other regular players. Intuitively, the mediator is a player who is overseeing the proper execution of a protocol.*

2. *The height h of the game tree is bounded.*

3. *Every player's utility is bounded.*

4. *Starting from the first or second action, the mediator plays every second action (and only these actions).*

5. *Every action played by the mediator shares the same action space:*

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

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

6. *The mediator's utility is binary (i.e. it is in $\{0,1\}$) and is described entirely by the notion of* allowed edges*. This is a set of edges in the game tree such that the mediator's utility is 1 iff it plays edges only in this set. Importantly, this set has the property that for every allowed edge, each grandchild subtree (i.e. subtree that represents the next mediator's action) must have at least one allowed edge from its root. Intuitively, these edges are the ones that follow the protocol the mediator is implementing.*

7. *A regular player's utility is continuous[2] in the action tuple.*

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

8. Allowed-edges-closedness*: given a convergent sequence of action tuples where the mediator plays only allowed edges, the mediator must play only allowed edges in the limit action tuple as well.*

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

We now give a series of definitions and lemmas that culminate in the main tool used in the proof of Theorem 9: all mediated games have an SPNE.

Definition 11. *A sequence of action tuples* $(a_1^i, ..., a_r^i) \mid_i$ *is said to be* consistent *if for every* j *the player who plays action* a_j^i *is constant throughout the sequence and, moreover, its action spaces are always subsets of* $[0,1]$ *or always the same subset of* $\{1, ..., h\}$ *throughout the sequence.*

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

Proof. Due to the finite number of players and bounded height of the game, we can find an infinite consistent subsequence $\mathbf{b}^i \mid_i = (b_1^i, ..., b_r^i) \mid_i$. It suffices to show this subsequence has a convergent subsequence of its own. It is fairly clear that we can find a convergent subsequence via compactness arguments, but there is a slight caveat: we must show that the limit action tuple is legal. That is, if the limit action tuple is $(a_1, ..., a_r)$ we must show that for every $i < r$ such that the mediator plays action i, action $i+1$ is played by the player prescribed by a_i, and within the bounds prescribed by it. We will prove this by induction.

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

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

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

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

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

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

Proof. We prove the result by induction on k.

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

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

Induction step ($k = m+1$)*:* Let $\mathbf{a}^i \mid_i = (a_1^i, ..., a_r^i) \mid_i$ be a convergent sequence of $(r-m)$-SPNE action tuples with the limit action tuple $(a_1, ..., a_r)$. We wish to show that if all actions before the last $m+1$ actions play their limit actions, then the remaining $m+1$ actions are SPNE — note that by Lemma 3 we know that the limit sequence is a valid action tuple.

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

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

$$\mathbf{b}^i \mid_i = (a_1^i, ..., a_{r-m-1}^i, \alpha^i, \tilde{a}_1^i, ..., \tilde{a}_m^i) \mid_i$$

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

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

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

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

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

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

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

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

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

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

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

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

With this machinery in hand, we are now ready to complete the proof of Theorem 9. Our main task is to make a formal connection between mediated games and (informed) GCC protocols.

of Theorem 9. Suppose we have a n-player GCC protocol $\mathcal{P}$ with a bounded number of steps and and set valuations of the players $V_1, ..., V_n$. Then we wish to prove that there exists an informed GCC protocol $\mathcal{P}'$ that is equivalent to $\mathcal{P}$ up to tie-breaking such that the game induced by $\mathcal{P}'$ and $V_1, ..., V_n$ has an SPNE.

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

1. This is by definition.

2. The height of the tree is twice the height of the GCC protocol.

3. The mediator's utility is bounded by 1 by definition, and all other player's utilities are bounded by 1 as that is their value of the entire cake.

4. This is by definition.

5. When the mediator wishes to ask a *Cut* query to player i in the interval $[a, b]$, it plays the action $(i, (a, b))$, whereas when it wishes to ask a *Choose* query to player i giving them the choice between the $x_1^{th}, ..., x_k^{th}$ pieces from the left, it plays the action $(i, \{x_1, ..., x_k\})$. This method of giving choose queries deviates slightly from the definition given in Section 4.2, but the two representations are clearly equivalent.

6. The allowed edges are ones that follow the rules of $\mathcal{P}$.

7. This property is only relevant when considering *Cut* nodes. To establish it, first consider the action in a single *Cut* node, and fix all the other actions. We claim that for every $\epsilon > 0$ there exists $\delta = \delta(\epsilon) > 0$ *that is independent of the choice of actions in other nodes* such that moving the cut by at most δ changes the values by at most ϵ. Indeed, let us examine how pieces change as the cut point moves. As long as the cut point moves without passing any other cut point, one piece shrinks

as another grows. As the cut point approaches another cut point, the induced piece — say k'th from the left — shrinks. When the cut point passes another cut point x, the k'th piece from the left grows larger, or it remains a singleton and another piece grows if there are multiple cut points at x. In any case, it is easy to verify that the sizes of various pieces received in *Choose* nodes change by at most δ if the cut point is moved by δ. Furthermore, note that the number of steps is bounded by r and — since the value density functions are continuous — there is an upper bound M on the value density functions such that if $y - x \leq \delta'$ then $V_i([x, y]) \leq M\delta'$ for all $i \in N$. Therefore, choosing $\delta \leq \epsilon/(Mr)$ is sufficient. Finally, $V_1, \ldots, V_n$ are continuous even in the actions taken in multiple *Cut* nodes, because we could move the cut points sequentially.

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

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

Now as M is a mediated game, it has an SPNE S by Lemma 5. Let $\mathcal{P}'$ be the informed GCC protocol equivalent to $\mathcal{P}$ up to tie-breaking such that for every point in the game tree of M that represents the mediator branching on an "*if* $x \leq y$ then A *else* B" statement in the original protocol $\mathcal{P}$, $\mathcal{P}'$ chooses the A or B that S takes in the event of a tie. Then the informedness of the tie-breaking is built into $\mathcal{P}'$ and we immediately see that the SPNE actions of the regular players in M correspond to SPNE actions in $\mathcal{P}'$. □

Chapter 5

A Dictatorship Theorem

As already observed in the previous chapters, the classical discrete protocols are not *strategy-proof* [28, 43, 84], i.e., there are scenarios (possible behaviors of the other players) in which it is possible for a player to get a piece of strictly larger value by misrepresenting its valuation function than by behaving truthfully. This begs the question of whether alternative strategy-proof protocols can be constructed. Addressing this question, Kurokawa *et al.* [84] showed a negative result: For any number of players $n \geq 2$, there is no Robertson-Webb protocol of complexity bounded only by a function of n (i.e., independent of the valuations) that is strategy-proof and computes an envy-free allocation.

The main results of this chapter are impossibility theorems closely releated to the result of Kurokawa *et al.*, but rather than stating that no fair allocation can computed, we essentially state that no reasonable allocation can be computed *at all*; thus, the "unfairness conclusions" of our theorems are stronger. Also, we do not need to make any assumption about the complexity of the protocols.

For two players, our result is particularly strong, with a conclusion similar to the classical *dictatorship* results of social choice theory, in particular the Gibbard-Satterthwaite theorem [66, 114], which is a cornerstone of social choice theory and mechanism design. The Gibbard-Satterthwaite theorem states that the only strategy-proof choice functions (i.e., *direct* revelation mechanisms without money) for settings where players have *general* tie-free preferences on at least three alternatives are dictatorships. Our theorem is similar in spirit, but applies to a particular setting of *restricted* preferences over allocations/alternatives (those induced by value density functions as described above) and a restricted class of *indirect* revelation mechanisms:

Theorem 11. *Suppose a deterministic cake cutting protocol for two players in the Robertson-Webb model is strategy-proof. Then, restricted to hungry players, the protocol is a dictatorship.*

Theorem 12. *Suppose a deterministic cake cutting protocol for $n \geq 3$ hungry players in the Robertson-Webb model is strategy-proof. Then, in every outcome associated with truthful reports by hungry players, there is at least one player that gets the empty piece (i.e., no cake).*

Recall we say that a player i is *hungry* if its value density function v_i is hungry, i.e., satisfies $v_i(x) > 0$ for all x. We say that a protocol is a *dictatorship* if there is a fixed player (the dictator) to whom the entire cake is allocated in all truthful executions of the protocol, no matter what the value density functions are[1].

5.1 Comments on the Impossibility Theorems

The theorems refer to the Robertson-Webb model as formalized originally by Woeginger and Sgall [126]. Recall that the alternative formulation given by Procaccia is more permissive, such as:

> "*Allocate* $[0, 0.5)$ *to player 1 and* $[0.5, 1]$ *to player 2*"

This protocol is clearly strategy-proof but not a dictatorship. The only difference between the two formalizations is that the Woeginger-Sgall version requires all cut points to be defined by the players rather than by the center. This property is essential for the theorems and their proofs.

Theorem 11 fails without the restriction to hungry players. The next protocol can be formalized in the Robertson-Webb model and is strategy proof but not a dictatorship (unless restricted to hungry players, in which case player 2 becomes the dictator):

> "*Ask player 1 to cut the cake in two pieces that are of equal value to him. If player 2 assigns strictly positive value to both of these pieces, give player 2 the entire cake, otherwise give player 1 the piece to which player 2 assigns 0 value and player 2 the other piece*"

The conclusion of Theorem 12 cannot be improved to the protocol being a dictatorship. Indeed, consider the following protocol: "*Player* 1 *cuts the cake in two pieces of equal value. Player* 2 *takes the piece it prefers. Player* 3 *takes the remaining piece.*" Player 1 never receives anything, so it has no incentive to misreport. Player 2 can always select its most preferrred piece, so it has no incentive to lie either. Finally, player 3 takes the remaining piece without

[1]This is consistent with the standard meaning of "dictatorship" in social choice theory: For all preference profiles, the social choice is the most preferred alternative of the dictator.

making any report; thus the protocol is strategyproof. However, it is not a dictatorship.

Our main theorem is of relevance to the general discussion about the merits of indirect revelation mechanisms versus direct revelation mechanisms. The *revelation principle* [102] informally states that any indirect revelation mechanism can be converted to an "equivalent" direct revelation mechanism. Concretely, if we had a strategy-proof Robertson-Webb protocol computing, say, an envy-free allocation, we would also have a strategy-proof direct revelation mechanism doing so. The revelation principle can be and often is used as motivation to focus research on direct revelation mechanisms. However, in the cake cutting scenario (and in many other scenarios), direct revelation would require submitting an infinite amount of information (in the cake cutting case, the value density function) which is not realistic, making models such a the Robertson-Webb model where players interact through a protocol in which information is revealed gradually the main object of study. But our main theorem shows that this can easily reduce drastically or even trivialize what can be done in a strategy-proof way (in particular, there are many strategy-proof direct revelation mechanisms for cake cutting with non-trivial fairness properties [95, 99**?**]).

5.2 The Robertson-Webb model

Recall the Robertson-Webb model (as formalized by Woeginger and Sgall [126]) – the protocol and the players interact through the following types of queries:

- Cut($i; \alpha$): Player i cuts the cake at a point y where $V_i([0, y]) = \alpha$. The point y becomes a *cut point*.
- Eval($i; y$): Player i returns $V_i([0, y])$ where y is a previously made cut point.

The queries made by the protocol may depend on the history (i.e. answers to previous queries). At termination, the cut points define a partition of the cake into a finite set of intervals that the protocol allocates to the players in some specified way.

To make the definition rigorous, we formally define a Robertson-Webb protocol as an *infinite decision tree* (see Figure 5.1) where each internal node $\mathcal{X}$ is labeled with the query made if $\mathcal{X}$ is reached. There is a directed outgoing edge e from such a node $\mathcal{X}$ for every possible answer to the given query (i.e., infinitely many), and the node $\mathcal{Y}$ reached through edge e is either an internal node, or a leaf containing the resulting allocation if the path to $\mathcal{Y}$ is taken. We require that the protocol does not ask for information it already knows and does not accept information from the players that is inconsistent with previous

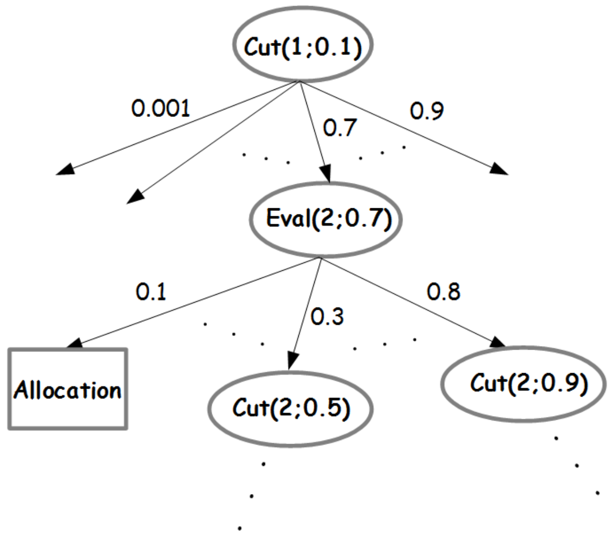

Figure 5.1: Representation of a Robertson-Webb protocol as an infinite decision tree

replies (e.g., reports that implies negatively valued subintervals). We require that the protocol terminates (reaches a leaf) for every profile of value density functions, if players report truthfully. If truthful reporting according to a value density function profile $\mathbf{v}$ makes the protocol reach a leaf u, we say that $\mathbf{v}$ *is associated with* u and vice versa.

Without loss of generality, protocols can be assumed to have alternating Cut and Eval queries: For any protocol $\mathcal{M}$, there is an equivalent protocol $\mathcal{M}'$ that after every cut asks for the values for all players of the newly generated subintervals and produces allocations identical to $\mathcal{M}$ on every instance. We say that a protocol is *strategy-proof* if for every profile of value density function it holds that truthful reporting is a dominant strategy for each player i, when the protocol is viewed as a complete and perfect information extensive form game where players choose strategically what to report.

Viewing the protocol as a complete and perfect information extensive form game this way in particular entails assuming that all communication between players and center is broadcast and accessible to all players. In addition, the transformation of a protocol $\mathcal{M}$ to a protocol $\mathcal{M}'$ with alternating Cut and Eval queries as described above preserves strategy-proofness: any extra Eval queries introduced in the transformation are payoff irrelevant "cheap talk" seen from the point of view of the players.

The following lemma will be useful.

Lemma 6. *Let $\mathcal{M}$ be a strategy-proof Robertson-Webb protocol for two players that is not dictatorial when restricted to hungry valuations. Then, in no leaf* *of $\mathcal{M}$ reached under truthful reporting for some profile of hungry valuations, is the entire cake given to a single player.*

Proof. Assume to the contrary that at a reachable leaf u the entire cake is allocated to a single player, say, player 1. Let $\mathbf{v} = (v_1, v_2)$ be the profile of hungry value density functions associated with u. Since the protocol is not dictatorial when restricted to hungry players, there is another reachable leaf u' where player 1 does not receive the entire cake. Let $\mathbf{v}' = (v_1', v_n')$ be a profile of hungry value density functions associated with u'. Consider now the outcome when players report according to the profile $w = (v_1', v_2)$. It must be the case that player 1 receives the entire cake in this outcome; otherwise the protocol is not strategy-proof, as player 1 could misrepresent his value density function as v_1 and get the entire cake, assuming that the other player reports according to v_2. But this means that when the protocol is played with profile w, player 2 would benefit from misrepresenting his value density function as v_2' rather than v_2, as he would then receive a non-empty piece rather than nothing at all. This contradicts the strategy-proofness of $\mathcal{M}$. □

5.3 Proof of the main theorems

For the proof of the theorems, it is convenient to define a restricted kind of protocols where the physical locations of the cut points do not matter; instead, the protocol is only concerned with the values that the players have for the generated pieces. We call such protocols *strictly mediated* and observe that in fact, all classical protocols in the Robertson-Webb model belong to this class. Strict mediation can be interpreted as the center not having direct access to the cake; instead, it can only see it through the eyes of the players.

Definition 12 (*Strictly Mediated Protocol*)**.** *A* strictly mediated protocol *for n players is an infinite decision tree containing two kinds of (internal) nodes –* Cut *nodes and* Eval *nodes – and leaves:*

- *An* Eval *node is labeled by a pair of natural numbers (i, j) and has a successor for each real number $\alpha \in (0, 1)$.*
- *A* Cut *node is labeled by a pair of natural numbers (i, j) and a real number $\alpha \in (0, 1)$ and has a* single *successor.*
- *Each leaf is labeled with a finite sequence of natural numbers in $\{1, \ldots, n\}$.*

The semantics is the following. At any point in the execution (at some node u in the tree), a set of cut points $x_0 = 0 < x_1 < x_2 < \cdots < x_k < x_{k+1} = 1$ has been defined (where k is the number of cut nodes above u):

- *When an* Eval *node X with labels (i, j) is reached, player i is asked for its value of interval $[x_j, x_{j+1}]$; given the player's answer, $\alpha \in (0, 1)$,* *execution moves to the successor node reached along the edge labeled with the value α.*

- *When a* Cut *node with labels (i, j, α) is reached, player i is asked to define a new cut point x' somewhere between x_j and x_{j+1} so that his value of the interval $[x_j, x']$ is an α-fraction of his value of the interval $[x_j, x_{j+1}]$.*

- *When a leaf node with labels $(i_0, i_1, \ldots, i_k)$ is reached, each interval (x_j, x_{j+1}) is allocated to player i_j, with the cut points themselves given arbitrarily.*

For convenience, we have defined strictly mediated protocols as a separate model rather than as a special case of Robertson-Webb protocols. However, given a strictly mediated protocol, it is easy to define a Robertson-Webb protocol that simulates it, so we shall also consider strictly mediated protocols as a special case of Robertson-Webb protocols.

To get some intuition, consider the following example.

Example 3. *Let $\mathcal{M}$ be some strictly mediated protocol that on an execution path reaches a leaf where the cut points discovered are $\{0.1, 0.7\}$, and the values of the players for each subinterval are:*

- *Player 1 has: $V_1([0, 0.1]) = v_1$, $V_1([0.1, 0.7]) = v_2$.*

- *Player 2 has: $V_2([0, 0.1]) = w_1$, $V_2([0.1, 0.7]) = w_2$.*

Say that $\mathcal{M}$ stopped after discovering these values and allocated the subintervals in the order $[1, 2, 1]$; that is, player 1 received $[0, 0.1] \cup [0.7, 1]$, while player 2 received $[0.1, 0.7]$. Then $\mathcal{M}$ has the property that if the answers of the players resulted instead in a different set of cut points, $\{x_1, x_2\}$, but the evaluate queries were answered in the same way (i.e. $V_1([0, x_1]) = v_1$, $V_1([x_1, x_2]) = v_2$, $V_2([0, x_1]) = w_1$, $V_2([x_1, x_2]) = w_2$), then $\mathcal{M}$ outputs the same allocation order (i.e. player 1 gets $[0, x_1] \cup [x_2, 1]$ and 2 gets $[x_1, x_2]$).

The relevance of the strictly mediated model is apparent from the following lemma.

Lemma 7. *Assume there exists a strategy-proof protocol $\mathcal{M}$ for $n \geq 2$ players with the property that there exists an outcome that is associated with a hungry valuation profile and where every player receives a non-empty piece. Then there exists a strategy-proof strictly mediated protocol $\mathcal{R}$ with the same property.*

Proof. Rather than formally describe the protocol $\mathcal{R}$ as a decision tree, we give an informal description, from which a formal (but probably less readable) description as a decision tree could easily be derived. First, we describe the idea of the construction.

The key constraint that a strictly mediated protocol has to satisfy is to not let the sequence of queries it makes nor its final allocation depend on the exact physical location of the cut points. It can only let these actions depend on the reports of the players. With this in mind, the idea of the protocol $\mathcal{R}$ is to directly simulate the protocol $\mathcal{M}$ step by step. But since the protocol $\mathcal{M}$ might have behavior that depends on the physical location of the cut points, we let $\mathcal{R}$ maintain a list of fictitious or pretend locations $y^*_t, t = 1, \ldots, k$ in $(0,1)$ that it feeds to $\mathcal{M}$ instead of the actual cut points $y^t, t = 1, \ldots, k$ made, preserving order, i.e, with the invariant maintained that $y^*_t < y^*_{t'}$ if and only if $y_t < y_{t'}$ for all t, t'. An alternative point of view is that $\mathcal{R}$, being strictly mediated, has no precise measuring device that can determine exactly where the players make the cut points, but that it makes its own primitive yardstick as it goes along, using the cut points actually made by the players as marks on its yard stick. However, we also have to make sure that we preserve the outcome of $\mathcal{M}$ where all players get a piece. Therefore, $\mathcal{R}$ has to be somewhat careful when defining the fictitious cut points.

Let X be some outcome (leaf) of $\mathcal{M}$ that is associated with a hungry value density function profile and in which all players receives a non-empty piece. Concretely, $\mathcal{R}$ simulates $\mathcal{M}$ as described in the next cases.

- ***Case 1:*** Whenever the protocol $\mathcal{M}$ wants to ask player i a cut query $\text{Cut}(i; \alpha)$, the protocol $\mathcal{R}$ computes numbers t, t', α' and by a Cut query asks player i to specify a point y_t in the subinterval $[y_{t'}; y_{t''}]$ between existing cut points $y_{t'}$ and $y_{t''}$ for which $V_i([y_{t'}, y_t]) = \alpha'$. The numbers t', t'', α' are computed so that a truthful player i will execute exactly the $\text{Cut}(i; \alpha)$ query. This computation can be performed by $\mathcal{R}$ for the following reason. As we explained when we defined the Roberson-Webb model, we maintain the invariant that all new subintervals are evaluated by all players after each Cut query in the original protocol $\mathcal{M}$. As $\mathcal{R}$ simulates $\mathcal{M}$ step by step, $\mathcal{R}$ also maintains this knowledge. When player i returns the new cut point y_t from the Cut query, the protocol $\mathcal{R}$ needs to find a suitable fictitious cut point y^*_t. There are two sub-cases:
  - In the execution of $\mathcal{M}$, it is still possible to reach X (i.e., X is a descendant of the $\text{Cut}(i; \alpha)$ node that $\mathcal{R}$ is simulating at the moment). In this case, there is a unique value for the cut point that will keep this possibility open by keeping the exeuction of $\mathcal{M}$ on the path to X. We let y^*_t be this unique value.
  - In the execution of $\mathcal{M}$, it is no longer possible to reach X. In this case, we let $y^*_t = (y^*_{t'} + y^*_{t''})/2$.

  In both sub-cases, we feed y^*_t back to $\mathcal{M}$ as the fictitious answer to the Cut query $\text{Cut}(i; \alpha)$.

- ***Case 2:*** Whenever $\mathcal{M}$ asks player i an Eval query Eval$(i; y^*_{t'})$ where $y_{t'}(y^*_{t'})$ is a new real (fictitious) cut point, $\mathcal{R}$ asks player i to evaluate $[y_{t''}, y_{t'}]$ where y'' is the largest cut point smaller than $y_{t'}$ (or 0, if no such cut point exists). As $y_{t''}$ is an older cut point, $\mathcal{R}$ already knows a report for $V_i([0, y_{t''}])$ and can return a report for $V_i([0, y_{t''}])$ as the sum of these reports to $\mathcal{M}$.
- ***Case 3:*** Finally, when $\mathcal{M}$ makes an allocation in the end, $\mathcal{R}$ allocates each subinterval $[y_{t'}, y_{t''}]$ to the player to which $\mathcal{M}$ allocates $(y_{t'}, y_{t''})$.

Now we check that $\mathcal{R}$ has the desired properties:

- (*Strict mediation*) By construction, $\mathcal{R}$ is strictly mediated.
- (*Shared leaf*) By construction, there is an outcome of $\mathcal{R}$ associated with a hungry valuation profile where all players get a piece of the cake, namely the hungy value density function profile where players answer Eval queries in the way that keeps the execution of $\mathcal{M}$ on the track to X.
- (*Truthfulness*) Finally, suppose that $\mathcal{R}$ is not truthful. That is, there is a scenario where a player, say player 1, has value density function v, and there is a strategies σ_i for players $i = 2, .., n$ in $\mathcal{R}$, so that truthful reporting is not an optimal strategy for player 1. Then some other strategy π is strictly better, yielding an increase in payoff $\delta > 0$. We claim that then there is a value density function v' for player 1 and strategies σ'_i for players $i = 2, .., n$ in $\mathcal{M}$ so that truthful reporting is not an optimal strategy for player 1. Hence $\mathcal{M}$ is also not truthful, contradicting the assumption on $\mathcal{M}$. We define:
  - v' simply to be any value density function that is consistent with the reports that $\mathcal{M}$ receives by $\mathcal{R}$ when $\mathcal{R}$ is given input v and
  - σ'_i to be the strategy of reporting to $\mathcal{M}$ the way $\mathcal{R}$ reports to $\mathcal{M}$ for player i, when player i plays according to σ_i in $\mathcal{R}$.

  Then truthful reporting of v' is not optimal for player 1 in $\mathcal{M}$, if all players $i = 2, .., n$ play according to σ_i, since player 1 would get an increase in payoff of δ by playing the strategy π' of reporting to $\mathcal{M}$ the way $\mathcal{R}$ reports to $\mathcal{M}$ for player 1, when player 1 plays using π in $\mathcal{R}$.

□

Our next lemma shows that strategy-proof strictly mediated protocols are very restricted in their behavior.

Lemma 8. *Let $\mathcal{M}$ be a strictly mediated protocol for $n \geq 2$ players that has some outcome, with an associated hungry valuation profile, in which each player receives a non-empty piece. Then $\mathcal{M}$ is not strategy-proof.*

Proof. Let X be a leaf of the protocol, associated with a hungry value density function profile $\mathbf{v}$, in which an allocation is made where all the players receive a non-empty piece. Denote by $x_1 < x_2 < \cdots < x_M$ the labels of the cut points defined on the path to X. (Note that they were not necessarily defined in that order on the path - the indices here indicate the order of the cut points according to usual ordering of real numbers. Also note that since the protocol is strictly mediated, $x_1, \ldots, x_M$ are symbolic label names rather than actual real numbers.) Without loss of generality, assume that $\mathcal{M}$ asks the first Cut query to player 1. Since the allocation at X is non-dictatorial, then we have the following:

- player 1 does not receive the entire cake
- player 1 receives at least one subinterval, say (x_{k-1}, x_k).

Suppose any concrete sequence $x_1^* < .. \cdots .. < x_M^*$ of real numbers strictly between 0 and 1 is given. By continously deforming $\mathbf{v}$, we can construct a hungry value density function profile $\mathbf{v}^*$ associated with the leaf X so that when the protocol is executed on $\mathbf{v}^*$, the actual cut point with label x_i becomes x_i^*. That is, the allocation order (from left to right) computed by $\mathcal{M}$ is the same for both valuations, $\mathbf{v}^*$ and $\mathbf{v}'$ – for example, if player 2 gets the piece $[0, x_1^*]$ on input $\mathbf{v}^*$, then player 2 also gets the piece $[0, x_1']$ on input $\mathbf{v}'$, and viceversa.

We shall define two such valuation profiles (see Figure 5.2), namely $\mathbf{v}^*$ (with actual cut points $x_1^* < .. \cdots .. < x_M^*$) and $\mathbf{v}'$ (with actual cut points $x_1' < \cdots < x_M'$). We choose $\mathbf{v}^*$ to be an arbitrary hungry profile associated with X. Let w^* be the valuation of player 1 for his piece in the outcome associated with $\mathbf{v}^*$; since player 1 does not receive the entire cake at X, we have that $w < 1$. The profile $\mathbf{v}'$ is constructed and its cut points are chosen so that $x_{k-1}' < x_1^* < \cdots < x_M^* < x_k'$. Moreover, x_{k-1}' is chosen sufficiently close to 0 and x_k' sufficiently close to 1, to ensure that the valuation of player 1 for (x_{k-1}, x_k) according to $\mathbf{v}^*$ is strictly larger than w. Since $\mathbf{v}'$ is associated with X, player 1 gets the subinterval (x_{k-1}', x_k') when both players report according to $\mathbf{v}'$.

