Fair Division of Land

Erel Segal-Halevi

Department of Computer Science

Ph.D. Thesis

Submitted to the Senate of Bar Ilan University

Ramat-Gan, Israel December 2016

This work was carried out under the supervision of

Prof. Yonatan Aumann

and

Prof. Avinatan Hassidim

Department of Computer Science,

Bar-Ilan University.

I would like to acknowledge
the generous sponsorship of

Drs. Mordecai & Monique Katz

and the

Katz Graduate fellowship program

for support of my
Doctoral degree studies and research

Acknowledgements

This research is a small step in a long journey, and now is a good time to acknowledge and thank all those who accompanied me along my share in this journey. This section is organized mostly in chronological order.

First and foremost, thank God who has given His guidance to us through His commandments in the Bible. Specifically, the commandment to divide the land equally (Numbers 26, Ezekiel 47) was the primary inspiration to this entire research.

Many thanks to my parents Dov and Dina, who raised me and taught me, since I was a child, to love both Bible and science. The education they instilled in me still drives me forward.

The current academic chapter in my life started when Sarit Kraus and Ido Dagan, of the BIU computer science department, accepted me as a research assistant. Later on, they encouraged me to start a Ph.D. research. Although I changed the topic along the way, I benefited a lot from their guidance and support. I also enjoyed a lot to work with colleagues in both labs, particularly: Inon Zuckerman, Avi Rosenfeld, Ariel Roth and Osnat Drein in Sarit's lab, and Meni Adler, Eden Erez, Ofer Bronstein, Yonatan Berant, Eyal Schnarch, Amnon Lotan, Oren Melamud, Lili Kotlerman, Chaya Liebeskind and Hila Weismann in Ido's lab.

The present research was initiated as a seminar project in Nitzotzot Shachar (sparks of dawn) — a multi-disciplinary seminar in Judaism and Science. The course leaders, Rav Shabtay Rappaport and Roly Belfer, encouraged me to pursue this research further, and connected me to Shmuel Nitzan.

Shmuel Nitzan of the BIU economics department introduced me to the classic literature on fairness, and helped me a lot to enter the community of economic theory. His guidance and support are still of great help to me.

I am especially grateful to Yonatan Aumann and Avinatan Hassidim, who agreed to become my Ph.D. advisors in an unusual schedule. They were very supportive of my half-baked ideas, and helped me make them into full-blown results. They guided me through new research topics in computational economics and game theory and introduced me to the community of researchers in these fields. They provided and still provide invaluable guidance and help for my first steps as a researcher.

During my research I met a lot of friends in the Bar-Ilan University. I gained a lot from participating in the BIU game theory seminar and presenting there some parts of this research. I got helpful comments from the participants, including Igal Milchtaich, Ron Peretz, Amnon Schreiber, Gilad Bavli, Ziv Hellman, Eyal Baharad, Shiri Alon-Eron and Gabi Gayer. I received helpful advice from people in the math department, including Ron Adin, Tahl Nowik, Mikhail Katz, Nathan Keller, Reuven Cohen and Simi Haber.

The secretaries, Hila Atia and Daphna Gad and Sylvie Baruch and Sigal Segev-Zrahia of the CS department and Michal Glazerman and Orit Nissim and Efrat Uzan of the economics department, provided all the help that I needed with administrative issues. Nachum Weil and Pinchas Weisberg and Yisrael Yedidya and Doron Kurzberg from the CS department helped me a lot with computational issues.

Many researchers inspired me through their beautiful papers before I met them personally, and even more so after meeting (some of) them. I learned a lot from Haris Aziz, Noga Alon, Reuven Bar-Yehuda, Julius Barbanel, Xiaohui Bei, Silvain Bouveret, Steve Brams, Simina Brânzei, Iannis Caragiannis, Katarina Cechlarova, Marco Dall'Aglio, Ulle Endriss, Shimon Even, Dorothea Herreiner, Michael Jones, Marc Kilgour, Christian Klamler, David Kurokawa, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Herve Moulin, Azaria Paz, Ariel D. Procaccia, Jack M. Robertson, Nisarg Shah, Walter Stromquist, Francis E. Su, Alan D. Taylor, and William A. Webb.

Herve Moulin kindly hosted me in the university of Glasgow and introduced me to interesting new topics in fair division. I enjoyed working with him, as well as with various other friends from abroad, particularly Balázs Sziklai, Arseniy Akopyan, Laszlo Kozma and Emel Ozturk.

I received hundreds of answers to questions in mathematics, economics, computer science, programming and many other issues through the Stack-Exchange websites. I did my best to mention all the people who kindly responded to my queries in the acknowledgments sections of the relevant chapters. I feel fortunate to live in a time where tools such as Wikipedia and Stack-Exchange make international collaboration so easy; not sure how I could have completed this research without them.

My research was funded by various funds, including the Doctoral Fellowships of Excellence Program in BIU, the ISF grants 1083/13 and 1241/12 and 1394/16, the Wolfson Chair and the Mordecai and Monique Katz Graduate Fellowship Program.

Last and most beloved — many thanks to my wife Galya. She gave me a lot of love, support and inspiration ever since we've met, listened and commented on all my practice presentations, provided genuine insights and ideas on my research, accepted all my culinary challenges, and baked wonderful cakes ☺

— Erel Segal-Halevi
— Hod-Hasharon, 2017-06

Contents

Abstract **i**

1 Introduction **1**
1.1 The Land Problem . . . 1
1.2 The Fair Division Solution . . . 1
1.3 Applying Fair Division to Land . . . 2
1.4 Related Work . . . 3

2 Model **6**
2.1 Cake . . . 6
2.2 Agents . . . 6
2.3 Queries . . . 6
2.4 Allocations . . . 7
2.5 Fairness . . . 7
2.6 Geometry . . . 8

3 Geometric Proportional Division **10**
3.1 Introduction . . . 10
3.1.1 Results . . . 11
3.1.2 Techniques . . . 14
3.1.3 Related work . . . 14
3.2 Model . . . 15
3.3 Impossibility Results . . . 16
3.3.1 Impossibility results for two, three and four walls . . . 17
3.3.2 Impossibility results for one wall . . . 18
3.3.3 Impossibility results for zero walls . . . 19
3.3.4 Impossibility results with fat rectangles . . . 22
3.4 Auctions and Covers . . . 24
3.4.1 Mark auction . . . 24
3.4.2 Eval auction . . . 25
3.4.3 Cover numbers . . . 26
3.5 Division procedures . . . 27
3.5.1 Four and three walls, guillotine cuts . . . 27
3.5.2 Two walls . . . 33
3.5.3 One and zero walls . . . 35
3.5.4 Three walls . . . 36
3.5.5 Four walls, guillotine cuts, identical valuations . . . 41
3.5.6 Compact cakes of any shape . . . 46
3.6 Conclusions and Future Work . . . 48
3.6.1 Open questions . . . 48
3.6.2 Different geometric constraints . . . 49
3.6.3 Web implementation . . . 49
3.7 Acknowledgments . . . 49
Chapter appendix . . . 51
3.A Staircase Lemma . . . 51
3.B Non-intersection of Squares in Fat Procedure . . . 52

3.C Existence of Best Pieces 53
3.D Non-Rectangular Pieces 54

4 Geometric Envy-Free Division **57**
4.1 Introduction 58
4.1.1 Results 59
4.1.2 Related Work 60
4.2 Model 61
4.3 Geometric Preliminaries 62
4.3.1 Geometric loss 62
4.3.2 Chooser Lemma 63
4.3.3 Cover Numbers and Cover Lemma 64
4.3.4 Knife functions 65
4.3.5 Geometric loss of knife functions 66
4.4 Envy-Free Division for Two Agents 67
4.4.1 Squares and rectangles 67
4.4.2 Cubes and archipelagos 68
4.4.3 Fat rectangles 69
4.4.4 Arbitrary fat objects 71
4.4.5 Between envy-freeness end proportionality 72
4.5 Envy-Free Division For n agents 73
4.5.1 The one-dimensional procedure 73
4.5.2 Knife tuples 74
4.5.3 Squares and fat rectangles 77
4.5.4 Arbitrary fat objects 78
4.6 Conclusions and Future Work 81
4.7 Acknowledgments 81
Chapter appendix 82
4.A Geometric conditions for S-good knife functions 82
4.A.1 S-smoothness 82
4.A.2 S-continuity 83
4.A.3 Examples 84
4.A.4 Conclusion 85
4.B Dividing a convex fat cake to convex fat pieces 85

5 Re-division and Price-of-Fairness **86**
5.1 Introduction 86
5.1.1 Results 87
5.1.2 Related Work 90
5.2 Model 92
5.2.1 Cake redivision 92
5.2.2 Social-welfare and Price-of-fairness 92
5.3 Arbitrary Cake and Arbitrary Pieces 93
5.4 Interval Cake and Interval Pieces 94
5.5 Polygonal Cake and Polygonal Pieces 96
5.6 Price-of-Fairness Bounds 97
5.7 Conclusions and Future Work 99
5.7.1 Handling other geometric constraints 99
5.7.2 Handling other fairness requirements 99
5.7.3 Improving the constants 100
5.7.4 Price-of-fairness 100
5.8 Acknowledgments 100
Chapter appendix 101
5.A Fair Division of a Rectilinear Polygon 101

6 Family Ownership **102**
6.1 Introduction 102
6.1.1 Results 103
6.1.2 Related Work 103
6.2 Model 104
6.2.1 Families and entitlements 105
6.2.2 Allocations and components 105
6.2.3 Three fairness criteria 106
6.3 Average fairness 106
6.4 Unanimous fairness 107
6.4.1 Exact division 108
6.4.2 UnanimousProp $\implies$ Exact 108
6.4.3 Exact $\implies$ UnanimousProp 109
6.4.4 Less components: equal entitlements 109
6.4.5 Less components: different entitlements 110
6.5 Democratic fairness 111
6.5.1 Two families: a division procedure 111
6.5.2 Three or more families: an impossibility result 112
6.5.3 Three or more families: positive results 112
6.6 Conclusions and Future Work 114
6.6.1 Open questions 114
6.6.2 Alternative fairness criteria 114
6.7 Acknowledgments 115

Bibliography **116**

Hebrew abstract

List of Tables

3.1 Summary of results for square cakes: upper and lower bounds on the level of attainable absolute proportionality. . . . 12
3.2 Summary of results for arbitrary compact cakes: upper and lower bounds on the level of attainable relative proportionality. . . . 13

4.1 Summary of results for geometric envy-free division: upper and lower bounds on the level of attainable proportionality. . . . 61

5.1 Summary of results for cake redivision: ownership and proportionality guarantees. . . . 89
5.2 Summary of price-of-fairness upper bounds 90

6.1 Summary of results for dividing a cake among families: upper and lower bounds on number of cuts 114

List of Figures

2.1 Fatness of several 2-dimensional geometric shapes 9

3.1 With geometric constraints, a proportional allocation might not exist. 11
3.2 A circular cake where all value is near the perimeter. 13
3.3 Impossibility results in a quarter-plane cake. 17
3.4 Impossibility results in a square cake. 17
3.5 Impossibility result for 3 agents on a half-plane. 18
3.6 Cover numbers of various polygons. 26
3.7 A staircase-shaped cake. 33
3.8 Dividing a staircase-shaped cake: the easy case. 34
3.9 Dividing a staircase-shaped cake: the hard case 35
3.10 A valley-shaped cake 37
3.11 The Staircase Lemma 51

4.1 With geometric constraints, a Pareto-efficient envy-free allocation might not exist. 58
4.2 Geometric loss factors relative to the family of squares. 63
4.3 Several knife functions. 65
4.4 Knife functions for a rectangle-union cake 69
4.5 Knife-function for a fat-rectangle cake 70
4.6 Dividing a general R-fat cake to two people. 72
4.7 An illustration of the Simmons-Su procedure. 73
4.8 A 3-knife-tuple on a square cake. 77
4.9 Dividing a general R-fat cake to 3 people. 79
4.10 A fat cake in which every proportional division must use slim pieces. 80
4.11 Square-continuity of a knife-function. 84
4.12 A knife-function that is not S-continuous. 84
4.13 Dividing a convex R-fat cake to two people. 85

5.1 With geometric constraints, a Pareto-efficient division might paradoxically have to discard some of the cake. 88
5.2 A rectilinear polygon with 4 reflex vertexes (circled). 89
5.3 The allocation-completion step: input and output. 96

Notation

The following table can serve as a quick reference for notation used throughout the thesis. For detailed formal definitions, see the model, in Chapter 2 on page 6.

Notation	Meaning
d	**d**imensions. Usually $d = 2$ since we divide a two-dim. resource.
C	**C**ake, the thing that should be divided fairly. A subset of $\mathbb{R}^d$.
R	upper bound on length/width **R**atio of a geometric object.
S	**S**et of pieces that are considered usable, e.g. squares, R-fat objects.
T	number of reflex ver**T**exes of a rectilinear polygon.
n	number of age**n**ts (people) entitled to receive a piece.
i	**i**ndex of an agent. $i \in \{1, \ldots, n\}$.
v_i	**v**alue-density of agent i; function from C to $\mathbb{R}$.
V_i	**V**alue-measure of agent i (integral of v_i).
V_i^S	Utility-function of agent i who can use only pieces from S.
X_i	piece of cake allocated to agent i.
X	cake allocation; n pairwise-disjoint pieces: $X = (X_1, \ldots, X_n)$.
Y, Z	alternative allocations.
$\text{Prop}(C, S, n)$	Proportionality function; see Chapter 3.
$\text{PropEF}(C, S, n)$	Envy-free-proportionality function; see Chapter 4.

Abstract

This research presents algorithms for fair division of land. The algorithms take as input a heterogeneous land-estate, and several people with different preferences over parts of the land-estate. They return as output a partition of the land-estate among the people, such that each person agrees that his/her share is "fair".

The baseline of this research is the classic problem of *fair cake-cutting*. There are many algorithms that take as input a heterogeneous cake and several people with different tastes, and give each person a piece the he/she considers "fair". However, these algorithms cannot be directly applied to fair division of land, since land is not a cake. There are several differences between land and cake, and they require new fair division algorithms.

The first difference is **geometry**. When a cake is divided, the geometric shape of the pieces is usually ignored. It is often assumed that the cake is a one-dimensional interval and that the pieces are sub-intervals or finite collections thereof. In contrast, when land is divided, the two-dimensional geometric shape of the pieces is of crucial importance. We present fair division algorithms that can handle multi-dimensional geometric constraints on the pieces. In particular, we present algorithms that guarantee that each piece is a square, a fat rectangle (a rectangle with a bounded length/width ratio) or an arbitrary fat object. We give upper and lower bounds on the degree of "fairness" (the value guarantee per agent) as a function of the geometric constraints.

The second difference is **redivision**. A cake is usually divided when it is fresh and new, so that no people have a previous claim on it. In contrast, many land-resources are already divided, and it is often required to re-divide them, as in a land-reform. We present algorithms for fair re-division, which balance the ownership rights of existing land-owners and the fairness claims of landless citizens. We first present a baseline algorithm for redivision without geometric constraints. Then, we combine the redivision model with the geometry model and present a redivision algorithm that can also handle one-dimensional and two-dimensional geometric constraints.

The redivision algorithms have implications on another important issue in fair division — the trade-off between fairness and efficiency. Our redivision algorithms allow us to prove upper bounds on the *price-of-fairness* — the loss of efficiency due to fairness considerations — with geometric constraints.

The third difference is **group ownership**. A piece of cake is usually eaten by a single person. In contrast, a plot of land is usually owned by a group, such as a family or a community. Different family members may have different preferences. We present algorithms for fair division that respect the different tastes of group-members.

While the primary focus of this research is land division, the concepts introduced herein are applicable in other division problems. Geometric considerations are relevant when dividing other two-dimensional resources, such as advertisement space in print or electronic media. Redivision considerations are relevant in other dynamic division problems, such as dividing computation resources among processes. Group ownership considerations are relevant also in the classic economic setting of dividing homogeneous resources among families. We believe the present research will enrich the general fair division literature by adding these new considerations.

Chapter 1

Introduction

1.1 The Land Problem

This research has been motivated by a pressing social problem — the rising prices of housing in Israel. Israeli youth find it more and more difficult to afford a house, and this problem is largely related to their inability to own land.[1]

Land division is not only a problem of the present. It has been an important issue since Biblical times. This is evident from the commandment to divide the land of Israel among the tribes in proportion to their size (Numbers 26:53-54), through the protests of the prophets against unfair land allocation (e.g. Isaiah 5:8), to the latter-day prophecies describing a futuristic fair land division (e.g. Ezekiel 47:14).

Land division is, of course, not only an Israeli problem. It has been an important issue all around the globe. Fair division of land has been the goal of numerous land reforms carried out in all five continents throughout history.[2] The earliest recorded land-reform attempts were done by Egyptian king Bakenranef in 8th century BC (Powelson, 1988). The latest such attempt was done by the Scottish government in 2016 AD.[3]

1.2 The Fair Division Solution

This research studies fair division of land from the perspective of a computer scientist. Its goal is to develop algorithms for fair division of land. The input to such an algorithm is a land-estate that has to be divided among several people. The goal is that all people agree that the division is “fair”. When I tell people about this goal, their immediate reply is:

> That’s impossible! Different people have different tastes. Some people might claim that a fair division should give each person access to the *road*; others might claim that that you must give each person the same area of *seashore*; yet others might claim that you should give each person the same probability of finding *oil*; there are as many opinions as there are people. How can you hope to find a division that will be conceived as fair by *everyone*?

They are quite surprised when I tell them the following 4-word algorithm:

I cut. You choose.

This algorithm is so simple, that it is even used by children to divide a birthday-cake. [4] It does not

[1] This is illustrated by the fact that people who own land can build a home in less than one year of labor (MarkerWeek 14.7.2013, http://www.themarker.com/markerweek/1.2069919), in contrast to over 8 years that are required to buy a house without owning land (Calcalist 26.02.15, http://www.calcalist.co.il/real_estate/articles/0,7340,L-3653354,00.html . See also BizPortal 4.4.2013 http://www1.bizportal.co.il/article/356140). Retrieved 21.11.16.

[2] See the Wikipedia page “Land reforms by country” for more details

[3] Land Reform (Scotland) Act 2016, http://www.gov.scot/Topics/Environment/land-reform retrieved 21.11.16

[4] This algorithm is already alluded to in the Bible (Genesis 13:9): when Abraham and Lot wanted to divide the land of Canaan between them, Abraham suggested a division of the land to two parts, and let Lot be the first to choose his part. See Isaac Dov Paris, "fairness and justice in dividing property" (Hebrew), http://www.daat.ac.il/mishpat-ivri/skirot/143-2.htm retrieved 21.11.16.

require any details about the land. It does not need to know the location of the road, nor the amount of seashore, nor the probability of finding oil, nor any other particular feature of the land. It can be implemented by any two people on their own — they do not have to employ an expert (and expensive) real-estate assessor.

Despite its simplicity, this algorithm can be called "fair". To see why it is fair, suppose first that the cutter divides the land to two pieces that are equal in his eyes, and that the chooser chooses the piece that is better in his eyes. The resulting division has two properties:

1. Each person receives a piece that is worth at least 1/2 of the total value, according to his own taste. This property is called **proportionality**.

2. Each person receives a piece that is at least as good as the other piece, according to his own taste. This property is called **envy-freeness**.[5]

What happens if one of the participants does not follow the rules? In that case, the division is still "fair" for the other participant. For example, if the cutter "breaks the rules" by cutting the cake to two unequal pieces, then the chooser can still follow the rules and pick the piece that is better in his eyes, so his piece is still worth at least 1/2 of the total value and at least as much as the other piece. Similarly, if the chooser breaks the rules by choosing the piece that is worse in his eyes, then the cutter (who followed the rules) still receives a piece that is worth exactly 1/2 of the total value and exactly the same as the other piece. When you use this algorithm, you receive a personal fairness guarantee, that does not depend on what the other person does: as long as *you* keep the rules, *you* are guaranteed a fair share.

Proportionality and envy-freeness are two very natural definitions of fairness, so it is nice that such a simple algorithm "I cut, you choose" can guarantee both of them.

The success of the "I cut, you choose" algorithm for two people naturally invokes the question: what happens when there are more than two people? The first person to ask this question was the Jewish-Polish mathematician Hugo Steinhaus. After World War II, he posed this question to two of his students, Banach and Knaster. They developed an algorithm that finds a *proportional* division for n people. Steinhaus published their algorithm (Steinhaus, 1948). This publication initiated a new field of research, that is now called: *fair division*.

Since then, the fair division problem has been studied by many researchers from different disciplines: mathematicians, economists, computer scientists and political scientists. Each discipline has brought its own questions and answers. Some of the interesting questions are: how can we find an *envy-free* division for n people? What is the *runtime complexity* of fair division (how many queries are required)? What happens if some of the items to divide are *indivisible*? What if the items have *negative value* (like a piece of lawn that has to be mowed)? What are the *strategic properties* of fair division algorithms, when they are viewed as competitive games? Is it possible to attain a fair division that is also *economically efficient*? What if people have *different entitlements*? What if there are *externalities* between the agents? What if the agents can form *coalitions*? And so on (see the Related Work section below for some references).

A particularly exciting recent development is the launching of several websites that let visitors apply state-of-the-art fair division algorithms to their own problems: FairOutcomes[6], The Fair Division Calculator[7] and Spliddit.[8] The latter website has been used by tens of thousands of visitors (Goldman and Procaccia, 2015), demonstrating the practical usefulness of fair division algorithms.

1.3 Applying Fair Division to Land

In the world of fair division, there are many different kinds of problems, depending on whether the resources to divide are homogeneous or heterogeneous, divisible or indivisible, and so on. We are interested of fair division of land, which is heterogeneous and divisible. The sub-field of fair division that handles heterogeneous and divisible resources is called *fair cake-cutting*. This uses the metaphor of Steinhaus (1948), of cutting a birthday cake among several siblings with different tastes. Indeed, many researchers explicitly mention land division as an important application of fair cake-cutting (e.g. Berliant and Raa, 1988; Berliant

[5]When there are two people and the entire land is divided, proportionality and envy-freeness are equivalent. This is not the case when there are more than two people or when some of the land is left undivided, so these two properties are independent.

[6]http://www.fairoutcomes.com/

[7]https://www.math.hmc.edu/s̃u/fairdivision/calc/

[8]http://www.Spliddit.org

et al., 1992; Legut et al., 1994; Chambers, 2005; Dall'Aglio and Maccheroni, 2009; Hüsseinov, 2011; Nicolò et al., 2012).

However, the cake metaphor might be misleading. Land is not a cake, and indeed existing cake-cutting algorithms have several shortcomings that make them impractical for division of land. This is the main motivation for the present research. Our goal is to develop new division algorithms, that handle the considerations that are important in land division.

We focus on three topics — three main differences between cakes and lands.

1. Geometry. When a cake is divided, the geometric shape of the pieces is usually ignored. It is often assumed that the cake is a one-dimensional interval and that the pieces are sub-intervals or finite collections thereof. In contrast, when land is divided, the two-dimensional geometric shape of the pieces is of crucial importance. This work extends the cake-cutting model to handle multi-dimensional cakes. It presents fair division algorithms that can handle multi-dimensional geometric constraints on the pieces. Due to its length, this part is divided to two chapters:

- Chapter 3 focuses on the simpler fairness requirement — *proportionality* — each agent is guaranteed a piece worth for him at least a given fraction of the total cake value. It presents algorithms that guarantee that the pieces are *squares or fat rectangles* (rectangles with a balanced length/width ratio).
- Chapter 4 adds the second fairness requirement — *envy-freeness* — each agent is guaranteed a piece worth for him at least as much as any other piece. It presents algorithms that guarantee that the pieces are squares or fat rectangles, but can also handle more general geometric constraints such as arbitrary fat pieces.

An interesting aspect of this work is the combination of different disciplines: computer science, geometry and economics. Indeed, parts of this work have appeared in preliminary forms in the AAAI 2015 conference (Segal-Halevi et al., 2015a) and EuroCG 2016 conference, and are now under revision for the Journal of Mathematical Economics.

2. Redivision. A cake is usually divided when it is fresh and new, so that no people have a previous claim on it. In contrast, many land-resources are already divided, and it is often required to re-divide them, as in a land-reform. In Chapter 5 we present algorithms for fair re-division, which balance the rights of existing land-owners and those of new landless citizens. We first present a baseline algorithm for redivision without geometric constraints. Then, we combine the redivision model with the geometry model and present a redivision algorithm that can also handle one-dimensional and two-dimensional geometric constraints.

The redivision algorithms have implications on another important issue in fair division — the trade-off between fairness and efficiency. Our redivision algorithms allow us to prove upper bounds on the *price-of-fairness* — the loss of efficiency due to fairness considerations — with geometric constraints.

3. Family ownership. A piece of cake is usually eaten by a single person. In contrast, a plot of land is usually owned by a group, such as a family or a community. Different family members may have different preferences. In Chapter 6 we present algorithms for fair division that respect the different tastes of group members.

While the primary focus of this research is land division, the concepts introduced herein are applicable in other division problems. Geometric considerations are relevant when dividing other two-dimensional resources, such as advertisement space in print or electronic media. Redivision considerations are relevant in other dynamic division problems, such as dividing computation resources among processes. Group ownership considerations are relevant also in the classic economic setting of dividing homogeneous resources among families. We believe the present research will enrich the general fair division literature by adding these new considerations.

1.4 Related Work

Fair division has greatly evolved since the days of Steinhaus, with hundreds of research papers and several books (Brams and Taylor, 1996; Robertson and Webb, 1998; Moulin, 2004; Barbanel, 2005; Brams, 2007;

Brânzei, 2015). A comprehensive survey of this literature is beyond the scope of this work, but to illustrate the diversity of the fair division research, we present some of its questions below.

1. A long-standing open question was how to find an *envy-free* cake-cutting for n people. The Banach-Knaster algorithm from the forties guarantees a division that is proportional but not necessarily envy-free. Envy-free division turned out to be much more difficult. It was solved only in the nineties. Three different algorithms find an envy-free division with *disconnected pieces* in *finite but unbounded time* (Brams and Taylor, 1995; Robertson and Webb, 1998; Pikhurko, 2000). A fourth algorithm finds an envy-free division with *connected pieces* but in *infinite time*. Brams et al. (2011) prove that the divide-and-conquer algorithm of Even and Paz (1984), while not guaranteeing envy-freeness, minimizes the maximum number of players that any single player can envy (the minimum taken over a family of algorithms for proportional cake-cutting).

2. Computer scientists have been mainly interested in the *computational complexity* of cake-cutting:

- How many queries are required to find a *proportional* cake-cutting? The Banach-Knaster algorithm uses $O(n^2)$ queries, but a later algorithm by Even and Paz (1984) requires only $O(n \log n)$ queries. Moreover, recent results by Edmonds and Pruhs (2006b); Woeginger and Sgall (2007) show that this is asymptotically optimal.

- How many queries are required for *envy-free* cake-cutting? Stromquist (2008) proved that an infinite number of queries may be required when the pieces are connected; Procaccia (2009) proved that $\Omega(n^2)$ queries may be required when pieces may be disconnected. Gasarch (2015) presented a comparison among the three unbounded procedures for envy-free fair division. Very recently, Aziz and Mackenzie (2016) published the first bounded-time algorithm for envy-free cake-cutting (with disconnected pieces). Their algorithm requires $n^{n^{n^{n^{n^{n}}}}}$ queries, much more than the upper bound of n^2, so there is still a lot of room for improvement. We made a modest contribution to this line of research by presenting quicker algorithms for envy-free cake-cutting, with either connected or disconnected pieces, when it is allowed to leave some cake unallocated (Segal-Halevi et al., 2015b).

3. The strategic aspects of cake-cutting have attracted the attention of researchers in algorithmic mechanism design:

- Are there fair cake-cutting algorithms that are also *truthful*, meaning that an agent always receives the highest possible value by playing according to his true valuations? "I cut, you choose" is not truthful, since an agent who knows the other agent's preferences may get a better piece by playing untruthfully. This will not damage the fairness guarantee to the other agent, but it might encourage the agents to spy on each other, which is undesirable. Recently, Chen et al. (2013); Aziz and Ye (2014) showed truthful algorithms for the special cases in which the agents' valuations are *piecewise-uniform* or *piecewise-constant*. In contrast, Brânzei and Miltersen (2015) showed that in the general case, every truthful query-based algorithm might give one of the agents a worthless piece, so there is a fundamental conflict between truthfulness and fairness.

- What happens when agents can not only lie about their preferences, but also create duplicates? Tsuruta et al. (2015) study the notion of *false-name-proof* mechanisms for cake-cutting.

- How do agents behave when they play a non-truthful cake-cutting algorithm? Is there a Nash equilibrium, and what are its properties? Brânzei et al. (2016) present a framework for studying this question and give some answers.

4. Economists have been mainly interested in the *economic efficiency* of cake-cutting:

- The fundamental definition of economic efficiency is *Pareto-efficiency* — there is no allocation which is better for one person and not worse for another one. Varian (1974) proved that, under fairly general conditions, there exists a Pareto-efficient and envy-free division of homogeneous resources. Weller (1985) proved a similar resource for a cake — a heterogeneous resource. Barbanel (2005) presented alternative proofs. Reijnierse and Potters (1998) showed how to (approximately) find a Pareto-efficient envy-free cake division. These results give each agent a disconnected piece. If each agent must get a connected piece, then an envy-free and Pareto-efficient division might not exist (Stromquist, 2007). What if each agent may get a union of two connected pieces? This question is still open.

- Another measure of economic efficiency is the *social welfare* — usually defined as the sum of utilities of all agents (Bentham, 1789; Mill, 1863). It is possible to calculate an allocation that approximates the maximum welfare (Aumann et al., 2013), but this allocation might not be fair. In contrast, a fair allocation might have a low social welfare. This invokes the question of what is the "price of fairness" — how much does society have to pay, in terms of welfare, for the different fairness requirements? Aumann and Dombb (2015) and Caragiannis et al. (2012) and Arzi (2012) study this question in various settings. Finally, Cohler et al. (2011); Bei et al. (2012) show how to calculate an allocation with optimal social welfare subject to fairness requirements.

5. What happens when there are indivisible items? The "I cut, you choose" algorithm assumes that the cake can always be divided to halves without losing value, but what if there are houses or trees, that cannot be divided? The papers on this topic are far too many to mention; see Bouveret et al. (2016) for a recent survey.

6. What happens when different agents have different entitlements? McAvaney et al. (1992) and Robertson and Webb (1995) present some solutions, but they require a large number of cuts; it is still open whether such weighted-fair divisions can be found using a smaller number of cuts.

7. What happens when there are externalities, i.e, the utility of an agent depends on the pieces allocated to other agents? See Brânzei et al. (2013). What happens when agents can cooperate and form coalitions before the division process? See Dall'Aglio et al. (2009).

In each of the following chapters, we present work that is more closely related to that chapter.

Chapter 2

Model

This chapter presents general definitions and notations that are applicable in all following chapters. Each of the following chapters will contain additional definitions and notations specific to that chapter.

2.1 Cake

The object that should be divided is called a *cake*, and denoted by C. In the cake-cutting literature, it is often assumed that C is a one-dimensional interval. In this work, we will usually assume that C is a polygon in the Euclidean plane $\mathbb{R}^2$, but we will also consider more general cakes that are objects in a d-dimensional Euclidean space $\mathbb{R}^d$.

Pieces of C are Borel subsets of C — the subsets that can be formed from open subsets through the operations of countable union, countable intersection, and relative complement.

2.2 Agents

The people among whom the cake should be divided are called *agents*. There are n agents, where $n \geq 1$.

Each agent $i \in \{1, \ldots, n\}$ has a *value-density* function v_i, which is a real, integrable, bounded and non-negative function on C. Value-density functions are common in real-estate assessments, for example, it is common to say that "in neighborhood X, the house prices are \$2000 per square meter." However, it should be emphasized that each agent has a personal value-density function, that need not be related to the market prices. Different agents may have different value-density functions; this is what makes the fair division problem interesting.

The *value* of a piece X_i to agent i is marked by $V_i(X_i)$ and it is the integral of its value-density:

$$V_i(X_i) = \int_{x \in X_i} v_i(x) dx$$

Even when C is unbounded, we assume that the v_i have finite support — they are nonzero only in a bounded subset of C. Hence the V_i are always finite.

The definition implies that the V_i are *measures*. In particular, they are countably additive: when a piece is divided to parts (even countably many parts), the value of the piece equals the sum of the values of its parts.

Moreover, the definition implies that the V_i measures are *absolutely continuous with respect to the Lebesgue measure*, i.e., any piece with zero Lebesgue measure (length, area, etc.) has zero value to all agents. This implies, in particular, that from any piece with a value of V, we can cut a sub-piece with a value of $\alpha \cdot V$, for any fraction $\alpha \in [0, 1]$. This assumption is already implicitly made by the "I cut, you choose" algorithm — it assumes that the cutter can cut the cake into two pieces with a value of exactly half the original value. Now we have a formal model that justifies this assumption.

2.3 Queries

The division protocols access the value measures via *queries* (Robertson and Webb, 1998): an *eval query* asks an agent to reveal its value for a specified piece of cake; a *mark query* asks an agent to mark a piece of cake with a specified value.

In this work, we ignore strategic considerations and assume that all agents answer truthfully. As usual in the cake-cutting literature since Steinhaus (1948), the fairness guarantees of our algorithms are valid for every *single* agent answering the queries truthfully, regardless of the behavior of the other agents. This is the common practice in the cake-cutting world.[1]

However, our protocols are not *dominant-strategy-truthful*, i.e, an agent may gain by answering untruthfully. Designing dominant-strategy-truthful mechanisms for cake-cutting is known to be a difficult problem even in one dimension (Brânzei and Miltersen, 2015).

As an alternative to queries, it may be more convenient to let agents submit their entire value-density function to a referee, who will calculate the division for them.[2] One way to do this is illustrated in a simple web-page that I built to demonstrate some of the algorithms for two-dimensional land division (Chapter 3).[3] The web-page lets the agents arrange points on the cake. Each point represents a certain amount of value. So, each agent can put more points in areas that are more valuable in his/her eyes. Once the points are arranged, the application automatically uses them to calculate the valuation functions and the fair division, without the need to query the agents.

2.4 Allocations

An *allocation* is a vector of n pieces, $X = (X_1, \ldots, X_n)$, one piece per agent, such that the X_i are pairwise-disjoint and contained in C. We express the latter two facts succinctly using the "disjoint union" operator, $\sqcup$:

$$X_1 \sqcup \cdots \sqcup X_n \subseteq C$$

The above definition implies that some cake may remain unallocated, i.e, *free disposal* is assumed. This is a reasonable assumption in land division: it is usually allowed, and often even desired, to leave some public lands unallocated.

2.5 Fairness

There are two common definitions of fairness. Both of them are natural generalizations of the guarantees of the "I cut, you choose" protocol described in the introduction.

1. Proportionality. Traditionally, a division X is called **proportional** if each agent receives at least $1/n$ of the total cake value, according to its personal valuation:

$$\forall i \in \{1, \ldots, n\} : \qquad V_i(X_i) \geq V_i(C)/n$$

In this work, we will often have to relax the proportionality requirement and require *partial-proportionality*.

Definition 2.5.1. The *proportionality* of a division X is defined as the largest value p such that:

$$\forall i \in \{1, \ldots, n\} : \qquad V_i(X_i) \geq p \cdot V_i(C)$$

Equivalently, the proportionality of X is:

$$Prop(X) := \min_{i=1}^{n} V_i(X_i)/V_i(C)$$

By this definition, X is proportional if-and-only-if $Prop(X) \geq 1/n$.

[1] In the words of Steinhaus (1948): "The greed, the ignorance, and the envy of other partners cannot deprive him of the part due to him in his estimation; he has only to keep to the methods described above. Even a conspiracy of all other partners with the only aim to wrong him, even against their own interests, could not damage him."

[2] We are grateful to a referee for suggesting this option.

[3] http://tora.us.fm/geometry/fair-division.html

2. Envy-freeness A division X is called **envy-free** if each agent receives at least as much as any other agent, according to its personal valuation. Formally:

Definition 2.5.2. A division X is called envy-free if

$$\forall i, j \in \{1, \ldots, n\} : \qquad V_i(X_i) \geq V_i(X_j)$$

2.6 Geometry

The geometric constraints on the pieces (if any) are represented by a set of *usable pieces*, which is denoted by S. For example, S may be the family of intervals, rectangles or squares. An element of S is called an S-piece. We assume that each agent can use only a single S-piece.[4] An allocation X where for every i, $X_i \in S$, is called an *S-allocation*.

Based on the value measures V_i and the geometric family S, the fair land division problem can be formulated in two equivalent ways.

1. Geometry is an external restriction. The division algorithms must return only S-allocations. So, for example, an envy-free land-division algorithm should give each agent i an S-piece X_i such that $\forall i, j$: $V_i(X_i) \geq V_i(X_j)$.

2. Geometry is a part of the agents' utility functions. An agent can derive utility only from an S-piece; when his allotted land-plot is not an S-piece, he selects the most valuable S-piece contained therein and utilizes it. For each agent i, we define the *S-value function*, which assigns to each piece X_i the value of the most valuable usable piece contained in it:

$$V_i^S(X_i) = \sup_{s \in S\ ,\ s \subseteq X_i} V_i(s)$$

Now, the division algorithms may return any allocation, but the fairness guarantees are judged according to the agents' S-value functions. So, for example, an envy-free land-division algorithm should give each agent i a piece X_i such that $\forall i, j : V_i^S(X_i) \geq V_i^S(X_j)$. Note that, in contrast to the V_i that are measures, the V_i^S are usually not measures since they are not additive. This means that cake-cutting algorithms that require additivity cannot be used.

The above two formulations are equivalent and we will use them interchangeably.

Fatness

In land division, it is often preferred that the pieces will have a balanced length/width ratio — not too long in one dimension and too short in another dimension. This preference is captured by the the concept of *fatness*, which we adapt from the computational geometry literature, (e.g. Agarwal et al., 1995; Katz, 1997):

Definition 2.6.1. Let $R \geq 1$ be a real number. A d-dimensional piece is called *R-fat*, if it contains a d-dimensional cube B^- and is contained in a parallel d-dimensional cube B^+, such that the ratio between the side-lengths of the cubes is at most R: $\text{len}(B^+)/\text{len}(B^-) \leq R$.