Consider now the following strategy σ_i for each other player $i \in \{2, \ldots, n\}$:

1. If player 1 answers the first cut query according to a valuation consistent with $\mathbf{v}'$, then player i answers for the remainder of the protocol as if his valuation is also consistent with $\mathbf{v}'$.
2. Otherwise, player i answers truthfully throughout the protocol.

Observe that since the first Cut query is addressed to player 1, the answer of the player will be different under valuations $\mathbf{v}^*$ and $\mathbf{v}'$ by choice of the two profiles.

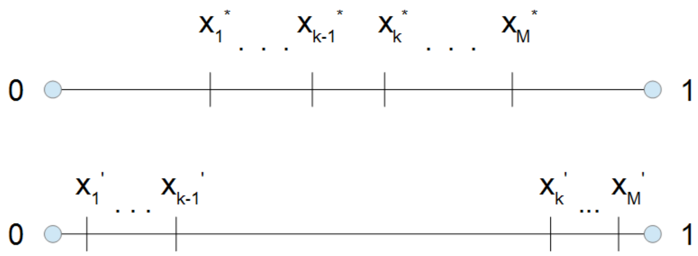

Figure 5.2: Valuation profile $\mathbf{v}^*$ at the top and $\mathbf{v}'$ at the bottom. The cut points $x_1^*, \ldots, x_M^*$ are completely contained in the interval (x_{k-1}^*, x_k').

Suppose the true type profile of the players is $\mathbf{v}^*$ and players $2, \ldots, n$ adopt strategies $\sigma_2, \ldots, \sigma_n$, respectively. Then, if player 1 answers truthfully, it gets in the end a piece worth w to the player. However, if player 1 lies by answering according to a valuation consistent with $\mathbf{v}'$, then in the end it gets a piece for which (x_{k-1}, x_k) is a subset. Since the value of player 1 for this interval alone is strictly larger than w^*, $\mathcal{M}$ is not strategy-proof. □

We are now ready to prove our two main theorems.

Proof. (of Theorem 11) Suppose we have a strategy-proof protocol $\mathcal{M}$ which is not a dictatorship when restricted to two hungry players. We have by Lemma 6 that it non-trivially shares the cake between the players in all outcomes corresponding to truthful reports of hungry value density functions. By Lemma 7, there is a strategy-proof strictly mediated protocol with an outcome in which the cake is shared. But this contradicts Lemma 8. □

Proof. (of Theorem 12) Suppose we have a strategy-proof protocol $\mathcal{M}$ with some outcome where $n \geq 3$ players get a non-empty piece. By Lemma 7, there exists a strategy-proof strictly mediated protocol $\mathcal{R}$ with the same property. This contradicts Lemma 8. □

5.4 Randomized Protocols

In this section we turn to randomized protocols in the Robertson-Webb model. A randomized protocol can formally be defined similar to the definition of deterministic protocols in Section 5.2, except that the decision tree now contains three types of internal nodes: cut nodes, evaluate nodes, and chance nodes. The cut and evaluate nodes are the same as for deterministic protocols, while each chance node X has some number of directed outgoing edges, each of which is labeled with the probability of being taken when the execution reaches the node X.

```
K ← ⌈2n(n−1)/ϵ⌉
for each player i ∈ {1, ..., n} do
    x_{i,0} ← 0
    x_{i,K+1} ← 1
    for each j ∈ {1, ..., K} do
        x_{i,j} ← Cut(i; j/K)
    end for
end for
X ← ⋃_{i=1}^{n} {x_{i,1}, ..., x_{i,K}}
for each subset Y ⊆ X, with |Y| ≤ n(n − 1) do
    for each allocation (A_1, ..., A_n) definable by cuts in Y do
        for each i, j ∈ {1, ..., n} do
            n_{i,j} ← #{k ∈ {0, ..., K} | (x_{i,k}, x_{i,k+1}) ⊆ A_j}
            w_{i,j} ← (1/K) · n_{i,j}
        end for
        if (1/n − 2/K ≤ w_{i,j}) and (w_{i,j} ≤ 1/n + 2/K), for all i, j then
            π ← RandomPermutation({1, ..., n})
            for each player i ∈ {1, ..., n} do
                W_{π_i} ← A_i // Player π(i) gets piece A_i
            end for
            return W
        end if
    end for
end for
```

Algorithm 6: Randomized Robertson-Webb protocol that is truthful in expectation and almost perfect

Recall that Mossel and Tamuz [99] showed a randomized *direct revelation* protocol that is truthful in expectation and computes a *perfect* allocation, that is, an allocation $A = (A_1, \ldots, A_n)$ where $V_i(A_j) = 1/n$, $\forall i, j \in N$:

> *Given as input valuations $V_1, \ldots, V_n$, find a perfect partition $A = (A_1, \ldots, A_n)$ and allocate it using a random permutation π over $\{1, \ldots, n\}$ (i.e. player i receives the piece A_{π_i}).*

This protocol is well-defined, but not constructive. Here, we observe that we can "discretize" the Mossel-Tamuz protocol to get an *explicit* Robertson-Webb protocol that is truthful in expectation and computes an "almost perfect" allocation.

Theorem 13. *Given $\epsilon > 0$, there is a randomized Robertson-Webb protocol $\mathcal{M}_\epsilon$ that asks at most $O(n^2/\epsilon)$ queries, is truthful in expectation and allocates* *to each player a piece of value between $1/n - \epsilon$ and $1/n + \epsilon$, according to the valuation functions of all players.*

The protocol is stated as Algorithm 2. Given ϵ, the protocol asks the players to divide the cake in many small cells (worth $\approx \frac{\epsilon}{2n(n-1)}$ each), and then uses the grids supplied by the players to search for an almost perfect allocation.

Proof. Our main tool is a lemma due to Alon [2].

Lemma 9. *Let $\mu_1, \mu_2, \dots, \mu_t$ be t continuous probability measures on the unit interval. Then it is possible to cut the interval in $(k-1) \cdot t$ places and partition the $(k-1) \cdot t + 1$ resulting intervals into k families $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k$ such that $\mu_i \left(\bigcup \mathcal{F}_j\right) = 1/k$, for all $1 \leq i \leq t$, $1 \leq j \leq k$. The number $(k-1) \cdot t$ is best possible.*

Given $\epsilon > 0$, let $\mathcal{M}_\epsilon$ be the protocol in Algorithm 5.4.

At a high level, protocol $\mathcal{M}_\epsilon$ asks each player to divide the cake in many small cells (K of them) of equal value $1/K$; then $\mathcal{M}_\epsilon$ exhaustively enumerates all subsets Y of size bounded by $n(n-1)$ from the cut points supplied by the players.

Given that a perfect partition is guaranteed to exist on the continuous cake within at most $n(n-1)$ cuts by Alon's lemma, one of the sets Y is guaranteed to work. That is, $\mathcal{M}_\epsilon$ finds a set of points Y and an allocation A that uses exclusively cut points in Y such that:

- every point in Y is close to a cut point of a perfect partition $\bar{A}$ on the continuous cake (within distance at most $1/K$ from the point of view of each player)
- the allocation order (from left to right) in A is the same as the one in $\bar{A}$.

Then for each contiguous piece $X \in A$, the value of a player i for X is the same as player i's value for the corresponding piece $\bar{X}$ in the perfect partition $\bar{A}$, except possibly for a gain or loss of $2/K$ due to estimation errors (at most $1/K$ at each endpoint of X) It follows that A approximates $\bar{A}$ within an error of at most ϵ. Finally, once $\mathcal{M}_\epsilon$ finds an appropriate partition, it allocates it using a random permutation π, and so the expected value of each player is *exactly* $1/n$, regardless of the strategies of the other players, as in the Mossel-Tamuz protocol. Thus $\mathcal{M}_\epsilon$ is truthful in expectation and ϵ-perfect. $\square$

Chapter 6

Simultaneous Cake Cutting

In this chapter, we introduce a novel computational model that, we believe, provides a fundamentally new perspective on cake cutting; we call it the *simultaneous model*. In our model, the players *simultaneously* report compact versions of their preferences, specifically, their values for specific pieces of cake; this information is used to compute a fair allocation, without further communication between the players. We define the *complexity* of a simultaneous protocol as the maximum number of pieces whose values a player may need to report.[1]

Cake cutting protocols in the simultaneous model have two advantages compared to their counterparts in the Robertson-Webb model:

1. Elicitation of preferences can be done in parallel.

2. The existence of computationally efficient simultaneous protocols would imply that players' valuation functions can be *sketched* in a way that preserves sufficient information for recovering a fair cake division (via the protocol).

On the other hand, the simultaneous model severely restricts the power of protocols. Is the restriction so severe that fair divisions, according to standard fairness properties, cannot be computed? Our research question is

> *... which fairness properties are computationally feasible in the simultaneous model, and what is the complexity of computing cake divisions satisfying those properties?*

Our simultaneous model of cake cutting is related to, and conceptually draws on, work on *communication complexity* [85] and *streaming algorithms* [101].

[1] This definition is better, formally and intuitively, than taking the *overall* amount of communication (summed over all players); it is also consistent with related work on communication complexity [82].

In particular, Kremer et al. [82] studied the relation between one-round communication complexity and simultaneous communication complexity. Similarly to our model, streaming algorithms deal with compact representations — called *sketches* — of data. Some papers focus specifically on sketching valuation functions or preferences in various contexts [11, 12, 39].

6.1 The Simultaneous Model

We define a *discretization* of the cake as a tuple $(\bar{x}, \bar{w})$, for which there exists $m \in \mathbb{N}$ such that:

- $\bar{x} = (x^0, x^1, \ldots, x^{m-1}, x^m)$ is a sequence of *cut points* with $0 = x^0 < x^1 < \cdots < x^{m-1} < x^m = 1$.
- $\bar{w} = (w^1, \ldots, w^m)$ is a sequence of values, such that w^i represents the value of the piece $[x^{i-1}, x^i]$ and $w^1 + \ldots + w^m = 1$.

Let $\mathcal{D}$ denote the space of all discretizations. Then a one-round protocol can be defined as follows:

Definition 13 (*Simultaneous protocol*)**.** *A* simultaneous protocol *is a function $\mathcal{F} : \mathcal{V} \to \mathcal{D}$, where $\mathcal{V}$ is the space of valuations, $\mathcal{D}$ is the space of discretizations of the cake, and $\mathcal{F}(V)$ is the discretization that a player is instructed to report when his valuation function is V.*

One could alternatively define a simultaneous protocol as reporting a set of (possibly overlapping) subintervals and their values. However, the two definitions are essentially equivalent for our purposes.

Note that in this section an allocation of the cake is denoted by $A = (A_1, \ldots, A_n)$. We focus on the Robertson-Webb model as formalized by Procaccia [109].

What does it mean for a simultaneous protocol to satisfy a *property*, such as envy-freeness, proportionality, or Pareto optimality? This question involves surprising subtleties even in the Robertson-Webb model, and so the definition must be carefully chosen. Very roughly speaking, the main difficulty (in both models) is that players could potentially use an injection from the space of valuation functions to $[0, 1]$ to encode their entire valuation function as a single number (e.g., the first cut point they make). For any given property, that would give enough information to compute an allocation satisfying the property (if one exists). The definition below circumvents this problem, by capturing the idea that reporting a value for an interval commits the player to a valuation function that actually assigns the reported value to that interval, and nothing else.

Definition 14 (*Property of a simultaneous protocol*)**.** *Let $\mathcal{P}$ be a property of cake allocations. A* protocol $\mathcal{F}$ satisfies property $\mathcal{P}$ *if the following holds for any tuple of valuation functions $\bar{V} = (V_1, \ldots, V_n)$:*

- *Whenever each player i follows the protocol by reporting its recommended discretization, $(\bar{x}_i, \bar{w}_i) := \mathcal{F}(V_i)$, there exists an allocation A that satisfies $\mathcal{P}$ with respect to any other valuations $\bar{V}' = (V_1', \ldots, V_n')$ consistent with the discretizations reported at $\bar{V}$ (i.e., $V_i'([x_i^{j-1}, x_i^j]) = w_i^j, \forall i, j$).*

For example, let us describe an envy-free simultaneous protocol $\mathcal{F}$ for two players. $\mathcal{F}(V_i)$ is the discretization $(\bar{x}_i, \bar{w}_i)$ where $\bar{x}_i = (x_i^0 = 0, x_i^1, x_i^2 = 1)$ and $\bar{w}_i = (\frac{1}{2}, \frac{1}{2})$, that is, each player essentially cuts the cake into two pieces worth $1/2$ using the cut point x_i^1. Now assume without loss of generality that $x_1^1 \leq x_2^1$, and consider the allocation $A_1 = [0, x_1^1]$, $A_2 = [x_1^1, 1]$. This allocation is clearly envy-free for the reported valuation functions, and, moreover, it is envy free for any valuation functions where $V_i'([0, x_i^1]) = V_i'([x_i^1, 1]) = 1/2$ for $i = 1, 2$.

In the Robertson-Webb model, the complexity of a protocol is the maximum number of cut and evaluation queries. We use an equivalent definition in the simultaneous model.

Definition 15 (*Complexity of a simultaneous protocol*)**.** *The* complexity *of a simultaneous protocol is the maximum number of intervals in the discretization $\mathcal{F}(V)$ taken over all $V \in \mathcal{V}$ (that is, the maximum number of cut points minus one). If the maximum does not exist, we say that the protocol is* unbounded*.*

For example, the complexity of the envy-free simultaneous protocol for two players is 2.

6.2 Proportionality

We start by examining proportionality in the simultaneous model. In the Robertson-Webb query model, the complexity of computing proportional allocations is $\Theta(n \log n)$: an $O(n \log n)$ upper bound is given by the Even-Paz [61] protocol, and a matching lower bound was established by Edmonds and Pruhs [58].

Similarly, the simultaneous model turns out to admit the computation of proportional allocations, but the complexity of proportionality in this model is only $\Theta(n)$. For the upper bound, we describe a protocol that is a simultaneous interpretation of a protocol designed in a different context by Manabe and Okamoto [91]. Importantly, this protocol requires $\Theta(n^2)$ cut queries in the Robertson-Webb model; but the simultaneous model allows us to implicitly parallelize the queries to the players, leading to a reduction in complexity. The simultaneous model captures the insight that the information elicited from one player does not need to rely on the information elicited from another.

Theorem 14. *There exists a proportional simultaneous protocol with complexity n.*

Proof. Consider the following simultaneous protocol:

- *Map the valuation function of each player to n disjoint contiguous intervals of value exactly $1/n$ each.*

Formally, the discretization is defined by $\bar{x}_i = (x_i^0, \ldots, x_i^n)$ and $w_i^j = 1/n$, for all $j = 1, \ldots, n$.

Given the intervals submitted by the players, we produce an allocation by scanning the cake from left to right until the first mark, $x_{i_1}^1$, of some player $i_1 \in N$ is encountered. Allocate the piece $[0, x_{i_1}^1]$ to player i_1. Then, scan to the right starting with the point $x_{i_1}^1$ while looking for the second mark $x_{i_2}^2$ of some player $i_2 \in N \setminus \{i\}$. Allocate the piece $[x_{i_1}^1, x_{i_2}^2]$ to player i_2 and continue in this fashion until the entire cake is allocated.

To see why the protocol is proportional, note that for player i_t that was allocated in round t, $x_{i_t}^{t-1} \geq x_{i_{t-1}}^{t-1}$, because i_t was not selected in round $t-1$. Thus, $[x_{i_t}^{t-1}, x_{i_t}^t] \subseteq [x_{i_{t-1}}^{t-1}, x_{i_t}^t]$. Moreover, $A_{i_t} = [x_{i_{t-1}}^{t-1}, x_{i_t}^t]$ and $V_{i_t}([x_{i_t}^{t-1}, x_{i_t}^t]) = 1/n$, thus $V_{i_t}(A_{i_t}) \geq 1/n$. □

Next, we show the bound given in Theorem 14 is tight.

Theorem 15. *Every proportional simultaneous protocol has complexity at least n.*

Proof. Assume by contradiction that there exists a proportional simultaneous protocol $\mathcal{F}$ with complexity less than n. Without loss of generality, let V_n be a valuation function such that $\mathcal{F}(V_n)$ reports the values of $n-1$ intervals with cut points $(x_n^0, \ldots, x_n^{n-1})$. (The case where the player reports fewer intervals is similar.) Then the valuations of the other players can be set such that for every player $i \in N \setminus \{n\}$, the entire value of the cake from the point of view of player i is concentrated in the interval $[x_n^{i-1}, x_n^i]$, that is, $V_i([x_n^{i-1}, x_n^i]) = 1$.

Let us now consider two (exhaustive) types of allocations. First, let A be an allocation such that for all $i \in \{1, \ldots, n-1\}$, player i gets a nonempty interval $I_i \subseteq [x_n^{i-1}, x_n^i]$. We can define the valuation function V_n' where $V_n'(I_i) = V_n([x_n^{i-1}, x_n^i])$ for all $i \in \{1, \ldots, n-1\}$. Then V_n' is consistent with player n's reported intervals, but $V_n'(A_n) = 0$, so the allocation is not proportional with respect to V_n'.

Second, let A be an allocation such that there exists a player $i \in \{1, \ldots, n-1\}$ that does not get a nonempty interval $I_i \subseteq [x_n^{i-1}, x_n^i]$. Then clearly $V_i(A_i) = 0$, and again the allocation is not proportional. □

6.3 (Approximate) Envy-Freeness, and Beyond

We have seen that simultaneous protocols can compute proportional allocations. For two players, proportionality and envy-freeness coincide, but for more players, envy-freeness is strictly stronger. It has long been known that envy-free allocations are guaranteed to exist, but it wasn't until the nineties that an envy-free protocol that can be simulated in the Robertson-Webb model was discovered [20].

The Brams-Taylor protocol is finite (i.e. terminates on every instance), but unbounded: its running time cannot be bounded by a function of the number of players, and so the execution can take arbitrarily long depending on the valuation functions themselves. It is an open problem whether a bounded envy-free protocol exists in the Robertson-Webb model for any number of players.

Our next result shows that no simultaneous protocol can be envy free. Interestingly, this impossibility result does not assume that the protocol is bounded: it says that there are valuation functions for which there is no discretization that is fine enough to guarantee envy-freeness in the simultaneous model.

Theorem 16. *For $n \geq 3$ there does not exist an envy-free simultaneous protocol.*

Proof. Let V_1 be the uniform valuation function (i.e., its value density function is $v(x) \equiv 1$), which yields a discretization $\mathcal{F}(V_1) = (\bar{x}_1, \bar{w}_1)$ under protocol $\mathcal{F}$. Let there be m reported intervals, and denote $X^i = [x_1^{i-1}, x_1^i]$ for $i = 1, \ldots, m$; then $w_1^i = |X^i| = x_1^i - x_1^{i-1}$. We will show that there exist valuation functions for the other players such that no envy-free allocation can be computed from these reported intervals.

Define a constant $c \in \left(\frac{1}{w_1^1+1}, 1\right)$ such that for all $i \in N \setminus \{1\}$, the value density function v_i of player i satisfies the following conditions:

(a) For all $j \in \{1, \ldots, m\}$, v_i is constant on X^j.

(b) $V_i(X^1) = c \cdot w_1^1 + 1 - c$

(c) $V_i(X^j) = c \cdot w_1^j$, for all $j \in \{2, \ldots, m\}$

(d) There do not exist distinct indices $a_1, \ldots, a_x \in \{1, \ldots, m\}$ such that the following identity holds:

$$w_1^{a_1} + \cdots + w_1^{a_x} = \frac{1}{c \cdot n}.$$

Note that any $c \in \left(\frac{1}{w_1^1+1}, 1\right)$ induces valid valuation functions that satisfy (b) and (c), because

$$V_i([0,1]) = \sum_{j=1}^{m} V_i(X^j) = c\left(\sum_{j=1}^{m} w_1^j\right) + (1-c) = 1.$$

Moreover, constraint (d) can be satisfied because there is an (uncountably) infinite number of possible values of c, and the constraint only rules out a finite number of them.

Let $A = (A_1, \ldots, A_n)$ be an allocation computed by the protocol. We consider two cases, depending on whether the interval X^1 is split or not among the players.

Case I: Interval X^1 is not split. We have several subcases:

(i) $|A_1| < \frac{1}{n}$: Then there exists another player i that receives a piece of length at least $\frac{1}{n}$ and player 1 envies i.

(ii) $|A_1| \geq \frac{1}{n}$ and player 1 receives X^1. Then the value of the other players for the piece received by 1 is:

$$\begin{aligned} c \cdot w_1^1 + 1 - c + c(|A_1| - w_1^1) &= c \cdot |A_1| + 1 - c \\ &\geq \frac{c}{n} + 1 - c. \end{aligned}$$

The length of the piece for all the other players is at most $\frac{n-1}{n}$. Since the remainder of the cake does not contain X^1, the minimum value $V_i(A_i)$ — taken over all $i \in \{2, \ldots, n\}$ — is at most $c\left(\frac{n-1}{n}\right)\left(\frac{1}{n-1}\right) = \frac{c}{n}$. It follows that there exists a player i that envies 1.

(iii) $|A_1| \geq \frac{1}{n}$ and a player $i \in N \setminus \{1\}$ receives X^1.

It must be the case that $|A_j| = |A_k|$ for all $j, k \in N \setminus \{i\}$ to prevent envy. For the same reason, we have $V_j(A_i) = V_j(A_j)$ for all $j \in \{2, \ldots, n\}$. Therefore, all the players, except player 1, value $A_1, \ldots, A_n$ equally. It follows that $V_j(A_1) = \frac{1}{n}$ for all $j \in \{2, \ldots, n\}$. This implies that $|A_1| = \frac{1}{c \cdot n}$, so by property (*d*), there exists a reported interval that is split between player 1 and at least one other player. Now we can define V_1' that is consistent with $\mathcal{F}(V_1)$, where player 1's value for its part(s) of the split interval(s) is zero; then player 1 would be envious.

Case II: Interval X^1 is split among at least two players. For each $i \in N$, let $A_i' = A_i \setminus X^1$. We have two subcases:

(i) There exists exactly one player $i \in N \setminus \{1\}$, such that $A_i \cap X^1 \neq \emptyset$. Then $A_1 \cap X^1 \neq \emptyset$. Consider another player $j \in N \setminus \{1, i\}$.

- If $|A_j| > |A'_1|$, let V'_1 be a valuation function consistent with $\mathcal{F}(V)$ such that player 1 has a value of zero for his portion of X^1. Then $V'_1(A_1) = V'_1(A'_1) < V'_1(A_j)$, violating envy-freeness.
- If $|A_j| \leq |A'_1|$, then $V_j(A'_1) \geq V_j(A_j)$. Moreover, $V_j(A_1 \setminus A'_1) > 0$. It follows that $V_j(A_1) > V_j(A_j)$, violating envy-freeness.

(ii) There exist distinct players $i, j \in N \setminus \{1\}$ such that $A_i \cap X^1 \neq \emptyset$, $A_j \cap X^1 \neq \emptyset$. Assume without loss of generality that player i's piece satisfies $|A_i \cap X^1| \leq \frac{|X^1|}{2}$. Then

$$V_i(A_i \cap X^1) \leq \frac{1}{2}(c \cdot w_1^1 + 1 - c)$$

and $V_i(A_1) \geq c|A'_1|$. It must also be the case that

$$|A'_i| \geq |A'_1| - \frac{1}{2}\left(w_1^1 + \frac{1}{c} - 1\right)$$

since otherwise

$$\begin{aligned} V_i(A_i) &\leq c|A'_i| + \frac{1}{2}\left(c \cdot w_1^1 + 1 - c\right) \\ &< c\left(|A'_1| - \frac{1}{2}\left(w_1^1 + \frac{1}{c} - 1\right)\right) + \frac{1}{2}(c \cdot w_1^1 + 1 - c) \\ &= c|A'_1|, \end{aligned}$$

which would imply that $V_i(A_1) > V_i(A_i)$.

Consider V'_1 consistent with $\mathcal{F}(V_1)$ such that $V'_1(A_1 \cap X^1) = 0$, and $V'_1(A_i \cap X^1) = w_1^1$. Then player 1's value for i's piece is:

$$\begin{aligned} V'_1(A_i) = w_1^1 + |A'_i| &\geq w_1^1 + |A'_1| - \frac{1}{2}\left(w_1^1 + \frac{1}{c} - 1\right) \\ &> w_1^1 + |A'_1| - \frac{1}{2}(w_1^1 + w_1^1) \\ &= |A'_1| = V'_1(A_1), \end{aligned}$$

where the third transition holds by the choice of $c \in \left(\frac{1}{w_1^1+1}, 1\right)$. Thus player 1 envies player i.

□

Theorem 16 tells us that we cannot hope to obtain envy-free allocations in the simultaneous model. However, it turns out that we can reach envy-free allocations arbitrarily close. Indeed, we say that an allocation is ϵ-envy free if for all $i, j \in N$, $V_i(A_i) \geq V_i(A_j) - \epsilon$. This notion of approximate envy-freeness has been studied in several previous papers [46, 50, 89]. We will show that there exists an ϵ-envy-free protocol of polynomial complexity in n and ϵ.

The main idea is to sketch the players' valuations using a very fine discretization, but then use a coarser discretization to partition the cake into *indivisible* goods. Then, each player's value for each indivisible good can be accurately estimated using the fine discretization. An allocation of the indivisible goods that is approximately envy-free with respect to the estimated values is therefore also approximately envy-free with respect to the real values (with a slightly worse additive approximation term).

Theorem 17. *For every $\epsilon > 0$ there exists an ϵ-envy-free simultaneous protocol with complexity $O(n/\epsilon^2)$.*

The proof uses the following lemma, which is a special case of a result by Lipton *et al.* [89], and deals with the allocation of *indivisible* goods. In this context, the valuation functions are said to be *additive* if the value of a bundle of goods is the sum of values of goods in the bundle.

Lemma 10 (Lipton et al. 2004)**.** *Let $V_1', \ldots, V_n'$ be additive valuation functions over a set G of indivisible goods. Assume that for all $i \in N$ and $g \in G$, $V_i(g) \leq \epsilon$. Then there exists an ϵ-envy-free allocation.*

Proof of Theorem 17. For every n and $\epsilon > 0$ we design a simultaneous protocol $\mathcal{F}^{n,\epsilon}$. Given a valuation V, $\mathcal{F}^{n,\epsilon}$ discretizes the cake as follows. First, the *coarse discretization* has $1/\delta$ subintervals of value δ each, for $1/\delta = \lceil 2/\epsilon \rceil$; note that $\delta \leq \epsilon/2$. Second, the *fine discretization* includes $1/\delta'$ intervals of value δ' each, for $1/\delta' = \lceil 16n/\epsilon^2 \rceil$; note that $\delta' \leq \epsilon^2/16n$. Formally speaking, $\mathcal{F}^{n,\epsilon}(V)$ contains the cut points of both discretizations, but we prefer to think of these two different discretizations for ease of exposition.

Given $\mathcal{F}^{n,\epsilon}(V_1), \ldots, \mathcal{F}^{n,\epsilon}(V_n)$, we wish to show that there is an allocation A that is ϵ-envy free with respect to any valuation functions that are consistent with these reported discretizations. Consider the partition of the cake obtained by ordering the cut points of all players' coarse discretizations, and treating the subinterval between two adjacent cut points as an *indivisible good.* Denote the set of indivisible goods by G.

This partition into indivisible goods has two properties:

1. For each indivisible good $g \in G$, $V_i(g) \leq \epsilon/2$ for all $i \in N$, because for each $i \in N$ there is a subinterval of the coarse discretization of V_i that contains g.