A 2-dimensional cube is a square. So, for example, the only 2-dimensional 1-fat shape is a square. An $L \times 1$ rectangle is L-fat, a right-angled isosceles triangle is 2-fat and a circle is $\sqrt{2}$-fat (see Figure 2.1).

Note that R is an upper bound, so if $R_2 \geq R_1$, every R_1-fat piece is also R_2-fat. So a square is also e.g. 2-fat and 3-fat, but a 10-by-20 rectangle is not 1-fat.

[4] If we want to allow each agent to use e.g. two squares, then we can just define S to be the family of all square-pairs. So the assumption of one S-piece per agent does not lose generality.

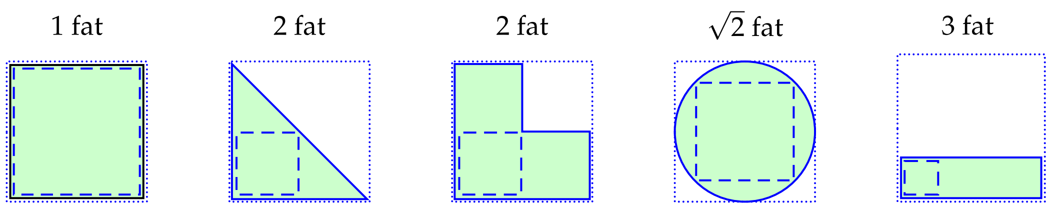

Figure 2.1: Fatness of several 2-dimensional geometric shapes. The dashed square is the largest contained cube; the dotted square is the smallest containing parallel cube. The shape is R-fat if the ratio of the side-lengths of these squares is at most R.

Chapter 3

Geometric Proportional Division

Contents

2.1 Cake 6
2.2 Agents 6
2.3 Queries 6
2.4 Allocations 7
2.5 Fairness 7
2.6 Geometry 8

This chapter was presented in the EuroCG 2016 conference and is now under revision for the Journal of Mathematical Economics.

3.1 Introduction

In many cake-cutting papers, it is assumed that the cake is a one-dimensional interval and the pieces are sub-intervals. This assumption is usually justified by the reasoning that higher-dimensional settings can always be projected onto one dimension, and hence fairness in one dimension implies fairness in higher dimensions.[1] However, projecting back from the one dimension, the resulting two-dimensional plots are thin rectangular slivers, of little use in most practical applications; it is hard to build a house on a $10 \times 1,000$ meter plot even though its area is a full hectare, and a thin 0.1-inch wide advertisement space would ill-serve most advertisers regardless of its height.

We claim that the *two-dimensional shape* of the allotted piece is of prime importance. Hence, we seek divisions in which the allotted pieces must be of some restricted family of "usable" two-dimensional shapes, e.g. squares or polygons of balanced length/width ratio.

Adding a two-dimensional geometric constraint re-opens most questions and challenges related to cake-cutting. Indeed, even the elementary proportionality criterion can no longer be guaranteed.

Example 3.1.1. A homogeneous square land-estate has to be divided between two heirs. Each heir wants to use his share for building a house with as large an area as possible, so the utility of each heir equals the area of the largest house that fits in his piece (see Figure 3.1). If the houses can be rectangular, then it is possible to give each heir 1/2 of the total utility (a); if the houses must be square, it is possible to give each heir 1/4 of the total utility (b) but impossible to give both heirs *more* than 1/4 the total utility (c). In particular, when the allotted pieces must be square, a proportional division does not exist.[2]

This example invokes several questions. What happens when the land-estate is heterogeneous and each agent has a different utility function? Is it always possible to give each agent a 2-by-1 rectangle worth for him at least 1/2 the total value? Is it always possible to give each agent a square worth for him at least 1/4 the total value? Is it even possible to guarantee a positive fraction of the total value? If it is possible, what

[1] In the words of Woodall (1980): "the cake is simply a compact interval which without loss of generality I shall take to be [0,1]. If you find this thought unappetizing, by all means think of a three-dimensional cake. Each point P of division of my cake will then define a plane of division of your cake: namely, the plane through P orthogonal to [0,1]".

[2] Berliant and Dunz (2004) use a very similar example to prove the nonexistence of a competitive equilibrium when the pieces must be square.

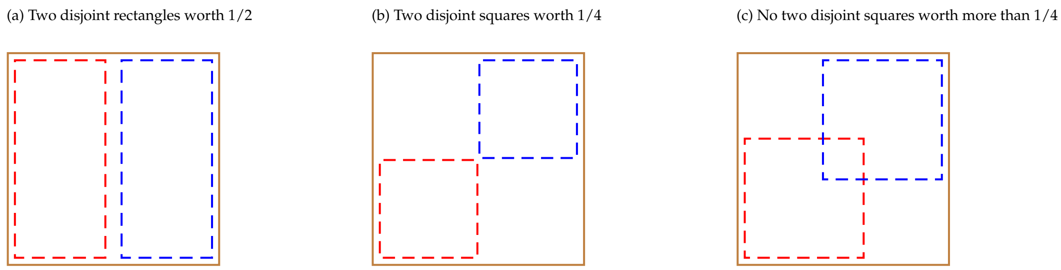

Figure 3.1: With geometric constraints, a proportional allocation might not exist.

division procedures can be used? How does the answer change when there are more than two agents? Such questions are the topic of the present chapter.

We use the term *proportionality* to describe the fraction that can be guaranteed to every agent. So when the shape of the pieces is unrestricted, the proportionality is always $1/n$, but when the shape is restricted, the proportionality might be smaller. Naturally, the attainable proportionality depends on both the shape of the cake and the desired shape of the allotted pieces. For every combination of cake shape and piece shape, one can prove *impossibility results* (for proportionality levels that cannot be guaranteed) and *possibility results* (for the proportionality that can be guaranteed). While we examined many such combinations, the present chapter focuses on several representative scenarios which, in our opinion, demonstrate the richness of the two-dimensional cake-cutting task.

Walls and unbounded cakes

In Example 3.1.1, the two pieces had to be contained in the square cake. One can think of this situation as dividing a square island surrounded in all directions by sea, or a square land-estate surrounded by 4 walls: no land-plot can overlap the sea or cross a wall.

In practical situations, land-estates often have less than 4 walls. For example, consider a square land-estate that is bounded by sea to the west and north but opens to a desert to the east and south. Allocated land-plots may not flow over the sea shore, but they may flow over the borders to the desert.

Cakes with less than 4 walls can also be considered as unbounded cakes. For example, the above-mentioned land-estate with 2 walls can be considered a quarter-plane. The total value of the cake is assumed to be finite even when the cake is unbounded. When considering unbounded cakes, the pieces are allowed to be "generalized squares" with an infinite side-length. For example, when the cake is a quarter-plane (a square with 2 walls), we allow the pieces to be squares or quarter-planes. When the cake is a half-plane (a square with 1 wall), we also allow the pieces to be half-planes, etc. The terms "square with 2 walls" and "quarter-plane" are used interchangeably throughout the chapter.

3.1.1 Results

Our results can be broadly summarized as follows.

- **Negative results:** when the pieces have to be squares or fat rectangles, a proportional division is usually [3] not guaranteed to exist. Moreover, there is a small constant $A > 1$ that depends on the shape of the cake and usable pieces, such that the largest value that can be guaranteed to all agents is $1/(A \cdot n)$.

- **Positive results:** when the pieces have to be squares or fat rectangles, a constant-factor approximation to a proportional division is usually guaranteed to exist: there is a small constant $B > 1$ that depends on the shape of the cake and usable pieces, such that all agents can be guaranteed a value of at least $1/(B \cdot n)$.

[3] We have proved this for most, but not all the cases that we have studied. The exception is when the cake is an unbounded plane and the pieces are non-parallel squares: in this case, we do not know whether a proportional division always exists. See Table 3.1 below.

Cake ↓		Impossibility		Possibility	
		Square pieces	R-Fat rects ($R \geq 2$)	Square pieces	R-Fat rects ($R \geq 2$)
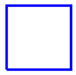	4 walls (Square)	$1/(2n)$ *	$1/(2n-1)$	$1/(4n-4)$ * *same*: $1/(2n)$ *	$1/(4n-5)$ *same*: $1/(2n-1)$
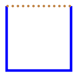	3 walls	$1/(2n-1)$		$1/(2n-1)$	
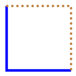	2 walls (quarter-plane)	$1/(2n-1)$		$1/(2n-1)$	
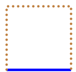	1 wall (half-plane)	$1/(\frac{3}{2}n-1)$		$1/(2n-2)$	
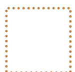	0 walls (plane)	*axes-parallel:* $1/(\frac{10}{9}n-1)$ *parallel:* $1/(\frac{30}{29}n-1)$ *general:* ?	?	$1/\max(2n-4, n)$	

Table 3.1: Summary of results for square cakes: upper and lower bounds on the level of attainable absolute proportionality.
All results assume that there are at least two agents ($n \geq 2$).
* means that the results are valid not only for square pieces but also for R-fat rectangles with $R < 2$.
? means that we do not have a non-trivial impossibility result for this case .

The constant A in our negative results is at most 2, and the constant B in our positive results is at least 2; this leads us to conjecture that the "real" constant is 2, i.e, a half-proportional division with square pieces always exists, and half-proportionality is the best that can be guaranteed. Currently we can prove this conjecture only in several restricted scenarios, that are presented below.

Square cakes bounded or unbounded

In the first set of results, the cake is a square bounded in zero or more sides. Table 3.1 summarizes our negative and positive results:

The **Impossibility** column shows upper bounds on the attainable proportionality. Each upper bound is proved by showing a specific scenario in which it is impossible to give all agents more than the mentioned fraction of their total value. The upper bound for a square with 4 walls and $n = 2$ is $1/(2n) = 1/4$, as was already seen in Example 3.1.1. The upper bounds for an unbounded plane are valid only when the pieces must be squares parallel to a pre-specified coordinate system, or parallel to each other (as is common in urban planning). The other upper bounds are valid even when the squares are allowed to be non-parallel.

The **Possibility** column shows our positive results. Each such result is proved constructively by an explicit division procedure that gives each agent at least the mentioned fraction of their total value. The *same* result means that there exists a different division procedure that guarantees a larger fraction per agent, but this procedure works only when all agents have the same valuations. We do not know whether there exists a division procedure that guarantees this larger fraction for agents with different valuations.

Note that all our impossibility results hold even for agents with the same valuations, and all our division procedures return axes-parallel pieces.

Intuitively, one may think that allowing rectangles instead of just squares should considerably increase the attainable proportionality level. But this is not the case if the pieces need to be fat. As seen in the table, most results for fat rectangles are almost the same as for squares. The only exception is the impossibility result for an unbounded plane, which we have not managed to extend to R-fat rectangles.

For $n = 2$, the proportionality levels in our possibility results are equal to the impossibility results. For a cake with two or three walls the guaranteed proportionality is equal to the impossibility result for every n. This means that in these cases, our procedures are optimal in their worst-case guarantee. For a cake with 4 walls, the guaranteed proportionality for agents with the same value measure is optimal. In the other cases, there is a multiplicative gap of at most 2 between the possibility and the impossibility result.

A secondary consideration in geometric division problems, in addition to value, is the type of cuts used for implementing the division. In some cases, *guillotine cuts* are preferred. Guillotine cuts are axis-parallel

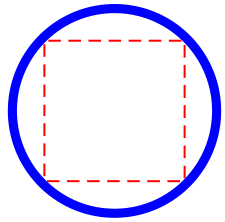

Figure 3.2: A circular cake where all value is near the perimeter. No positive value can be guaranteed to an agent who wants a square piece.

Pieces ↓		Impossibility	Possibility
Parallel squares	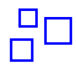	$1/(2n)$	$1/(8n-6)$ *same:* $1/(2n)$
General squares	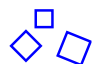	$1/(2n)$	$1/(16n-14)$ *same:* $1/(2n)$
Parallel R-fat rectangles	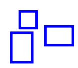	$1/(2n-1)$	$1/([4R+4][n-1]+2)$ *same:* $1/(2n)$

Table 3.2: Summary of results for arbitrary compact cakes: upper and lower bounds on the level of attainable relative proportionality.

cuts running from one end to the opposite end of an already cut piece. They are considered easier to implement (e.g. Alvarez-Valdés et al., 2002; Cui et al., 2008; Hifi et al., 2011). In the industry, guillotine cuts are used for cutting stock such as plates of glass. In the context of land division, guillotine cuts may be desired because they may make it easier to build fences between land-plots. Our procedures for a cake with 4 walls find divisions that can be implemented using guillotine cuts. The other procedures use general cuts, and we do not know if it is possible to attain the same value guarantees using guillotine cuts.

Bounded cakes of any shape

While some states in the USA are rectangular (e.g. Colorado or Wyoming), most land-estates have irregular shapes. In such cases, it may be impossible to guarantee any positive proportionality. For example, consider Robinson Crusoe arriving at a circular island. Assume that Robinson's value measure is such that all value is concentrated in a very thin strip along the shore, as in Figure 3.2. The value contained in any single square might be arbitrarily small. Clearly, no division procedure for n agents can guarantee a better fraction of the total value.

Therefore, for arbitrary cakes we use a *relative* rather than absolute fairness measure. For each agent, we calculate the maximum value that this agent can attain in a square piece if he doesn't need to share the cake with other agents. We guarantee the agent a certain fraction of this value, rather than a certain fraction of the entire cake value. This fairness criterion is similar to the *uniform preference externalities* criterion suggested by Moulin (1990b). Similar criteria have been recently studied in the context of indivisible item assignment (Budish, 2011; Procaccia and Wang, 2014; Bouveret and Lemaître, 2015).

Table 3.2 summarizes our bounds on relative proportionality. The impossibility results follow trivially from those for square cakes. The possibility results require new division procedures. They are valid for any cake that is a compact (closed and bounded) subset of the plane. The guarantees are better when the pieces are required to be axis-parallel. This is in accordance with the common practice in urban planning, in which axis-parallel plots are usually preferred.

3.1.2 Techniques

Most of our division procedures can be presented as sequences of *auctions*.[4] The general process is as follows. Initially, each of the n agents receives a ticket with an entitlement to share a certain cake, C, in a group of n agents. Then, the divider performs a well-designed sequence of auctions. In each auction, the winning agents exchange their ticket for another ticket with an entitlement to share a smaller cake $C' \subset C$ in a smaller group of $n' < n$ agents. This goes on until finally each agent holds a private entitlement for a single piece of the cake. Note that there are no monetary payments: the winners 'pay' only by giving away their tickets.

We use auctions of two types: *mark auction* and *eval auction*.[5] They are presented briefly below; formal definitions and detailed examples are given in Section 3.4.

- In a *mark auction*, each agent bids by marking a piece of cake. All bids must satisfy a given geometric constraint (such as "mark a square at the bottom-left corner"). An agent bidding a piece X_i is interpreted as saying "I am willing to give my ticket in exchange for X_i". The agent bidding the smallest piece is the winner. The winner receives his bid and goes home, while the remaining agents continue to divide the remaining cake.

- In an *eval auction*, the divider specifies a piece $C' \subset C$, and each agent bids by declaring his/her evaluation of C'. An agent bidding a value V is interpreted as saying "I am willing to give my ticket for sharing C in a group of n, in exchange for a ticket for sharing C' in a group of $n'(V)$.". Here n' is some weakly-increasing function of V that depends on the situation. The agent or agents bidding the highest values are the winners, since they are willing to share C' with the largest number of other agents. The number of winners is determined as the largest value n' such that the n' highest winners are willing to share C' in a group of n'. These winners go on and divide C' among them, while the remaining $n - n'$ agents continue to compete on $C \setminus C'$.

The geometric constraints are carefully designed in order to guarantee that the final pieces are usable. A key geometric concept here is the *cover number* — the minimum number of squares required to cover a given region. By making sure that all sub-pieces have a sufficiently small cover-number, we ensure that they can be divided effectively. See Section 3.4 for details.

3.1.3 Related work

Many authors regard land division as an important application of division procedures (e.g. Berliant and Raa, 1988; Berliant et al., 1992; Legut et al., 1994; Chambers, 2005; Dall'Aglio and Maccheroni, 2009; Hüsseinov, 2011; Nicolò et al., 2012). Hence, they note the importance of imposing some geometric constraints on the pieces allotted to the agents.

The most well-studied constraint is *connectivity* — each agent should receive a single connected piece. The cake is usually assumed to be the one-dimensional interval $[0,1]$ and the allotted pieces are sub-intervals (e.g. Stromquist, 1980; Su, 1999; Nicolò and Yu, 2008; Azrieli and Shmaya, 2014)). Several authors studied a circular cake (Thomson, 2007; Brams et al., 2008; Barbanel et al., 2009), but it is still a one-dimensional circle and the pieces are one-dimensional arcs.

The importance of the multi-dimensional geometric shape of the plots was noted by several authors.

Hill (1983); Beck (1987); Webb (1990); Berliant et al. (1992) study the problem of dividing a disputed territory between several bordering countries, with the constraint that each country should get a piece that is adjacent to its border.

Berliant et al. (1992); Ichiishi and Idzik (1999); Dall'Aglio and Maccheroni (2009) acknowledge the importance of having nicely-shaped pieces in resolving land disputes. They prove that, if the cake is a simplex in any number of dimensions, then there exists an envy-free and proportional partition of the cake into polytopes. However, this proof is purely existential when the cake has two or more dimensions. Additionally, there are no restrictions on the fatness of the allocated polytopes and apparently these can be arbitrarily thin triangles. Berliant and Dunz (2004) studies the existence of competitive equilibrium with

[4] The relation between division procedures and auctions has already been mentioned by Brams and Taylor (1996).

[5] The two auction types are analogous to the two query types — *mark query* and *eval query* — used in the cake-cutting literature in computer science, e.g. Robertson and Webb (1998); Woeginger and Sgall (2007). In fact, each mark/eval auction can be implemented by n mark/eval queries. Therefore, all our division procedures require $O(poly(n))$ queries. We prefer to use auctions because their economic meaning is clearer.

utility functions that may depend on geometric shape; their *nonwasteful partitions assumption* explicitly excludes fat shapes such as squares. Devulapalli (2014) studies a two-dimensional division problem in which the geometric constraints are connectivity, simple-connectivity and convexity.

Iyer and Huhns (2009) describe a procedure for giving each agent a rectangular plot with an aspect ratio determined by the agent. Their procedure asks each of the n agents to draw n disjoint rectangles on the map of the two-dimensional cake. These rectangles are supposed to represent the "desired areas" of the agent. The procedure tries to give each agent one of his n desired areas. However, it does not succeed unless each rectangle proposed by an individual intersects at most one other rectangle drawn by any other agent. If even a single rectangle of Alice intersects two rectangles of George (for example), then the procedure fails and no agent gets any piece.