2. The number of indivisible goods is given by the number of "internal" (not 0 or 1) cut points plus one, i.e.,

$$|G| \leq n\left(\left\lceil \frac{1}{\delta} \right\rceil - 1\right) + 1 \leq \frac{4n}{\epsilon}.$$

Let us create additive valuation functions $V'_1, \dots, V'_n$ over the indivisible goods in G. For $g \in G$, let $H_i(g)$ be the set of intervals in the fine partition of V_i that are contained inside g. We define $V'_i(g) = \delta' \cdot |H_i(g)|$.

We claim that

$$V'_i(g) \le V_i(g) \le V'_i(g) + 2\delta'. \tag{6.1}$$

Indeed, the left inequality is trivial. For the right inequality, let I be the interval obtained by taking $H_i(g)$ and adding one subinterval to the left and one to the right. It holds that $g \subseteq I$, hence

$$V_i(g) \le V_i(H_i(g)) = \delta' \cdot (|H_i(g)| + 2) = V'_i(g) + 2\delta'.$$

Note that for all $i \in N$ and $g \in G$, $V'_i(g) \le V_i(g) \le \epsilon/2$. We can therefore use Lemma 10 to create an allocation A of the goods G such that for all $i, j \in N$, $V'_i(A_i) \ge V'_i(A_j) - \epsilon/2$. We claim that A is ϵ-envy free with respect to the original valuation functions (and any other valuations that are consistent with the reported discretizations). Indeed,

$$\begin{aligned}
V_i(A_i) &\ge V'_i(A_i) \ge V'_i(A_j) - \frac{\epsilon}{2} \\
&= \left(\sum_{g \in A_j} V'_i(g) \right) - \frac{\epsilon}{2} \ge \left(\sum_{g \in A_j} (V_i(g) - 2\delta') \right) - \frac{\epsilon}{2} \\
&= V_i(A_j) - 2\delta'|A_j| - \frac{\epsilon}{2} \ge V_i(A_j) - 2\delta'|G| - \frac{\epsilon}{2} \\
&\ge V_i(A_j) - 2 \cdot \frac{\epsilon^2}{16n} \cdot \frac{4n}{\epsilon} - \frac{\epsilon}{2} = V_i(A_j) - \epsilon.
\end{aligned}$$

where the first and fourth transitions follow from Equation (6.1). □

Envy-freeness and proportionality are examples of what we call *linear* properties, in the sense that they are specified by linear constraints involving the players' valuations for pieces. Another example of a linear property is *equitability*, which requires that $V_i(A_i) = V_j(A_j)$ for all $i, j \in N$, that is, players must have identical values for their own pieces. We formally define linear properties using the matrix form, as is common in linear programs.

Definition 16 (*Linear property*)**.** *A property of allocations is* linear *if there exist $m \in \mathbb{N}$, matrix $B \in \mathbb{R}^{m \times n^2}$ such that $\sum_{j=1}^{n^2} |B_{ij}| \le 1$ for $i = 1, \dots, m$, and vector $c \in \mathbb{R}^m$, such that an allocation A satisfies the property if it satisfies the constraints: $B \cdot \alpha \ge c$, where $\alpha_k = V_i(A_j)$, with $i = \lceil \frac{k}{n} \rceil$, and $j = k \mod n$ if $n \nmid k$ and $j = n$ otherwise.*

To illustrate the definition of α_k, note that $\alpha_1 = V_1(A_1)$, $\alpha_n = V_1(A_n)$, and $\alpha_{n+1} = V_2(A_1)$. Importantly, this representation captures equality constraints, as they can be represented using two inequalities. Furthermore, the assumption that $\sum_{j=1}^{n^2} |B_{i,j}| \le 1$ is without loss of generality: we just divide each entry in the matrix B and vector c by the maximum sum of absolute values of any row of B, which is a constant in the context of the properties we are interested in.

As an example, we explicitly represent envy-freeness as a linear property for the case of three players. Let $m := n(n-1) = 6$ and define:

$$B = \frac{1}{2} \cdot \begin{pmatrix} 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \end{pmatrix}$$

and

$$c = (0, 0, 0, 0, 0, 0) .$$

Then the constraint $B \cdot \alpha \geq c$ is equivalent to requiring that $V_i(A_i) - V_i(A_j) \geq 0$, for each $i \neq j$.

Every linear property P (defined by a matrix B and a vector c) naturally admits an approximate version $\mathcal{P}_\epsilon$, which requires each linear constraint of $\mathcal{P}$ to hold up to an error of ϵ; formally, $B \cdot \alpha \geq c - \epsilon \cdot \mathbf{1}$. Using this new notion, we can establish a more general version of Theorem 17.

Theorem 18. *For every $\epsilon > 0$ and every bounded protocol in the Robertson-Webb model that allocates the entire cake and guarantees some linear property $\mathcal{P}$ with complexity $f(n)$, there exists a simultaneous protocol that guarantees the property $\mathcal{P}_\epsilon$ with complexity $O(f(n)/\epsilon)$.*

The theorem's proof appears in the appendix, which was submitted as supplementary material (it also contains a formal definition of properties in the Robertson-Webb model). Theorem 18 implies Theorem 17 because ϵ-envy free allocations can be computed in the Robertson-Webb model using $O(n/\epsilon)$ queries. And while exact equitability is impossible to achieve in the Robertson-Webb model [42], ϵ-equitability can also be achieved with with $O(n/\epsilon)$ queries, leading to an ϵ-equitable simultaneous protocol with complexity $O(n/\epsilon^2)$.

We note that a technique for approximating general density functions with piecewise constant density functions [46, Lemma 8] can be leveraged to obtain a strictly weaker version of Theorems 17 and 18, requiring the assumption that the value density functions are piecewise K-Lipschitz continuous, and giving a bound that also depends on K.

6.4 Discussion

In some ways, simultaneous protocols are weaker than their counterparts in the Robertson-Webb model: players cannot interact, but rather are allowed to send one message only. However, in other ways simultaneous protocols are stronger. Indeed, under the Robertson-Webb model, information is elicited via cut and evaluation queries, without ever seeing the full valuations. This means that properties such as Pareto optimality are impossible to achieve in this model, even when the value density functions are restricted to be piecewise constant and the protocol is allowed to have unbounded complexity [84]. Intuitively, the reason is that a Pareto optimal allocation cannot allocate to player i a subinterval I such that $V_i(I) = 0$ and $V_j(I) > 0$. But in the Robertson-Webb model, it is impossible to exactly identify the boundaries of subintervals that are worthless to a player.

In contrast, in the simultaneous model players can observe their full valuation functions before deciding which subintervals to report, which allows them to exactly mark worthless intervals. Now, suppose for simplicity that the players' value density functions are piecewise constant, so each has a finite number of intervals on which its density is zero. Each player reports a discretization that pinpoints the zero-density intervals. Then we can allocate the intervals using *serial dictatorship*: in stage i, allocate to player i all unclaimed intervals on which its density is positive. This allocation is clearly Pareto-optimal.

Unfortunately, the protocol just described is not formally Pareto optimal according to Definition 2, because the allocation is not guaranteed to be Pareto optimal with respect to all valuation functions consistent with the reports (some may have additional worthless subintervals). In fact, Pareto optimality cannot be guaranteed in the simultaneous model — as can be shown using an argument that is similar to the proof of the equivalent result in Robertson-Webb [84, Theorem 5]. However, this negative result can be circumvented via a slight augmentation of the simultaneous model, which allows players to mark intervals on which their density is strictly positive.

It is therefore natural to consider a relaxed model that allows protocols to enjoy the best of both worlds: multi-round interaction à la Robertson-Webb, and allowing players to report discretizations by observing their own valuation function (and information previously communicated by others) à la the simultaneous model. This hybrid model gives rise to intriguing questions. Most importantly: does it admit *bounded* envy-free protocols? We view this question as a natural, compelling relaxation of what is perhaps the most enigmatic open problem in computational fair division [109]: settling the existence of bounded envy-free protocols in the Robertson-Webb model.

6.5 Proof of Linear Properties Theorem

In this section we prove Theorem 18. But we first need to formally introduce the notion of property for cake cutting protocols in the Robertson-Webb model. For ease of exposition, we restrict attention to protocols that allocate the entire cake and only use cut points discovered through queries. However, the proof carries over to the case where the protocol can use arbitrary cuts and discard portions of the cake.

Definition 17 (*Property of a Robertson-Webb protocol*)**.** *Let $\mathcal{P}$ be a property of cake allocations. A protocol $\mathcal{F}$ in the Robertson-Webb model satisfies property $\mathcal{P}$ if the following holds for any tuple of valuation functions $\bar{V} = (V_1, \ldots, V_n)$:*

- *Whenever each player i answers the cut and evaluate queries addressed by $\mathcal{F}$ correctly (i.e. according to V_i), the protocol outputs an allocation A that satisfies $\mathcal{P}$ with respect to any valuations $\bar{V}' = (V_1', \ldots, V_n')$ consistent with the answers given by the players during the execution of $\mathcal{F}$ on $\bar{V}$.*

We are now ready to restate and prove Theorem 18.

Theorem 18. *For every $\epsilon > 0$ and every bounded protocol in the Robertson-Webb model that allocates the entire cake and guarantees some linear property $\mathcal{P}$ with complexity $f(n)$, there exists a simultaneous protocol that guarantees the property $\mathcal{P}_\epsilon$ with complexity $O(f(n)/\epsilon)$.*

Proof. Let $\mathcal{M}$ be a bounded protocol in the Robertson-Webb model that guarantees a linear property $\mathcal{P}$, where $\mathcal{P}$ is given by $B \cdot \alpha \geq c$, for some $m \in \mathbb{N}$, $B \in \mathbb{R}^{m \cdot n^2}$, and $c \in \mathbb{R}^m$. Moreover, let $f(n)$ be the maximum number of steps that $\mathcal{M}$ takes on an instance with n players. Each query makes two "marks": $\text{Eval}_i(x, y)$ makes marks at x and y, and $\text{Cut}_i(x, \alpha) = y$ makes marks at x and the point y such that $V_i([x, y]) = \alpha$. Overall, $\mathcal{M}$ makes at most $2f(n)$ marks.

For every $\epsilon > 0$, let $\mathcal{F}_{\mathcal{P}}^{\epsilon}$ be the simultaneous protocol stated as Algorithm 7.

Protocol $\mathcal{F}_{\mathcal{P}}^{\epsilon}$ asks each player i to submit a discretization of the cake containing very small cells of equal value according to i. Then $\mathcal{F}_{\mathcal{P}}^{\epsilon}$ guesses (by trying all possibilities) the number of contiguous intervals used by $\mathcal{M}$, and then approximates the pieces discovered by $\mathcal{M}$ using the discretizations provided by the players. Next we show that one of these guesses is guaranteed to work.

Given an arbitrary tuple of valuation functions $\bar{V} = (V_1, \ldots, V_n)$, let $Y = \{y_0, y_1, \ldots, y_{M-1}, y_M\}$ be the marks made during the execution of $\mathcal{M}$ when the valuations of the players are $\bar{V}$, where $y_0 = 0$, $y_M = 1$, and $M \leq 2f(n)+1$. Denote by $I = (I_1, \ldots, I_M)$ the resulting disjoint, consecutive contiguous intervals with $I_j = (y_{j-1}, y_j)$. Let $A = (A_1, \ldots, A_n)$ be the allocation computed by protocol $\mathcal{M}$. We can assume without loss of generality that each piece A_i is a union of intervals from I [108], $A_i \cap A_j = \emptyset$, $\forall i, j$ and $\bigcup_{i=1}^{n} A_i = [0, 1]$.

For each mark $y_j \in Y$, let z_j be the rightmost point in X with the property that $z_j \leq y_j$ (recall that X is the collection of points submitted by all players

Map the valuation of each player i to a discretization $(\bar{x}_i, \bar{w}_i)$ consisting of $T = \left\lceil \frac{4f(n)+2}{\epsilon} \right\rceil$ cells, each worth $1/T$ to player i.
$X \leftarrow \bigcup_{i=1}^{n} \bigcup_{j=1}^{T} \{x_{i,j}\}$
for $M = 1$ *to* $f(n)+1$ **do**
  for each subset $Y \subseteq X$, where $|Y| = M+1$ **do**
    for each allocation A demarcated only by points in Y **do**
      for each player $i \in N$ and piece $A_j \in A$ **do**
        $n_{i,j} \leftarrow$ # intact cells in A_j from $\bar{x}_i$
        $k \leftarrow (i-1) \cdot n + j$
        $\tilde{\alpha}_k \leftarrow n_{i,j} \cdot \left(\frac{1}{T}\right)$
      end for
      if $B \cdot \tilde{\alpha} \geq c - \epsilon \cdot \mathbf{1}$ **then**
        return A
      end if
    end for
  end for
end for

Algorithm 7: Simultaneous protocol $\mathcal{F}_{\mathcal{P}}^{\epsilon}$

under $\mathcal{F}_{\mathcal{P}}^{\epsilon}$) Observe that for each player i, we have that $V_i(z_j, y_j) \leq 1/T$. Then we can construct an approximate version $\tilde{I}_j$, of each interval I_j such that the endpoints of $\tilde{I}_j$ belong to the set $\{0, z_1, \ldots, z_{M-1}, 1\}$. More formally, we find the intervals $\tilde{I} = (\tilde{I}_1, \ldots, \tilde{I}_M)$ by scanning the cake from left to right as follows:

1. Let $z_1 \in X$ be maximum such that $z_1 \leq y_1$.
2. $\tilde{I}_1 \leftarrow [0, z_1]$.
3. *For each* $j \in \{2, \ldots, M-1\}$:
   - Let $z_j \in X$ be maximum such that $z_j \leq y_j$.
   - *If* $(z_j = z_{j-1})$ *then*:
     $I_j \leftarrow \emptyset$
   - *Else*:
     $\tilde{I}_j \leftarrow [z_{j-1}, z_j]$
4. $\tilde{I}_M \leftarrow [z_{M-1}, 1]$.

By construction, for each player i and interval $\tilde{I}_j$ we have that $|V_i(\tilde{I}_j) - V_i(I_j)| \leq \frac{2}{T}$; intuitively, player i views $\tilde{I}_j$ as identical to I_j, except possibly for the two endpoints of the interval, where the player might have lost or gained a cell of value $1/T$ on each side.

Define an allocation $\tilde{A} = (\tilde{A}_1, \ldots, \tilde{A}_n)$, such that $\tilde{I}_j \in \tilde{A}_i$ if and only if $I_j \in A_i$, for all $i \in N$ and $j \in [M]$. Then since each piece $\tilde{A}_k$ contains at most M contiguous intervals from $\tilde{I}$, we have that $\tilde{A}_k$ is an approximation of A_k within an additive error term of $M \cdot \left(\frac{2}{T}\right)$, from the point of view of each player. More formally,

$$|V_i(A_k) - V_i(\tilde{A}_k)| \leq \frac{2M}{T} \leq \frac{2(2(f(n)+1)}{\left\lceil \frac{4f(n)+2}{\epsilon} \right\rceil} \leq \epsilon,$$

for all $i \in N$.

Next we show that allocation $\tilde{A}$ approximately satisfies property $\mathcal{P}$. Recall that $\mathcal{P}$ is defined as $B \cdot \alpha \geq c$, where α is the vector with the values of each player for every piece in A.

For each row $i \in [m]$, allocation A satisfies the constraint: $\sum_{j=1}^{n^2} B_{i,j}\alpha_j \geq c_i$, where $\alpha_j = V_k(A_l)$ and

- $k = \lceil \frac{j}{n} \rceil$
- $l = j \mod n$ if $n \nmid j$ and $l = n$ otherwise.

Let $\tilde{\alpha}_j = V_k(\tilde{A}_l)$. We have shown that $|\tilde{\alpha}_j - \alpha_j| \leq \epsilon$. By definition, $\sum_{j=1}^{n^2} |B_{i,j}| \leq 1$, and therefore we have:

$$\sum_{j=1}^{n^2} B_{i,j}\tilde{\alpha}_j \geq \sum_{j=1}^{n^2} B_{i,j}\alpha_j - \epsilon \sum_{j=1}^{n^2} |B_{i,j}| \geq c_i - \epsilon.$$

It follows that $B \cdot \tilde{\alpha} \geq c - \epsilon \cdot \mathbf{1}$, and so the allocation $\tilde{A}$ approximately satisfies property $\mathcal{P}$. The simultaneous protocol $\mathcal{F}_{\mathcal{P}}^{\epsilon}$ checks all the possible allocations that can be formed with the cut points submitted by the players, and one of these (i.e. $\tilde{A}$) is guaranteed to work; thus the allocation computed by $\mathcal{F}_{\mathcal{P}}^{\epsilon}$ ϵ-satisfies $\mathcal{P}$ whenever the valuations of the players are consistent with the discretizations $(\bar{x}, \bar{w})$. □

Chapter 7

Externalities

Recall that two of the most prominent notions of fairness are *proportionality* and *envy-freeness*. Informally, proportionality requires that each of the n players involved in the division of the resource receive at least $1/n$ of the total value. Envy-freeness is a much stronger notion, which stipulates that no player prefer another player's allocation to their own. On a closer look, it becomes clear that the two notions of fairness are fundamentally different. While proportionality requires each player to only evaluate the quality of their own allocation (compared to their best possible), the very idea of envy assumes that players naturally compare their own allocations with those of others. This latter notion is derived from psychology research and conveys the more general concept that players are influenced not only by their own state, but also by the states of other players. Such influences are called *externalities*.

Generally speaking, externalities are costs or benefits that are not transmitted through prices, and may be incurred by a party that was not involved in a transaction. For example, vaccination reduces the risk of illness not only for the individual receiving the vaccine, but for all others around them. In network formation games externalities are known as network effects, and play an important role during the adoption of new technologies [56]. For example, when the phone was introduced, the value of the phone for a potential customer depended on how many other people were also using a phone.

Externalities play a role in resource allocation settings, where the allocation of one player can affect the others. These circumstances are particularly relevant in the context of social networks, where players derive value from the allocations of others due to the existence of synergies. For example, consider the scenario in which each player is trying to carry out an online project and is allocated slots of working time on a server. The players may be able to use portions of their collaborators' idle time to run additional experiments and improve the quality of the project. Similarly, the exploitation of land (e.g. crop harvesting or road construction) can be done more efficiently by the players with the most advanced equipment, and their efforts can benefit everyone else. Our goal here is to model externalities in cake cutting; in particular, addressing the conceptual challenge of defining fairness and understanding the existence and computability of fair allocations in this model.

7.1 Related Work

Theories of externalities are widely studied in economics [7, 80], but recently have also been receiving increasing attention in the computer science literature. Such studies include the analysis of externalities in coalitional games [96], auctions [70, 83], voting [3], and matchings [35].

Velez [124] considers externalities in the fair division of indivisible goods and money (e.g., tasks and salary). On the conceptual side, among other contributions he (independently) introduces the notion of swap envy-freeness, which we discuss below. On the technical side his intriguing results can be mapped to the cake cutting setting, but the outcome is rather restricted. Specifically, in the cake cutting context his results only capture contiguous allocations (a piece is specified by its "position" and size), and only externalities that are "anonymous", that is, each player cares about allocations to others only insofar as they affect its own allocation, and is indifferent to the identities of the other players that receive various pieces.

7.2 Model

We introduce a general model for cake cutting with externalities, in which each player i has multiple valuation functions, to reflect the influence of every other player j on player i. We naturally extend the notion of proportionality to the setting with externalities and formalize two notions of envy-freeness, namely swap envy-freeness and swap stability. Under the former notion, a player cannot benefit by swapping its allocation with another player; under the latter notion, no player is better off when any two players swap their allocations.

Formally, the cake is represented by the interval $[0, 1]$; there is also a set $N = \{1, \ldots, n\}$ of players. A piece of cake X is a set of disjoint intervals of $[0, 1]$. In the context of externalities, we will sometimes discuss the existence of infinite allocations, in which a piece of cake is a countable union of intervals.[1] Each player i has n integrable, non-negative value density functions, such that $v_{i,j}(x)$ defines the value that i receives when x is allocated to player j. The value that player i derives from a piece X that is allocated to player j is $V_{i,j}(X) = \int_X v_{i,j}(x)dx$. This definition assigns zero value to singleton intervals, therefore we allow "disjoint" pieces to intersect at boundaries of

[1] Such allocations can also appear in the classical cake cutting model, for example when dividing a cake among two players to achieve an irrational ratio [**?**].

intervals. In the classical model of cake cutting, $V_{i,j}(X) = 0$ for all pieces X and players $i \neq j$.

An *allocation* $A = (A_1, \ldots, A_n)$ is an assignment of a piece of cake A_i to each player i, such that the pieces are disjoint and $\bigcup_{i \in N} A_i = [0, 1]$. Moreover, each piece A_i is a possibly infinite set of disjoint intervals of $[0, 1]$. The value of player i under allocation A is: $V_i(A) = \sum_{j=1}^{n} V_{i,j}(A_j)$.

Similarly to the classical model, utilities are normalized so that all the players have equal weight. That is, for each player i, $V_i(\tilde{\mathcal{A}}_i) = 1$, where $\tilde{\mathcal{A}}_i$ is the best possible allocation for player i (note that in general this may not be giving the whole cake to i). For our results this assumption is merely for ease of exposition and without loss of generality.

Even before generalizing the classical fairness criteria it is immediately apparent that our model is fundamentally different from the standard model. Indeed, we note that computing the optimal allocation for a single player can require infinitely many cuts, as the following example shows. In contrast, in the standard model, the optimal allocation for any given player requires no cuts and can be obtained by giving the entire cake to that player.

Example 4. *For every player $i \in N$, let: $v_{i,1}(x) = \frac{x}{4}$ and $v_{i,2}(x) = x \sin\left(\frac{1}{x}\right)$, $\forall x \in \left[0, \frac{1}{n}\right]$, $v_{i,2}(x) = \frac{n(1-w)}{n-1}$, $\forall x \in \left(\frac{1}{n}, 1\right]$, where $w = \int_0^{\frac{1}{n}} \max\left(\frac{x}{4}, x \sin\left(\frac{1}{x}\right)\right) dx$. For every player i, the optimal allocation requires giving alternating pieces of cake in the interval $\left[0, \frac{1}{n}\right]$ to players 1 and 2, respectively. However $v_{1,1}(x)$ and $v_{1,2}(x)$ intersect infinitely many times on this interval, and so the optimal allocation for player i requires infinitely many cuts.*

7.3 Fairness Criteria

As noted above, the two most commonly used fairness criteria are proportionality and envy-freeness. Proportionality has a very natural interpretation in our model.

Definition 18 (**Proportionality**)**.** *An allocation A is* proportional *if for every player $i \in N$, $V_i(A) \geq \frac{1}{n}$.*

In words, each player must receive at least $1/n$ of the value it receives under the optimal allocation from its point of view. Note that this definition directly generalized the classical definition: when there are no externalities, each player simply receives a piece of cake that it values at $1/n$ of the whole cake.

In contrast, the notion of envy-freeness lends itself to several possible interpretations.

Definition 19 (**Swap Envy-Freeness, see also [124]**)**.** *An allocation* $A = (A_1, \ldots, A_n)$ *is* swap envy-free *if for any two players* $i, j \in N$, $V_{i,i}(A_i) + V_{i,j}(A_j) \geq V_{i,i}(A_j) + V_{i,j}(A_i)$.

That is, a player cannot improve by swapping its allocation with that of another player. This definition generalizes and implies the classical definition of envy-freeness when there are no externalities. We also define an even stronger version of swap envy-freeness, in which a player cannot benefit from a swap between *any* pair of players.

Definition 20 (**Swap Stability**)**.** *An allocation* $A = (A_1, \ldots, A_n)$ *is* swap stable *if for every three players* $i, j, k \in N$, $V_{i,j}(A_j) + V_{i,k}(A_k) \geq V_{i,j}(A_k) + V_{i,k}(A_j)$.

Note that swap stable allocations are always swap envy-free, but the converse may not be true.

7.4 Relationship Between Fairness Properties

In the classical cake cutting model, proportionality coincides with envy-freeness when $n = 2$, and envy-freeness is strictly stronger than proportionality when $n > 2$. Of course, implications that do not hold in the classical model are also false in our more general model (as our notions of fairness reduce to the classical notions). However, it may be the case that some classical implications are no longer true.

Focusing first on the case of two players, we immediately see that proportionality and swap envy-freeness are no longer equivalent. Indeed, the following example constructs an allocation that is proportional but not swap envy-free (and, therefore, not swap stable).

Example 5. *Consider the value density functions:* $v_{1,1}(x) = v_{2,2}(x) = v_{2,1}(x) = 1$, $\forall x \in [0, 1]$; $v_{1,2}(x) = \frac{1}{3}$, $\forall x \in \left[0, \frac{1}{2}\right]$, *and* $v_{1,2}(x) = \frac{1}{4}$, $\forall x \in \left[\frac{1}{2}, 1\right]$. *The allocation* $A = (A_1, A_2)$, *where* $A_1 = \left[0, \frac{1}{2}\right]$ *and* $A_2 = \left[\frac{1}{2}, 1\right]$ *is proportional, but not swap envy-free, since player* 1 *would improve by swapping its piece with that of player* 2.

In addition, swap envy-freeness does not imply proportionality when $n > 2$, as the next example shows.

Example 6. *Let* $N = \{1, 2, 3\}$ *and define the intervals* $I_1 = \left[0, \frac{1}{3}\right]$, $I_2 = \left[\frac{1}{3}, \frac{2}{3}\right]$, *and* $I_3 = \left[\frac{2}{3}, 1\right]$. *Let* $v_{1,2}(x) = \frac{3}{2}$, $\forall x \in I_3$; $v_{1,3}(x) = \frac{3}{2}$, $\forall x \in I_2$; $v_{2,2}(x) = 3$, $\forall x \in I_2$; *and* $v_{3,3}(x) = 3$, $\forall x \in I_3$. *All the other densities are set to zero. Then the allocation* $A = (I_1, I_2, I_3)$, *where player* i *receives the interval* I_i, *has utilities:* $V_1(A) = V_{1,1}(I_1) = 0$, *while* $V_2(A) = V_{2,2}(I_2) = 1$ *and*

$V_3(A) = V_{3,3}(I_3) = 1$. *The allocation is swap envy-free, but not proportional, as player* 1 *only receives a value of zero.*

So far we have not determined whether swap envy-freeness implies proportionality in the case of two players. Our main positive result in this section establishes a much stronger statement: swap stability implies proportionality for any number of players whenever the entire cake is allocated (this assumption is also required for the classical implication). In particular, for only two players (where our two notions of envy-freeness coincide), swap envy-freeness does imply proportionality.

Theorem 19. *Every swap stable allocation that contains the entire cake is proportional.*

Proof. Let $A = (A_1, \ldots, A_n)$ be any swap stable allocation that contains the entire cake. By definition of swap stability, we have that for all $i, j, k \in N$:

$$V_{i,j}(A_j) + V_{i,k}(A_k) \geq V_{i,j}(A_k) + V_{i,k}(A_j)$$

By summing over all $j \in N$, we obtain:

$$\sum_{j=1}^{n} V_{i,j}(A_j) + \sum_{j=1}^{n} V_{i,k}(A_k) \geq \sum_{j=1}^{n} V_{i,j}(A_k) + \sum_{j=1}^{n} V_{i,k}(A_j)$$

Since $V_i(A) = \sum_{j=1}^{n} V_{i,j}(A_j)$, we have:

$$V_i(A) + nV_{i,k}(A_k) \geq \sum_{j=1}^{n} V_{i,j}(A_k) + V_{i,k}([0,1]) \tag{7.1}$$

By summing Inequality (7.1) over all $k \in N$, we get:

$$\begin{aligned} & \sum_{k=1}^{n} V_i(A) + n \sum_{k=1}^{n} V_{i,k}(A_k) \\ \geq \quad & \sum_{k=1}^{n} \sum_{j=1}^{n} V_{i,j}(A_k) + \sum_{k=1}^{n} V_{i,k}([0,1]) \\ = \quad & \sum_{j=1}^{n} \sum_{k=1}^{n} V_{i,j}(A_k) + \sum_{k=1}^{n} V_{i,k}([0,1]) \\ = \quad & \sum_{j=1}^{n} V_{i,j}([0,1]) + \sum_{k=1}^{n} V_{i,k}([0,1]) \end{aligned}$$

Equivalently,

$$\begin{aligned} 2nV_i(A) \quad &= \quad nV_i(A) + nV_i(A) \\ &\geq \quad \sum_{j=1}^{n} V_{i,j}([0,1]) + \sum_{k=1}^{n} V_{i,k}([0,1]) \geq 1 + 1 \end{aligned}$$

Thus $V_i(A) \geq \frac{1}{n}$, and so A is proportional. □

As noted above, swap stability also implies swap envy-freeness by definition. In contrast, the next example shows that proportionality and swap envy-freeness, even combined, do not imply swap stability, that is, there exist proportional and swap envy-free allocations that are not swap stable.

Example 7. *Consider the value density functions:* $v_{2,2}(x) = v_{3,3}(x) = 1$, $\forall x \in [0,1]$*;* $v_{1,1}(x) = 1$, $\forall x \in \left[0, \frac{1}{3}\right]$*;* $v_{1,3}(x) = 1$, $\forall x \in \left(\frac{1}{3}, \frac{2}{3}\right)$*; and* $v_{1,2}(x) = 1$, $\forall x \in \left[\frac{2}{3}, 1\right]$*; all remaining densities are zero. Let* $A = (A_1, A_2, A_3)$*, where* $A_1 = \left[0, \frac{1}{3}\right]$, $A_2 = \left[\frac{1}{3}, \frac{2}{3}\right]$*, and* $A_3 = \left[\frac{2}{3}, 1\right]$*. Each player receives a value of at least* $\frac{1}{3}$ *under* A*, and the allocation is also swap envy-free. However,* A *is not swap stable, since player* 1 *would prefer that players* 2 *and* 3 *swap pieces, which would bring player* 1*'s utility to* 1 *(compared to* $\frac{1}{3}$ *under* A*).*

7.5 Existence of Fair Allocations

In the classical model, the case of two players trivially admits an envy-free (and therefore proportional) allocation: simply divide the cake into two pieces that player 1 values equally, and let player 2 choose its favorite piece. It turns out that the analogous result also holds in the presence of externalities.[2]

Theorem 20. *Let* $n = 2$*. Then there exists a proportional and swap envy-free allocation that requires a single cut.*

Proof. Define $\mathcal{D} : [0,1] \to \mathbb{R}$ such that for all $x \in [0,1]$:

$$\mathcal{D}(x) = V_{1,1}([0,x]) + V_{1,2}([x,1]) - V_{1,1}([x,1]) - V_{1,2}([0,x]).$$

Note that:

$$\mathcal{D}(0) = V_{1,2}([0,1]) - V_{1,1}([0,1])$$

and

$$\mathcal{D}(1) = V_{1,1}([0,1]) - V_{1,2}([0,1]).$$

It holds that $\mathcal{D}(0) + \mathcal{D}(1) = 0$, and since $\mathcal{D}$ is continuous it follows from the intermediate value theorem that there exists $\tilde{x} \in [0,1]$ such that $\mathcal{D}(\tilde{x}) = 0$. We claim that the allocation in which player 2 takes its favorite piece among $\{[0,\tilde{x}], [\tilde{x},1]\}$ — giving the other piece to player 1 – is proportional and swap envy-free.

Without loss of generality, assume player 2 chooses the piece $[\tilde{x}, 1]$. Then the resulting allocation is $A = (A_1, A_2)$, where $A_1 = [0, \tilde{x}]$ and $A_2 = [\tilde{x}, 1]$. By optimality of player 2's choice, we have:

$$V_{2,2}([\tilde{x},1]) + V_{2,1}([0,\tilde{x}]) \geq V_{2,2}([0,\tilde{x}]) + V_{2,1}([\tilde{x},1])$$

[2]The proof is excluded due to space constraints and can be found in the full version of the paper, available on: http://www.cs.cmu.edu/ arielpro/papers.html.

Thus player 2 is not swap-envious. Assume for contradiction that player 2 obtains less than $\frac{1}{2}$. Then we have

$$\begin{aligned}\frac{1}{2} &> V_2(A) = V_{2,2}([\tilde{x},1]) + V_{2,1}([0,\tilde{x}]) \\ &\geq V_{2,2}([0,\tilde{x}]) + V_{2,1}([\tilde{x},1])\end{aligned}$$

and so

$$\begin{aligned}1 &> V_{2,2}([\tilde{x},1]) + V_{2,1}([0,\tilde{x}]) + V_{2,2}([0,\tilde{x}]) + V_{2,1}([\tilde{x},1]) \\ &= V_{2,1}([0,1]) + V_{2,2}([0,1]) \geq 1\end{aligned}$$

This is a contradiction, thus $V_2(A) \geq \frac{1}{2}$.

We next show that A also satisfies fairness for player 1. By the choice of $\tilde{x}$, $V_{1,1}([0,\tilde{x}]) + V_{1,2}([\tilde{x},1]) = V_{1,1}([\tilde{x},1]) + V_{1,2}([0,\tilde{x}])$, and so player 1 is not swap-envious. Moreover,

$$\begin{aligned}&2V_1(A) \\ &= V_{1,1}([0,\tilde{x}]) + V_{1,2}([\tilde{x},1]) + V_{1,1}([\tilde{x},1]) + V_{1,2}([0,\tilde{x}]) \\ &= V_{1,1}([0,1]) + V_{1,2}([0,1]) \geq 1,\end{aligned}$$

and so $V_1(A) \geq \frac{1}{2}$. Thus A is proportional, swap envy-free, and requires one cut. □

In the classical cake cutting model envy-free (and hence proportional) allocations that require only $n-1$ cuts are guaranteed to exist [**?**]. Of course, at least that many cuts are required because each player must receive a piece. In stark contrast, in our model there are instances where zero cuts are needed to achieve a swap stable allocation of the whole cake! To see this, simply consider an instance where all players derive value only from allocating the cake to player 1.

On the other hand, a proportional and swap envy-free allocation can require strictly more than $n-1$ cuts. Note that swap stability implies both proportionality and swap envy-freeness, hence this lower bound also holds for swap stability.

Theorem 21. *A proportional and swap envy-free allocation may require strictly more than $n-1$ cuts.*

Proof. Informally, we consider an instance where each player has exactly one "representative" player. The idea is that each player can obtain a value of approximately 1 only by giving the entire cake to their representative. In addition, different players require different regions of the cake. Formally, for each $i \in N$, let r_i be the representative of i, where $r_i = 1$ if i is odd and $r_i = 2$ if i is even. Define the value density functions as follows:

$$v_{i,r_i}(x) = \begin{cases} n(1-\epsilon) & x \in \left[\frac{i-1}{n}, \frac{i}{n}\right] \\ \frac{n\epsilon}{n-1} & x \in \left[0, \frac{i-1}{n}\right) \cup \left(\frac{i}{n}, 1\right] \end{cases}$$

and for all $x \in [0,1]$,

$$v_{i,j}(x) = \begin{cases} \epsilon & j \in N \setminus \{r_1, r_2\} \\ 0 & j \in \{r_1, r_2\} \setminus \{r_i\} \end{cases}$$

Note that $v_{i,r_2} = 0$ for all odd i, and $v_{i,r_1} = 0$ for all even i. That is, a player does not receive utility from both representatives. Any proportional allocation of the cake requires at least $n-1$ cuts, since it would have to give player r_1 a piece in each of the intervals $\left[\frac{i-1}{n}, \frac{i}{n}\right]$, where i is odd, and player r_2 a piece in each of the intervals $\left[\frac{i-1}{n}, \frac{i}{n}\right]$, where i is even. However, an allocation with $n-1$ cuts cannot be swap envy-free in this example, since every player $i \in N \setminus \{r_1, r_2\}$ will want to swap with the other representative. Thus each player $i \in N \setminus \{r_1, r_2\}$, where i is odd, requires a piece of length equal to that of r_2, and each player $i \in N \setminus \{r_1, r_2\}$, where i is even, requires a piece of length equal to that of r_1. We conclude that any swap envy-free and proportional allocation requires at least n cuts. □

In contrast, our main result for this section shows that a swap stable allocation (which is in particular swap envy-free and proportional) necessarily exists under mild assumptions, and also gives an upper bound on the number of required cuts.

Theorem 22. *Assume that the value density functions are continuous. Then a swap stable allocation is guaranteed to exist and requires at most $(n-1)n^2$ cuts.*

Our main tool is the following lemma that is due to Alon [**?**].

Lemma 11 (Alon 1987)**.** *Let $\mu_1, \mu_2, \ldots, \mu_t$ be t continuous probability measures on the unit interval. Then it is possible to cut the interval in $(k-1)\cdot t$ places and partition the $(k-1)\cdot t+1$ resulting intervals into k families $\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_k$ such that $\mu_i\left(\bigcup \mathcal{F}_j\right) = 1/k$, for all $1 \le i \le t$, $1 \le j \le k$. The number $(k-1)\cdot t$ is best possible.*

Proof of Theorem 22. Let

$$\Psi = \{(i,j) \in N \times N \mid v_{i,j} \neq \mathbf{0}\}.$$

Define a normalized instance of each value density function: $v'_{i,j}(x) = \frac{v_{i,j}(x)}{V_{i,j}([0,1])}$, $\forall (i,j) \in \Psi$. Note that the denominator is strictly positive for all $(i,j) \in \Psi$. Then the functions $v'_{i,j}(x)$ are continuous probability measures on the unit interval. By Lemma 11, there exists a partition of the cake into n pieces, $A = (A_1, \ldots, A_n)$, where the number of cuts is bounded by $(|\Psi|-1)n^2 \le (n-1)n^2$, such that $V'_{i,j}(A_k) = 1/n$ for all $(i,j) \in \Psi$ and $k \in N$.

Consider the allocation given by A, where player i receives the piece A_i, $\forall i \in N$. By construction of A, we have that: $V_{i,j}(A_k) = \frac{V_{i,j}([0,1])}{n}$, for all $i, j, k \in N$ (the identity trivially holds for all $(i, j) \notin \Psi$), and so:

$$\begin{aligned} V_{i,j}(A_j) + V_{i,k}(A_k) &= \frac{V_{i,j}([0,1])}{n} + \frac{V_{i,k}([0,1])}{n} \\ &= V_{i,j}(A_k) + V_{i,k}(A_j) \end{aligned}$$

Thus A is swap stable, with at most $(n-1)n^2$ cuts. □

Even more generally, it can be shown that fair allocations are guaranteed to exist when the value density functions are piecewise continuous.

7.6 Complexity Considerations

An important question in cake cutting is how protocols operate and what can be achieved depending on the type of operations allowed. In the presence of externalities, the Robertson-Webb query model (as formalized by Procaccia [109] naturally generalizes to the following types of queries:

1. ***Evaluate***$_{i,j}(x, y)$:
   Player i outputs α such that $V_{i,j}([x, y]) = \alpha$.

2. ***Cut***$_{i,j}(x, \alpha)$:
   Player i outputs y such that $V_{i,j}([x, y]) = \alpha$.

We can show that under this extended form of the Robertson-Webb communication model, it is possible to guarantee a value of $\frac{1}{n^2}$ to all the players. This relies on the observation that for each player i, there exists a "representative" that holds at least $\frac{1}{n}$ of the value for player i. Then by running any of the classical proportional protocols while querying only the representatives, we obtain an allocation that gives at least $\frac{1}{n^2}$ to each player.

Theorem 23. *An allocation in which every player receives utility at least $\frac{1}{n^2}$ can be computed with $O(n^2)$ queries in the extended Robertson-Webb model.*

Proof. For every player $i \in N$, let t_i be the player which brings i the highest value in the optimal allocation for i, $\tilde{\mathcal{A}}_i$. Then i receives a value of at least $\frac{1}{n}$ from t_i in $\tilde{\mathcal{A}}_i$, and thus giving the entire cake to player t_i guarantees i a value of at least $\frac{1}{n}$. We refer to t_i as the representative of player i. Run the Dubins-Spanier procedure, where a player i calls cut whenever the piece to the left of the knife is worth $\frac{1}{n}$ according to the valuation function V_{i,t_i}. Let $A = (A_1, \ldots, A_n)$ be the resulting allocation, where piece A_i is given to player

t_i. Then the utility of player i is:

$$\begin{aligned} u_i(A) &\geq V_{i,t_i}(A_i) \geq \left(\frac{1}{n}\right) V_{i,t_i}([0,1]) \\ &\geq \left(\frac{1}{n}\right)\left(\frac{1}{n}\right) = \frac{1}{n^2} \end{aligned}$$

Thus there exists an allocation with $n-1$ cuts which gives utility at least $\frac{1}{n^2}$ to each player. □

However, one cannot significantly improve this result. Specifically, we show that no proportional protocol can be obtained even for two players under the extended Robertson-Webb communication model. The proof idea is reminiscent of the technique used to show that no finite protocol can compute an exact allocation in the standard cake cutting model [112].

Theorem 24. *There exists no finite protocol that can compute a proportional allocation of the entire cake even for two players in the extended Robertson-Webb model.*

Proof. Consider an instance where the two players have symmetric valuations. That is, $v_{1,1}(x) = v_{2,2}(x)$ and $v_{1,2}(x) = v_{2,1}(x)$, $\forall x \in [0,1]$. Moreover, let $V_{1,1}([0,1]) = \frac{2}{3}$ and $V_{1,2}([0,1]) = \frac{1}{3}$. Note that it is possible to set the value density functions such that each player still obtains a value of 1 in the optimal allocation over $[0,1]$. However, by giving the entire cake only to player 1 or player 2, player 1 obtains $\frac{2}{3}$ or $\frac{1}{3}$, respectively.

We first claim that it is sufficient to restrict attention to cut and evaluate queries to player 1. Indeed, let $A = (A_1, A_2)$ be any proportional allocation that contains the entire cake. Then it must be the case that:

$$V_{1,1}(A_1) + V_{1,2}(A_2) \geq \frac{1}{2}$$

and

$$V_{2,2}(A_2) + V_{2,1}(A_1) \geq \frac{1}{2}.$$

By choice of the valuations, we have:

$$V_{1,1}(A_1) + V_{1,1}(A_2) + V_{1,2}(A_1) + V_{1,2}(A_2) = 1$$

and

$$V_{2,2}(A_1) + V_{2,2}(A_2) + V_{2,1}(A_1) + V_{2,1}(A_2) = 1$$

The inequalities can be rewritten as:

$$\begin{aligned} V_{1,1}(A_1) + V_{1,2}(A_2) &\geq \frac{1}{2}(V_{1,1}(A_1) + V_{1,1}(A_2) \\ &+ V_{1,2}(A_1) + V_{1,2}(A_2)) \end{aligned}$$

and

$$V_{2,2}(A_2) + V_{2,1}(A_1) \geq \frac{1}{2}(V_{2,2}(A_1) + V_{2,2}(A_2) + V_{2,1}(A_1) + V_{2,1}(A_2)) \tag{7.2}$$

Equivalently,

$$V_{1,1}(A_1) + V_{1,2}(A_2) \geq V_{1,1}(A_2) + V_{1,2}(A_1) \tag{7.3}$$

and

$$V_{1,1}(A_2) + V_{1,2}(A_1) \geq V_{1,1}(A_1) + V_{1,2}(A_2) \tag{7.4}$$

where Inequality (7.4) is obtained from (7.2) by symmetry of the valuations. From Inequality (7.3) and (7.4) we get:

$$V_{1,1}(A_1) + V_{1,2}(A_2) = V_{1,1}(A_2) + V_{1,2}(A_1) \tag{7.5}$$

By definition of the valuations, we have: $V_{1,1}(A_1) + V_{1,1}(A_2) = V_1([0,1]) = \frac{2}{3}$ and $V_{1,2}(A_1)+V_{1,2}(A_2) = V_{1,2}([0,1]) = \frac{1}{3}$, thus Equation (7.5) can be rewritten as:

$$V_{1,1}(A_1) - V_{1,2}(A_1) = \left(\frac{2}{3} - V_{1,1}(A_1)\right) - \left(\frac{1}{3} - V_{1,2}(A_1)\right)$$
$$= \frac{1}{3} - V_{1,1}(A_1) + V_{1,2}(A_1)$$

Thus to achieve proportionality it must hold that $V_{1,1}(A_1) - V_{1,2}(A_1) = \frac{1}{6}$. By symmetry, the allocation of player 2 must also verify: $V_{2,2}(A_2) - V_{2,1}(A_2) = \frac{1}{6}$.

We prove the theorem by tracing an infinite path through the algorithm tree and proceed by induction on the number of *Cut* queries. Note that the given instance requires at least two pieces, since giving the entire cake to either player results in a utility of $\frac{1}{3}$ for the other one. Assume that after $k-1$ steps we arrived at a non-terminating vertex, where the pieces $W_1, \ldots, W_k$ have been cut and the values $V_{i,j}(W_l)$ have been provided, $\forall i, j\{1,2\}$, $\forall l \in \{1, \ldots, k\}$. This is all that is known about the value density functions at this stage. Based on this information, the protocol decides which piece is cut next, according to which valuation, and the sizes of the pieces that should be produced. Recall that since the valuations are symmetric, it is sufficient to query player 1.

By the induction hypothesis, a proportional and swap envy-free allocation cannot be obtained with the pieces $W_1, \ldots, W_k$. That is, for any allocation A_1^i of player 1, which can be obtained from the set of already demarcated pieces, we have:

$$V_{1,1}(A_1^i) - V_{1,2}(A_1^i) = \frac{1}{6} + \delta_i, \text{ where } \delta_i \neq 0, \ \forall i.$$

Assume the protocol can query inside some interval W_j such that a proportional allocation is obtained in the next step. We illustrate the case where the query is made with respect to $V_{1,1}$. The other case, when the query is made with respect to $V_{1,2}$, is similar.

Let α denote the value of the query with respect to the left interval of W_j. That is, W_j is divided into two pieces, W_j^1 and W_j^2, such that $V_{1,1}(W_j^1) = \alpha$ and $V_{1,1}(W_j^2) = V_1(W_j) - \alpha$. In order for a proportional allocation to be obtained in the next step, it should be the case that one of the allocations of player 1 from the previous step A_1^i, which does not contain piece W_j, becomes proportional when player 1 obtains the piece W_j^1 and player 2 obtains the piece W_j^2, or vice versa. That is,

$$V_{1,1}(A_1^i \cup W_j^1) - V_{1,2}(A_1^i \cup W_j^1) = \frac{1}{6}$$

or

$$V_{1,1}(A_1^i \cup W_j^2) - V_{1,2}(A_1^i \cup W_j^2) = \frac{1}{6}$$

The identities are equivalent to:

$$\begin{aligned}\Big((V_{1,1}(A_1^i) \quad &- \quad V_{1,2}(A_1^i)\Big) + V_{1,1}(W_j^1) - V_{1,2}(W_j^1) \\ &= \quad \left(\frac{1}{6} + \delta_i\right) + V_{1,1}(W_j^1) - V_{1,2}(W_j^1) = \frac{1}{6}\end{aligned}$$

or

$$\begin{aligned}\Big(V_{1,1}(A_1^i) \quad &- \quad V_{1,2}(A_1^i)\Big) + V_{1,1}(W_j^2) - V_{1,2}(W_j^2) \\ &= \quad \left(\frac{1}{6} + \delta_i\right) + V_{1,1}(W_j^2) - V_{1,2}(W_j^2) = \frac{1}{6}\end{aligned}$$

Recall that $V_{1,1}(W_j^1) = \alpha$, $V_{1,1}(W_j^2) = V_{1,1}(W_j) - \alpha$, $V_{1,2}(W_j^2) = V_{1,2}(W_j) - V_{1,2}(W_j^1)$. Rewriting, we get:

$$V_{1,2}(W_j^1) = \delta_i + \alpha \tag{7.6}$$

or

$$V_{1,2}(W_j^1) = V_{1,2}(W_j) - V_{1,1}(W_j) + \alpha - \delta_i \tag{7.7}$$

However, there exist at most 2^k different values of δ_i (which correspond to different allocations), and so an adversary can report a value of $V_{1,2}(W_j^1)$ such that all the equalities (7.6) and (7.7) fail simultaneously, for every value of i. That is, there exists w, where $0 \leq w \leq V_{1,2}(W_j)$, such that by setting $V_{1,2}(W_j^1) = w$, we have that for each i,

$$V_{1,2}(W_j^1) \neq \delta_i + \alpha$$

and

$$V_{1,2}(W_j^1) \neq V_{1,2}(W_j) - V_{1,1}(W_j) + \alpha - \delta_i.$$

Thus the protocol requires at least one more step before terminating, which shows the existence of an infinite path in the algorithm tree.

Note that at the k-th step, the values of the demarcated pieces sum up to $2/3$ with respect to $V_{1,1}$ and $1/3$ with respect to $V_{1,2}$. Thus at the k-th cut, the adversary must respect the condition that the valuations for the two subsets of W_j sum up to $V_{1,1}(W_j)$ and $V_{1,2}(W_j)$, respectively. This can be done by having interleaved value density functions, such that $v_{1,1}(x) > 0 \Rightarrow v_{1,2}(x) = 0$, and vice versa. We can partition any interval whose values are known into two such disjoint subintervals and set the densities to recover the known values. □

Intuitively, the protocol is severely restricted if valuations can only be accessed one at a time. However, by allowing simultaneous access, it becomes possible to obtain proportional allocations in finite time. The communication model we consider instead is the following:

1. ***Evaluate Optimal**$_i(x, y)$:* Player i outputs a pair $(\alpha, \tilde{\mathcal{A}}_\alpha)$ such that $\tilde{\mathcal{A}}_\alpha$ is an optimal allocation for i on the interval $[x, y]$ and gives the player exactly α: $V_i(\tilde{\mathcal{A}}_\alpha) = \alpha$.

2. ***Cut Optimal**$_i(x, \alpha)$:* Player i outputs y such that i's optimal allocation on $[x, y]$, $\tilde{\mathcal{A}}_\alpha$, gives the player exactly α: $V_i(\tilde{\mathcal{A}}_\alpha) = \alpha$.

The queries reduce to *Cut* and *Evaluate* from Robertson-Webb in the absence of externalities. Note that the optimal allocation may contain an unbounded number of cuts, and so it is not known apriori how much information the player may send. However, this is also true of the classical Robertson-Webb model; there, the players can communicate infinitely long strings in $O(1)$ (for example, if the value returned by an evaluate query is an irrational number).

Theorem 25. *Every proportional protocol from the standard cake cutting model translates to a proportional protocol with externalities when the Cut and Evaluate queries are replaced by Cut Optimal and Evaluate Optimal, respectively.*

7.7 Piecewise Constant Valuations

When the representation of the value density functions is succinct (such as piecewise constant), then swap envy-free, proportional, and welfare-maximizing allocations can be computed efficiently. For the standard cake cutting model, a linear program for computing such allocations was given by Cohler *et al.* [46].

Let the cake be given as a set of intervals $I = (I_1, \ldots, I_m)$, such that $\forall i, j \in N$, the influence of player j on player i in interval I_k is given by a value density function constant on I_k: $v_{i,j}(x) = c_{i,j,k}$, $\forall x \in I_k$.

Proposition 8. *Consider a cake cutting instance with externalities, where the value density functions are piecewise constant. Then optimal allocation for any given player requires at most* $mn-1$ *cuts and can be computed in time* $\Theta(mn)$*, where* m *is the number of intervals in the representation.*

Proof. It can be immediately observed that a player does not receive additional utility by fractional allocations of any piece I_k. That is, the best allocation of any piece I_k for player i is to give the entire piece to the player j which maximizes the value $V_{i,j}(I_k)$. □

Definition 21 (Uniform Allocation)**.** *Given a cake cutting problem with piecewise constant valuations over intervals* $I = (I_1, \ldots, I_m)$*, an allocation is* uniform *if it gives each player a contiguous piece of length* $|I_j|/n$ *of each interval* I_j*.*

Proposition 9. *Consider a cake cutting instance with externalities, where the value density functions are piecewise constant. Then the uniform allocation is proportional and swap envy-free.*

Proof. The uniform allocation is trivially swap envy-free, since all the players have identical pieces. Consider now a player i and some interval I_k. In the allocation which is optimal for player i, interval I_k is given to the player k_i which maximizes $V_{i,k_i}(I_k)$. In the uniform allocation, player k_i receives $1/n$ of interval I_k, and so i gets a value of at least $\left(\frac{1}{n}\right) V_{i,k_i}(I_k)$ from I_k. Let u_i^U denote the utility of i in the uniform allocation. Then:

$$u_i^U = \sum_{k=1}^{m} \sum_{j=1}^{n} \left(\frac{1}{n}\right) V_{i,j}(I_k) \geq \sum_{k=1}^{m} \left(\frac{1}{n}\right) V_{i,k_i}(I_k) = \frac{1}{n}$$

Thus the uniform allocation is proportional. □

Theorem 26. *Consider a cake cutting instance with externalities, where the value density functions are piecewise constant,* m *is the number of intervals in the representation, and* n *is the number of players. Then Algorithm 8 returns an optimal swap envy-free and proportional allocation in time polynomial in* m *and* n*.*

7.8 Discussion and Future Work

This work lays the foundations of externalities in cake cutting. One of the main open questions is the design of a query model and computationally efficient protocols for the computation of swap envy-free and swap stable allocations for any number of players. The existence result of Theorem 22 relies on a non-constructive result (Lemma 11), and so it does not give a bounded algorithm.

Solve the following linear program, where $x_{i,k} \in [0,1]$ is the percentage of piece I_k that gets allocated to player i:

$$\max \sum_{i,j=1}^{n} \sum_{k=1}^{m} x_{j,k} V_{i,j}(I_k) \tag{7.8}$$

$$\text{s.t. } \sum_{i=1}^{n} x_{i,k} = 1, \forall k \in \{1, \ldots, m\} \tag{7.9}$$

$$\sum_{k=1}^{m} \sum_{j=1}^{n} x_{j,k} V_{i,j}(I_k) \geq \frac{1}{n}, \forall i \in N \tag{7.10}$$

$$\sum_{k=1}^{m} x_{i,k} V_i(I_k) + x_{j,k} V_{i,j}(I_k) \geq \sum_{k=1}^{m} x_{j,k} V_i(I_k) + x_{i,k} V_{i,j}(I_k), \forall i,j \in N \tag{7.11}$$

$$x_{i,k} \geq 0, \forall i \in N, \forall k \in \{1, \ldots, m\} \tag{7.12}$$

Algorithm 8: Optimal Swap Envy-Free and Proportional Allocation

In addition, we conjecture that both proportionality and swap envy-freeness can be computed with at most $n-1$ cuts when required separately.

A separate direction for future work is the study of negative externalities. One can certainly imagine relevant settings where externalities are negative; for example, when allocating time slots for advertising, it hurts Coca Cola if Pepsi is allocated the best slots. Negative externalities invalidate some of our positive results, and present a nice challenge for future work.

Chapter 8

Notes on Envy-Free Cake Cutting

8.1 Approximate Envy-Freeness

In this section we design an ϵ-envy-free protocol that runs in $O(n^2/\epsilon)$ and is simpler conceptually than the previously known ones. These include the ϵ-envy-free protocol designed by Lipton *et al.* [89] design in the context of allocating indivisible resources and the ϵ-envy-free algorithms derived from the computation of approximately fair partitions (Robertson and Webb [112]) in the context of approximating exact allocations (in ratios such as 1:1).