In our model (see Section 3.2), the utility functions depend on geometry, which makes them non-additive. They are not even sub-additive like in the models of Maccheroni and Marinacci (2003); Dall'Aglio and Maccheroni (2005, 2009). [6] Previous papers about cake-cutting with non-additive utilities can be roughly divided to two kinds: some (Berliant and Dunz, 2004; Sagara and Vlach, 2005; Hüsseinov and Sagara, 2013) handle general non-additive utilities but provide only pure existence results. Others (Su, 1999; Caragiannis et al., 2011; Mirchandani, 2013) provide constructive division procedures but only for a 1-dimensional cake. Our approach is a middle ground between these extremes. Our utility functions are more general than the 1-dimensional model but less general than the arbitrary utility model; for this class of utility functions, we provide both existence results and constructive division procedures.

Besides fair division problems, geometric methods have been used in many other economics problems,[7] such as voting (Plott, 1967), trade theory and growth theory (e.g. Johnson, 1971), tax burdens (Hines et al., 1995), social choice (Cantillon and Rangel, 2002), mechanism design (Goeree and Kushnir, 2011), public good/bad allocation (e.g. Öztürk et al., 2013, 2014; Chatterjee et al., 2016), utility theory (Abe, 2012) and general economics models (Michaelides, 2006).

With square pieces a proportional allocation may not exist, so we have to settle for partial-proportionality. Other goals that justify partial-proportionality are speed of computation (Edmonds and Pruhs, 2006a; Edmonds et al., 2008), improving the social welfare (Zivan, 2011; Arzi, 2012) and guaranteeing a minimum-length constraint of a 1-dimensional piece (Caragiannis et al., 2011).

3.2 Model

We briefly recall some terminology from Chapter 2 (see there for formal definitions).

- C is the cake to be divided. In this chapter it will usually be a polygon or a polygonal domain in the Euclidean plane $\mathbb{R}^2$.
- S is the family of pieces that are considered *usable*. An *S-piece* is an element of S. In this chapter it will usually be the family of squares or fat rectangles.
- For each agent $i \in \{1, \ldots, n\}$, $V_i(X_i)$ is agent i's value-measure of the piece X_i.
- For each agent $i \in \{1, \ldots, n\}$, $V_i^S(X_i)$ is agent i's utility of the piece X_i. It is the value-measure of the most valuable S-piece contained in X_i.

The fairness of an allocation is determined by the agents' *normalized* values. Values can be normalized in two ways:

- either divide them by the *absolute* cake value for the agent and get $V_i(X_i)/V_i(C)$,
- or divide them by the *relative* cake utility for the agent and get $V_i(X_i)/V_i^S(C)$.

[6] Dall'Aglio and Maccheroni (2009) do not explicitly require sub-additivity, but they require *preference for concentration*: if an agent is indifferent between two pieces X and Y, then he prefers 100% of X to 50% of X plus 50% of Y. This axiom is incompatible with geometric constraints: an agent who wants square pieces will give away 100% of a 20×10 rectangle, in exchange for 50% of a 20×20 square that is the union of two such rectangles. We are grateful to Marco Dall'Aglio for his help in clarifying this issue.

[7] We are thankful to Steven Landsburg, Michael Greinecker, Kenny LJ, Alecos Papadopoulos, B Kay and Martin van der Linden for contributing these references in economics.stackexchange.com website (http://economics.stackexchange.com/q/6254/385).

Throughout the chapter absolute normalization is used, except in Subsection 3.5.6 where relative normalization is used.

An allocation is called *proportional* if the normalized value of every agent is at least $1/n$. Example 3.1.1 shows that a proportional allocation is not always attainable (whether absolute or relative normalization is used). Hence, we define:

Definition 3.2.1. *(Absolute proportionality)* For a cake C, a family of usable pieces S and an integer $n \geq 1$:

(a) The *proportionality level* of C, S and n, marked $\text{Prop}(C, S, n)$, is the largest fraction $r \in [0, 1]$ such that, for every n value measures $(V_i, ..., V_n)$, there exists an S-allocation $(X_1, ..., X_n)$ for which $\forall i: \ V_i(X_i)/V_i(C) \geq r$.[8]

(b) The *same-value proportionality level* of C, S and n, marked $\text{PropSame}(C, S, n)$, is the largest fraction $r \in [0, 1]$ such that, for every single value measure V, there exists an S-allocation $(X_1, ..., X_n)$ for which $\forall i: \ V(X_i)/V(C) \geq r$.

The analogous definition for relative proportionality is given in Subsection 3.5.6.

Obviously, for every C, S and n: $\text{Prop}(C, S, n) \leq \text{PropSame}(C, S, n) \leq 1/n$.

Applying this notation, classic cake-cutting results (e.g. Steinhaus, 1948) imply that for every cake C

$$\text{Prop}(C, \ All, \ n) = \text{PropSame}(C, \ All, \ n) = 1/n$$

where "All" is the collection of all pieces. That is: when there are no geometric constraints on the pieces, for every cake C and every combination of n continuous value measures there is a division in which each agent receives a utility of $1/n$, which is the best that can be guaranteed. One-dimensional procedures with contiguous pieces (e.g. Even and Paz, 1984) prove that $\text{Prop}(Interval, \ intervals, \ n) = 1/n$ and when translated to two dimensions they yield:

$$\text{Prop}(Rectangle, \ rectangles, \ n) = \text{PropSame}(Rectangle, \ Rectangles, \ n) = 1/n$$

However, these procedures do not consider constraints that are two-dimensional in nature, such as squareness. Such two-dimensional constraints are the focus of the present chapter.

Our challenge in the rest of this chapter will be to establish bounds on $\text{Prop}(C, S, n)$ and PropSame (C, S, n) for various cake shapes C and piece families S. Two types of bounds are provided:

- Impossibility results (upper bounds), of the form $\text{Prop}(C, S, n) \leq f(n)$ where $f(n) \in [0, 1]$, are proved by showing a set of n value measures on C, such that in any S-allocation, the value of one or more agents is *at most* $f(n)$. Such bounds are established in Section 3.3.

- Positive results (lower bounds), of the form $\text{Prop}(C, S, n) \geq g(n)$ where $g(n) \in [0, 1]$, are proved by describing a division procedure which finds, for every set of n value measures on C, an S-allocation in which the value of every agent is *at least* $g(n)$. Such bounds are established in Sections 3.4-3.5.

3.3 Impossibility Results

Our impossibility results are based on the following scenario.

- The cake C is a desert with only k water-pools; the set of pools is denoted P_k.

- Each pool in P_k is a square with side-length $\epsilon > 0$ containing 1 unit of water.

- There are n agents with the same value measure: the value of a piece equals the total amount of water in the piece. So the value of each pool in P_k is 1 and the total cake value is k.

- We say that a piece X_i is **supported by** P_k if X_i contains strictly more than 1 unit of water from P_k. This implies that X_i touches at least two pools of P_k.

- We say that P_k **supports** m **squares** if there exists a collection of m pairwise-disjoint squares each of which contains strictly more than one unit of water from P_k.

[8] Shortly: $\text{Prop}(C, S, n) = \inf_V \sup_X \min_i V_i(X_i)/V_i(C)$, where the infimum is on all combinations of n value measures $(V_1, ..., V_n)$, the supremum is on all S-allocations $(X_1, ..., X_n)$ and the minimum is on all agents $i \in \{1, ..., n\}$.

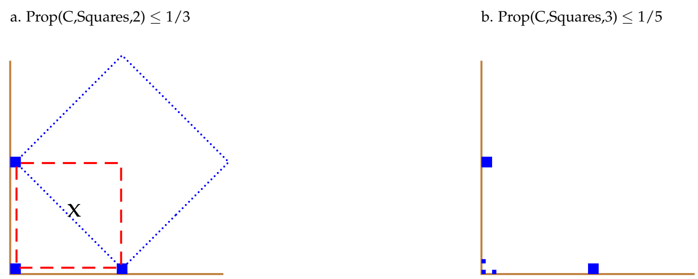

Figure 3.3: Impossibility results in a quarter-plane cake.

Figure 3.4: Impossibility results in a square cake.

The latter definition implies the following lemma:

Lemma 3.3.1. *A collection of k pools supports at most $k-1$ squares.*

Proof. Let P_k be a collection of k pools and suppose that it supports m squares. This means that there exists a collection of m pairwise-disjoint squares, each of which contains more than one unit of water from P_k. So the union of these squares contains strictly more than m units of water from P_k. Since each pool in P_k contains exactly one unit of water, necessarily $k \geq m+1$ so $m \leq k-1$. □

In each impossibility result, we present a set P_k and prove that it supports at most $n-1$ squares. This implies that, in every allocation of n pairwise-disjoint squares, at least one agent receives a piece not supported by P_k — a piece with at most 1 unit of water. The value of this agent is at most a fraction $1/k$ of the total cake value. This implies that $\text{PropSame}(C, \mathit{Squares}, n) \leq 1/k$, which implies that $\text{Prop}(C, \mathit{Squares}, n) \leq 1/k$.

3.3.1 Impossibility results for two, three and four walls

We start with impossibility results for two agents.

Claim 3.3.1.

$$\text{PropSame}(\mathit{Quarter\ plane}, \mathit{Squares}, 2) \leq 1/3$$

Proof. Let P_3 be the set of 3 pools shown in Figure 3.3/a, where the bottom-left corners of the pools are in $(0,0)$, $(10,0)$, $(0,10)$. Every square in C touching two pools of P_3 must contain e.g. the point $(6,6)$ in its interior (marked by x in the figure). Hence, every two squares touching two pools of P_3 must overlap. Hence, P_3 supports at most one square. Hence, in any allocation of squares to two agents, at least one square touches at most one pool of P_3; the agent receiving such a square has at most $1/3$ of the total value. □

Claim 3.3.2.

$$\text{PropSame}(\mathit{Square}, \mathit{Squares}, 2) \leq 1/4$$

Proof. Analogous to the previous claim, based on the set P_4 shown in Figure 3.4/a. □

To extend these results to $n > 2$ agents, we construct new sets of pools by shrinking existing sets into pools of other sets.

As an example, consider P_3 from the proof of Claim 3.3.1. Suppose the entire plane is shrunk (deflated) towards the origin. If the deflation factor is sufficiently large, all three pools of the shrunk P_3 are contained in $[0,\epsilon] \times [0,\epsilon]$, which is a pool of the original P_3. The cake itself (the quarter-plane) is not changed by the deflation. By adding the other two pools of P_3, namely $(10,0)$ and $(0,10)$, we get a larger pool set, P_5, which is depicted in Figure 3.3/b. We already know that the shrunk P_3 supports at most one square. The

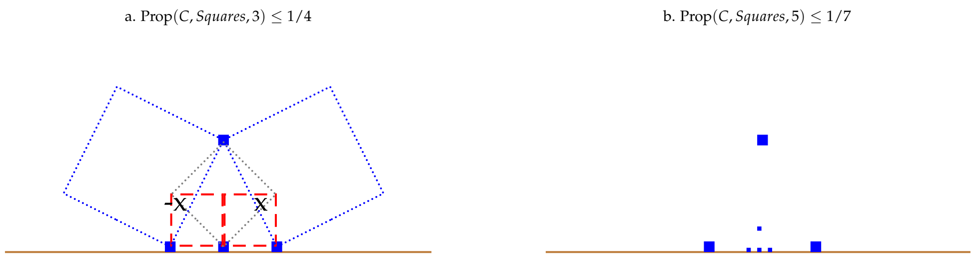

Figure 3.5: Impossibility result for 3 agents on a half-plane. See Claims 3.3.5-3.3.6.

additional two pools support at most one additional square, since there is at most one square touching two new pools or a new pool and a shrunk pool. Hence, P_5 supports at most two squares. This proves that $\text{PropSame}(\textit{Quarter plane}, \textit{Squares}, 3) \leq 1/5$. The following claim generalizes this construction.

Claim 3.3.3. For every $n \geq 1$:

$$\text{PropSame}(\textit{Quarter plane}, \textit{Squares}, n) \leq \frac{1}{2n-1}$$

Proof. [9]It is sufficient to prove that for every n there is an arrangement of $2n-1$ pools in C that supports at most $n-1$ squares. The proof is by induction on n. The base case $n = 1$ is trivial (and the case $n = 2$ is Claim 3.3.1). For $n > 2$, assume there is an arrangement of $2(n-1)-1$ pools that supports at most $n-2$ squares. Deflate the entire arrangement towards the origin until it is contained in $[0,\epsilon] \times [0,\epsilon]$, where $\epsilon > 0$ is a sufficiently small constant.

Add two new pools with side-length ϵ cornered at $(10,0)$ and $(0,10)$. We now have an arrangement of $2n-1$ pools. Every square touching a new pool and another pool (either new or old), must contain e.g. the point $(6,6)$ in its interior, so every two such squares must overlap. Hence, the additional pools support at most one additional square. All in all, the new arrangement of $2n-1$ pools supports at most $(n-2)+1 = n-1$ squares. □

The upper bound for two walls is also trivially true when the cake is a square with three walls, since adding walls cannot increase the proportionality:

$$\text{PropSame}(\textit{Square with 3 walls}, \textit{Squares}, n) \leq \frac{1}{2n-1}$$

The bound also holds for a square with 4 walls, but in this case a slightly tighter bound is true:

Claim 3.3.4. For every $n \geq 2$,

$$\text{PropSame}(\textit{Square with 4 walls}, \textit{Squares}, n) \leq \frac{1}{2n}$$

Proof. W.l.o.g. assume C is the square $[0,10+\epsilon] \times [0,10+\epsilon]$. Create the arrangement of $2(n-1)-1$ pools from the induction step of Claim 3.3.3. Deflate it into to $[0,\epsilon] \times [0,\epsilon]$. The shrunk collection supports at most $n-2$ squares. Add *three* new pools with side-length ϵ cornered at $(10,0)$, $(0,10)$ and $(10,10)$, as in Figure 3.4/b. Every square in C touching a new pool and another pool must contain $(5,5)$ in its interior. Hence, the three additional pools allow us to support at most one additional square. All in all, the new arrangement of $2n$ pools supports at most $n-1$ squares. □

3.3.2 Impossibility results for one wall

Claim 3.3.5.

$$\text{PropSame}(\textit{Half plane}, \textit{Squares}, n = 3) \leq 1/4$$

Proof. Let P_4 be the set of 4 pools shown in Figure 3.5/a. Assume the side-length of each pool is $\epsilon \leq 0.01$ and that their bottom-left corner is in $(-5,0)$, $(0,0)$, $(0,10)$, $(5,0)$. We prove that P_4 supports at most 2 squares. Examine the squares in C that touch two pools of P_4:

[9]We are grateful to Boris Bukh for the idea underlying this proof.

- Every square touching $(5,0)$ and another pool must contain the point x $(4,4.5)$ in its interior.
- Every square touching $(-5,0)$ and another pool must contain the point -x $(-4,4.5)$.
- Every square touching $(0,0)$ and another pool must touch either x or -x or both.

Hence, in every set of three squares, each of which touches two pools of P_4, at least two squares must overlap. Hence, P_4 supports at most two squares. Hence, in any allocation to three agents, at least one of them receives at most $1/4$ of the total value. □

Claim 3.3.6. For every $n \geq 2$:

$$\text{PropSame}(\textit{Half plane}, \textit{Squares}, n) \leq \frac{1}{(3/2)n - 1}$$

Proof. The proof is analogous to that of Claim 3.3.3. With each induction step, the current arrangement of pools is shrunk towards the central pool at the origin, *three* new pools are added, but only *two* new squares are supported. Hence the coefficient of n is $3/2$. The -1 ensures that the right-hand side is a correct upper bound for every $n \geq 2$. □

Figure 3.5/b shows the set of 7 pools for the case $n = 5$.

3.3.3 Impossibility results for zero walls

Finding an impossibility result for an unbounded cake is a challenging task. The main difficulty is that, when there are no walls, any arrangement of pools can be rotated arbitrarily, as will be explained shortly.

We begin with impossibility results for the restricted case in which the squares must be parallel to a specific coordinate system. Such a restriction may be meaningful, for example, in the installation of solar power-plants or the building of houses with electric solar panels, where the positioning relative to the sun is important.

Claim 3.3.7. Given a fixed coordinate system in the plane:

$$\text{PropSame}(\textit{Plane}, \textit{Axes Parallel Squares}, n = 5) \leq 1/6$$

Proof. Let P_6 be the set of 6 pools: A(0,2.5), B(-3,0), C(-1,0), C'(1,0), B'(3,0), A'(0,-2.5). We prove that P_6 supports at most 4 axes-parallel squares. First, consider the squares that touch two pools of P_6:

(a) P_6 Pools:

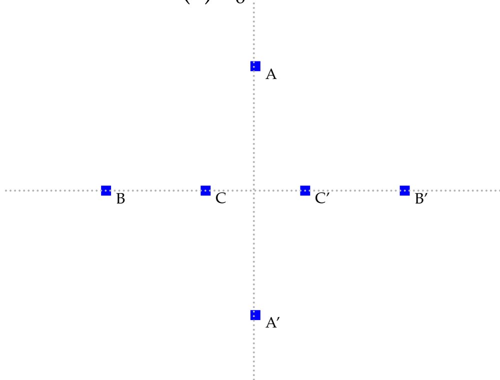

(b) Potential squares:

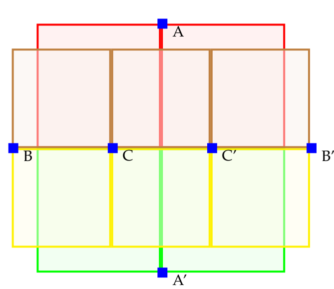

We can ignore squares that contain other squares or that contain pools in their interior, since such squares can be shrunk without interfering with other squares. Hence, any set of supported squares must contain a subset of the following:

- At most two disjoint "top squares" (squares touching pool A) and two disjoint "bottom squares" (touching pool A'). Each such square has a side-length of 2.5.

- At most one "left square" (touching pools B and C), one "right square" (touching pools B′ and C′) and one "central square" (touching C and C′). Each such square has a side-length of 2 and can be located anywhere between $y = -2$ and $y = 2$. [10]

We prove that at most four of these squares can be supported simultaneously. There are two cases:

Case #1: there are no bottom squares. The pool A′ is not used, so only 5 pools are used. By Lemma 3.3.1, these pools can support at most 4 squares. The situation is similar if there are no top squares, since in this case the pool A is not used.

Case #2: there is at least one bottom square (e.g, a square supported by A′ and C′) and at least one top square (e.g, supported by A and C). These two squares leave no room for a central square. Hence, there is room for at most two additional squares: one above the x axis (e.g, supported by A and C′, or C′ and B′), and one below the x axis (e.g, supported by A′ and C, or C and B).

In all cases, P_6 supports at most 4 axes-parallel squares. □

Claim 3.3.8. Given a fixed coordinate system in the plane, for every $k \geq 0$:

$$\text{PropSame}(\textit{Plane},\ \textit{Axes Parallel Squares},\ n = 5 + 9k) \leq 1/(6 + 10k)$$

Proof. We prove that for every $k \geq 0$, there exists an arrangement of $6 + 10k$ pools that supports at most $4 + 9k$ axes-parallel squares. The proof is by induction on k. The base $k = 0$ is proved by P_6 from Claim 3.3.7. Assume that there exists an arrangement $P_{6+10(k-1)}$ which supports at most $9 + 4k$ squares. Construct a new arrangement P_{6+10k} in the following way. Take P_6, replace the pool A a with shrunk copy of P_6 and the pool A′ with a shrunk copy of $P_{6+10(k-1)}$. The following illustration shows P_{16}, the arrangement for $k = 1$ (the shrunk copies are enlarged for the sake of clarity):

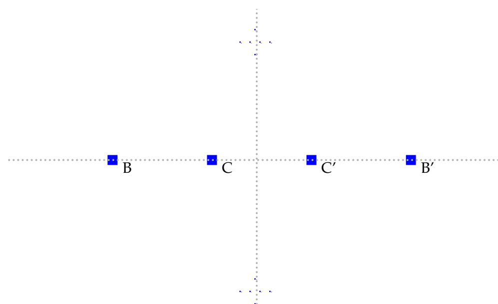

The number of pools in the new arrangement is $6 + 4 + 6 + 10(k - 1) = 6 + 10k$. We claim that it supports at most $4 + 9k$ squares:

- The shrunk copy of P_6 supports at most 4 squares;
- The shrunk copy of $P_{6+10(k-1)}$ supports at most $4 + 9(k - 1)$ squares, by the induction assumption;
- The four pools B C C′ B′ in the large P_6 support at most 3 large squares;
- If there is a top large square then there is at most one additional large square above the x axis, and if there is a bottom large square then there is at most one additional large square below the x axis.

All in all, at most $4 + 4 + 9(k - 1)$ squares are supported by the shrunk copies and at most 3+2=5 additional large squares are supported by the outer arrangement, so the total number of supported squares is at most $4 + 9k$. □

In general, every 10 additional pools support at most 9 additional squares. Hence:

$$\text{PropSame}(\textit{Plane},\ \textit{Axes Parallel Squares},\ n) \leq \frac{1}{(10/9)n - 1} \approx \frac{9}{10} \cdot \frac{1}{n}$$

This implies that any division procedure which works in a pre-specified coordinate system cannot guarantee a proportional division of the plane with square pieces.

[10]While there can two disjoint squares touching pools B+C, Lemma 3.3.1 implies that the pools B+C can support at most one square. The same is true for the pools B′+C′ and C+C′.

In our next results, we relax the axes-parallel restriction and only require that the squares be parallel to each other. While this is still not the most general setting, it is natural e.g. in urban planning. Equivalently, we still require that the squares be parallel to the axes, but allow the arrangement of pools to rotate.

Note that the proof of Claim 3.3.7 (Case 2) relies on the fact that any pair of a top-square and a bottom-square leaves no room for a central square. This follows from the facts that A and A′ lie horizontally between C and C′, and the horizontal distance between C and C′ is larger than the vertical distance between B and B′. These facts are still true if the entire arrangement is rotated by at most 18° to either direction:[11]

(a) P_6 rotated 18°:

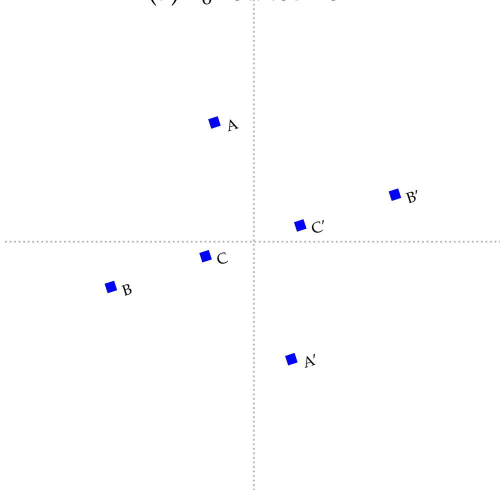

(b) Potential squares:

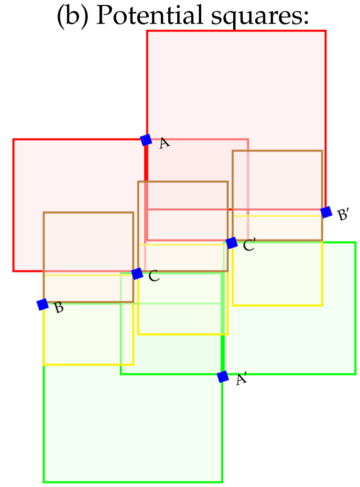

For every angle θ, define $ParallelSquares[\theta]$ as the family of squares rotated at exactly θ degrees (counter-clockwise) relative to the axes. Then, the proofs of Claim 3.3.7 and 3.3.8 and the above explanation imply:

Claim 3.3.9. For every $\theta \in [-18°, +18°]$ and every $k \geq 0$:

$$\text{PropSame}(Plane, ParallelSquares[\theta], n = 5 + 9k) \leq 1/(6 + 10k)$$

The arrangement P_{6+10k} "covers" a range of rotation-angles of size 36°. By using three copies of P_{6+10k} rotated in different angles, we can cover the entire range of relevant rotation angles. We use this idea to prove an impossibility result for rotated parallel squares.

Claim 3.3.10. For every $k \geq 0$:

$$\text{PropSame}(Plane, ParallelSquares, n = 18 + 29k) \leq 1/(18 + 30k)$$

Proof. Construct an arrangement P_{18+30k} from three copies of P_{6+10k}:

- A *leftmost copy* — rotated by $-27°$ and translated by $(-300, 0)$;
- A *central copy* — not rotated;
- A *rightmost copy* — rotated by $+27°$ and translated by $(+300, 0)$.

The following illustration shows P_{18} (the construction for $k = 0$) with the three copies enlarged for the sake of clarity:

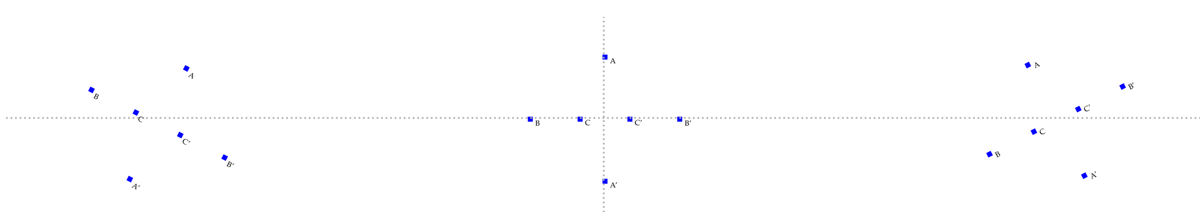

[11]The calculation was done using Geogebra (Hohenwarter, 2002; Hohenwarter et al., 2013). The worksheet is available here: https://tube.geogebra.org/m/zzNY3ag4

We claim that if P_{18+30k} is rotated by any angle $\theta \in [-45°, 45°]$, then the rotated arrangement supports at most $18 + 29k$ axes-parallel squares. Consider three cases:

(a) P_{18+30k} is rotated by $\theta \in [-45°, -9°]$. Then, the rightmost copy is P_{6+10k} rotated by $\theta + 27° \in [-18°, 18°]$, so it supports at most $4 + 9k$ squares.

(b) P_{18+30k} is rotated by $\theta \in [-18°, +18°]$. Then the central copy supports at most $4 + 9k$ squares.

(c) P_{18+30k} is rotated by $\theta \in [+9°, +45°]$. Then the leftmost copy is P_{6+10k} rotated by $\theta - 27° \in [-18°, 18°]$, so it supports at most $4 + 9k$ squares.

In all cases, one of the copies supports at most $4 + 9k$ squares. Each of the other two copies has $6 + 10k$ pools, so by Lemma 3.3.1 it supports at most $5 + 10k$ squares. Additionally, between the three copies there can be at most four (huge) pairwise-disjoint squares: two above and two below the x axis. All in all, the number of supported squares is at most $(4 + 9k) + (5 + 10k) + (5 + 10k) + 4 = 18 + 29k$.

Therefore, for any angle $\theta \in [-45°, 45°]$, if the family S of usable pieces is the family of squares rotated by θ, then P_{18+30k} supports at most $18 + 29k$ S-pieces. But, *any* square is identical to a square rotated by $\theta \in [-45°, 45°]$. Therefore, the existence of P_{18+30k} proves the claim. □

In Claim 3.3.10, for every 30 new pools, at most 29 new squares can be supported. Therefore,

Claim 3.3.11. For every $n \geq 1$:

$$\text{PropSame}(Plane,\ Parallel\ Squares,\ n) \leq \frac{1}{(30/29)n - 1} \approx \frac{29}{30} \cdot \frac{1}{n}$$

3.3.4 Impossibility results with fat rectangles

Our impossibility results so far have assumed that S is the family of squares. One could think that allowing fat rectangles, instead of just squares, can overcome these impossibility results. But this is not necessarily true. Claim 3.3.1 holds as-is for R-fat rectangles:

Claim 3.3.12. For every finite $R \geq 1$:

$$\text{PropSame}(Quarter\ plane,\ R\ fat\ rectangles,\ 2) \leq 1/3$$

Proof. Let P_3 be the arrangement of 3 pools from the proof of Claim 3.3.1:

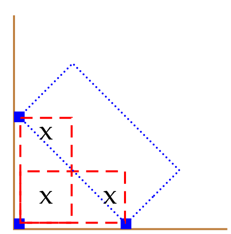

The side-length of each pool is $\epsilon > 0$. Every R-fat rectangle touching the two bottom pools must have a height of at least $(10 - 2\epsilon)/R$ and thus, when ϵ is sufficiently small, it must contain the point $(5/R, 5/R)$ and the point $(10 - 10/R,\ 5/R)$. Every R-fat rectangle touching the two left pools must contain the point $(5/R, 5/R)$ and the point $(5/R,\ 10 - 10/R)$. Every R-fat rectangle touching the top-left and the bottom-right pools must contain $(10 - 10/R,\ 5/R)$ and $(5/R,\ 10 - 10/R)$. Hence, in every allocation of disjoint R-fat rectangles, at most one rectangle touches two or more pools. □

Claim 3.3.3 is based on Claim 3.3.1, so it holds as-is for R-fat rectangles. The same is true for the 3-walls result. The 1-wall claims 3.3.5 and 3.3.6 can be generalized in a similar way:

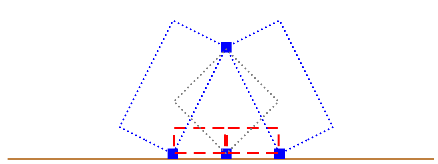

We omit the details. We obtain:

Claim 3.3.13. For every $R \geq 1$:

$$\text{PropSame}(Square\,with\,1\,wall,\,R\,fat\,rectangles,\,n) \leq \frac{1}{(3/2)n-1}$$
$$\text{PropSame}(Square\,with\,2\,walls,\,R\,fat\,rectangles,\,n) \leq \frac{1}{2n-1}$$
$$\text{PropSame}(Square\,with\,3\,walls,\,R\,fat\,rectangles,\,n) \leq \frac{1}{2n-1}$$

Claims 3.3.2 and 3.3.4 hold whenever $R < 2$, since in this case, every R-fat rectangle touching one of the corner-pools must contain the central point of the cake in its interior, as shown below:

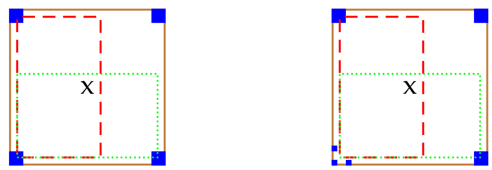

This gives:

Claim 3.3.14. For every R such that $1 \leq R < 2$:

$$\text{PropSame}(Square\,with\,4\,walls,\,R\,fat\,rectangles,\,n) \leq \frac{1}{2n}$$

When $R \geq 2$, the following slightly weaker result follows immediately from Claim 3.3.13 (since adding walls cannot increase the proportionality):

Claim 3.3.15. For every $R \geq 2$: [12]

$$\text{PropSame}(Square\,with\,4\,walls,\,R\,fat\,rectangles,\,n) \leq \frac{1}{2n-1}$$

The impossibility results for an unbounded plane are different for R-fat rectangles. Consider first Claim 3.3.7, which assumes that the pieces must be axes-parallel. When the pieces have to be squares, the set P_6 supports at most 2 pieces above the x axis and 2 pieces below the x axis. But when the pieces may be R-fat rectangles and $R \geq 2.5$, it is possible to support 3 pieces above or below the x axis, e.g:

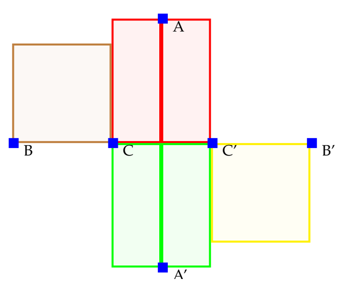

The impossibility result can be maintained by locating the pool A at $(2.5R, 0)$ instead of $(2.5, 0)$, and the pool A' at $(-2.5R, 0)$ instead of $(-2.5, 0)$:

[12]By classic cake-cutting protocols, PropSame$(Square,\,\infty\,fat\,rectangles,\,n) = 1/n$ (an ∞-fat rectangle is just an arbitrary rectangle). The PropSame function is thus discontinuous at $R = \infty$. If the agents agree to use any rectangular piece, they can receive their proportional share of $1/n$, but if they insist on using R-fat rectangles, even when R is very large, they might have to settle for about half of this share.

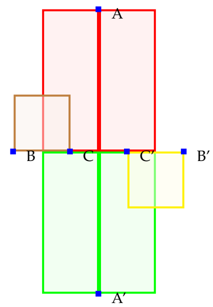

So Claim 3.3.7, and hence Claim 3.3.8, are valid for R-fat rectangles, and we obtain:

Claim 3.3.16. Given a fixed coordinate system in the plane, for every $R \geq 1$:

$$\text{PropSame}(Plane,\ Axes\,Parallel\,R\,fat\,rectangles,\ n) \leq \frac{1}{(10/9)n - 1}$$

However, the angle-range in which Claim 3.3.16 holds is no longer $[-18°, 18°]$ — the range becomes smaller as a (complicated) function of R. This means that more copies may be needed to "cover" the entire range of $[-45°, 45°]$. Therefore, the upper bound for parallel squares will probably be a complicated function of R. We leave this issue for future work.

3.4 Auctions and Covers

Our cake-cutting procedures are composed of two types of auctions. In a *mark auction*, each agent bids by marking a piece of the cake; the winner is the agent marking the smallest piece. In an *eval auction*, each agent bids by declaring a value for a pre-specified piece of cake; the winners are the agents declaring the highest value. As usual in the cake-cutting literature, no monetary transfers are involved; the agents effectively 'pay' with their entitlements for a share of the cake. Below we explain each auction type in detail.

3.4.1 Mark auction

In a mark auction, the divider specifies a geometric constraint and a value v. Each agent has to mark a piece of the cake which satisfies the geometric constraint and is worth for him exactly v. The geometric constraint guarantees that the marked pieces are totally ordered by containment (i.e. for every two agents i, j, the bid of i either contains or is contained in the bid of j). Hence, there is a *smallest bid* — a bid contained in all other bids. There can be more than one smallest bid; in this case, one smallest bid is selected arbitrarily. The agent making the selected smallest bid is the *winner*; he is allocated his bid and goes home. The remaining cake is divided among the remaining $n-1$ agents.

Example 3.4.1. Dividing a rectangle to rectangles. The cake C is a rectangle and S is the family of rectangles. We normalize the valuations of all agents such that the value of the entire cake is n. We show how a sequence of mark auctions can be used to give each agent a rectangle with a value of at least 1.

The proof is by induction on the number of agents n. When $n = 1$, C can just be given to the single agent. Suppose we already know how to divide a rectangle to $n-1$ agents who value it as $n-1$. Now we are given n agents who value the cake as n. We do a mark auction with the following geometric constraint: *mark a rectangle whose rightmost edge coincides with the rightmost edge of* C. The auction value is $v = 1$. The continuity of the valuations guarantees that all agents can indeed bid as required, and the geometric constraint guarantees that the bids are totally ordered by containment. An example is illustrated below, where there are four bids marked by dotted lines:

the winning bid — the smallest rectangle — is marked by a thick dotted line. The winner is given his bid, so he now has a rectangle with a value of exactly 1, as required (recall that our guarantees are valid for every agent bidding truthfully, regardless of what the other agents do). Since the $n-1$ losing bids contain the winning bid, the $n-1$ losers value the winning bid as at most 1. By additivity, they value the remaining cake as at least $n-1$. Hence, by the induction assumption we can divide the remaining cake among them in a similar way, finally giving each agent a rectangle with a value of at least 1.[13] □

A mark auction has the following interpretation. Initially, each agent holds an entitlement for a piece of cake. An agent bidding a piece X_i is interpreted as saying "I am willing to give my entitlement in exchange for piece X_i". The agent marking the smallest piece is effectively offering the highest "price" per unit area; hence this agent is the winner. He pays for the win by giving up his entitlement and leaving the remaining cake to the remaining agents.

3.4.2 Eval auction

In an eval auction, the divider specifies a piece of cake $C' \subset C$. Each agent i has to declare the value $V_i(C')$. The agents are ordered in a descending order of their bids, such that $V_1(C') \geq V_2(C') \geq \cdots \geq V_n(C')$. The procedure calculates the number of winners n' (we explain shortly how this number is calculated). The n' highest bidders, $1, \ldots, n'$, are the winners. The remaining $n - n'$ agents are the losers. The procedure then divides C' among the winners and $C \setminus C'$ among the losers.

To calculate the number of winners n', we should already have a plan for dividing C' among each possible number of winners $n' \leq n$. Specifically, we should have a procedure for dividing C' among n' agents, each of whom values C' as at least $F(n')$ (where F is some increasing function), such that each agent is guaranteed a piece with a value of at least 1. Assuming that we have such a procedure, the number of winners is defined as the largest integer n' such that:

$$V_{n'}(C') \geq F(n')$$

or 0 if already $V_1(C') < F(1)$. Since $V_{n'}(C')$ is a decreasing sequence, the definition implies that:

- For every winner $i \in \{1, \ldots, n'\}$: $V_i(C') \geq F(n')$
- For every loser $i \in \{n'+1, \ldots, n\}$: $V_i(C') < F(n'+1)$

(this is true even when $n' = 0$). Hence, the set of winners is a largest set of agents for whom we can divide C' in a way which guarantees each of them a value of at least 1.

Example 3.4.2. Dividing an archipelago to rectangles. The cake C is an "archipelago" — a union of m disjoint rectangular "islands". S is the family of rectangles. We normalize the valuations of all agents such that the value of the entire archipelago is $n + m - 1$. We show how a sequence of eval auctions can be used to give each agent a rectangle, contained in one of the islands, with a value of at least 1.

The proof is by induction on the number of islands m. When $m = 1$, C is a single rectangle and all agents value it as at least n, so the procedure of Example 3.4.1 can be used to give each agent a rectangle with a value of at least 1. Suppose we already know how to divide an archipelago of $m-1$ islands. Given an archipelago of m islands, pick one island arbitrarily and call it C'. Do an eval auction on C'. Order the bids in descending order, and let n' be the largest index such that:

$$V_{n'}(C') \geq n'$$

or 0 if already $V_1(C') < 1$. If $n' = 0$ then just discard C'; otherwise use the procedure of Example 3.4.1 to divide C' among the n' winners. By definition, each winner values C' as at least n' so he is guaranteed a rectangular piece of C' with a value of at least 1.