Theorem 27. *For each $\epsilon > 0$, an ϵ-envy-free allocation that contains the entire cake can be computed in $O\left(\frac{n^2}{\epsilon}\right)$.*

Proof. The allocation produced by Algorithm 9 is ϵ-envy-free and contains the entire cake. At a high level, the players cut the cake into $O(n^2/\epsilon)$ pieces (intervals between adjacent cuts), and the pieces are allocated in a round robin fashion. Taking the point of view of player i, this player can make the initial cuts so that it values each interval between its own adjacent cuts at most at ϵ; it follows that it values any of the $O(n^2/\epsilon)$ pieces induced by everyone's cuts at most at ϵ. Partition the choices into phases, where in each phase, i chooses first, followed by players $i+1, \ldots, n, 1, \ldots, i-1$. In each phase, i prefers its own piece to the piece selected by any other player. Player i may envy the choices made by players $1, \ldots, i-1$ before the beginning of the first phase, but its value for each of these pieces is at most ϵ. □

8.2 Polynomial Valuations

Returning to the question of exact envy-free cake cutting, we show that there exists a protocol in the Robertson-Webb model that is guaranteed to output an

for each $i \in N$ **do**
    Player i makes $\lceil n/\epsilon \rceil$ cuts in $[0, 1]$
end for
$i = 1$
while there exist available pieces between two adjacent cuts **do**
    Player i takes his (remaining) favorite piece
    $i = (i \bmod n) + 1$
end while

Algorithm 9: An ϵ-envy-free protocol for n players

envy-free allocation for the family of polynomial valuations and is much simpler conceptually than the general protocol; moreover, the number of queries required by the protocol is bounded by the sum of the degrees of the polynomials. Another recent result in this area was given by Kurokawa *et al.* [84], who designed an exact envy-free protocol for *piecewise linear* value density functions. Their protocol is guaranteed to produce an envy-free allocation within $O(n^6 k \log k)$ queries on any given instance, where n is thenumber of players and kis the number of break points in the valuations (break points are discontinuities in the derivatives of the valuation functions). Kurokawa *et al.* also showed that if a protocol can compute envy-free allocations for the class of piecewise uniform valuations, then it can also solve the envy-free cake cutting problem for general valuations. The result of Kurokawa *et al.* suggests that the main difficulty is detecting the break points in the value density function (and possibly its derivative). Polynomial valuations are interesting from the point of view of envy-free cake cutting because no such discontinuities exist, yet no bounded protocol is known for this class.

Theorem 28. *There exists a protocol in the Robertson-Webb communication model such that on every n-player cake cutting instance with value density functions given by polynomials, the protocol is guaranteed to terminate with an envy-free allocation using $O(d \cdot n^2)$ queries, where d is the maximum degree of any polynomial in the representation.*

Proof. Consider Algorithm 5.4 and assume the value density functions of the players can be expressed as polynomials. That is, $v_i(x) = \sum_{j=0}^{d_i} a_{i,j} x^j$, for some $d_i \in \mathbb{N}$ and $a_{i,j} \in \mathbb{R}$, for all $j \in \{0, \ldots, d_i\}$ and $i \in N$. Recall that value density functions are always non-negative and normalized to give equal weight to all the players. That is, $v_i(x) \geq 0$, for all $x \in [0, 1]$ and $\int_0^1 v_i(x)dx = 1$, for each player $i \in N$.

Define polynomials P_i for each player i as follows:

$$\begin{aligned} P_i(x) &= V_i([0,x]) = \int_0^x v_i(y)dy = \int_0^x \left(\sum_{j=0}^{d_i} a_{i,j} y^j \right) dy = \sum_{j=0}^{d_i} a_{i,j} \int_0^x y^j dy \\ &= \sum_{j=0}^{d_i} \left(\frac{a_{i,j}}{j+1} \right) x^{j+1} \end{aligned}$$

Then the polynomial P_i has the property that $P_i(0) = 0$, $P_i(1) = 1$, and $(P_i(x))' = v_i(x)$, for each player $i \in N$.

Algorithm 1 starts by assuming that the players have valuations given by polynomials of degree zero (i.e. constant) and increases the degrees with every iteration. Consider the iteration in which the correct upper bound has been reached: $d = \max(d_1, \ldots, d_n)$. Then the answers of player i to the evaluate queries on the intervals:

$$\{[0,0], [0, \frac{1}{d+1}], [0, \frac{1}{d}], \ldots, [0,1]\}$$

can be used to obtain $d+2$ values for the unique interpolating polynomial, of maximum degree $d+2$. That is, the protocol has obtained the following values:

$$P_i(0), P_i(\frac{1}{d+1}), P_i(\frac{1}{d}), \ldots, P_i(1).$$

By taking the derivative of the interpolating polynomial (Line 7), the protocol can find the exact value density function of player i. Since d is an upper bound on the degrees of all the players, it follows that all the value density functions have been guessed correctly, and so the allocation X computed in this iteration (Line 9) is guaranteed to be envy-free.

It is immediate that the protocol terminates after at most d iterations, and the number of *Evaluate* queries asked in each iteration is $n^2 + 1$. Thus the total number of queries required to output an envy-free allocation when the maximum degree is d is bounded by $d(n^2 + 1)$. □

An interesting open question is whether there exists a bounded algorithm for envy-free cake cutting with polynomial valuations, where the runtime of the protocol is only a function of the number of players. A negative result for this class would also answer the existence question for general valuations.

```
d ← 0 // running upper bound on the degrees of the polynomials
while (true) do
    for (each i ∈ N) do
        x_{i,d} ← 1/(d+1)
        y_{i,d} ← Evaluate_i([0, x_{i,d}])
        P_i(x) ← Polynomial-Interpolate ({(0,0)} ∪_{j=0}^{d} {(x_{i,j}, y_{i,j})})
        w_i(x) ← (P_i(x))' // Player i's value density function assuming it's a polynomial of max degree d
    end for
    // Compute a contiguous envy-free allocation w.r.t. {w_1, ..., w_n}
    X ← Contiguous-EF-Allocation ({w_1, ..., w_n})
    // Ask the players if X is envy-free
    for (all (i, j) ∈ N^2) do
        W_{i,j} ← Evaluate_i(X_j)
    end for
    if (Envy-Free(W)) then
        return X // Output and exit
    else
        d ← d + 1 // Increase the maximum degree and try again
    end if
end while
```

Algorithm 10: Algorithm for envy-free cake cutting with polynomial valuations

Part II

Multiple Goods

Chapter 9

Background

In this part we move to the fair allocation of multiple goods, both divisible and indivisible. Again there is a set of players $N = \{1, \ldots, n\}$ and a set of goods $M = \{1, \ldots, m\}$, which have to be allocated in a way that is fair; the players typically do not have monetary endowments. An allocation $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n) \in [0,1]^{n \times m}$ is a matrix, such that $\mathbf{x}_i = (x_{i,1}, \ldots, x_{i,m}) \in [0,1]^m$ denotes the bundle received by player i and $x_{i,j}$ is the fraction received by player i from good j.

When the valuations are additive, the problem of fairly allocating multiple divisible goods can be viewed as a subset of cake cutting. However, non-additive families of valuations (such as Leontief utilities) are incomparable with the standard cake cutting model.

Next we briefly overview some of the important recent developments in the allocation of multiple divisible and indivisible goods. The fairness notions often studied for multiple divisible goods include proportionality and envy-freeness (defined as in the cake cutting problem), as well as others such as proportional fairness [47] and the competitive equilibrium from equal incomes [123].

The proportional fairness solution aims to give a good tradeoff between fairness and efficiency and was studied recently by Gkatzelis *et al.* [47], who designed a truthful mechanism that approximates the proportional fairness guarantees within a factor of $1/e$.

Guo and Conitzer [68] study the truthful allocation of multiple divisible goods among two players with additive valuations such that social welfare is optimized, and consider a family of mechanisms called "linear increasing price", in which the players are given equal amounts of artificial currency that they can use to purchase the goods. Han *et al.* [71] study the setting of multiple players and provide several negative results for the social welfare attainable by truthful allocation mechanisms.

Leontief utilities are another important class for multiple divisible goods and have been studied in a body of literature, such as in work by Dolev *et* *al.* [54], who examined the notion of "bottleneck-based fairness", Ghodsi *et al.* [65], who designed the "dominant resource fairness mechanism", with very good theoretical properties. Parkes *et al.* [107] extend the results of Ghodsi *et al.* to capture weighted players and indivisibilities. Gutman and Nisan [69] generalize the fairness notions of Ghodsi *et al* [65] and Dolev *et al.* [54] and design polynomial time algorithms to compute fair allocations for a larger family of utilities.

In the realm of allocating multiple indivisible (or discrete) goods, fairness notions of interest include proportionality, envy-freeness, maximin fairness, and the competitive equilibrium from equal incomes. Recently, Bouveret and Lemaître [18] investigated a scale of criteria for the allocation of indivisible goods, which included all the fairness notions above, as well as a new notion called min-max fair share. Procaccia and Wang [110] studied the maximin fairness solution and showed that while such allocations do not always exist, there exists an algorithm that guarantees each player 2/3 of their maximin value and runs in polynomial time when the number of players is constant. Lipton *et al.* [89] design an algorithm that computes approximately envy-free allocations.

An interesting version of proportionality for indivisible goods was shown to always exist by Hill [75] for additive valuations. Given any such fair division problem with n players, there exists an allocation that guarantees to each player i a value $F_n(\alpha)$, where α is the maximum value of a player i for any good j. Surprisingly, Markakis and Psomas [92] showed that there exists a polynomial time algorithm that computes allocations guaranteeing this type of value to every player (but in fact, the guarantee is even stronger: $F_n(\alpha_i)$, such that the minimum value of player i only depends on its maximum valuation for any good, namely α_i). In follow-up work, Gourvès *et al.* [67] extend the algorithm of Markakis and Psomas for the generalized problem on matroids.

Approximation algorithms for the maximin fairness solution (also known as the Santa Claus problem) were studied by Bezáková and Dani [16], Bansal and Sviridenko [14], Asadpour and Saberi [5]. Feige [62] studies the restricted assignment version of the maximin fair allocation problem, showing that the approximation algorithm of Bansal and Sviridenko approximates the optimal solution within a constant factor.

Recent work investigated the competitive equilibrium from equal incomes for the allocation of discrete goods (Othman *et al.* [106], Budish [38]); this background is discussed in more detail in Chapter 11.

Chapter 10

The Adjusted Winner Procedure

The *Adjusted Winner* procedure was introduced by Brams and Taylor ([22]) as a highly desirable mechanism for allocating *multiple divisible resources* among two parties. The procedure requires the participants to declare their preferences over the items and the outcome satisfies strong fairness and efficiency properties. Adjusted Winner has been advocated as a fair division rule for divorce settlements [22], international border conflicts [120], political issues [52, 94], real estate disputes [88], water disputes [90], deciding debate formats [87] and various negotiation settings [23, 111]. For example, it has been shown that the agreement reached during Jimmy Carter's presidency between Israel and Egypt is very close to what Adjusted Winner would have predicted [24]. Adjusted Winner has been patented by New York University and licensed to the law firm Fair Outcomes, Inc [78].

Although the merits of Adjusted Winner have been discussed in a large body of literature, the procedure is still not fully understood theoretically. We provide two novel characterizations, together with an alternative interpretation that turns out to be very useful for analyzing the procedure.

In addition, as observed already in [21], the procedure is susceptible to manipulation. However, fairness and efficiency are only guaranteed when the participants declare their preferences honestly. In a review of a well-known book on Adjusted Winner by Brams and Taylor [23], Nalebuff [103] highlights the need for research in this direction:

> *..thus we have to hypothesize how they (the players) would have played the game and where they would have ended up.*

In this chapter, we answer these questions by studying the existence, structure, and properties of pure Nash equilibria of the procedure. Until now, our understanding of the strategic aspects has been limited to the case of two items [21] and experimental predictions [49]; our work identifies conditions

Continuous Procedure	*Lexicographic tie-breaking*	*Informed tie-breaking*
pure Nash	✗	✓
ϵ-Nash	✓	✓
Discrete Procedure	*Lexicographic tie-breaking*	*Informed tie-breaking*
pure Nash	✗	✓
ϵ-Nash	✓(*)	✓

Table 10.1: Existence of pure Nash equilibria in Adjusted Winner. The (*) result holds when the number of points is chosen appropriately.

under which Nash equilibria exist and provides theoretical guarantees for the performance of the procedure in equilibrium.

10.1 Contributions

We start by presenting the first characterizations of Adjusted Winner. We show that among all protocols that split at most one item, it is the only one that satisfies Pareto-efficiency and equitability. Under the same condition, we further show that it is equivalent to the protocol that always outputs a maxmin allocation.

Next, we obtain a complete picture for the existence of pure Nash equilibria in Adjusted Winner. We find the following: neither the discrete nor the continuous variants of the procedure are guaranteed to have pure Nash equilibria, but they do have ϵ-Nash equilibria, for every $\epsilon > 0$. Additionally, under *informed* tie-breaking, pure Nash equilibria always exist for both variants of the procedure.

Finally, we prove that the pure Nash equilibria of Adjusted Winner are envy-free and Pareto optimal with respect to the true valuations and that their social welfare is at least 3/4 of the optimal. Our results concerning the existence or non-existence of pure Nash equilibria are summarized in Table 10.1.

10.2 Background

We begin by introducing the classical fair division model for which the Adjusted Winner procedure was developed [21]. Let there be two players, Alice and Bob, that are trying to split a set $M = \{1, \ldots, m\}$ of divisible items. The players have preferences over the items given by numerical values that express their level of satisfaction. Formally, let $\mathbf{a} = (a_1, a_2, \ldots, a_m)$ and

$\mathbf{b} = (b_1, \ldots, b_m)$ denote their *valuation vectors*, where a_j and b_j are the values assigned by Alice and Bob to item j, respectively.

An *allocation* $W = (W_A, W_B)$ is an assignment of fractions of items (or *bundles*) to the players, where $W_A = (w_A^1, \ldots, w_A^m) \in [0,1]^m$ and $W_B = (w_B^1, \ldots, w_B^m) \in [0,1]^m$ are the allocations of Alice and Bob, respectively.

The players have additive *utility* over the items. Alice's utility for a bundle W_A, given that her valuation is $\mathbf{a}$, is: $u_{\mathbf{a}}(W_A) = \sum_{j \in M} a_j \cdot w_A^j$. Bob's utility is defined similarly. The players are weighted equally, such that their utility for receiving all the resources is the same: $\sum_{i \in M} a_i = \sum_{i \in M} b_i$.

There are two main settings studied in this context: *discrete* and *continuous* valuations. In the discrete setting, valuations are positive natural numbers that add up to some integer P and can be interpreted as *points* (or coins of equal size) that the players use to acquire the items. For ease of notation, we will consider the equivalent interpretation of valuations as rationals with common denominator P, where the valuations sum to 1. In the *continuous* setting, the valuations are positive real numbers, which are without loss of generality normalized to sum to 1. These normalizations make procedures invariant to any rescaling of the bids [27, 78].

10.3 The Adjusted Winner Procedure

The Adjusted Winner procedure works as follows. Alice and Bob are asked by a mediator to state their valuations $\mathbf{a}$ and $\mathbf{b}$, after which the next two phases are executed.

> ***Phase 1:*** For every item i, if $a_i > b_i$ then give the item to Alice; otherwise give it to Bob. The resulting allocation is (W_A, W_B) and without loss of generality, $u_{\mathbf{a}}(W_A) \geq u_{\mathbf{b}}(W_B)$.
>
> ***Phase 2:*** Order the items won by Alice increasingly by the *ratio* a_i/b_i: $\frac{a_{k_1}}{b_{k_1}} \leq \ldots \leq \frac{a_{k_r}}{b_{k_r}}$. From left to right, continuously transfer fractions of items from Alice to Bob, until an allocation (W'_A, W'_B) where both players have the same utility is produced: $u_{\mathbf{a}}(W'_A) = u_{\mathbf{b}}(W'_B)$.

Let $AW(\mathbf{a}, \mathbf{b})$ denote the allocation produced by Adjusted Winner on inputs $(\mathbf{a}, \mathbf{b})$, where $AW_A(\mathbf{a}, \mathbf{b})$ and $AW_B(\mathbf{a}, \mathbf{b})$ are the bundles received by Alice and Bob. Note that the procedure is defined for strictly positive valuations, so the ratios are finite and strictly positive numbers. Examples can be found on the Adjusted Winner website[1] as well as in [21].

Adjusted Winner produces allocations that are *envy-free*, *equitable*, *Pareto optimal*, and *minimally fractional*. An allocation W is said to be Pareto optimal if there is no other allocation that strictly improves one player's utility

[1] http://www.nyu.edu/projects/adjustedwinner/.

without degrading the other player. Allocation W is equitable if the utilities of the players are equal: $u_{\mathbf{a}}(W_A) = u_{\mathbf{b}}(W_B)$, envy-free if no player would prefer the other player's bundle, and minimally fractional if at most one item is split.

Envy-freeness of the procedure implies *proportionality*, where an allocation is proportional if each player receives a bundle worth at least half of its utility for all the items. A procedure is called envy-free if it always outputs an envy-free allocation (similarly for the other properties).

10.4 Characterizations

In this section, we provide two characterizations of Adjusted Winner[2] for both the discrete and continuous variants. We begin with a different interpretation of the procedure that is useful for analyzing its properties.

An allocation is *ordered* if it can be produced by sorting the items in decreasing order of the valuation ratios a_i/b_i and placing a boundary line somewhere (possibly splitting an item), such that Alice gets the entire bundle to the left of the line and Bob gets the remainder:

$$\underbrace{\frac{a_{k_1}}{b_{k_2}} \geq \frac{a_{k_2}}{b_{k_2}} \geq \cdots \geq \frac{a_{k_i}}{b_{k_i}} \geq}_{\text{Alice's allocation}} \Bigg| \geq \underbrace{\frac{a_{k_{i+1}}}{b_{k_{i+1}}} \geq \cdots \geq \frac{a_{k_m}}{b_{k_m}}}_{\text{Bob's allocation}}$$

The placement of the boundary line could lead either to an integral or a minimally fractional allocation. Note that the allocation that gives all the items to Alice is also ordered (but admittedly unfair).

It is clear to see that Adjusted Winner produces an ordered allocation (using some tie-breaking rule for items with equal ratios) with the property that the boundary line is appropriately placed to guarantee equitability. This is the way we will be interpreting the procedure for the remainder of the paper. We start by characterizing Pareto optimal allocations.

Lemma 12. *For any valuations* $(\mathbf{a}, \mathbf{b})$ *and any tie-breaking rule, an allocation* W *is* not *Pareto optimal if and only if there exist items* i *and* j *such that Alice gets a non-zero fraction (possibly whole) of* j*, Bob gets a non-zero fraction (possibly whole) of* i*, and* $a_i b_j > a_j b_i$*.*

Proof. ($\Longleftarrow$) If such items i, j exist, then consider the exchange in which Bob gives $\lambda_i > 0$ of item i to Alice and Alice gives $\lambda_j > 0$ of item j to Bob, where:

$$\frac{b_i}{b_j}\lambda_i < \lambda_j < \frac{a_i}{a_j}\lambda_i$$

[2] The results here refer to the case when the players report their true valuations to the mediator. We discuss the strategic aspects of the procedure in Section 10.5.

Since $a_i/a_j > b_i/b_j$, such λ_i and λ_j do exist. Then Alice's net change in utility is:

$$a_i\lambda_i - a_j\lambda_j > a_i\lambda_i - a_j\frac{a_i}{a_j}\lambda_i = 0,$$

while Bob's net change is:

$$b_j\lambda_j - b_i\lambda_i > b_j\lambda_j - b_i(\lambda_j\frac{b_j}{b_i}) > b_j\lambda_j - b_j\lambda_j = 0.$$

Thus the allocation is not Pareto optimal.

($\Longrightarrow$) If the allocation W is not Pareto optimal, then Alice and Bob can exchange positive fractions of items to get a Pareto improvement.

Consider such an exchange and let S_A be the set of items for which positive fractions are given by Alice to Bob. Let S_B be defined similarly for Bob. Without loss of generality, S_A and S_B are disjoint; otherwise we could just consider the net transfer of any items that are in both S_A and S_B. Let $j \in S_A$ be the item with the lowest ratio a_j/b_j, and $i \in S_B$ with the highest ratio a_i/b_i.

If $a_ib_j > a_jb_i$ then we are done. Otherwise, assume by contradiction that for each item $k \in S_A$ and $l \in S_B$ it holds $a_kb_l \geq a_lb_k$. Then $a_k/b_k \geq a_l/b_l$; but then any Pareto improving exchange involving the transfer of items from S_A and S_B is only possible if at least one player gets a larger fraction of items without the other player getting a smaller fraction, which is impossible. □

By Lemma 12, a Pareto optimal allocation can be obtained by sorting the items by the ratios of the valuations and drawing a boundary line somewhere. No matter where the boundary line is, the allocation is Pareto optimal (even if not equitable); thus an allocation is Pareto optimal and splits at most one item if and only if it is ordered. From this we obtain our first characterization.

Theorem 29. *Adjusted Winner is the only Pareto optimal, equitable, and minimally fractional procedure. Any ordered equitable allocation can be produced by Adjusted Winner under some tie-breaking rule.*

Note that both Pareto optimality and equitability are necessary for the characterization. By restricting to Pareto optimal allocations only, then even the allocation that gives all the items to one player is Pareto optimal, while by restricting to equitable allocations only, even an allocation that throws away all the items is equitable. Similarly when the players have identical utilities for some items, then there exist Pareto optimal and equitable allocations that split more than one item. For example, if the two players have identical utilities over all items, then the allocation that gives half of each item to each player is equitable and Pareto optimal. However, in the case that the valuation are such that $a_i/b_i \neq a_j/b_j$ for all items $i \neq j$, then Adjusted Winner is exactly characterized by Pareto optimality and equitability.

Theorem 30. *If the valuations satisfy $a_i/b_i \neq a_j/b_j$ for all items $i \neq j$, then the only Pareto optimal and equitable allocation is the result of Adjusted Winner.*

An allocation is *maxmin* if it maximizes the minimum utility over both players. From Lemma 3.3 [48], an allocation is maxmin if and only if it is Pareto optimal and equitable. Together with Theorem 29, this leads to another characterization.

Theorem 31. *Adjusted Winner is equivalent to the procedure that always outputs a maxmin and minimally fractional allocation.*

10.5 Equilibrium Existence

In this section, we study Adjusted Winner when the players are *strategic*, that is, their reported valuations are not necessarily the same as their actual valuations. Let $\mathbf{x} = (x_1, x_2, \ldots, x_m)$ and $\mathbf{y} = (y_1, y_2, \ldots, x_m)$ be the *strategies* (i.e. declared valuations) of Alice and Bob respectively. Call $(\mathbf{x}, \mathbf{y})$ a *strategy profile*. We will refer to $\mathbf{a}$ and $\mathbf{b}$ as the *true values* of Alice and Bob. Note that since strategies are reported valuations they are positive numbers that sum to 1.

Since the input to Adjusted Winner is now a strategy profile $(\mathbf{x}, \mathbf{y})$ instead of $(\mathbf{a}, \mathbf{b})$, this means that the properties of the procedure are only guaranteed to hold with respect to the *declared* valuations, and not necessarily the true ones[3].

A strategy profile $(\mathbf{x}, \mathbf{y})$ is an ϵ-*Nash equilibrium* if no player can increase its utility by more than ϵ by deviating to a different (pure) strategy. For $\epsilon = 0$, we obtain a *pure Nash equilibrium*.

The main result of this section is that ϵ-Nash equilibria always exist. Furthermore, using an appropriate rule for settling ties between items with equal ratios x_i/y_i, the procedure also has exact pure Nash equilibria. We start our investigations from simple tie-breaking rules.

The main result of this section is that Adjusted Winner is only guaranteed to have ϵ-Nash equilibria when $\epsilon > 0$ using standard tie-breaking. For the discrete case, this is achieved by the center setting the number of points or equivalently the denominator large enough. Furthermore, we prove that when using an appropriate rule for settling ties between items with equal ratios x_i/y_i, the procedure does admit pure Nash equilibria. We start our investigations from the standard tie-breaking rules.

[3] We will show that in the equilibrium, the procedure guarantees some of the properties with respect to the true values as well.

10.6 Lexicographic Tie-Breaking

The classical formulation of Adjusted Winner resolves ties in an arbitrary deterministic way, for example by ordering the items lexicographically, such that items with lower indices come first.

10.7 Continuous Strategies

First, we consider the case of continuous strategies. We start with the following theorem.

Theorem 32. *Adjusted Winner with continuous strategies is not guaranteed to have pure Nash equilibria.*

Proof. Take an instance with two items and valuations $(\mathbf{a}, \mathbf{b})$, where $b_1 > a_1 > a_2 > b_2 > 0$. Assume by contradiction there is a pure Nash equilibrium at strategies $(\mathbf{x}, \mathbf{y})$, where $\mathbf{x} = (x, 1-x)$ and $\mathbf{y} = (y, 1-y)$. We study a few cases and show the players can always improve.

Case 1: $(x \neq y)$. Without loss of generality $x > y$ (the case $x < y$ is similar). Then there exists $\delta \in \mathbb{R}$ with $x - \delta > y \Rightarrow 1 - x + \delta < 1 - y$, and Alice can improve by playing $\mathbf{x}' = (x - \delta, 1 - x + \delta)$, as the boundary line moves to the left of its former position.

Case 2: $(x = y < 1/2)$. Here both players report higher values on the item they like less; Alice's allocation is $(1, \lambda)$ while Bob's is $(0, 1-\lambda)$, for some $\lambda \in (0, 1)$. Then $\exists\, \delta \in \mathbb{R}$ with $x + \delta < 1/2$. By playing $\mathbf{y}' = (x+\delta, 1-x-\delta)$, Bob gets $(1, 1-\lambda')$, for some $\lambda' \in (0,1)$. This is a strict improvement since $a_1 > a_2$.

Case 3: $(x = y > 1/2)$. Both players report higher values on the item they like more. Bob gets $(1 - \frac{1}{2x}, 1)$ and Alice gets $(\frac{1}{2x}, 0)$, with utilities:

$$u_{\mathbf{a}}(AW(\mathbf{x}, \mathbf{y})) = \frac{a_1}{2x}$$

and

$$u_{\mathbf{b}}(AW(\mathbf{x}, \mathbf{y})) = \left(1 - \frac{1}{2x}\right) b_1 + b_2.$$

Let $\delta \in (0, \min(1-x, 2x-1))$ such that:

$$\delta < \max\left\{\frac{4x(x-a_1)}{2x-a_1}, \frac{4x(b_1-x)}{2x-b_1}\right\}.$$

Observe that since $b_1 > a_1$ and $2x - a_1$ and $2x - b_1$ are positive, at least one of $x - a_1$ and $b_1 - x$ is strictly positive and by continuity of the strategy space, such a δ exists. Now consider alternative profiles $(\mathbf{x}', \mathbf{y}) = ((x-\delta, 1-x+\delta), (x, 1-x))$ and $(\mathbf{x}, \mathbf{y}') = ((x, 1-x), (x+\delta, 1-x-\delta))$. Since $\delta < 2x-1$, the first item is still the item that gets split in the new profile. Using the identities $a_1 + a_2 = b_1 + b_2 = 1$ and the assumption that $(\mathbf{x}, \mathbf{y})$ is a pure Nash equilibrium, we have that

$$\begin{cases} a_1\left(1 - \frac{1}{2x} - \frac{1}{2x-\delta}\right) + a_2 \leq 0 \implies \delta \geq \frac{4x(x-a_1)}{2x-a_1} \\ b_1\left(1 - \frac{1}{2x} - \frac{1}{2x+\delta}\right) + b_2 \geq 0 \implies \delta \geq \frac{4x(b_1-x)}{2x-b_1} \end{cases}$$

and we obtain a contradiction.