All $n - n'$ losers value C' as less than $n' + 1$, so they value the remaining archipelago $C \setminus C'$ as more than $(n + m - 1) - (n' + 1) = (n - n') + (m - 1) - 1$. This is an archipelago of $m - 1$ islands, so by the induction assumption we can divide it among the remaining $n - n'$ agents giving each agent a rectangle with a value of at least 1. Note that this is true even when $n' = 0$.[14] □

[13] Example 3.4.1 shows that $\text{Prop}(Rectangle, Rectangles, n) = 1/n$. This result is not new since it follows immediately from known results on 1-dimensional cake-cutting. It is presented here to show that it fits well into the auction framework.

[14] Example 3.4.2 shows that $\text{Prop}(m \text{ } disjoint\ rectangles,\ Rectangles,\ n) \geq 1/(n + m - 1)$. It is easy to construct an arrangement of pools, analogous to the ones in Section 3.3, proving that this is the best proportionality that can be guaranteed.

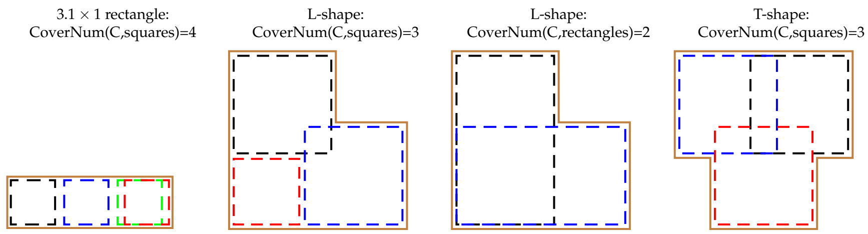

Figure 3.6: Cover numbers of various polygons.

An eval auction has the following interpretation. Initially, each agent has an entitlement to share the entire cake C with n agents (including the agent himself). An agent bidding a value V is interpreted as saying "I am willing to give my entitlement in exchange for an entitlement to share C' with at most n' agents, where n′ is the largest integer such that $V \geq F(n')$." The agents with the highest bids are actually offering a higher "price" for C', since they are willing to share C' with a larger number of other agents. Hence, the highest bidders are the winners. They pay for their win by giving up their entitlement to $C \setminus C'$ and leaving it to the remaining agents.

3.4.3 Cover numbers

The last ingredient we need for our division procedures, in addition to the two auction types, is the *cover number*. It is a well-known concept in computational geometry (see Keil (2000) for a survey).

Definition 3.4.3. Let C be a cake and S a family of pieces.
(a) An *S-cover of C* is a set of S-pieces, possibly overlapping, whose union equals C.
(b) The *S-cover number of C*, CoverNum(C, S), is the minimum cardinality of an S-cover of C.

Some examples are depicted in Figure 3.6.

The cover number is related to the utility that a single agent can derive from a given cake:

Lemma 3.4.4. *(Covering Lemma) For every cake C and family S:*

$$Prop(C, S, n = 1) \geq \frac{1}{CoverNum(C, S)}$$

Proof. Let $k = \text{CoverNum}(C, S)$ and let $\{C_1, ..., C_k\}$ be an S-cover of C. By definition $\cup_{j=1}^{k} C_j = C$. By additivity, if an agent's valuation function is V, then:

$$\sum_{j=1}^{k} V(C_j) \geq V(C)$$

so the average value of the left-hand side is at least $V(C)/k$. By the properties of the average, at least one summand must be weakly larger than the average value, i.e, there exists j for which $V(C_j) \geq V(C)/k$. This C_j, which is an S-piece, gives the single agent a utility of at least $1/k$ of the total cake value. □

The next example combines an eval auction, a mark auction and the Covering Lemma.

Example 3.4.5. Dividing a square between two agents who want square pieces. The cake C is a square, S is the family of squares and there are $n = 2$ agents. Example 3.1.1 shows that the maximum utility that can be guaranteed to both agents is $1/4$ of the total value. We now present a division procedure that guarantees this utility. We normalize the valuations of both agents such that their value of C is 4 and give each agent a square with a value of at least 1.

Partition the cake to a 2×2 grid. Denote one of the four quarters as C', e.g.:

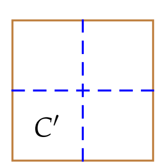

Do an **eval auction** on C'. Let n' be the number of agents whose bid is at least 1.

Case #1: $n' = 0$ (both agents value C' as less than 1). Denote another quarter as C' and do an eval auction again. Because the total cake value is 4, this can happen at most three times; eventually one of the other cases must happen.

Case #2: $n' = 1$. The single agent who values C' as at least 1 wins C' and goes home. The losing agent values C' as less then 1 so he values $C \setminus C'$ as more than 3. $C \setminus C'$ is a union of 3 squares, so by the Covering Lemma the losing agent can get from it a square with a value of at least 1.

Case #3: $n' = 2$. Do a **mark auction** with the following constraint: *mark a square with a value of 1 contained in C' and adjacent to a corner of C*. Both agents can bid as required, since they value C' as at least 1 so they have a square with a value of exactly 1 inside C'. An example is illustrated below, where the two bids are marked by dotted lines:

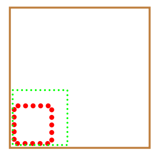

The winning bid (the smallest square) is marked with thicker dots. It is given to the winner, who walks home with a square worth 1. The remaining cake is an L-shape similar to the one in Figure 3.6. Its cover number is 3 and its value for the loser is at least 3. By the Covering Lemma, it contains a square whose value to the loser is at least 1.[15] The final allocation may look like:

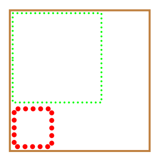

The fairness of this allocation is evident: both agents agree that the south-west is the most valuable district, so the agent who has to go to a less valuable district is compensated by a larger plot.

Note that some land remains unallocated. This is unavoidable if the pieces have to be square. Moreover, in realistic land-division scenarios it is common to leave some land unallocated and available for public use. □

3.5 Division procedures

In this section we use the building-blocks developed in Section 3.4 to create various division procedures.

3.5.1 Four and three walls, guillotine cuts

We develop simultaneously a pair of division procedures. Both procedures accept a cake C which is assumed to be the rectangle $[0, L] \times [0, 1]$, and return n disjoint square pieces $\{X_i\}_{i=1}^{n}$ such that for every agent i: $V_i(X_i) \geq 1$.

The two procedures differ in their requirement on L (the length/width ratio of the cake) and in the number of "walls" (bounded sides) they assume on the cake:

- The *3-walls procedure* requires that $L \in [0, 1]$ and it guarantees that the allocated squares are contained in $[0, \infty] \times [0, 1]$ (in other words, there is no wall in the rightmost edge of the cake).
- The *4-walls procedure* requires that $L \in [1, 2]$ (i.e, the cake is a 2-fat rectangle) and it guarantees that all allocated squares are contained in C.

Additionally, the two procedures differ in their requirement on the total cake value:

- The 3-walls procedure requires that for every agent i: $V_i(C) \geq \max(1, 4n - 5)$.
- The 4-walls procedure requires that for every agent i: $V_i(C) \geq \max(2, 4n - 4)$.

[15]Combining the lower bound proved by Example 3.4.5 with the upper bound proved by Claim 3.3.2 gives a tight result for two agents: Prop$(Square, Squares, n = 2) = 1/4$.

The procedures are developed by induction on the number of agents. We first consider the base case in which there is a single agent ($n = 1$).

In the 3-walls procedure, the single agent values C as at least 1. The square $[0,1] \times [0,1]$ contains all the value of C and it is contained within its three walls, so it can be given to the single agent:

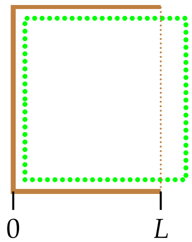

In the 4-walls procedure, the single agent values C as at least 2. The requirement on L guarantees that the cake can be covered by at most 2 squares:

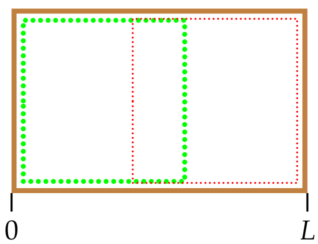

Hence, by the Covering Lemma, the single agent can be given a square with a value of at least 1.

We now assume that we can handle any number of agents less than n. Given n agents ($n \geq 2$), we proceed as follows.

3 Walls procedure

At this point, there are $n \geq 2$ agents who value the cake as at least $4n - 5$.

(1) **Mark auction**. Ask each agent to mark a rectangle with a value of exactly 1 adjacent to the rightmost edge of the cake (the edge without the wall):

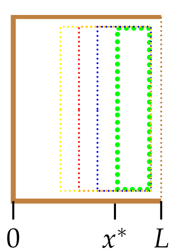

The winning bid (marked by thicker dots above) is a rectangle $[x^*, L] \times [0,1]$. There are two cases:

- *Easy case*: $x^* \geq 1/2$. Make a vertical guillotine cut at x^*. Give to the winner the square $[x^*, x^* + 1] \times [0,1]$. This square contains the winning bid, so its value for the winner is at least 1. The remaining cake is a 2-fat rectangle and its value for the remaining $n - 1$ agents is at least $V(C) - 1 \geq 4n - 6 \geq \max(2, 4(n-1) - 4)$. Use the *4 walls procedure* to divide the remainder among the losers.

- *Hard case*: $x^* < 1/2$. Now we cannot let the winner have the winning bid, since the remainder will be too thin for the remaining agents. Our solution relies on the following observation: the fact that $x^* < 1/2$ means that all agents value the rectangle $[1/2, L] \times [0,1]$ as less than 1. Therefore, they value the rectangle $[0, 1/2] \times [0,1]$ as at least $4n - 6$. Since all agents believe that this "far left" rectangle is so valuable, we are going to do an *eval auction* inside it.

(2) **Eval auction**. Let $C' = [0, 1/2] \times [1/2, 1]$ and $C'' = [0, 1/2] \times [0, 1/2]$:

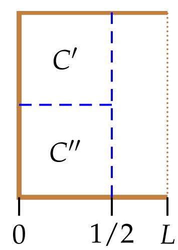

Do an eval auction on C'. Order the agents in a descending order of their bid, $V_1(C') \geq \cdots \geq V_n(C')$, and let n' be the largest integer with:

$$V_{n'}(C') \geq \max(4n' - 5, 1)$$

If $n' = n$ then all agents value C' as the entire cake, so the other parts of the cake can be discarded and the division procedure can start again with C' as the cake. Hence, we assume that $n' < n$. There are two main cases to consider:

- *Easy case*: $1 \leq n' \leq n - 2$. Make a horizontal guillotine cut between C' and C''. Use the 3-walls procedure to divide C' among the n' winners.

  The losers value C' as less than $\max(4(n'+1) - 5, 1) = 4n' - 1$. At this point all agents value the rectangle $C' \cup C''$ as at least $4n - 6$; hence, all losers value C'' as at least $(4n - 6) - (4n' - 1) = 4(n - n') - 5$. Since $n - n' \geq 2$, this value is also larger than 1, so we can use the 3-walls procedure to divide C'' among the $n - n'$ losers.

  Note that no square is allocated to the right of the line $x = 1/2$, so we can assume that the rightmost border of both C' and C'' is open and use the 3-walls procedure to divide them.

- *Hard case*: $n' = 0$. This means that all agents value C' as less than 1, so they value C'' as at least $4n - 7$. Now we have a problem: we cannot give C' even to a single agent since it is not sufficiently valuable, but we also cannot divide C'' among all n agents since it too is not sufficiently valuable.

  Our solution is to shrink C'' towards the corner, until one of the agents decides that it is better to take a piece outside C'' and leave C'' to the remaining $n - 1$ agents. This solution is implemented using a *mark auction*, which is described in detail in step (3) below. But before proceeding there is one more case that must be handled:

- *Mixed case*: $n' = n - 1$. This is handled according to the bid of the single losing agent (agent n): if $V_n(C') < 4n - 7$, then the losing agent values C'' as at least 1, so we can proceed as in the Easy case (the winning agents receive C' and the losing agent receives C''). Otherwise, $V_n(C') \geq 4n - 7$, so all agents value C' as at least $4n - 7$ (because the agents are ordered in descending order of their bid). Switch the roles of C' and C'' (e.g. by reflecting the cake about the line $y = 1/2$), and proceed as in the hard case to the next auction.

(3) **Mark auction.** Ask each agent to mark an L-shape with a value of exactly 2, the complement of which is a square inside C'' with a value of $4n - 7$ cornered at the corner of C, like this:

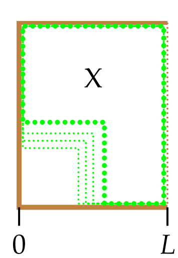

Let X be the winning bid. X can be covered by two overlapping pieces: a square near the top-left corner of C (denoted by Y below) and a square overlapping the right edge of C (denoted by Z below):

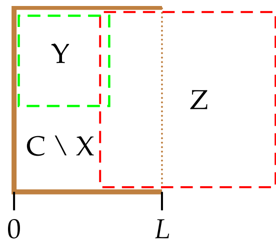

At least one of these squares must have a value of at least 1 to the winner. If Y has value 1 then give Y to the winner and leave Z unallocated; otherwise, give Z to the winner, leave Y unallocated and rotate C clockwise 90°. In both cases, $C \setminus X$ can be separated from the piece given to the winner using a horizontal guillotine cut. Moreover, in both cases the cake to the right of $C \setminus X$ is unallocated. The remaining $n-1$ agents value $C \setminus X$ as at least $(4n-5)-2$, which is more than $\max(1, 4(n-1)-5)$. Use the 3 walls procedure to divide $C \setminus X$ among them. □

4 Walls procedure

At this point, there are $n \geq 2$ agents who value the cake as at least $4n-4$.

The 4-walls procedure is similar to the 3-walls procedure except that it has one additional *eval* auction at the beginning. If this auction succeeds, then it effectively cuts the cake to two halves each of which is a 2-fat rectangle, so each half can be divided recursively using the 4-walls procedure. If this auction fails (as will be explained below), then the situation is similar to the 3-walls procedure and we can use a similar sequence of three auctions.

(0) **Eval auction.** Let $C' = [L/2, 1] \times [0, 1]$ = the rightmost half of C. Note that both C' and its complement are 2-fat rectangles:

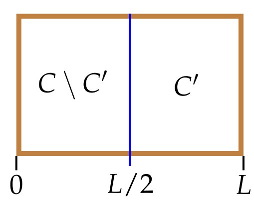

Do an eval auction on C'. Order the agents in a descending order of their bid, $V_1(C') \geq \cdots \geq V_n(C')$, and let n' be the largest integer with:

$$V_{n'}(C') \geq \max(4n'-4, 2)$$

If $n' = n$ then for all agents $V_i(C') = V_i(C)$, so $C \setminus C'$ can be ignored and the procedure can be restarted with C' as the entire cake. Hence, there are two non-trivial cases to consider:

- *Easy case*: $1 \leq n' \leq n-2$. Make a vertical guillotine cut between C' and $C \setminus C'$. Use the 4-walls procedure to divide C' among the n' winners. This is possible since C' is a 2-fat rectangle and all winners value it as at least $\max(4n'-4, 2)$.

  The losers value C' as less than $\max(4(n'+1)-4, 2) = 4n'$, so they value the remaining half $C \setminus C'$ as more than $(4n-4)-4n' = 4(n-n')-4$. Since $n-n' \geq 2$, this value is also larger than 2. Use the 4-walls procedure to divide $C \setminus C'$ among the $n-n'$ losers; this is possible since $C \setminus C'$ is a 2-fat rectangle and all losers value it as at least $\max(4(n-n')-4, 2)$.

- *Hard case*: $n' = 0$. This means that all agents value C' as less than 2 so they value the remainder $C \setminus C'$ as at least $4n-6$. We are going to enlarge C' leftwards, until it becomes sufficiently valuable such that some agent is willing to accept it. We implement this solution using a *mark auction*, described in step (1) below. But beforehand, one more case must be handled:

- *Mixed case*: $n' = n-1$. This case is handled according to the bid of the losing agent: if $V_n(C') < 4n-6$, then the losing agent values $C \setminus C'$ as at least 2, so we can proceed as in the Easy case (the winning agents receive C' and the losing agent receives $C \setminus C'$). Otherwise, $V_n(C') \geq 4n-6$, so all agents value C' as at least $4n-6$. Switch the roles of C' and $C \setminus C'$ (e.g. by reflecting the cake C about the line $x = L/2$), and proceed as in the hard case to the next auction.

(1) **Mark auction.** Ask each agent to mark a rectangle with a value of exactly 2 adjacent to the rightmost edge of C:

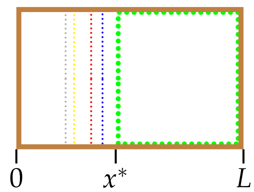

The smallest rectangle wins. Let x^* be the x coordinate of its leftmost edge, so the winning bid is $[x^*, L] \times [0,1]$. Since all agents value C' as less than 2, all bids must contain C', so $x^* \le L/2$. There are two cases:

- *Easy case*: $x^* \ge 1/2$. Make a vertical guillotine cut at x^*. Both the winning bid and its complement are 2-fat rectangles. By the Covering Lemma, the winner can be allocated from its bid a square with a value of at least 1. The $n-1$ losers value the remaining cake, $[0, x^*] \times [0,1]$, as at least $4n-6$, which is at least $\max(2, 4(n-1)-4)$. Hence, the 4-walls procedure can be used to divide the remainder among the losers.

- *Hard case*: $x^* < 1/2$. Now we cannot let the winner have the winning bid, since the remainder will be too thin for the remaining agents. But we know that all agents value the rectangle $[1/2, L] \times [0,1]$ as less than 2 so they value the rectangle $[0,1/2] \times [0,1]$ as at least $4n-6$. Since all agents believe that this rectangle is so valuable, we are going to do an *eval auction* inside it.

(2) **Eval auction**. Let $C' = [0,1/2] \times [1/2,1]$ and $C'' = [0,1/2] \times [0,1/2]$:

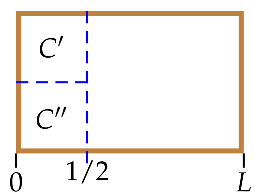

Do an eval auction on C' and let n' be the largest integer with:

$$V_{n'}(C') \ge \max(4n'-5, 1)$$

As in step (0), the case $n' = n$ is trivial and can be ignored. There are two non-trivial cases:

- *Easy case*: $1 \le n' \le n-2$. Make a horizontal guillotine cut between C' and C''. Use the *3-walls procedure* to divide C' among the n' winners. The 3-walls procedure might allocate pieces that flow over the right boundary of C' (the line $x = 1/2$). This does not cause any problem because the side-length of these rectangles is at most $1/2$, so they are still contained in the original cake C.

  The losers value C' as less than $\max(4(n'+1)-5, 1) = 4n'-1$. At this point of the procedure, all agents value the rectangle $C' \cup C''$ as at least $4n-6$; hence, all losers value C'' as at least $(4n-6) - (4n'-1) = 4(n-n')-5$. Since $n - n' \ge 2$, this value is also larger than 1, so we can use the 3-walls procedure to divide C'' among the $n-n'$ losers.

- *Hard case*: $n' = 0$. This means that all agents value C' as less than 1 and value C'' as at least $4n-7$. We are going to "shrink" C'' using a mark-auction in step (3). But beforehand we handle the remaining case:

- *Mixed case*: $n' = n-1$. Proceed according to the bid of the losing agent: if $V_n(C') < 4n-7$, then the losing agent values C'' as at least 1, so we can proceed as in the Easy case (the winning agents receive C' and the losing agent receives C''). Otherwise, $V_n(C') \ge 4n-7$, so all agents value C' as at least $4n-7$. Switch the roles of C' and C'', and proceed as in the hard case to the next auction.

(3) **Mark auction.** Ask each agent to mark an L-shape with a value of 3, whose complement is a square inside C'' cornered at the corner of C, like this:

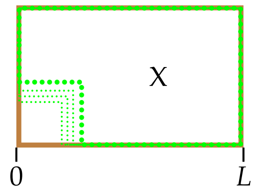

Since all agents value C'' as at least $4n-7=(4n-4)-3$ they can indeed bid as required. Let X be the winning bid. X is an L-shape that can be covered by two overlapping pieces: a square near the top-left corner of C (denoted by Y below) and a rectangle near the right edge of C (denoted by Z below):

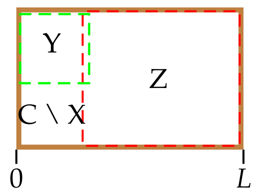

Since the winner values X as 3, at least one of the following must hold:

- The winner values Y as at least 1; if this is the case then the winner receives Y, and Z remains unallocated.
- The winner values Z as at least 2; if this is the case then the winner selects a square from Z with a value of at least 1 (this is possible by the Covering Lemma since Z is a 2-fat rectangle), and Y remains unallocated. If this is the case, then rotate C clockwise 90°.

In both cases, $C \setminus X$ can be separated from the piece given to the winner using a horizontal guillotine cut. In both cases, the cake to the right of $C \setminus X$ is unallocated. The $n-1$ losers value X as at most 3 so they value $C \setminus X$ as at least $(4n-4)-3$, which is at least $\max(1, 4(n-1)-5)$. Therefore, the 3 walls procedure can be used to divide $C \setminus X$ among them. □

The above pair of procedures prove the following pair of positive results $\forall n \geq 2$:

$$\text{Prop}(2\,fat\,rectangle\,with\,all\,sides\,bounded,\,Squares,\,n) \geq \frac{1}{4n-4}$$
$$\text{Prop}(Rectangle\,with\,a\,long\,side\,unbounded,\,Squares,\,n) \geq \frac{1}{4n-5}$$

Since a square is a 2-fat rectangle:

$$\text{Prop}(Square\,with\,4\,walls,\,Squares,\,n) \geq \frac{1}{4n-4}$$
$$\text{Prop}(Square\,with\,3\,walls,\,Squares,\,n) \geq \frac{1}{4n-5}$$

Fat rectangle pieces

When the pieces are allowed to be R-fat rectangles, the above lower bounds are of course still true, since a square is an R-fat rectangle. But when $R \geq 2$, the 4-walls division procedure can give slightly stronger guarantees — the required value is $\max(1, 4n-5)$ instead of $\max(2, 4n-4)$ (this is analogous to the fact that in Subsection 3.3.4, when the pieces are allowed to be R-fat rectangles with $R \geq 2$, our upper bound for a cake with 4 walls is slightly weaker — the denominator is $2n-1$ instead of $2n$). The required modifications are briefly explained below:

- In the base case ($n=1$), since the cake is 2-fat, the single agent can have it all, so it is sufficient that its value be 1.

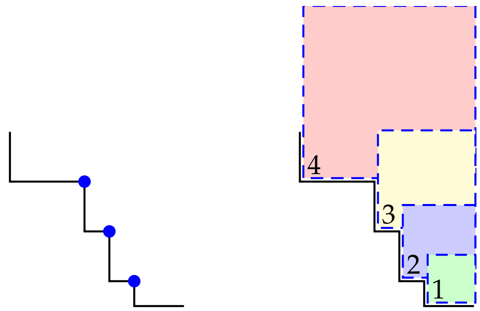

Figure 3.7: A staircase with $T = 3$ teeth marked by discs (Left). It has $T + 1 = 4$ corners and can be covered by 4 squares (Right).

- In step (0), after the Eval auction, n' is the largest integer with $V_{n'}(C') \geq \max\boldsymbol{(4n' - 5, 1)}$. In the easy case, the n' winners value their share C' as at least $\max(4n' - 5, 1)$ and the $n - n'$ losers value their share $C \setminus C'$ as at least $\max(4(n - n') - 5, 1)$, so each part can be divided recursively using the 4-walls procedure. In the hard case, all agents value C' as less than 1 so they value the remainder $C \setminus C'$ as at least $4n - 6$; proceed to the next step.
- In step (1), the Mark auction asks each agent to mark a rectangle with a value of exactly **1** adjacent to the rightmost edge of C. In the easy case, both the winning bid and its complement are 2-fat rectangles. The winning bid can be given entirely to the winner; the $n - 1$ losers value the remaining cake as at least $4n - 6$, which is at least $\max(1, 4(n - 1) - 5)$, so the 4-walls procedure can be used to divide the remainder among them. In the hard case, all agents value the rectangle $[1/2, L] \times [0, 1]$ as less than 1 so they value the rectangle $[0, 1/2] \times [0, 1]$ as at least $4n - 6$; proceed to the next step.
- In step (2), the Eval auction proceeds exactly as in the case of square pieces. The values are sufficient for using the 3-walls procedure.
- In step (3), the Mark auction asks each agent to mark an L-shape with a value of exactly **2**. Let X be the winning bid. Since the winner values X as 2, he values either its topmost part or its rightmost part as at least 1; both these parts are 2-fat rectangles so the winner can pick one of them and get a fair share. In both cases, $C \setminus X$ (which is a square) can be separated from the piece given to the winner using a horizontal guillotine cut. In both cases, the $n - 1$ losers value X as at most 2 so they value $C \setminus X$ as at least $(4n - 5) - 2$, which is at least $\max(1, 4(n - 1) - 5)$. Therefore, the 3 walls procedure can be used to divide $C \setminus X$ among them.
- The 3-walls procedure remains unchanged.

So for every $n \geq 2$ and $R \geq 2$:

$$\text{Prop}(2\, fat\, rectangle\, with\, all\, sides\, bounded,\ R\, fat\, rectangles,\ n) \geq \frac{1}{4n - 5}$$

3.5.2 Two walls

We present a division procedure for dividing the top-right quarter-plane, i.e, the cake is a square with two walls and two unbounded sides. We would like to do a *mark auction* in which each agent is asked to mark a square adjacent to the bottom-left corner. Then, the smallest square should be allocated to its bidder and the remaining cake should be divided among the remaining agents. However, when we try to do this we run into trouble, as the remaining cake is no longer a quarter-plane.

As it often happens, the solution is to generalize the problem. Instead of dividing a quarter-plane, we divide a *rectilinear polygonal domain unbounded in two directions*, which for brevity we call “staircase” because of its shape (see Figure 3.7).

Each staircase has vertexes with inner angle 90° and vertexes with inner angle 270°; we call the former *corners* and the latter *teeth*.[16] A staircase with T teeth has $T + 1$ corners. A quarter-plane is a staircase with $T = 0$ teeth.

[16]Other common names are *convex vertexes* vs. *concave/reflex vertexes*, or *inner corners* vs. *outer corners*.

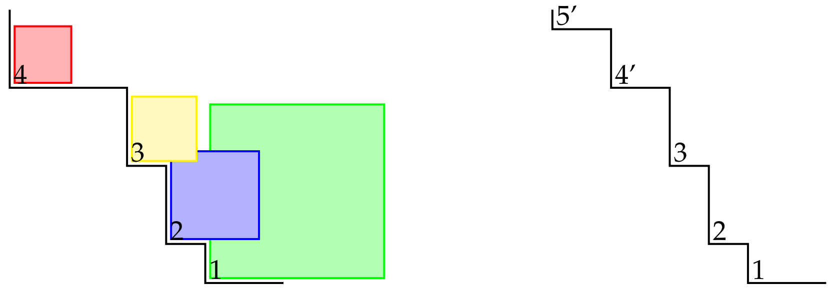

Figure 3.8: The square at corner 4 is entirely contained in the corner (left). After it is allocated, the remaining cake is a staircase with 4 teeth and 5 corners (right).

By putting the arrangement of Claim 3.3.3 in one of the corners and adding a pool in each of the other T corners, the following upper bound is obtained:

$$\text{Prop}(T\,staircase, Squares, n) \leq \frac{1}{2n-1+T}$$

We normalize the valuations of all agents such that the value of the entire cake is $2n-1+T$. We use a sequence of mark auctions to give each agent a square with a value of at least 1.

We proceed by induction on the number of agents n. When $n=1$, the cake value for the single agent is at least $T+1$. The cake can be covered by $T+1$ sufficiently large squares — one square per corner (see Figure 3.7/Right). By the Covering Lemma, the agent can get a square with a value of at least 1.

Suppose we already know how to divide a T-staircase to $n-1$ agents, for every integer $T \geq 0$. Now there are n agents. Start by doing $T+1$ **mark auctions:** for each corner $j \in \{1, \ldots, T+1\}$, ask each agent to mark a square with a value of exactly 1 adjacent to corner j. If the total value of the agent in corner j is less than 1, then the agent is allowed to not participate in that auction, or equivalently mark a square with an infinite side-length. By the Covering Lemma, each agent can mark at least one finite square.

In each corner, the "corner-winning-bid" is the smallest square (contained in all other bids in that corner). We now have $T+1$ corner-winners, and we have to select a single global-winner. There are two cases.

Easy case: there is a $j \in \{1, \ldots, T+1\}$ such that the corner-j winning-bid is smaller than the two edges of C adjacent to corner j. An example is the square in corner 4 in Figure 3.8. Select one such square arbitrarily as the global "winning bid". Give the winning bid to its bidder. The remaining cake is a staircase with $T+1$ teeth (see Figure 3.8). The $n-1$ losing agents value the allocated square as at most 1, so they value the remaining staircase as at least $(2n-1+T)-1 = 2(n-1)-1+(T+1)$. Hence, by induction we can divide the remainder among the losers.

Hard case: all corner-winning-bids are larger than the edges adjacent to their corners, as in Figure 3.9. Now, when a square is allocated, the remainder is no longer a staircase. In order to restore the staircase shape, we have to remove an additional part of C. We do this by cutting, from the top-right corner of the allocated square, a straight line downwards to the bottom boundary of C, and a straight line leftwards to the leftmost boundary of C. The parts of C that are removed besides the allocated square are called the *shadows* of the square. An example is illustrated in Figure 3.9, where the square at corner 2 has two shadows denoted by dotted lines.

We now need the following geometric lemma, which is formally stated and proved in Appendix 3.A:

Lemma 3.5.1. *(Staircase Lemma) Given a staircase in which a square is located in each corner, there exists a square whose shadows are contained in the union of the other squares.*

Based on the Staircase Lemma, we proceed as follows. From the $T+1$ corner-winning-bids, select one square whose shadows are contained in the other squares (e.g. the square in corner 2 in Figure 3.9). Declare this square as the global winning square, give it to its bidder, and remove its shadows from C.

We have to prove that the remaining cake is sufficiently valuable for each losing agent. The number of agents changes by $\Delta n = -1$ since the winning agent leaves. The cake value for a losing agent changes

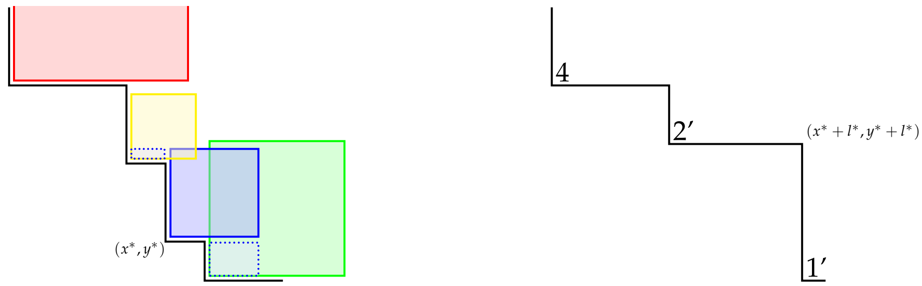

Figure 3.9: The square at corner 2 (second from the bottom-right) satisfies the Staircase Lemma, since its "shadows" (dotted) are contained in the other squares. After it is allocated, the remaining cake is a staircase with 2 teeth and 3 corners (right).

by ΔV (a negative quantity). The number of teeth changes by ΔT which may be positive or negative. Looking at the value requirement $V \geq 2n + T - 1$, we see that in order to use the induction assumption, it is sufficient to prove that for every loser:

$$\Delta V \geq 2\Delta n + \Delta T = \Delta T - 2$$

I.e, the value of the remaining agents should drop by at most two units, plus one unit for each removed tooth.

The shadows of the winning square can be partitioned to m disjoint rectangular components, to its top-left and to its bottom-right, such that each component is located in a different corner (e.g. in Figure 3.9, $m = 2$). After the shadows are removed, m teeth disappear. One tooth is added at the top-right of the winning square. Hence, $\Delta T = 1 - m$.

The winning square is worth at most 1 for the remaining agents, since it is contained in all other squares in its corner. By the selection of the global-winning-bid, each of the m shadows is contained in a corner-winning-bid, so its value for the losing agents is at most 1. Hence, the total value of the removed region to the $n - 1$ losers is at most $m + 1$, so $\Delta V \geq -1 - m = \Delta T - 2$, as required. Hence, by the induction assumption we can proceed and divide the remainder among the losers. [17] □

The above procedure proves that, for every $n \geq 1, T \geq 0$:

$$\text{Prop}(T\,staircase, Squares, n) = \frac{1}{2n - 1 + T}$$

By letting $T = 0$ we get:

$$\text{Prop}(Quarter\ plane,\ Squares,\ n) = \frac{1}{2n - 1}$$

3.5.3 One and zero walls

A half-plane can be divided by partitioning it to two quarter-planes:

Claim 3.5.1. For every $n \geq 2$:

$$\text{Prop}(Half\ plane,\ Squares,\ n) \geq \frac{1}{2n - 2}$$

Proof. Assume the cake is the half-plane $y \geq 0$ and there are n agents who value it as $2n - 2$. Do the following *mark auction*: ask each agent to mark a quarter-plane open to the top-left, whose bottom edge is adjacent to the bottom edge of C and its value is exactly 1. An example is illustrated below, where the winning bid is — as usual — marked by thicker dots:

[17] The easy case is, in fact, contained in the hard case, since a square smaller than the edges adjacent to its corner has an empty shadow (so $m = 0$). The split to easy and hard cases is done for presentation purposes only.

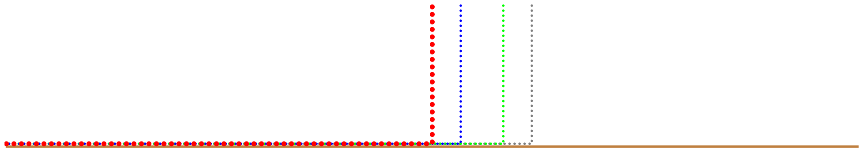

After the winning bid is allocated to its winner, the $n-1$ losers value the remaining quarter-plane as at least $(2n-2)-1 = 2(n-1)-1$; divide it among them using the procedure of Subsection 3.5.2. □

An unbounded plane can be divided by partitioning it to two half-planes:

Claim 3.5.2. For every $n \geq 4$:

$$\text{Prop}(Plane, Squares, n) \geq \frac{1}{2n-4}$$

Proof. Normalize the cake value to $2n-4$. Do the following *mark auction*: ask each agent to mark a half-plane bounded at its bottom, with a value of exactly 2 (so each agent i marks a half-plane $Y_i = [-\infty, \infty] \times [y_i, \infty]$). Order the bids by containment, so that $Y_1 \subseteq Y_2 \subseteq \cdots \subseteq Y_n$. Select two winners — the agents with the two smallest bids (Y_1 and Y_2). Both winners value Y_2 as at least 2; divide it among them using cut-and-choose. Each of them receives a quarter-plane with a value of at least 1. The $n-2$ losers value the remaining half-plane as at least $(2n-4)-2 = 2(n-2)-2$; divide it among them using the procedure of Claim 3.5.1. □

The lower bounds for one and zero walls do not match the upper bounds proved in Section 3.3: the proportionality coefficient (the coefficient of n in the denominator) is 2 in both cases, while the coefficients in the upper bounds are 3/2 for a half-plane and almost 1 for an unbounded plane. We believe that the procedures presented above are tight and the "real" coefficient is 2. The reason is that, whenever a plane is cut by even a single straight line, the remainder is a half-plane, and when a half-plane is cut, the remainder is a quarter-plane, and for a quarter-plane we know that the proportionality coefficient is 2. In future work we plan to look for tighter impossibility results showing that the proportionality coefficient is indeed 2 in half-planes and unbounded planes, too.

3.5.4 Three walls

Our next goal is to divide a square bounded by three walls. We already presented a procedure for a square with three walls in Subsection 3.5.1, but the value guarantee of the present procedure is better and it matches the upper bound of $1/(2n-1)$. On the other hand, the present procedure uses general (non-guillotine) cuts.

Similarly to the two-walls case, we have to generalize the problem and divide a *rectilinear polygonal domain unbounded in one direction*, which for brevity we call a "valley". Again the number of teeth is denoted by T; see Figure 3.10.

We require the valley to have the *Sunlight property*, which means that light coming from the top can reach all parts of the bottom border. In other words: no part of the valley lies below a wall; the bottom border of a valley goes from the left wall (at $x=0$) to the right wall (at $x=1$) in stairs climbing to the top-right or bottom-right, but never back to the left. Hence a valley can be represented as a sequence of $T+1$ *levels* $\{[x^{\min}, x_i^{\max}] \times y_i\}_{i=1}^{T+1}$, where (see Figure 3.10):

$$0 = x_1^{\min} < x_1^{\max} = x_2^{\min} < x_2^{\max} \cdots < x_T^{\max} = x_{T+1}^{\min} < x_{T+1}^{\max} = 1$$

Our valley-division procedure is essentially similar to the staircase-division procedure: a mark-auction is performed in each "corner" of the valley; the smallest bid in each corner is the corner-winning-bid; and a global winning-bid is selected such that its "shadows" are contained in all other bids. We have to carefully define the "corners" and the "shadows", and this requires several definitions.