Case 4: $(x = y = 1/2)$. Alice and Bob get allocations $(1, 0)$ and $(0, 1)$, respectively. Let $0 < \delta < \frac{(b_1-b_2)}{b_2}$ and consider the strategy $\mathbf{y}' = (x+\delta, 1-x-\delta)$ of Bob. Using $\mathbf{y}'$, Bob gets the allocation $(\frac{1}{\delta+1}, 0)$, which is better than $(0, 1)$. Since $b_1 > b_2$, such δ exists.

As none of the cases $1 - 4$ are stable, the procedure has no pure Nash equilibrium. □

However, we show that Adjusted Winner admits approximate Nash equilibria.

Theorem 33. *Each instance of Adjusted Winner has an ϵ-Nash equilibrium, for every $\epsilon > 0$.*

Proof. Let $(\mathbf{a}, \mathbf{b})$ be any instance. We show there exists an ϵ-Nash equilibrium in which Alice plays her true valuations and Bob plays a small perturbation of Alice's valuations. More formally, we show there exist $\epsilon_1, \ldots, \epsilon_m$, such that an ϵ-equilibrium is obtained when Alice plays $\mathbf{a} = (a_1, \ldots, a_m)$ and Bob plays $\tilde{\mathbf{a}} = (\tilde{a}_1, \ldots, \tilde{a}_m)$, where $\tilde{a}_i = a_i + \epsilon_i$ for each item $i \in [m]$ and $\sum_{i=1}^{m} \epsilon_i = 0$. The theorem will follow from the next two lemmas. □

Lemma 13. *For any pair of strategies $(\mathbf{a}, \tilde{\mathbf{a}})$, where $|a_i - \tilde{a}_i| < \epsilon/m$ for all $i \in [m]$, Alice's strategy is an ϵ-best response.*

Proof. Since the procedure is envy-free, Alice gets at least half of the total value by being truthful regardless of Bob's strategy, and so $u_{\mathbf{a}}(AW_A(\mathbf{a}, \tilde{\mathbf{a}})) \geq 1/2$. The allocation must also be envy-free according to Bob's declared valuation profile $\tilde{\mathbf{a}}$, and so $u_{\tilde{\mathbf{a}}}(AW_B(\mathbf{a}, \tilde{\mathbf{a}})) \geq 1/2$.

Since strategies $\mathbf{a}$ and $\tilde{\mathbf{a}}$ are ϵ-close, that is $\sum_i |a_i - \tilde{a}_i| < \epsilon$, then their evaluations of the same allocation, namely $AW_B(\mathbf{a}, \tilde{\mathbf{a}})$, are also close:

$$u_{\mathbf{a}}(AW_B(\mathbf{a}, \tilde{\mathbf{a}})) \geq u_{\tilde{\mathbf{a}}}(AW_B(\mathbf{a}, \tilde{\mathbf{a}})) - \epsilon \geq 1/2 - \epsilon$$

It follows that $1/2 \leq u_{\mathbf{a}}(AW_A(\mathbf{a}, \tilde{\mathbf{a}})) \leq 1/2 + \epsilon$. Moreover, Alice cannot use some other strategy $\mathbf{a}'$ to force an allocation that gives her more than $1/2 + \epsilon$; otherwise, Bob's utility as measured by $\tilde{\mathbf{a}}$ under strategy profiles $(\mathbf{a}', \tilde{\mathbf{a}})$ would be strictly less than $1/2 - \epsilon$, contradicting the envy-freeness of the procedure.

Thus when Bob's strategy is ϵ-close to Alice's truthful strategy $\mathbf{a}$, Alice has an ϵ-best response at her truthful strategy $\mathbf{a}$, which completes the proof of the lemma. □

Lemma 14. *When Alice plays* $\mathbf{a}$*, Bob has an* ϵ*-best response that is* ϵ*-close to Alice's strategy.*

Proof. Let $\pi = (\pi_1, \ldots, \pi_m)$ be a fixed permutation of the items. Then there exist uniquely defined index $l \in \{1, \ldots, m\}$ and $\lambda \in [0, 1)$ such that

$$a_{\pi_1} + \ldots a_{\pi_{l-1}} + \lambda a_{\pi_l} = \frac{1}{2} = (1 - \lambda) a_{\pi_l} + a_{\pi_{l+1}} + \ldots + a_{\pi_m} \qquad (10.1)$$

Note that Adjusted Winner uses lexicographic tie breaking to sort the items when there exist equal ratios $x_i/y_i = x_j/y_j$, for some $i \neq j$. Thus the order π may never appear in an outcome of the procedure when the players use the same strategies.

However, we show that Bob can approximate the outcome of Equation (10.1) arbitrarily well. We have two cases:

Case 1: $\lambda \in (0, 1)$. Then there exist $\epsilon_1, \ldots, \epsilon_m$ such that the following conditions hold:

(i) $|\epsilon_j| < \min\left(\frac{\epsilon}{m}, \frac{2\lambda a_{\pi_l}}{m}\right)$, for all $j \in [m]$,

(ii) the items are strictly ordered by π: $\frac{a_{\pi_1}}{a_{\pi_1} + \epsilon_{\pi_1}} > \ldots > \frac{a_{\pi_m}}{a_{\pi_m} + \epsilon_{\pi_m}}$,

(iii) $\sum_{j=1}^{m} \epsilon_j = 0$, and

(iv) it's still item π_l that gets split, in a fraction $\delta \in (0, 1)$ close to λ; that is, $|\lambda - \delta| < \frac{\epsilon}{b_{\pi_l}}$.

Informally, Bob plays a perturbation of Alice's truthful strategy inducing ordering π on the items (with no ties) and splits item π_l in a fraction close to λ.

Case 2: $\lambda = 0$. Again, there exist $\epsilon_1, \ldots, \epsilon_m$ such that the following conditions are met:

(i) $\epsilon_j < \min\left(\frac{\epsilon}{m}, \frac{a_{\pi_l}}{m}\right)$ for all $j \in [m]$,

(ii) the item order is π: $\frac{a_{\pi_1}}{a_{\pi_1} + \epsilon_{\pi_1}} > \ldots > \frac{a_{\pi_m}}{a_{\pi_m} + \epsilon_{\pi_m}}$,

(iii) $\sum_{j=1}^{m} \epsilon_j = 0$, and

(iv) item π_l is split in a ratio δ close to zero: $|\delta| < \frac{\epsilon}{b_{\pi_l}}$.

Thus Bob can approximate the outcome of Equation (10.1).

Now consider any ϵ-best response $\mathbf{y}$ of Bob; this induces some permutation of the items according to the ratios. If $\mathbf{y}$ is ϵ-close to the strategy of Alice we are done. Otherwise, Bob could change his strategy to be ϵ-close to the strategy of Alice while inducing the same permutation. This will only improve his utility as the boundary line moves to the left. □

It can be observed that there is at least one other ϵ-Nash equilibrium, at strategies $(\mathbf{b}, \tilde{\mathbf{b}})$, where $\tilde{\mathbf{b}}$ is a perturbation of Bob's truthful profile.

10.8 Discrete Strategies

Even though the continuous procedure is not guaranteed to have pure Nash equilibria, this does not imply that the discrete variant should also fail to have pure Nash equilibria. However we do find that this is indeed the case.

Theorem 34. *Adjusted Winner with discrete strategies is not guaranteed to have pure Nash equilibria.*

Proof. Consider a game with 4 items and 7 points, where Alice and Bob have valuations $(1, 1, 2, 3)$ and $(2, 3, 1, 1)$, respectively. This game does not admit a pure Nash equilibrium; this fact can be verified with a program that checks all possible configurations. □

Our next theorem shows that an ϵ-Nash equilibrium always exists in the discrete case if the number of points is set adequately, such that the players can approximately represent their true valuations.

Theorem 35. *For any profile* $(\mathbf{a}, \mathbf{b})$ *and any* $\epsilon > 0$*, there exists* P' *such that the procedure has an* ϵ*-Nash equilibrium when the players are given* P' *points.*

Proof. Let $\epsilon > 0$, and consider any profile $(\mathbf{a}, \mathbf{b})$ with denominator P. Then if we interpret $(\mathbf{a}, \mathbf{b})$ as a profile for the continuous setting, we get a $\epsilon/2$-Nash equilibrium $(\mathbf{a}, \tilde{\mathbf{a}})$ from Theorem 33, where $\tilde{a}_j = a_j + \epsilon_j$, for all $j \in [m]$.

Recall that $a_j, b_j \in \mathbb{Q}$; where $a_j = \frac{s_j}{P}$ and $b_j = \frac{t_j}{P}$, for some $s_j, t_j, \in \mathbb{N}$. We can find a rational number $\epsilon'_j = \frac{q_j}{r_j}$ (with $q_j, r_j \in \mathbb{N}$) that approximates ϵ_j within $\frac{\epsilon}{2m}$ for each $j \in [m]$, and such that the ordering of the items induced by the ratios $\frac{a_j}{a_j+\epsilon_j}$ is the same as the one given by $\frac{a_j}{a_j+\epsilon'_j}$. Define $\tilde{\mathbf{a}}'$ such that $\tilde{a}'_j = a_j + \epsilon'_j$.

It follows that $(\mathbf{a}, \tilde{\mathbf{a}}')$ is an ϵ-Nash equilibrium with $a_j, \tilde{a}'_j \in \mathbb{Q}$, for all $j \in [m]$. Thus whenever the players have a denominator of $P' = P \cdot \prod_{j=1}^{m} r_j$, the strategy profiles $(\mathbf{a}, \tilde{\mathbf{a}}')$ can be represented in the discrete procedure, so by giving P' points to the players, there exists an ϵ-Nash equilibrium. □

10.9 Informed Tie-Breaking

If the tie-breaking rule is not independent of the valuations, then both the discrete and continuous variants of Adjusted Winner have exact pure Nash equilibria. The deterministic tie-breaking rule under which this is possible is the one in which one of the players, for example Bob, is allowed to resolve ties by sorting them in the best possible order for him. Bob can compute the optimal order as outlined in the next definition.

Definition 22 (Informed Tie-Breaking)**.** *Let there be a fixed player, for example Bob. Given any strategies* $(\mathbf{x}, \mathbf{y})$*, for each permutation* π*, let* $l_\pi \in [m]$

and $\lambda_\pi \in [0, 1)$ *be the uniquely defined item and fraction for which:*

$$x_{\pi_1} + \dots x_{\pi_{l-1}} + \lambda x_{\pi_l} = (1 - \lambda)y_{\pi_l} + y_{\pi_{l+1}} + \dots + y_{\pi_m}$$

Let π^* *be an optimal permutation with respect to* $(\mathbf{x}, \mathbf{y})$, *namely* $\pi^* \in \arg\max_\pi (1-\lambda)y_{\pi_l} + y_{\pi_{l+1}} + \dots + y_{\pi_m}$. *Then under* informed tie-breaking, *the procedure resolves ties in the order given by* π^*.

Note that there might be more than one choice of π^* and Bob picks any fixed one. Now we can state the equilibrium existence theorems.

Theorem 36. *Adjusted Winner with continuous strategies and informed tie-breaking is guaranteed to have a pure Nash equilibrium.*

Proof. We show that the profile $(\mathbf{a}, \mathbf{a})$ is an exact equilibrium. By envy-freeness of the procedure, Alice gets at least half of the points at this strategy profile. Moreover, she cannot get strictly above half, since that would violate envy-freeness from the point of view of Bob's declared valuation, which is also $\mathbf{a}$. Thus Alice's strategy is a best response. As argued in Theorem 33 and 35, there exists an optimal permutation π^* such that by playing $\mathbf{a}$ and sorting the items in the order π^*, Bob can obtain the best possible utility (and as mentioned in Lemma 14, this value is achievable at these strategies). □

Similarly, it can be shown that the strategy profile $(\mathbf{a}, \mathbf{a})$ is a pure Nash equilibrium in the discrete procedure.

Theorem 37. *Adjusted Winner with discrete strategies and informed tie-breaking is guaranteed to have a pure Nash equilibrium.*

10.10 Efficiency and Fairness of Equilibria

Having examined the existence of pure Nash equilibria in Adjusted Winner, we now study their fairness and efficiency. For fairness, we observe that following.

Theorem 38. *All the pure Nash equilibria of Adjusted Winner are envy-free with respect to true valuations of the players.*

For efficiency, we use the well known measure of the *Price of Anarchy* [81, 105].

First, the *social welfare* of an allocation W is defined as the sum of the players' utilities: $SW(W) = u_A(W_A) + u_B(W_B)$. Then the Price of Anarchy is defined as the ratio between the maximum social welfare and the welfare of the worst-case pure Nash equilibrium. Our main findings are that when the procedure is equipped with an informed tie-breaking rule (i) all the pure Nash equilibria are Pareto optimal with respect to the true valuations and (ii) the price of anarchy is constant; that is, each pure Nash equilibrium achieves at least 75% of the optimal social welfare. We start with a lemma.

Lemma 15. *Let* $(\mathbf{x}, \mathbf{x})$ *be a pure Nash equilibrium of Adjusted Winner with informed tie-breaking and let* π^* *be the permutation that Bob chooses. Then, among all possible permutations,* π^* *maximizes Alice's utility.*

Proof. Assume by contradiction that there exists a permutation π that gives Alice a strictly larger utility; let α be her marginal increase from π^* to π. As discussed in Section 10.5, Alice can find appropriate constants $\epsilon_1, \ldots, \epsilon_m$ such that $AW(\mathbf{x}', \mathbf{x})$ with $\mathbf{x}' = (x_1 + \epsilon_1, \ldots, x_m + \epsilon_m)$ orders the items by π and the allocations $AW(\mathbf{x}, \mathbf{x})$ and $AW(\mathbf{x}', \mathbf{x})$ differ only in the allocation of the split item by by δ. Moreover, by continuity of the strategies, for each α, there exist ϵ_i's such that δ is small enough for $AW(\mathbf{x}', \mathbf{x})$ to be better for Alice than $AW(\mathbf{x}, \mathbf{x})$. □

Next we show that all equilibria are Pareto optimal.

Theorem 39. *All the pure Nash equilibria of Adjusted Winner with informed tie-breaking are Pareto optimal with respect to the true valuations* $\mathbf{a}$ *and* $\mathbf{b}$*.*

Proof. Let $(\mathbf{x}, \mathbf{x})$ be a pure Nash equilibrium of Adjusted Winner under informed tie-breaking and let l be the item that gets split (if any, otherwise the item to the left of the boundary line). Order Alice's items decreasing order of ratios a_i/x_i and Bob's items in increasing order of ratios b_i/x_i. Since $(\mathbf{x}, \mathbf{x})$ is a pure Nash equilibrium, by Lemma 15, both players are getting their maximum utility over all possible tie-breaking orderings of items. This means that for every item $i \le l$ and every item $j \ge l$ with $i \neq j$, it holds that

$$\frac{a_j}{x_j} \ge \frac{a_i}{x_i} \text{ and } \frac{b_i}{x_i} \ge \frac{b_j}{x_j} \Rightarrow \frac{a_i}{x_i} \cdot \frac{b_j}{x_j} \le \frac{a_j}{x_j} \cdot \frac{b_i}{x_i},$$

which by Lemma 12, implies that $AW(\mathbf{x}, \mathbf{x})$ is Pareto optimal. □

The Pareto optimality of a strategy profile has a direct implication on the social welfare achieved at that profile.

Theorem 40. *The Price of Anarchy of Adjusted Winner is* $4/3$*.*

Proof. Let $(\mathbf{x}, \mathbf{y})$ be any pure Nash equilibrium and let OPT_A and OPT_B be the utilities of Alice and Bob respectively in the optimal allocation. Since $AW(\mathbf{x}, \mathbf{y})$ is Pareto optimal by Theorem 39, the allocation for at least one of the players, (e.g. Alice), is at least as good as that of the optimal allocation. In other words, $u_A(AW(\mathbf{x}, \mathbf{y})) \ge OPT_A$. On the other hand, since $AW(\mathbf{x}, \mathbf{y})$ is envy-free, Bob's utility from $AW(\mathbf{x}, \mathbf{y})$ is at least $1/2$ which is at least $\frac{1}{2}OPT_B$. Overall, the social welfare of $AW(\mathbf{x}, \mathbf{y})$ is at least $OPT_A + \frac{1}{2}OPT_B$ and the ratio is minimized when OPT_A and OPT_B are minimum. Since $OPT_A \ge OPT_B \ge 1/2$, the ratio is at least $4/3$.

The bound is (almost) tight, given by the following simple instance with two items. Let $\mathbf{a} = (1 - \epsilon, \epsilon)$ and $\mathbf{b} = (\epsilon, 1 - \epsilon)$ and consider the strategy profile $\mathbf{x} = (\epsilon, 1-\epsilon)$ and $\mathbf{y} = (\epsilon, 1-\epsilon)$. It is not hard to see that $\mathbf{x}, \mathbf{y}$ is a pure Nash equilibrium for Alice breaking ties. The social welfare of the optimal allocation is $2-2\epsilon$. In the allocation of Adjusted Winner, Alice wins the first item and the second item is split (almost) in half. The social welfare of the mechanism is $1+\frac{1}{2}+o(\epsilon)$ and the approximation ratio is (almost) 4/3. As ϵ grows smaller, the ratio becomes closer to 4/3. □

10.11 Future Work

According to Foley [63], the quintessential characteristics of fairness are envy-freeness and Pareto optimality. We show that Adjusted Winner is guaranteed to have pure Nash equilibria, which satisfy both of these fairness notions. This attests to the usefulness and theoretical robustness of the procedure. A very interesting direction for future work is to study the *imperfect information* setting, as the Nash equilibria studied here require the players to have full information of each other's preferences.

Chapter 11

Characterization and Computation of Equilibria for Indivisible Goods

The systematic study of economic mechanisms began in the 19th century with the pioneering work of Irving Fisher [19] and Léon Walras [125], who proposed the Fisher market and the exchange economy as answers to the question: "How does one allocate scarce resources among the participants of an economic system?". These models of a competitive economy are central in mathematical economics and have been studied ever since in an extensive body of literature [105].

The high level scenario is that of several economic players arriving at the market with an initial endowment of resources and a utility function for consuming goods. The problem is to compute prices and an allocation for which an optimal exchange takes place: each player is maximally satisfied with the bundle acquired, given the prices and his initial endowment. Such allocation and prices form a *market equilibrium* and, remarkably, are guaranteed to exist under mild assumptions when goods are divisible [4].

In real scenarios, however, goods often come in discrete quantities; for example, clothes, furniture, houses, or cars may exist in multiple copies, but cannot be infinitely divided. Scarce resources, such as antique items or art collection pieces are even rarer – often unique (and thus *indivisible*). The problem of allocating discrete or indivisible resources is much more challenging because the theoretical guarantees from the divisible case do not always carry over; however, it can be tackled as well using market mechanisms [18, 38, 51, 106]. In this chapter, we are concerned with the question of allocating indivisible resources using the leading fairness concept from economics: the *competitive equilibrium from equal incomes* (CEEI).

The competitive equilibrium from equal incomes solution embodies the ideal notion of fairness [63, 76, 106, 123] and is a special case of the Fisher market model [122]. Informally, there are m goods to be allocated among n buyers, each of which is endowed with one unit of an artificial currency that they can use to acquire goods. The buyers declare their preferences over the goods, after which the equilibrium prices and allocation are computed. When the goods are divisible, a competitive equilibrium from equal incomes is guaranteed to exist for very general conditions and each equilibrium allocation satisfies the desirable properties of envy-freeness and efficiency.

In recent years, the competitive equilibrium from equal incomes has been studied for the allocation of discrete and indivisible resources in a series of papers. Bouveret and Lemaître [18] considered it for allocating indivisible goods, together with several other notions of fairness such as proportionality, envy-freeness, and maximin fairness. Budish [38] analyzed the allocation of multiple discrete goods for the course assignment problem[1] and designed an approximate variant of CEEI that is guaranteed to exist for any instance. In this variant, buyers have permissible bundles of goods and the approximation notion requires randomization to perturb the budgets of the buyers while relaxing the market clearing condition. In follow-up work, Othman, Papadimitriou, and Rubinstein [106] analyzed the computational complexity of this variant, showing that computing the approximate solution proposed by Budish is PPAD-complete, and that it is NP-hard to distinguish between an instance where an exact CEEI exists and the one in which there is no approximate-CEEI tighter than guaranteed in Budish [38].

In this chapter, we study the competitive equilibrium from equal incomes for two major classes of valuations, namely *perfect substitutes* and *perfect complements*. Perfect substitutes represent goods that can replace each other in consumption, such as Pepsi and Coca-Cola, and are modeled mathematically through *additive utilities*. This is the setting examined by Bouveret and Lemaître [18] as well. Perfect complements represent goods that have to be consumed together, such as a left shoe and a right shoe, and are modeled mathematically through *Leontief utilities*. For indivisible goods, Leontief utilities are in fact equivalent to the class of single-minded buyers, which have been studied extensively in the context of auctions [105].

We study the computation of competitive equilibria for indivisible goods and establish polynomial time algorithms and hardness results (where applicable). Our algorithm for Leontief utilities gives a very succinct characterization of markets that admit a competitive equilibrium from equal incomes for indivisible resources. The computational results of Othman, Papadimitriou, and Rubinstein [106] are orthogonal to our setting since they refer to combinatorial valuations.

[1] Given a set of students and courses to be offered at a university, how should the courses be scheduled given that the students have preferences over their schedules and the courses have capacity constraints on enrollment?

11.1 Competitive Equilibrium from Equal Incomes

We begin by formally introducing the competitive equilibrium from equal incomes. Formally, there is a set $N = \{1, \ldots, n\}$ of buyers and a set $M = \{1, \ldots, m\}$ of goods which are brought by a seller. In general, the goods can be either infinitely divisible or discrete and, without loss of generality, there is exactly one unit from every good $j \in M$. Each buyer i is endowed with:

- A utility function $u_i : [0, 1]^m \to \mathbb{R}_{\geq 0}$ for consuming the goods, which maps each vector $\mathbf{x} = \langle x_1, \ldots, x_m \rangle$ of resources to a real value, where $u_i(\mathbf{x})$ represents the buyer's utility for bundle $\mathbf{x}$; note that $\mathbf{x}_j$ is the amount received by the buyer from good j.

- An initial budget $B_i = 1$, which can be viewed as (artificial) currency to acquire goods, but has no intrinsic value to the buyer. However, the currency does have intrinsic value to the seller.

Each buyer in the market wants to spend its entire budget to acquire a bundle of items that maximizes its utility, while the seller aims to sell all the goods (which it has no intrinsic value for) and extract the money from the buyers.

A market outcome is defined as a tuple $(\mathbf{x}, \mathbf{p})$, where $\mathbf{p}$ is a vector of prices for the m items, and $\mathbf{x} = \langle \mathbf{x}_1, \ldots, \mathbf{x}_n \rangle$ is an allocation of the m items, with p_j denoting the price of item j and x_{ij} representing the amount of item j received by buyer i. A market outcome that maximizes the utility of each buyer subject to its budget constraint and clears the market is called a *market equilibrium* [105]. Formally, $(\mathbf{x}, \mathbf{p})$ is a market equilibrium if and only if:

- For each buyer $i \in N$, the bundle $\mathbf{x}_i$ maximizes buyer i's utility given the prices $\mathbf{p}$ and budget $B_i = 1$.

- Each item $j \in M$ is completely sold or has price zero. That is:

$$\left(\sum_{i=1}^{n} x_{ij} - 1 \right) p_j = 0.$$

- All the buyers exhaust their budgets; that is, $\sum_{j=1}^{m} p_j \cdot x_{ij} = 1$, for all $i \in N$.

Every competitive :equilibrium from equal incomes $(\mathbf{x}, \mathbf{p})$ is envy-free; if buyer i would strictly prefer another buyer j's bundle $\mathbf{x}_j$, then i could simply purchase $\mathbf{x}_j$ instead of $\mathbf{x}_i$ since they have the same buying power, which is in contradiction with the equilibrium property.

A market with divisible goods is guaranteed to have a competitive equilibrium under mild conditions [4]. Moreover, for the family of *Constant Elasticity* *of Substitution* valuations, the equilibrium can be computed using a remarkable convex program due to Eisenberg and Gale [60], which is one of the few algorithmic results in general equilibrium theory and stated in Figure 11.1. The Eisenberg-Gale program computes an equilibrium for the more general Fisher model, where the budgets B_i of the buyers are not necessarily equal.

Figure 11.1: The Eisenberg-Gale convex program for Fisher markets

$$
\begin{aligned}
\max \quad & \sum_{i=1}^{n} B_i \cdot \log(u_i) \\
s.t. \quad & u_i = \left(\sum_{j=1}^{m} a_{ij} \cdot x_{ij}^{\rho} \right)^{\frac{1}{\rho}}, \quad \forall\, i \in \{1, \ldots, n\} \\
& \sum_{i=1}^{n} x_{ij} \leq 1, \quad \forall\, j \in \{1, \ldots, m\} \\
& x_{ij} \geq 0, \quad \forall\, i \in \{1, \ldots, n\}, j \in \{1, \ldots, m\}
\end{aligned}
$$

The classes of valuations studied in this paper – perfect complements and perfect substitutes – belong to the constant elasticity of substitution family. In the following sections we study these classes in detail in the context of allocating indivisible resources.

11.2 Perfect Complements

Let $\mathcal{M} = (N, M, \mathbf{v})$ denote a market with perfect complements, represented through Leontief utilities; recall N is the set of buyers, M the set of items, and $\mathbf{v}$ a matrix of constants, such that $v_{i,j}$ is the value of buyer i for consuming one unit of good j. The utility of buyer i for bundle $\mathbf{x} = \langle x_1, \ldots, x_m \rangle \in [0,1]^m$ is given by:

$$u_i(\mathbf{x}) = \min_{j=1}^{m} \left(\frac{x_j}{v_{i,j}} \right) \tag{11.1}$$

In our model the goods are indivisible, and so $x_{i,j} \in \{0, 1\}$, for all i, j. By examining Equation 11.1, it can be observed that buyer i's utility for a bundle depends solely on whether the buyer gets all the items that it values positively (or not). To capture this we define the notion of *demand set*.

Definition 23 (Demand Set)**.** *Given a CEEI market with indivisible goods and Leontief utilities, let the* demand set *of buyer i be the set of items that i has a strictly positive value for; that is, $D_i = \{j \in M \mid v_{i,j} > 0\}$.*

Now we can introduce the precise utility equation for indivisible goods with Leontief valuations.

Definition 24 (Leontief Utility for Indivisible Goods)**.** *Given a market with Leontief utilities and indivisible goods, the utility of a buyer i for a bundle $\mathbf{x} = \langle x_1, \ldots, x_m \rangle \in [0, 1]^m$ is:*

$$u_i(\mathbf{x}) = \begin{cases} \min_{j \in D_i} \left(\frac{1}{v_{i,j}} \right), & \text{if } D_i \subseteq \mathbf{x}, \\ 0, & \text{otherwise} \end{cases}$$

where D_i represents buyer i's demand set.

We illustrate this utility class with an example. Note that valuations are not necessarily normalized.

Example 8. *Let $\mathcal{M}$ be a market with buyers $N = \{1, 2, 3\}$, items $M = \{1, 2, 3, 4\}$, and values: $v_{1,1} = 1$, $v_{2,2} = 2$, $v_{2,4} = 3$, $v_{3,1} = 0.5$, $v_{3,2} = 2.5$, $v_{3,3} = 5$, and $v_{i,j} = 0$, for all other i, j. Recall the demand set of each buyer consists of the items it values strictly positively, and so: $D_1 = \{1\}$, $D_2 = \{2, 4\}$, $D_3 = \{1, 2, 3\}$. Then the utility of buyer 1 for a bundle $S \subseteq M$ is: $u_1(S) = 0$ if $D_1 \not\subseteq S$, and $u_1(S) = \min_{j \in D_1} \left(\frac{1}{v_{1,j}} \right) = \frac{1}{v_{1,1}} = 1$ otherwise. Similarly, $u_2(S) = 0$ if $D_2 \not\subseteq M$ and $u_2(S) = \min \left(\frac{1}{v_{2,2}}, \frac{1}{v_{2,4}} \right) = \frac{1}{3}$ otherwise.*

Next, we examine the computation of allocations that are fair according to the CEEI solution concept. The main computational problems that we consider are : *Given a market, determine whether a competitive equilibrium exists and compute it when possible.* Depending on the scenario at hand, an allocation of the resources to the buyers may have already been made (or the seller may have already set prices for the items). The questions then are to determine whether an equilibrium exists at those prices or allocations. Our algorithm for computing a competitive equilibrium for Leontief utilities with indivisible goods yields a characterization of when a market equilibrium is guaranteed to exist.

Theorem 41. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with Leontief utilities, indivisible goods, and a tuple $(\mathbf{x}, \mathbf{p})$, where $\mathbf{x}$ is an allocation and $\mathbf{p}$ a price vector, it can be decided in polynomial time if $(\mathbf{x}, \mathbf{p})$ is a market equilibrium for $\mathcal{M}$.*

Proof. It is sufficient to verify that these conditions hold:

- Each buyer i exhausts their budget: $\sum_{j \in \mathbf{x}_i} p_j = 1$.
- Each item is either allocated or has a price of zero.
- No buyer i can afford a better bundle; that is, if $u_i(\mathbf{x}, \mathbf{p}) = 0$, then $\sum_{j \in D_i} p_j > 1$.

Clearly all three conditions can be verified in polynomial time, namely $O(mn)$, which concludes the proof. □

Theorem 42. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with Leontief utilities, indivisible goods, and a price vector $\mathbf{p}$, it is co-NP-complete to decide if there exists an allocation $\mathbf{x}$ such that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium for $\mathcal{M}$.*

Proof. From Theorem 41, given an allocation $\mathbf{x}$ for $(\mathcal{M}, \mathbf{p})$, it can be verified in polynomial time if there exists a market equilibrium for $\mathcal{M}$ at $(\mathbf{x}, \mathbf{p})$. To show hardness, we use the NP-complete problem PARTITION:

> *Given a set of positive integers $S = \{s_1, \ldots, s_m\}$, are there subsets $A, B \subset S$ such that $A \cup B = S$, $A \cap B = \emptyset$, and $\sum_{a \in A} a = \sum_{b \in B} b$?*

Given partition input S, construct a market $\mathcal{M} = (N, M, \mathbf{v})$ and price vector $\mathbf{p}$ as follows:

- Set $N = \{1, 2, 3\}$ of buyers.
- Set $M = \{0, 1, \ldots, m\}$ of items.
- Price vector $\mathbf{p} = (p_0, \ldots, p_m)$, such that $p_0 = 1$ and $p_j = \frac{2 \cdot s_j}{\sum_{l=1}^{m} s_l}$, for all $j \in \{1, \ldots, m\}$.
- Demand sets: $D_1 = \{0\}$ and $D_2 = D_3 = \{1, \ldots, m\}$. Clearly these demand sets can be expressed through Leontief valuations – for example, let $v_{1,0} = 1$ and $v_{1,j} = 0$, for all $j \in \{1, \ldots, m\}$, and $v_{2,k} = v_{3,k} = \frac{1}{m+1}$, for all $k \in \{0, \ldots, m\}$.

Note that the total price of the items in $\{1, \ldots, m\}$ is:

$$p(\{1, \ldots, m\}) = \sum_{j=1}^{m} p_j = \sum_{j=1}^{m} \frac{2 \cdot s_j}{\sum_{l=1}^{m} s_l} = 2$$

.

($\Longrightarrow$) If there is a partition (A, B) of $\mathcal{S}$, then we show that the allocation $\mathbf{x}$ given by $\mathbf{x}_1 = \{0\}$, $\mathbf{x}_2 = A$, $\mathbf{x}_3 = B$ is a market equilibrium:

- All the items are sold, since $\mathbf{x}_1 \cup \mathbf{x}_2 \cup \mathbf{x}_3 = \{0, 1, \ldots, m\}$, and the bundles are disjoint, since $\mathbf{x}_i \cap \mathbf{x}_j = \emptyset$, for all $i, j \in N$, $i \neq j$.
- Each buyer gets an optimal bundle at prices $\mathbf{p}$, since buyer 1 gets the best possible bundle: $u_1(\mathbf{x}_1, \mathbf{p}) = u_1(D_1, \mathbf{p}) = 1$, while buyers 2 and 3 cannot afford anything better, as the price of their demanded bundle is higher than their budget: $\mathbf{p}(D_2) = \mathbf{p}(D_3) = \sum_{j=0}^{m} p_j = 3 > 1$.
- The amount of money spent by each buyer is equal to their budget: $\mathbf{p}(\mathbf{x}_1) = 1$, $\mathbf{p}(\mathbf{x}_2) = \mathbf{p}(A) = 1$, and $\mathbf{p}(\mathbf{x}_3) = \mathbf{p}(B) = 1$.

($\Longleftarrow$) On the other hand, if there is an allocation $\mathbf{x}$ such that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium for the market $\mathcal{M}$, we claim that $A = \mathbf{x}_2$ and $B = \mathbf{x}_3$ is a correct partition of $\mathcal{S}$. First, note that if $\mathbf{x}$ is such that buyer 1 does not get item 0, then the optimality condition fails for buyer 1 because $p_0 = 1$ and so the buyer could always afford it. Thus buyer 1 must get item 0 and moreover, it cannot get anything else. Thus $\mathbf{x}_1 = \{0\}$ and $\mathbf{x}_2, \mathbf{x}_3 \subseteq \{1, \ldots, m\}$.

- Since all the items are sold at $(\mathbf{x}, \mathbf{p})$, we have that $A \cup B = S$.
- Buyers 2 and 3 spend their entire budgets at $(\mathbf{x}, \mathbf{p})$, and so $\mathbf{p}(\mathbf{x}_2) = 1$ and $\mathbf{p}(\mathbf{x}_3) = 1$. Then $\mathbf{p}(A) = \mathbf{p}(B) = 1$ and:

$$\begin{aligned} \sum_{s_j \in A} \left(\frac{2 \cdot s_j}{\sum_{k=1}^m s_k} \right) &= \sum_{s_j \in B} \left(\frac{2 \cdot s_j}{\sum_{k=1}^m s_k} \right) \iff \\ \sum_{s_j \in A} s_j &= \sum_{s_j \in B} s_j \end{aligned}$$

Thus S has a partition if and only if the corresponding market and price vector admit a market clearing allocation, which completes the proof. □

Theorem 43. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with Leontief utilities and an allocation $\mathbf{x}$, it can be decided in polynomial time if there exists price vector $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium for $\mathcal{M}$.*

Proof. This problem can be solved using linear programming (see Algorithm 5.4). At a high level, one needs to check that the allocation $\mathbf{x}$ is feasible, that each item is either sold or has a price of zero, and that (i) each buyer spends all their money and (ii) whenever a buyer does not get their demand set, the bundle is too expensive. Since the number of constraints is polynomial in the number of buyers and items, the algorithm runs in polynomial time. □

Finally, we investigate the problem of computing both market equilibrium allocation and prices given an instance. We will later also discuss improving the efficiency of the computed equilibria.

Theorem 44. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with Leontief utilities and indivisible items, a competitive equilibrium from equal incomes exists if and only if the following hold:*

- *There are at least as many items as buyers ($m \geq n$)*
- *No two buyers have identical demand sets of size one.*

Moreover, an equilibrium $(\mathbf{x}, \mathbf{p})$ of the market $\mathcal{M}$ can be computed in polynomial time if it exists.

input: Market $\mathcal{M}$ with Leontief valuations; allocation $\mathbf{x}$
output: price vector $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium for $\mathcal{M}$, or NULL if none exists
$\mathcal{A} \leftarrow \emptyset$ *// Set of items allocated under* $\mathbf{x}$
// Check that $\mathbf{x}$ *is feasible*
for $i = 1$ ***to*** n **do**
  $\mathcal{A} \leftarrow \mathcal{A} \cup \mathbf{x}_i$
  for $j = i + 1$ ***to*** n **do**
    if $(\mathbf{x}_i \cap \mathbf{x}_j \neq \emptyset)$ **then**
      return NULL
    end if
  end for
end for
$\mathcal{C} \leftarrow \emptyset$ *// Initialize the set of constraints*
for $j \in \mathcal{A} \setminus M$ **do**
  $\mathcal{C} \leftarrow \mathcal{C} \cup \{p_j \leq 0\}$ *// Price the unsold items at zero*
end for
for $i \in \{1, \ldots, n\}$ **do**
  $\mathcal{C} \leftarrow \mathcal{C} \cup \left\{\sum_{j \in \mathbf{x}_i} p_j \leq 1, -\sum_{j \in \mathbf{x}_i} p_j \leq -1\right\}$
  if $(D_i \not\subseteq \mathbf{x}_i)$ **then**
    $\mathcal{C} \leftarrow \mathcal{C} \cup \left\{-\sum_{j \in \mathbf{x}_i} p_j \leq -1 - \epsilon\right\}$ *// If buyer i does not get its demand set, then it's because the bundle is too expensive*
  end if
end for
return SOLVE$(\max \epsilon, \mathcal{C}, \mathbf{p} \geq 0)$ *// Linear program solver*

Algorithm 11: COMPUTE-EQUILIBRIUM-PRICES$(\mathcal{M}, \mathbf{x})$

Proof. Clearly the two conditions are necessary; if there are fewer items than buyers, then the budgets can never be exhausted, while if there exist two buyers whose demand sets are identical and consist of exactly the same item, at least one of them will be envious under any pair of feasible allocation and prices.

To see that the conditions are also sufficient, consider the allocation produced by Algorithm 5.4. At a high level, the algorithm first sorts the buyers in increasing order by the sizes of their demand sets, breaking ties lexicographically. Then each buyer i in this order is given one item k_i, where k_i is selected from the unallocated items in the buyer's demand set (if possible), and otherwise it represents an arbitrary un-allocated item. Finally, the last buyer (i.e. with the largest demand set) additionally gets all the items that remained unallocated at the end of this iteration (if any). For each buyer i, the items in its bundle, $\mathbf{x}_i$, are priced equally, at $1/|\mathbf{x}_i|$.

Now we verify that whenever the market $\mathcal{M}$ has an equilibrium, the allo-

input: Market $\mathcal{M}$ with Leontief valuations
output: Equilibrium allocation and prices $(\mathbf{x}, \mathbf{p})$, or NULL if none exist
if $(m < n)$ **then**
    return NULL *// No equilibrium : too few items*
end if
for $(\ i, j \in N)$ **do**
    if $(i \neq j$ ***and*** $D_i = D_j$ ***and*** $|D_i| = 1)$ **then**
        return NULL *// No equilibrium : buyers i and j have identical singleton demand sets*
    end if
end for
$\mathcal{A} \leftarrow \emptyset$ *// Items allocated so far*
for (buyer $i \in N$ *in increasing order by* $|D_i|$) **do**
    if $(|D_i \setminus \mathcal{A}| \geq 1)$ **then**
        $k_i \leftarrow \text{argmin}_{\ell \in D_i \setminus \mathcal{A}} \ell$ *// If not all the items in buyer i's demand set have been allocated, give the buyer one of them*
    else
        $k_i \leftarrow \text{argmin}_{\ell \in M \setminus \mathcal{A}} \ell$ *// Otherwise, i gets an arbitrary unallocated item*
    end if
    $\mathbf{x}_i \leftarrow \{k_i\}$
    $\mathcal{A} \leftarrow \mathcal{A} \cup \{k_i\}$
end for
$L \leftarrow \text{argmax}_{i \in N} |D_i|$ *// The buyer with the largest demand also gets all the unallocated items (if any)*
$\mathbf{x}_L \leftarrow \mathbf{x}_L \cup (M \setminus \mathcal{A})$
for $(i \in N)$ **do**
    for $(j \in \mathbf{x}_i)$ **do**
        $p_j \leftarrow 1/|\mathbf{x}_i|$
    end for
end for
return $(\mathbf{x}, \mathbf{p})$

Algorithm 12: COMPUTE-EQUILIBRIUM($\mathcal{M}$)

cation and prices $(\mathbf{x}, \mathbf{p})$ computed by Algorithm 5.4 represent indeed a market equilibrium:

- *Budgets exhausted*: Each buyer i gets a non-empty bundle $\mathbf{x}_i$ priced at:

$$\mathbf{p}(\mathbf{x}_i) = \left(\frac{1}{|\mathbf{x}_i|}\right) \cdot |\mathbf{x}_i| = 1.$$

Thus the buyer spends all its money.

- *Items sold*: Each item is allocated by the algorithm.

• *Optimality for each buyer*: We show that each buyer i either gets its demand set or cannot afford it using a few cases:

∘ *Case 1*: ($|D_i| = 1$). Since there are no two identical demand sets with the size of one, buyer i gets the unique item in its demand set, and this allocation maximizes i's utility.

∘ *Case 2*: ($|D_i| \geq 2$) and i is not the last buyer. Then if i gets an item from its demand set, since $|D_i| \geq 2$ and all items are positively priced, the bundle D_i is too expensive: $\mathbf{p}(D_i) > 1$. Otherwise, i gets an item outside of its demand set. Then all the items in D_i must have been allocated to the previous buyers. Since $|D_i| \geq 2$ and each previously allocated item has price 1, D_i is too expensive: $\mathbf{p}(D_i) > 1$.

∘ *Case 3*: ($|D_i| \geq 2$) and i is the last buyer. If i does not get all its demand, then some item in D_i was given to an earlier buyer at price 1. From $|D_i| \geq 2$, there is at least one other desired item in D_i positively priced, thus $\mathbf{p}(D_i) > 1$.

Thus, Algorithm 5.4 computes an equilibrium. □

To gain more intuition about the model, we illustrate the execution of Algorithm 5.4 on an example.

Example 9. *Consider a market with buyers* $N = \{1, \ldots, 6\}$, *items* $M = \{1, \ldots, 8\}$, *and demands:* $D_1 = \{1\}$, $D_2 = \{2\}$, $D_3 = \{2, 3\}$, $D_4 = \{2, 3\}$, $D_5 = \{4, 5, 6\}$, $D_6 = \{6, 7, 8\}$.

Algorithm 5.4 sorts the buyers in increasing order of the number of items in their demand sets, breaking ties lexicographically. The order is: $(1, 2, 3, 4, 5, 6)$.

- *Step 1 : Buyer* 1 *gets item* 1 *at price* 1*:* $\mathbf{x}_1 = \{1\}$, $p_1 = 1$.
- *Step 2 : Buyer* 2 *gets item* 2 *at price* 1 *:* $\mathbf{x}_2 = \{2\}$, $p_2 = 1$.
- *Step 3 : There is one unallocated item left from buyer* 3*'s demand set, and so* 3 *gets it:* $\mathbf{x}_3 = \{3\}$ *and* $p_3 = 1$.
- *Step 4 : Buyer* 4*'s demand set has been completely allocated, thus* 4 *gets the free item (outside of its demand) with smallest index:* $\mathbf{x}_4 = \{4\}$ *and* $p_4 = 1$.
- *Step 5 : There are two items (*5 *and* 6*) left unallocated in buyer* 5*'s demand set. Thus :* $\mathbf{x}_5 = \{5\}$ *and* $p_5 = 1$.
- *Step 6 : Buyer* 6 *gets all the leftover items :* $\mathbf{x}_6 = \{6, 7, 8\}$ *at equal prices:* $p_6 = p_7 = p_8 = 1/3$.

The characterization obtained through Algorithm 5.4 raises several important questions. For example, not only do fair division procedures typically guarantee fairness (according to a given solution concept), but also they improve some measure of efficiency when possible.

The *social welfare* of an allocation $\mathbf{x}$ is defined as the sum of the buyers' utilities: $\mathrm{SW}(\mathbf{x}) = \sum_{i=1}^{n} u_i(\mathbf{x}_i)$. For measuring social welfare, valuations must be normalized such that players are weighted equally (since their rights over the goods are equal), and so: $u_i(M) = 1$, for each buyer i. This can also be interpreted as the number of buyers that receive their demand sets (possibly in addition to other items). As the next example illustrates, the allocation computed by Algorithm 5.4 can be the worst possible among all market equilibria.

Example 10. *Given $n \in \mathbb{N}$, let $N = \{1, \ldots, n\}$ be the set of buyers, $M = \{1, \ldots, 2n\}$ the set of items, and the demand sets given by: $D_i = \{2i-1, 2i\}$, for each $i \in N$. Then Algorithm 5.4 computes the allocation: $\mathbf{x}_1 = \{1\}$, $\mathbf{x}_2 = \{2\}$, ..., $\mathbf{x}_{n-1} = \{n-1\}$, $\mathbf{x}_n = \{n, \ldots, 2n\}$, with a social welfare of $\mathrm{SW}(\mathbf{x}) = 1$. The optimal allocation that can be supported in a competitive equilibrium is: $\mathbf{x}_i^* = \{2i-1, 2i\}$, for each $i \in N$, with a social welfare of $\mathrm{SW}(\mathbf{x}^*) = n$.*

These observations give rise to the question: *Is there an efficient algorithm for computing a competitive equilibrium from equal incomes with optimal social welfare (among all equilibria) for perfect complements with indivisible goods?*

It is important to note that the allocation that maximizes social welfare among all possible allocations cannot always be supported in a competitive equilibrium. We illustrate this phenomenon in Example 11.

Example 11. *Consider a market with buyers: $N = \{1, 2\}$ and items: $M = \{1, 2, 3\}$, where the demand sets are: $D_1 = D_2 = \{1, 2\}$. Concretely, let these demands be induced by the valuations: $v_{1,1} = v_{1,2} = 1$, $v_{1,3} = 0$ and $v_{2,1} = v_{2,2} = 1$, $v_{2,3} = 0$. The optimal social welfare is 1 and can be achieved by giving one of the buyers its entire demand set and the other buyer the remaining item; for example, let $\mathbf{x}_1^* = \{1, 2\}$ and $\mathbf{x}_2^* = \{3\}$, with $p_1 = p_2 = 1/2$ and $p_3 = 1$. Clearly no such allocation can be supported in an equilibrium, because whenever a buyer gets their full demand, the other buyer does not get its own demand but can afford it (their initial budgets are equal). Thus every competitive equilibrium for this instance has a social welfare of zero, such as $\mathbf{x}_1 = \{1\}$, $\mathbf{x}_2 = \{2, 3\}$, with $p_1 = 1$, $p_2 = 1$, $p_3 = 0$.*

The next result implies that equilibria with optimal social welfare cannot be computed efficiently in the worst case.

Theorem 45. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with Leontief valuations, indivisible goods, and an integer $K \in \mathbb{N}$, it is NP-complete to decide if $\mathcal{M}$ has a competitive equilibrium from equal incomes with social welfare at least K.*

Proof. (sketch) We use a reduction from the NP-complete problem SET PACKING:

> *Given a collection $\mathcal{C} = \langle C_1, \ldots, C_n \rangle$ of finite sets and a positive integer $K \leq n$, does $\mathcal{C}$ contain at least K mutually disjoint sets?*

Given collection $\mathcal{C}$ and integer K, let $\mathcal{M}$ be a market with buyers $N = \{1, \ldots, n\}$, items $M = \{1, \ldots, m+n\}$, and demands $D_i = C_i \cup \{m+i\}$, for all $i \in N$. It can be checked that $\mathcal{M}$ has a competitive equilibrium with social welfare at least K if and only if $\mathcal{C}$ has a disjoint collection of at least K sets. □

In the full version of the paper (see [37]) we show that there exists a 1/n approximation for the social welfare maximization problem (even for weighted valuations) and this is close to optimal.

Another important notion of efficiency is known as *Pareto efficiency*. Informally, we say that a market equilibrium $(\mathbf{x}, \mathbf{p})$ is Pareto efficient (with respect to the set of all the possible market equilibria) if there is no other equilibrium $(\mathbf{x}', \mathbf{p}')$ that strictly improves the utility of at least one buyer without degrading the other buyers. Clearly, the equilibrium that maximizes social welfare is Pareto efficient. However, this particular equilibrium cannot be computed efficiently in the worst case (by Theorem 45), which leads to the next question that we leave open: *Is there a polynomial time algorithm that computes a Pareto efficient market equilibrium for indivisible goods with Leontief valuations?*

11.3 Perfect Substitutes

We begin by introducing the utility function in a market with perfect substitutes, represented through additive valuations.

Definition 25 (Additive Utility for Indivisible Goods)**.** *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with additive utilities and indivisible goods, the utility of a buyer i for a bundle $\mathbf{x} = \langle x_1, \ldots, x_m \rangle \in \{0,1\}^m$ is:*

$$u_i(\mathbf{x}) = \sum_{j=1}^{m} v_{i,j} \cdot x_{i,j} \tag{11.2}$$

where $v_{i,j}$ are constants and represent the value of buyer i for consuming one unit of good j, while $x_{i,j} = 1$ if buyer i gets good j, and $x_{i,j} = 0$, otherwise.

Next we investigate the computation of competitive equilibria from equal incomes with indivisible goods and additive utilities. Note that if a market $\mathcal{M}$ has a competitive equilibrium at some allocation and prices $(\mathbf{x}, \mathbf{p})$, then $\mathcal{M}$ is guaranteed to have an equilibrium at the same allocation $\mathbf{x}$ where all the prices are rational numbers, $(\mathbf{x}, \mathbf{p}^*)$; this aspect appears implicitly in some of the following proofs.

Theorem 46. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with additive valuations, indivisible goods, and tuple $(\mathbf{x}, \mathbf{p})$, where $\mathbf{x}$ is an allocation and $\mathbf{p}$ is a price vector, it is coNP-complete to determine whether $(\mathbf{x}, \mathbf{p})$ is a competitive equilibrium for $\mathcal{M}$.*

Proof. The problem admits efficiently verifiable "no" instances: it can be checked in polynomial time if the allocation is not feasible, or the budgets are not exhausted, or not all the items are sold. Otherwise, if $(\mathbf{x}, \mathbf{p})$ is not a market equilibrium for $\mathcal{M}$, then there exists a buyer k with a suboptimal bundle. In other words, the certificate that $(\mathbf{x}, \mathbf{p})$ is not a market equilibrium is given by a tuple (k, D), where k is a buyer that strictly prefers bundle $D \subseteq M$ to $\mathbf{x}_k$ and can also afford it; that is, $u_k(\mathbf{x}_k) < u_k(D)$ and $p(D) \leq 1$.

We show hardness using the SUBSET-SUM problem:

> *Given a set of positive integers $\mathcal{W} = \{w_1, \ldots, w_n\}$ and a target number K, is there a subset $S \subseteq \mathcal{W}$ that adds up to exactly K?*

Given $\langle \mathcal{W}, K \rangle$, construct market $\mathcal{M} = (N, M, \mathbf{v})$ and tuple $(\mathbf{x}, \mathbf{p})$, with buyers $N = \{0, 1, \ldots, n\}$, items $M = \{0, 1, \ldots, 2n\}$, and values:

- Buyer 0: $v_{0,0} = K - 1$; $v_{0,j} = w_j$, for all $j \in \{1, \ldots, n\}$; $v_{0,j} = 0$, for all $j \in \{n+1, \ldots, 2n\}$.
- Buyer $i \in \{1, \ldots, n\}$: $v_{i,n+i} = 1$; $v_{i,j} = 0$, for all $j \in M \setminus \{n+i\}$.

Let $\mathbf{x}_0 = \{0\}$ and $\mathbf{x}_i = \{i, n+i\}$, for all $i \in \{1, \ldots, n\}$. Define prices: $p_0 = 1$, $p_j = \frac{w_j}{K}$ and $p_{n+j} = 1 - p_j$, for all $j \in \{1, \ldots, n\}$ (Note that if there exist items with $w_j > K$, those items can be thrown away from the beginning).

($\Longrightarrow$) If there is a solution $S \subseteq U$ to $\mathcal{W}$, then we claim that $\mathcal{M}$ does not have an equilibrium at $(\mathbf{x}, \mathbf{p})$ since buyer 0 can acquire a better bundle, namely S:

- Buyer 0 can afford S:

$$\sum_{j \in S} p_j = \sum_{j \in S} \frac{w_j}{K} = \frac{K}{K} = 1.$$

- Bundle S is strictly better than $\mathbf{x}_0$:

$$\sum_{j \in S} v_{0,j} = \sum_{j \in S} w_j = K > K - 1 = v_{0,0}.$$

($\Longleftarrow$) If $(\mathbf{x}, p)$ is not a market equilibrium, then it must be that buyer 0 can get a better bundle (since all budgets are spent, all items are sold, and the other buyers already have their unique valuable item).

Thus there is bundle S such that: (i) $\sum_{j \in S} v_{0,j} > v_{0,0}$ and (ii) $\sum_{j \in S} p_j \leq 1$. From $v_{0,j} = 0$, for all $j \in \{n+1, \ldots, 2n\}$, it follows that $S \subset \{1, \ldots, n\}$ (otherwise, just take $S' = S \cap \{1, \ldots, n\}$). Condition (i) is equivalent to: $\sum_{j \in S} w_j \geq K$ and condition (ii) can be rewritten as:

$$\sum_{j \in S} \frac{w_j}{K} \leq 1 \iff \sum_{j \in S} w_j \leq K$$

Then S is a subset-sum solution; this completes the proof. □

Theorem 47. *Given a market with indivisible goods and additive valuations, $\mathcal{M} = (N, M, \mathbf{v})$, it is NP-hard to decide if $\mathcal{M}$ has a competitive equilibrium.*

Proof. We reduce from the NP-complete problem EXACT COVER BY 3-SETS (X3C):

> *Given universe $\mathcal{U} = \{1, \ldots, 3n\}$ of elements and family of subsets $\mathcal{F} = \{S_1, \ldots, S_k\}$, with $|S_i| = 3$, $\forall i$, decide if there is collection $S \subseteq \mathcal{F}$ such that each element of $\mathcal{U}$ occurs exactly once in S.*

Given X3C instance $\langle \mathcal{U}, \mathcal{F} \rangle$, define $N = \{1, \ldots, k\}$, $M = \{1, \ldots, 3n, 3n + 1, \ldots, 2n + k\}$ (note this assumes that $k \leq n$, since otherwise the answer to the X3C instance is trivially "no"), and valuations for each buyer $i \in N$:

- $v_{i,j} = \frac{1}{3}$, for all $j \in S_i$.
- $v_{i,j} = 1$, for all $j \in \{3n + 1, \ldots, 2n + k\}$.
- $v_{i,j} = 0$, otherwise.

If the market has some competitive equilibrium $(\mathbf{x}, \mathbf{p})$, then the following conditions hold:

- Each buyer gets a bundle worth at least 1, since the items in $\{3n + 1, \ldots, 2n + k\}$ are each worth 1 to every buyer and each of their prices is at most 1 (since all items get sold).
- No buyer can get a bundle worth more than 1.

Then each buyer gets a bundle worth exactly 1, and so the items in $\{3n + 1, \ldots, 2n+k\}$ are priced at 1 each. The remaining n buyers get a bundle worth 1 each from the items $\{1, \ldots, 3n\}$, which can only happen if their allocations form a solution to the X3C instance.

If X3C has a solution S, then a market equilibrium is obtained immediately by giving the sets in S to the buyers that want them, and the leftover items, in $\{3n + 1, \ldots, 2n + k\}$, to the remaining buyers. □

The next question, of computing an equilibrium allocation given a market $\mathcal{M}$ and a price vector $\mathbf{p}$ was raised by Bouveret and Lemaître ([18]). In a recent note, Aziz ([9]) also studied the hardness of this problem. We include our proof as well, which uses the PARTITION problem.

Theorem 48. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with indivisible goods, additive valuations, and price vector $\mathbf{p}$, it is coNP-hard to decide if there is an allocation $\mathbf{x}$ such that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium.*

Proof. We use a reduction from the NP-complete problem PARTITION. Given a set $S = \{s_1, \ldots, s_m\}$, where $\sum_{j=1}^{m} s_j = 2V$ and $s_j \in \mathbb{N}$, $\forall j \in \{1, \ldots, m\}$, construct the following market with indivisible goods and additive valuations:

- Let $N = \{1, 2\}$ and $M = \{1, \ldots, m + 2\}$.
- Buyer 1's valuations: $v_{1,j} = s_j$, $\forall j \in \{1, \ldots, m\}$, $v_{1,m+1} = 3V$ and $v_{1,m+2} = V - 1$.
- Buyer 2's valuations: $v_{2,j} = 1, \forall j \in \{1, \ldots, m\}$ and $v_{2,m+1} = v_{2,m+2} = 0$.

Consider the price vector given by $p_j = \frac{s_j}{2V}$, $\forall j \in \{1, \ldots, m\}$, $p_{m+1} = \frac{1}{2} = p_{m+2}$.

We claim that S has a partition if and only if $\mathcal{M}$ does not have an equilibrium at $\mathbf{p}$. First, note that buyer 2 can afford to buy all the items it has a strictly positive value for – i.e. the set $M' = \{1, \ldots, m\}$ – since:

$$\mathbf{p}(M') = \sum_{j=1}^{m} p_j = \sum_{j=1}^{m} \frac{s_j}{2V} = \frac{2V}{2V} = 1$$

Thus any equilibrium allocation $\mathbf{x}$ has the property that $M' \subseteq \mathbf{x}_2$. In addition, buyer 2 cannot afford any other item, and so it must be the case that $\mathbf{x}_1 = M \setminus M' = \{m + 1, m + 2\}$ and $\mathbf{x}_2 = M'$ in any equilibrium.

($\Longrightarrow$) If there is a partition $\langle A, B\rangle$ of S, then buyer 1 can afford a better bundle at these prices, namely $Y = \{m + 1\} \cup A$, since:

$$\begin{aligned} v_1(Y) &= v_{1,m+1} + \sum_{j\in A} v_{1,j} = 3V + \sum_{j\in A} s_j \\ &= 4V > 4V - 1 = v_1(\mathbf{x}_1) \end{aligned}$$

and

$$\mathbf{p}(Y) = p_{m+1} + \sum_{j\in A} p_j = \frac{1}{2} + \sum_{j\in A} \frac{s_j}{2V} = \frac{1}{2} + \frac{V}{2V} = 1$$

Thus the market cannot have a competitive equilibrium at $\mathbf{p}$.

($\Longleftarrow$) If the market does not have an equilibrium at $\mathbf{p}$, then it must be the case that in any feasible allocation there exists an improving deviation. Consider the allocation $\mathbf{x}_1 = \{m + 1, m + 2\}$ and $\mathbf{x}_2 = \{1, \ldots, m\}$. Since buyer 1 is already getting its optimal bundle, it follows that buyer 1 has an improving deviation. We consider a few cases:

- If buyer 1 replaces both items $m + 1$ and $m + 2$, then the only bundle it can afford by doing this is $\mathbf{x}_2$ and $v_1(\mathbf{x}_2) < v_1(\mathbf{x}_1)$; thus buyer 1 does not have an improving deviation of this type.

- If buyer 1 replaces item $m+1$ with some subset C of M' then again its utility decreases since:

$$\begin{aligned} v_1(\{m+2\} \cup C) &= V - 1 + v_1(C) \\ &= V - 1 + \sum_{j \in C} s_j \\ &< V - 1 + 2V < 4V - 1 \\ &= v_1(\mathbf{x}_1) \end{aligned}$$

- The only type of deviation left is the one where buyer 1 replaces item $m+2$ with some subset C of M'. Then the only improvements in value can come from bundles worth at least V. That is, there must exist a subset $C \subset M'$ with the property that:

$$\begin{aligned} v_1(C) &= \sum_{j \in C} s_j > v_1(\{m+2\}) = V - 1 \\ &\iff \sum_{j \in C} s_j \geq V \end{aligned}$$

 and

$$\mathbf{p}(C) = \sum_{j \in C} \frac{s_j}{2V} \leq \frac{1}{2} \iff \sum_{j \in C} s_j \leq V$$

 It follows that $\sum_{j \in C} s_j = V$ and so $\langle C, M' \setminus C \rangle$ are a partition of S.

This completes the proof of the theorem. □

Our final proof is the most subtle and included next.

Theorem 49. *Given a market $\mathcal{M} = (N, M, \mathbf{v})$ with indivisible goods, additive valuations, and allocation $\mathbf{x}$, it is coNP-hard to decide if there is a price vector $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is a market equilibrium.*

Proof. We use the NP-complete problem SUBSET-SUM. Given set of positive integers $\mathcal{W} = \{w_1, \ldots, w_m\}$ and integer K, we construct a market $\mathcal{M} = (N, M, \mathbf{v})$ and an allocation $\mathbf{x}$, such that an equilibrium price vector exists at $\mathbf{x}$ if and only if the subset-sum problem does not have a solution.

Let $N = \{1, 2\}$, $M = \{1, \ldots, m+2\}$, allocation $\mathbf{x}$ given by $\mathbf{x}_1 = \{m+1, m+2\}$, $\mathbf{x}_2 = \{1, \ldots, m\}$, and values:

- Buyer 1: $v_{1,m+1} = K - 1$; $v_{1,m+2} = 4\left(\sum_{j=1}^{m} w_j\right)^2$; $v_{1,j} = w_j$, for all $j \in \{1, \ldots, m\}$.
- Buyer 2: $v_{2,m+1} = K + 1$; $v_{2,m+2} = 0$; $v_{2,j} = w_j$, for all $j \in \{1, \ldots, m\}$.

Note we can assume the sum of the numbers in $\mathcal{W}$ is at least K and none is greater than K.

($\Longrightarrow$) If there is a solution S to $\langle \mathcal{W}, K \rangle$, then we claim there can be no market equilibrium. Let $\mathbf{p}$ be any feasible price vector. Then the utility of buyer 1 for bundle S is: $u_1(S) = \sum_{j \in S} v_{1,j} = \sum_{j \in S} w_j = K > K-1 = v_{1,m+1}$.

By the equilibrium property, it must be the case that buyer 1 cannot afford to swap pay for item $m+1$ instead of the set S, and so: $\mathbf{p}(S) > \mathbf{p}_{m+1}$. However, this implies buyer 2 can afford to swap the set S with item $m+1$, and moreover, this is an improving deviation since: $v_{2,m+1} = K + 1 > K = \sum_{j \in S} w_j$. Thus there can be no equilibrium prices.

($\Longleftarrow$) If there is no market equilibrium, then we claim there is a subset-sum solution. To this end, we show that whenever there is no set $S \subseteq U$ such that $\sum_{j \in S} w_j = K$, then a market equilibrium exists. For example, define the next price vector (at which all the budgets are spent):

- $p_j = \frac{w_j}{\sum_{k=1}^m w_k}$, for all $j \in \{1, \ldots, m\}$
- $p_{m+1} = \frac{K-1+\epsilon}{\sum_{k=1}^m w_k}$, where $\epsilon = \frac{1}{4(m+1)^2}$
- $p_{m+2} = 1 - p_{m+1}$

First we claim that buyer 1 does not have a deviation. Note that item $m + 2$ is very valuable, i.e. buyer 2 would never exchange it for any subset of $\{1, \ldots, m\}$. Thus the only remaining type of deviation is one in which buyer 1 exchanges item $m + 1$ for a subset $S \subseteq \{1, \ldots, m\}$. Then it hold that $u_1(S) > v_{1,m+1} = K - 1$, that is, $u_1(S) \geq K$. We have:

$$\begin{aligned} \mathbf{p}(S) - p_{m+1} = \left(\sum_{j \in S} p_j \right) - p_{m+1} \\ = \quad \sum_{j \in S} \left(\frac{w_j}{\sum_{k=1}^m w_k} \right) - \frac{K - 1 + \epsilon}{\sum_{k=1}^m w_k} > 0 \\ \Longleftrightarrow \quad \sum_{j \in S} w_j > K - 1 + \epsilon \end{aligned}$$

The last inequality holds since $u_1(S) \geq K \iff \sum_{j \in S} w_j \geq K > K-1+\epsilon$. Thus bundle S is too expensive for buyer 1 to afford it with the price of item $m + 1$.

Second, the only type of improving deviation of buyer 2 is one in which a set $S \subseteq \{1, \ldots, m\}$ is exchanged for item $m+1$. For this to be an improvement, it must hold that:

$$v_{2,m+1} > u_2(S) \iff u_2(S) \leq K \iff \sum_{j \in S} w_j \leq K$$

Input \ Valuations	**Perfect Complements**	**Perfect Substitutes**
Market $\mathcal{M}$	$\mathcal{P}$	$\mathcal{NP}$-hard
Market $\mathcal{M}$, allocation $\mathbf{x}$	$\mathcal{P}$	co-$\mathcal{NP}$-hard
Market $\mathcal{M}$, prices $\mathbf{p}$	co-$\mathcal{NP}$-complete	co-$\mathcal{NP}$-hard
Market $\mathcal{M}$, allocation $\mathbf{x}$, prices $\mathbf{p}$	$\mathcal{P}$	co-$\mathcal{NP}$-complete

Table 11.1: Summary of the computational results. The market instance is denoted by a tuple $\mathcal{M} = \langle N, M, \mathbf{v} \rangle$, where N is a set of buyers, M a set of indivisible items, and $\mathbf{v}$ the values of the buyers for the items; $\mathbf{x}$ is an allocation of the items to the buyers and $\mathbf{p}$ a price vector.

Since there is no subset-sum solution, we have: $\sum_{j \in S} w_j < K$, and so: $\sum_{j \in S} w_j \leq K - 1$. Equivalently:

$$\mathbf{p}(S) = \sum_{j \in S} \left(\frac{w_j}{\sum_{k=1}^{m} w_k} \right) \leq \frac{K-1}{\sum_{k=1}^{m} w_k} < \frac{K-1+\epsilon}{\sum_{k=1}^{m} w_k} = p_{m+1}$$

It follows that $p_{m+1} > \mathbf{p}(S)$, and so buyer 2 cannot afford to exchange the set S for item $m+1$. Thus neither buyer 1 nor buyer 2 have a deviation, and so $(\mathbf{x}, \mathbf{p})$ is an equilibrium. Then there is a subset-sum solution if and only if the market does not have an equilibrium, which completes the proof. □

Our findings on the complexity of computing a competitive equilibrium from equal incomes for indivisible goods are summarized in Table 11.1.

11.4 Discussion and Future Work

This work leaves several interesting directions for the future. In the case of indivisible resources, it would be interesting to understand the computation of Pareto efficient equilibria and if (or when) the equilibrium outcomes computed by Algorithm 2 can be improved. Moreover, what can be said about the more general Fisher market model with indivisible goods?

Moving to the setting of multiple discrete goods (that come in several copies), it remains to be determined whether and when efficient algorithms can be designed. For the class of Leontief utilities, our characterization from the indivisible setting essentially carries over for two players: "*Given a CEEI market with two buyers and multiple discrete goods, a competitive equilibrium from equal incomes is guaranteed to exist if and only if: there are at least two items* and *the two buyers do not have identical demand sets of size one from an item that comes in an odd number of copies*". However, does there exist a succinct characterization for more than two buyers?

Finally, an interesting open problem is investigating the existence of truthful equilibria in markets with perfect complements and, possibly, other special classes of preferences.

Bibliography

[1] Margareta Ackerman and Simina Brânzei. The authorship dilemma: Alphabetical or contribution? In *Proceedings of the 2014 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1487–1488, 2014. 4

[2] N. Alon. Splitting necklaces. *Advances in Mathematics*, 63:241–253, 1987. 12, 22, 74

[3] Noga Alon, Moshe Babaioff, Ron Karidi, Ron Lavi, and Moshe Tennenholtz. Sequential Voting with Externalities: Herding in Social Networks. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM EC)*, page 36, 2012. 90

[4] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954. 127, 129

[5] A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 114–121, 2007. 112

[6] Y. Aumann, Y. Dombb, and A. Hassidim. Computing socially-efficient cake divisions. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 343–350, 2013. 7, 41

[7] Robert Ayres and Allen Kneese. Production, Consumption, and Externalities. *The American Economic Review*, 59(3):282–297, 1969. 90

[8] H. Aziz and C. Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, 2014. Forthcoming. 22

[9] Haris Aziz. Competitive equilibrium with equal incomes for allocation of indivisible objects. 2015. http://arxiv.org/pdf/1501.06627v1.pdf. 140

[10] Haris Aziz, Simina Brânzei, Aris Filos-Ratsikas, and Søren Kristoffer Stiil Frederiksen. The Adjusted Winner Procedure: Characterizations and Equilibria. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015. 3

[11] Y. Bachrach, E. Porat, and J. S. Rosenschein. Sketching techniques for collaborative filtering. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2016–2021, 2009. 76

[12] A. Badanidiyuru, S. Dobzinski, H. Fu, R. Kleinberg, N. Nisan, and T. Roughgarden. Sketching valuation functions. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1025–1035, 2012. 76

[13] E. Balkanski, S. Brânzei, D. Kurokawa, and A. D. Procaccia. Simultaneous cake cutting. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 566–572, 2014. 3

[14] N. Bansal and M. Sviridenko. The Santa Claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 31–40, 2006. 112

[15] X. Bei, N. Chen, X. Hua, B. Tao, and E. Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1263–1269, 2012. 7, 22

[16] I. Bezáková and V. Dani. Allocating indivisible goods. *SIGecom Exchanges*, 5(3):11–18, 2005. 112

[17] L. Blumrosen and N. Nisan. On the computational power of demand queries. *SIAM Journal on Computing*, 39(4):1372—-1391, 2009. 41

[18] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1321–1328, 2014. 112, 127, 128, 140

[19] William Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, 2000. 127

[20] S. J. Brams and A. D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995. 14, 43, 79

[21] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution.* Cambridge University Press, 1996. 7, 14, 113, 114, 115

[22] S. J. Brams and A. D. Taylor. A procedure for divorce settlements. *Issue Mediation Quarterly Mediation Quarterly*, 13(3):191–205, 1996. 113

[23] S. J. Brams and A. D. Taylor. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody.* Norton, 2000. 113

[24] S. J. Brams and J. M. Togman. Camp david: Was the agreement fair? *Conflict Management and Peace Science*, 15(1):99–112, 1996. 113

[25] S. J. Brams, A. D. Taylor, and W. S. Zwicker. A moving-knife solution to the four-person envy-free cake-division problem. *Proceedings of the American Mathematical Society*, 125(2):547–554, 1997. 35, 43

[26] S. J. Brams, M. A. Jones, and C. Klamler. Better ways to cut a cake. *Notices of the AMS*, 53(11):1314–1321, 2006. 19

[27] S. J. Brams, M. Feldman, J. Morgenstern, J. K. Lai, and A. D. Procaccia. On maxsum fair cake divisions. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1285–1291, 2012. 7, 22, 115

[28] S. Brânzei and P. B. Miltersen. Equilibrium analysis in cake cutting. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 327–334, 2013. 1, 3, 44, 50, 63

[29] S. Brânzei and Peter Bro Miltersen. A Dictatorship Theorem for Cake Cutting. In *Proceedings of the24stInternational Joint Conference on Artificial Intelligence (IJCAI)*, 2015. 3

[30] S. Brânzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013. Forthcoming. 4

[31] S. Brânzei, A. D. Procaccia, and J. Zhang. Externalities in cake cutting. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 55–61, 2013. 3, 7

[32] S. Brânzei, I. Caragiannis, D. Kurokawa, and A. D. Procaccia. An algorithmic framework for strategic fair division. In *Proceedings of the30stInternational Joint Conference on Artificial Intelligence (IJCAI)*, 2016. 1, 3

[33] Simina Brânzei. A note on envy-free cake cutting with polynomial valuations. *Information Processing Letters*, 115(2):93–95, 2015. 3

[34] Simina Brânzei and Ariel D. Procaccia. Verifiably truthful mechanisms. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 297–306, 2015. 3

[35] Simina Brânzei, Tomasz Michalak, Talal Rahwan, Kate Larson, and Nicholas Jennings. Matchings with Externalities and Attitudes. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2013. Forthcoming. 90

[36] Simina Brânzei, Yiling Chen, Xiaotie Deng, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. The fisher market game: Equilibrium and welfare. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI)*, pages 587–593, 2014. 3

[37] Simina Branzei, Hadi Hosseini, and Peter Bro Miltersen. Characterization and Computation of Equilibria for Indivisible Goods". In *Proceedings of the8thInternational Symposium on Algorithmic Game Theory (SAGT)*, 2015. 3, 138

[38] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6), 2011. 112, 127, 128

[39] I. Caragiannis and A. D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9–10):1655–1671, 2011. 76

[40] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, pages 475–482, 2009. 22

[41] I. Caragiannis, J. K. Lai, and A. D. Procaccia. Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 127–132, 2011. 7, 22

[42] K. Chechlárová and E. Pillárová. On the computability of equitable divisions. *Discrete Optimization*, 9(4):249–?257, 2012. 84

[43] Y. Chen, J. K. Lai, D. C. Parkes, and A. D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77:284–297, 2013. 7, 12, 19, 20, 21, 63

[44] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006. 7

[45] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006. 7

[46] 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)*, pages 626–631, 2011. 7, 22, 82, 84, 101

[47] R. Cole, V. Gkatzelis, and G. Goel. Mechanism design for fair division: Allocating divisible items without payments. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 251–268, 2013. 111

[48] M. Dall'Aglio and R. Mosca. How to allocate hard candies fairly. *Mathematical Social Sciences*, 54(3):218—-237, 2007. 118

[49] T. E. Daniel and J. E. Parco. Fair, efficient and envy-free bargaining: An experimental test of the brams-taylor adjusted winner mechanism. *Group Decision and Negotiation*, 14(3):241–264, 2005. 113

[50] X. Deng, Q. Qi, and A. Saberi. Algorithmic solutions for envy-free cake cutting. *Operations Research*, 60(6):1461–1476, 2012. 41, 82

[51] Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. On the complexity of price equilibria. *J. Comput. Syst. Sci.*, 67(2):311–324, 2003. 127

[52] D. B. H. Denoon and S. J. Brams. Fair division: A new approach to the spratly islands controversy. *International Relations*, 2(2):303–329, 1997. 113

[53] S. Dobzinski. Better mechanisms for combinatorial auctions via maximal-in-range algorithms? *SIGecom Exchanges*, 7(1):30–33, 2007. 41

[54] D. Dolev, D. G. Feitelson, J. Y. Halpern, R. Kupferman, and N. Linial. No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, pages 68–75, 2012. 112

[55] L. E. Dubins and E. H. Spanier. How to cut a cake fairly. *American Mathematical Monthly*, 68(1):1–17, 1961. 40

[56] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World.* Cambridge University Press, 2010. 89

[57] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the Generalized Second-Price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007. 44

[58] 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)*, pages 271–278, 2006. 18, 41, 77

[59] J. Edmonds and K. Pruhs. Balanced allocations of cake. In *Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS)*, pages 623–634, 2006. 18, 41

[60] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. volume 30, pages 165–168, 1959. 130

[61] S. Even and A. Paz. A note on cake-cutting. *Discrete Applied Mathematics*, 7:285–296, 1984. 14, 40, 77

[62] U. Feige. On allocations that maximize fairness. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 287–293, 2008. 112

[63] D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967. 125, 127

[64] L. Fortnow. Program equilibria and discounted computation time. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 128–133, 2009. 41

[65] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, pages 24–37, 2011. 112

[66] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973. 2, 63

[67] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. A matroid approach to the worst case allocation of indivisible goods. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, pages 136–142, 2013. 112

[68] M. Guo and V. Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 881–888, 2010. 111

[69] A. Gutman and N. Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 719–728, 2012. 7, 112

[70] Nima Haghpanah, Nicole Immorlica, Vahab Mirrokni, and Kamesh Munagala. Optimal Auctions with Positive Network Externalities. In *Proceedings of the 12th ACM Conference on Electronic Commerce (ACM EC)*, pages 11–20, 2011. 90

[71] L. Han, C. Su, L. Tang, and H. Zhang. On strategy-proof allocation without payments or priors. In *WINE*, pages 182–193, 2011. 111

[72] C. Harris. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica*, 53(3):613–628, 1985. 51

[73] S. Hart and N. Nisan. The query complexity of correlated equilibria. CoRR abs/1305.4874, 2013. 41

[74] H. Hellwig and W. Leininger. On the existence of subgame-perfect equilibrium in infinite-action games of perfect information. *Journal of Economic Theory*, 45:55–75, 1987. 51

[75] T. Hill. Partitioning general probability measures. *Annals of Probability*, 15(2):804—-813, 1987. 112

[76] Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, pages 293–314, 1979. 127

[77] Egor Ianovski. Cake Cutting Mechanisms. *http://arxiv.org/abs/1203.0100*, 2012. 24

[78] J. Karp, A. M. Kazachkov, and A. D. Procaccia. Envy-free division of sellable goods. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 728–734. AAAI Press, 2014. 113, 115

[79] I. Kash, A. D. Procaccia, and N. Shah. No agent left behind: Dynamic fair division of multiple resources. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 351–358, 2013. 7

[80] Michael Katz and Carl Shapiro. Network Externalities, Competition, and Compatibility. *The American Economic Review*, 75(3):424–440, 1985. 90

[81] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999. 123

[82] I. Kremer, N. Nisan, and D. Ron. On randomized one-round communication complexity. *Computational Complexity*, 8(1):21–49, 1999. 75, 76

[83] Piotr Krysta, Tomasz Michalak, Tuomas Sandholm, and Michael Wooldridge. Combinatorial Auctions with Externalities (extended abstract). In *Proceedings of AAMAS*, pages 1471–1472, 2010. 90

[84] D. Kurokawa, J. K. 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)*, pages 555–561, 2013. 7, 18, 20, 41, 63, 85, 106

[85] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1996. 75

[86] Rida Laraki, Eilon Solan, and Nicolas Vieille. Continuous-time games of timing. *Journal of Economic Theory*, 129:206–238, 2005. 23

[87] J. R. Lax. Fair division: A format for the debate on the format of debates. *Political Science and Politics*, 32(1):45–52, 1999. 113

[88] G. M. Levy. Resolving real estate disputes. *Real Estate Issues*, 1999. 113

[89] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004. 82, 105, 112

[90] K. Madani. Game theory and water resources. *Journal of Hydrology*, 381:225—-238, 2010. 113

[91] D. F. Manlove and G. O'Malley. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. In *Proceedings of the 11th International Symposium on Experimental Algorithms (SEA)*, pages 271–282, 2012. 77

[92] E. Markakis and C.-A. Psomas. On worst-case allocations in the presence of indivisible goods. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, pages 278–289, 2011. 112

[93] E. Maskin. Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66:23–38, 1999. First circulated in 1977. 55

[94] T. G. Massoud. Fair division, adjusted winner procedure (aw), and the israeli-palestinian conflict. *Journal of Conflict Resolution*, 44(3): 333–358, 2000. 113

[95] A. Maya and N. Nisan. Incentive compatible two player cake cutting. In *Proceedings of the 8th Conference on Web and Internet Economics (WINE)*, pages 170–183, 2012. 12, 19, 65

[96] Tomasz Michalak, Talal Rahwan, Jacek Sroka, Ann Dowell, Michael Wooldridge, Peter McBurney, and Nicholas R. Jennings. On Representing Coalitional Games with Externalities. In *Proceedings of the 10th ACM Conference on Electronic Commerce (ACM EC)*, pages 11–20, 2009. 90

[97] Tomasz P. Michalak, Piotr L. Szczepanski, Talal Rahwan, Agata Chrobak, Simina Brânzei, Michael Wooldridge, and Nicholas R. Jennings. Implementation and computation of a value for generalized characteristic function games. *ACM Trans. Economics and Comput.*, 2(4): 16, 2014. 4

[98] D. Monderer and M. Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, 2009. 51

[99] E. Mossel and O. Tamuz. Truthful fair division. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 288–299, 2010. 22, 65, 73

[100] H. Moulin. *Fair Division and Collective Welfare.* MIT Press, 2003. 7

[101] S. Muthukrishnan. *Data Streams: Algorithms and Applications.* Now Publishers, 2005. 75

[102] Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58—-73, 1981. 65

[103] B. Nalebuff. Review of the win-win solution: Guaranteeing fair shares to everybody. *Journal of Economic Literature*, 39:125–127, 2001. 113

[104] A. Nicolò and Y. Yu. Strategic divide and choose. *Games and Economic Behavior*, 64(1):268–289, 2008. 41, 44

[105] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, (editors) 2007. 123, 127, 128, 129

[106] A. Othman, C. H. Papadimitriou, and A. Rubinstein. The complexity of fairness through equilibrium. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 209–226, 2014. 112, 127, 128

[107] D. C. Parkes, A. D. Procaccia, and N. Shah. Beyond Dominant Resource Fairness: Extensions, limitations, and indivisibilities. In *Proceedings of the 13th ACM Conference on Economics and Computation (EC)*, pages 808–825, 2012. 112

[108] A. D. Procaccia. Thou shalt covet thy neighbor's cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 239–244, 2009. 7, 18, 41, 86

[109] A. D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56(7):78–87, 2013. 7, 9, 12, 13, 17, 41, 42, 76, 85, 97

[110] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 675–692, 2014. 112

[111] M. G. Raith. Fair-negotiation procedures. *Mathematical Social Sciences*, 39:303—-322, 2000. 113

[112] J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be Fair If You Can.* A. K. Peters, 1998. 7, 9, 12, 13, 14, 35, 40, 41, 43, 98, 105

[113] 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)*, pages 292–300, 2009. 17

[114] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975. 2, 63

[115] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* Cambridge University Press, 2008. 44

[116] F. W. Simmons. *Private communication to Michael Starbird*, 1980. 9

[117] W. Stromquist. How to cut a cake fairly. *American Mathematical Monthly*, 87(8):640–644, 1980. 9, 52

[118] Walter Stromquist. How to cut a cake fairly. *American Mathematical Monthly*, (8):640–644, 1980. Addendum, vol. 88, no. 8 (1981). 613-614. 30

[119] F. E. Su. Rental harmony: Sperner's lemma in fair division. *American Mathematical Monthly*, 106(10):930–942, 1999. 9

[120] A. D. Taylor and A. M. Pacelli. *Mathematics and Politics.* Springer, 2008. 113

[121] M. Tennenholtz. Program equilibrium. *Games and Economic Behavior*, 49(2):363–373, 2004. 41

[122] William Thomson and Hal Varian. Theories of justice based on symmetry. *Social Goals and Social Organizations: Essays in Memory of Elisha Pazner*, 1985. 128

[123] H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9: 63–91, 1974. 111, 127

[124] Rodrigo Velez. Fairness and Externalities. *Working Paper*, 2011. 90, 92

[125] Léon Walras. Elements d'economie politique pure, ou theorie de la richesse sociale (in french). 1874. English translation: Elements of pure economics; or, the theory of social wealth. American Economic Association and the Royal Economic Society, 1954. 127

[126] G. J. Woeginger and J. Sgall. On the complexity of cake cutting. *Discrete Optimization*, 4:213–220, 2007. 17, 18, 41, 64, 65

[127] P. H. Young. *Equity in theory and practice.* Princeton University Press, 1994. 7

[128] R. Zivan. Can trust increase the efficiency of cake cutting algorithms? In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1145–1146, 2011. 22