The structure of a valley

For every level $i \in \{1, \ldots, T+1\}$, when we look from $(x_i^{\min}, y_i)$ leftwards, we see a wall. Let x_i^{left} be the x coordinate of that wall and y_i^{left} be the y coordinate of the level at the top of the wall ($y_i^{\text{left}} > y_i$). If $(x_i^{\min}, y_i)$

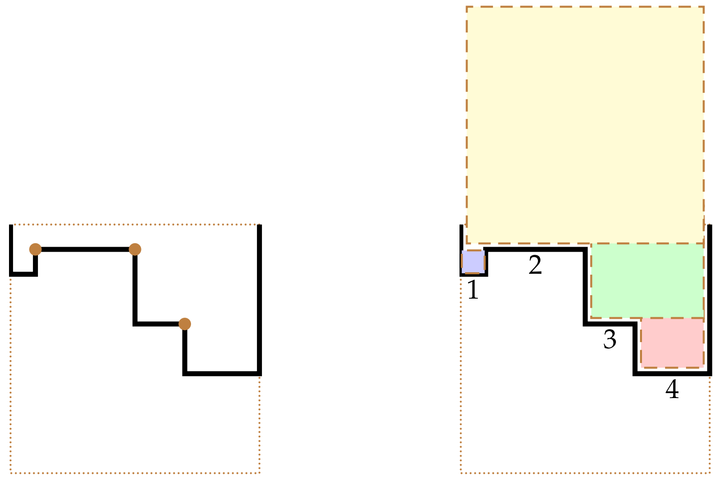

Figure 3.10: A valley with $T = 3$ teeth marked by discs (Left). It has $T + 1 = 4$ levels and can be covered by 4 squares (Right). The levels coordinates are: $[0, .1] \times .8$, $[.1, .5] \times .9$, $[.5, .7] \times .6$, $[.7, 1.0] \times .4$. The levels are covered from bottom to top: 4, then 3, then 1, then 2. In each level, the bottom rectangle, which is not overlapped by higher squares, is the *covering rectangle* of that level.

is a bottom-left corner (such as in levels 1 and 3 and 4 in Figure 3.10), then $x_i^{\text{left}} = x_i^{\min}$ and $y_i^{\text{left}} = y_{i-1}$ (if $x_i^{\text{left}} = 0$, i.e. we hit the left boundary, then we define $y_i^{\text{left}} = 1$). Otherwise (as in level 2), $x_i^{\text{left}} < x_i^{\min}$.

Similarly, define x_i^{right} as the x coordinate of the wall we see at the right and y_i^{right} as the y coordinate of the level at the top of the wall ($y_i^{\text{right}} > y_i$). If $(x_i^{\max}, y_i)$ is a bottom-right corner (such as in levels 1 and 4 in the figure), then $x_i^{\text{right}} = x_i^{\max}$ and $y_i^{\text{right}} = y_{i+1}$ (if $x_i^{\text{right}} = 1$, i.e. we hit the right boundary, then we define $y_i^{\text{right}} = 1$). Otherwise (as in levels 2 and 3), $x_i^{\text{right}} > x_i^{\max}$.

The horizontal distance between the two walls surrounding a level is denoted:

$$dx_i := x_i^{\text{right}} - x_i^{\text{left}}$$

In the figure, the values of dx_i for the 4 levels are: $0.1, 1.0, 0.5, 0.3$. The vertical depth of a level is denoted by:

$$dy_i := \min(y_i^{\text{right}}, y_i^{\text{left}}) - y_i$$

It is the height to which one has to climb in order to move to another level, or to exit the unit square. In the figure, the values of dy_i for the 4 levels (from left to right) are: $0.1, 0.1, 0.3, 0.2$.

Initially we handle the case of a single agent. This requires a bound on the square-cover-number of the valley, as a function of T. In general, the square-cover-number of a valley can be arbitrarily large, e.g, if the valley has a single level $[0, 1/m] \times 0$, then m squares are required to cover it, for every integer m. For our purposes, we can restrict our attention to valleys that do not have such deep levels. Formally, we require the valley to have the *Shallowness property*, which means that for every level i:

$$dy_i \le dx_i$$

This property guarantees that the valley can be covered by at most $T + 1$ squares, as we show in the following subsection.

Covering a valley with squares

Lemma 3.5.2. *If C is a valley with T teeth satisfying the Shallowness property, then:*

$$CoverNum(C, Squares) \le T + 1$$

Proof. Consider the lowest level — the level i with the smallest y_i. Consider the square:

$$S_i := [x_i^{\text{left}}, x_i^{\text{left}} + dx_i] \times [y_i, y_i + dx_i]$$

Because this is the lowest level, both its endpoints are inner corners, so $x_i^{\text{left}} = x_i^{\min}$ and $x_i^{\text{left}} + dx_i = x_i^{\text{right}} = x_i^{\max}$.

The Shallowness property guarantees that $dx_i \geq dy_i$. Hence, $y_i + dx_i \geq y_i + dy_i = \min(y_i^{\text{right}}, y_i^{\text{left}})$. Hence, S_i contains the rectangle:

$$R_i := [x_i^{\text{left}}, x_i^{\text{right}}] \times [y_i, \min(y_i^{\text{left}}, y_i^{\text{right}})]$$

Call R_i the *covering rectangle of level i* (see Figure 3.10/Right). If we remove from the valley the covering rectangle of i (the lowest level), then at least one of the teeth adjacent to it (from the left or from the right) is flattened, and we remain with at most $T-1$ teeth. In some remaining levels j, the $x_j^{\min}$ and $x_j^{\max}$ values might change, but the x_j^{left} and x_j^{right} do not change since the removed level was lower than all surrounding levels. Hence, dx_j and dy_j do not change, the Shallowness property is preserved, and we can continue this process iteratively until all the valley is covered. The number of squares in the covering is at most the number of levels, $T+1$. □

The division procedure

We are now ready to present the valley-division procedure.

We normalize the valuations of all agents such that the value of the entire valley for each agent is $2n - 1 + T$. We use a sequence of mark auctions to give each agent a square with a value of at least 1.

We proceed by induction on the number of agents n. When $n = 1$, the value for the single agent is at least $T + 1$. By Lemma 3.5.2 the valley can be covered by $T + 1$ squares, so by the Covering Lemma the agent can get a square with a value of at least 1.

Suppose we already know how to divide a T-valley to $n-1$ agents, for every integer $T \geq 0$. Now there are n agents. Start by doing $2(T+1)$ **mark auctions**. There are two auctions per level: one on the left and one on the right of its covering rectangle. For every level $i \in \{1, \ldots, T+1\}$, ask each agent to mark two squares with a value of exactly 1: a square with its bottom-left corner at the bottom-left corner of R_i (x_i^{left}, y_i) and a square with its bottom-right corner at the bottom-right corner of R_i $(x_i^{\text{right}}, y_i)$. The squares may overlap. An agent can refrain from participating in an auction if the largest square he can mark at this corner has a value of less than 1. By the Covering Lemma, each agent can participate in at least one auction.

In each corner, there are at most n squares. From these, we select a smallest square as the "corner-winning-bid". Now we have at most $2(T+1)$ corner-winners. The global-winner is **the square with a lowest top side**. I.e, if the side-length of the i-level winning-bid is l_i, then the global winner is a square with a smallest $y_i + l_i$.

In the illustration below, the index of each level is written below the level. There are squares only in 7 out of 10 corners, since no agents participated in the auction for the corner (x_3^{left}, y_3) (marked with x) and for level 5. The global-winner (marked with thicker dots) is the corner-winner at the corner $(x_1^{\text{right}}, y_1)$:

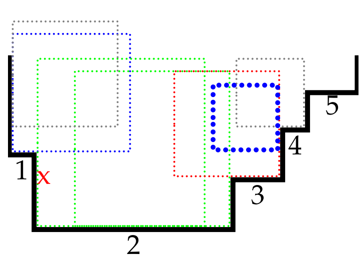

In addition to the winning square, we may have to remove some other parts of the valley, in order to ensure that the remaining valley satisfies the two properties defined above: the *Sunlight property* and the *Shallowness property*. We have to prove that this allocation leaves a sufficiently high value for the losing agents.

After all the removals, the number of agents changes by $\Delta n = -1$ since one agent leaves; the cake value for a losing agent changes by ΔV (a negative quantity); and the number of teeth changes by ΔT which may be positive or negative. Looking at the value requirement $V \geq 2n + T - 1$, we see that in order to use the induction assumption, it is sufficient to prove that for every loser:

$$\Delta V \geq 2\Delta n + \Delta T = \Delta T - 2$$

so the value of each loser should drop by at most two units, plus one unit for each removed tooth.

The following analysis depends on whether the winning square is adjacent to a right corner $(x_i^{\text{right}}, y_i)$ as in the illustration above, or a left corner (x_i^{left}, y_i). The two cases are entirely symmetric; henceforth we assume that the winning square is adjacent to a right corner.

First, we handle the *Sunlight property* by cutting from the left edge of the winning square down to the bottom border of C:

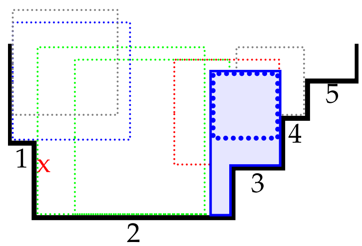

The winning square casts a shadow on $m \geq 0$ teeth below it, which are all removed. In the illustration above, $m = 1$. Additionally, a new tooth is added at the top-left of the winning square. Additionally, if the winning square is higher than the tooth at its right (as in the figure), then that tooth is removed and a new tooth is added at the top-right of the winning square. All in all, $\Delta T = 1 - m$.

The winning square casts a shadow on $1 + m$ levels. All squares of the losing agents in these levels are higher than the winning square; hence, the shadows of the winning square are contained in the losers' squares, and the total value of the shadows is at most $1 + m$. All in all, $\Delta V \geq -1 - m = \Delta T - 2$, as required.

Next, we have to handle the *Shallowness property* by removing *deep levels* — levels for which $dy_j > dx_j$ or equivalently:

$$\min(y_j^{\text{right}}, y_j^{\text{left}}) - y_j > x_j^{\text{right}} - x_j^{\text{left}} \tag{3.1}$$

This is done separately to the left and to the right of the winning square:

- A level to the left of the winning square ($j < i$) may become deep if the left edge of the winning square, and the cut from that edge downwards, becomes its rightmost wall:

$$x_j^{\text{right}} \leftarrow x_i^{\text{left}} \qquad\qquad y_j^{\text{right}} \leftarrow y_i + l_i$$
$$(y_i + l_i) - y_j > x_i^{\text{left}} - x_j^{\text{left}}$$

- A level to the right of the winning square ($j > i$) may become deep if the right edge of the winning square becomes its leftmost wall:

$$x_j^{\text{left}} \leftarrow x_i^{\text{right}} \qquad\qquad y_j^{\text{left}} \leftarrow y_i + l_i$$
$$(y_i + l_i) - y_j > x_j^{\text{right}} - x_i^{\text{right}}$$

In each side, we remove the highest deep level, and all levels below it. In the illustration below, only level 4 (to the right of the winning square) is removed:

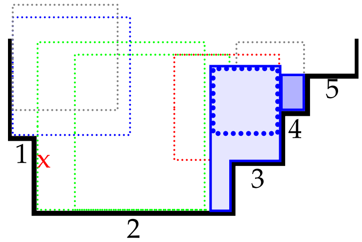

By selection of the global winner: $y_j + l_j \geq y_i + l_i$, which implies:

$$l_j \geq (y_i + l_i) - y_j \tag{3.2}$$

If a level $j < i$ becomes deep, then (3.2) implies:

$$l_j > x_i^{\text{left}} - x_j^{\text{left}}$$
$$\implies x_j^{\text{left}} + l_j > x_i^{\text{left}}.$$

In addition to $y_j + l_j \geq y_i + l_i$, this implies that the removed rectangle $[x_j^{\text{left}}, x_i^{\text{left}}] \times [y_j, \min(y_i + l_i, y_j^{\text{left}})]$ is contained in the corner-winner: $[x_j^{\text{left}}, x_j^{\text{left}} + l_j] \times [y_j, y_j + l_j]$. Hence, the value of the removed rectangle is at most 1. At most one unit of value is removed, and one tooth is removed. Hence, the balance between ΔV and ΔT is maintained.

If a level $j > i$ becomes deep, then (3.2) implies:

$$l_j > x_j^{\text{right}} - x_i^{\text{right}}$$
$$\implies x_j^{\text{right}} - l_j < x_i^{\text{right}}.$$

In addition to $y_j + l_j \geq y_i + l_i$, this implies that the removed rectangle $[x_i^{\text{right}}, x_j^{\text{right}}] \times [y_j, \min(y_i + l_i, y_j^{\text{right}})]$ is contained in the corner-winner: $[x_j^{\text{right}} - l_j, x_j^{\text{right}}] \times [y_j, y_j + l_j]$. Hence, the value of the removed rectangle is at most 1. At most one unit of value is removed, and one tooth is removed. The balance between ΔV and ΔT is maintained.

Finally, we have to handle the *Sunlight property* again by removing all levels below the levels removed in the previous step. We now prove that in all such levels, no agent marked any square. Indeed, let j be a level that became deep, and k be a level shadowed by it. So $y_k < y_j$ and $x_k^{\text{left}} > x_j^{\text{left}}$ and $x_k^{\text{right}} < x_j^{\text{right}}$. The side-length of any square marked in level k is at most $x_k^{\text{right}} - x_k^{\text{left}}$, so $l_k < x_k^{\text{right}} - x_k^{\text{left}} < x_j^{\text{right}} - x_j^{\text{left}}$ and:

$$y_k + l_k < y_j + (x_j^{\text{right}} - x_j^{\text{left}})$$

Combining this with (3.1) gives:

$$y_k + l_k < \min(y_j^{\text{right}}, y_j^{\text{left}}) \leq y_i + l_i$$

but this contradicts the assumption that i is the global-winning-square. Therefore, all levels below a deep level have a value of less than 1 to all agents. At most one unit of value is removed per level, so the balance between ΔV and ΔT is maintained.

To summarize: after allocating the winning square to the winner and removing some parts of the valley, we have a new valley with $T + \Delta T$ teeth satisfying the Sunlight and the Shallowness properties, and each losing agent values it as at least $((2n - 1 + T) + \Delta V) \geq ((2n - 1 + T) + (\Delta T - 2)) = 2(n - 1) - 1 + (T + \Delta T)$. Therefore, by the induction assumption we can continue to divide it among the $n - 1$ losers. □

The above procedure proves that, for every $n \geq 1, T \geq 1$:

$$\text{Prop}(T\, valley, Squares, n) = \frac{1}{2n - 1 + T}$$

A square with 3 walls is a valley with no teeth. It obviously satisfies the Sunlight property and the

Shallowness property. Letting $T = 0$ in the above formula yields:

$$\text{Prop}(\textit{Square with three walls}, \textit{Squares}, n) = \frac{1}{2n-1}$$

matching the upper bound.

Remark

We divided a 2-walls square by generalizing it to a "staircase", and divided a 3-walls square by generalizing it to a "valley". The natural next step is to divide a 4-walls square by generalizing it to a rectilinear polygon. This is a much more challenging task even for a single agent. The algorithmic problem of finding a minimal square-covering for a rectilinear polygon has been solved by Bar-Yehuda and Ben-Hanoch (1996), and we believe that their algorithm can be used for developing a rectilinear polygon division procedure. However, this algorithm is much more complicated than our covering algorithm of Subsection 3.5.4, so the division procedure will probably also be much more complicated.

In the next subsection we present a procedure for dividing a square using a different approach, which works only when the value measures are identical.

3.5.5 Four walls, guillotine cuts, identical valuations

Our procedures for identical valuations differ from the other procedures in that they do not use auctions, since all agents would make the same bids anyway.

We develop simultaneously a pair of division procedures. Both procedures accept a cake C which is assumed to be the rectangle $[0, 1] \times [0, L]$, and a *single* continuous value measure V. They return some disjoint square pieces $\{X_i\}$ such that for every i: $V(X_i) \geq 1$.

The two procedures differ in their requirement on L (the height/length ratio of the cake) and in the number of “walls” (bounded sides) they assume on the cake:

- The *fat-procedure* requires that $L \in [1, 2]$ (i.e, the cake is a 2-fat rectangle) and it guarantees that all allocated squares are contained in C;
- The *thin-procedure* requires that $L \in [2, \infty)$ (i.e, the cake is a "2-thin" rectangle) and it returns one of the following two outcomes:
  1. $n - 1$ squares contained in C (i.e, bounded by the 4 walls of the cake), or -
  2. n squares contained in $[0, \infty] \times [0, L]$, i.e, bounded by only 3 walls but may flow over the rightmost border. Every square that flows over the rightmost border is guaranteed to have its leftmost edge adjacent to the leftmost edge of C and its side-length at most $L - 1$ (the longer side of the cake minus its shorter side), so that all squares are contained in $[0, L - 1] \times [0, L]$.

Additionally, the two procedures differ in their requirement on the total cake value:

- The fat-procedure requires that $V(C) \geq 2n$.
- The thin-procedure requires that $V(C) \geq 2n - 2$.

The procedures are developed by induction on n. We first consider the base case $n = 1$:

- In the fat-procedure, the cake value is 2 and the cake is 2-fat, so by the Covering Lemma it contains a square with a value of at least 1.
- The thin-procedure can just return an empty set. This is an instance of the first outcome — $n - 1$ squares contained in C.

We now assume that both procedures work well for any number less than n. Given $n \geq 2$, we proceed as in the following subsections.

Henceforth, we make the following **positivity assumption**: every piece with positive area has positive value. This assumption is only for convenience: it simplifies the presentation and reduces the number of cases to consider. It can be dropped by adding sub-cases to each case in the procedures.

Fat procedure

At this point, the cake is a 2-fat rectangle with width 1 and height $L \in [1,2]$. Its total value is $2n$, and $n \geq 2$.

For every integer $u \in [0,2n]$, let y_u be the value $y \in [0,L]$ such that the cake below y has value u: $V([0,1] \times [0,y_u]) = u$. By the positivity assumption, y_u is unique, $y_0 = 0$ and $y_{2n} = L \geq 1$. Therefore, there exists a smallest $k \in [1,n]$ such that: $y_{2k} \geq 1/2$. Let $Bottom := [0,1] \times [0,y_{2k}] =$ the cake below y_{2k}; note that it is a 2-fat rectangle. Let $Top := C \setminus Bottom = [0,1] \times [y_{2k}, L] =$ the cake above y_{2k}. We have $V(Bottom) = 2k$ and $V(Top) = 2(n-k)$. Now there are two cases:

Case A: $L - y_{2k} \geq 1/2$ (this implies $k < n$). Thus $Bottom$ and Top are both 2-fat rectangles:

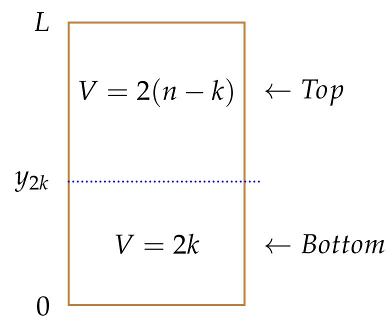

Apply the fat procedure to $Bottom$ and Top and get $k + (n-k) = n$ squares contained in C.

Case B: $L - y_{2k} < \frac{1}{2}$, so $Bottom$ is 2-fat and Top is 2-thin. Now consider y_{2k-2}. By definition of k, $y_{2k-2} < \frac{1}{2}$. Let $Bottom' := [0,1] \times [0, y_{2k-2}]$ and $Top' := C \setminus Bottom' = [0,1] \times [y_{2k-2}, L]$, so $V(Bottom') = 2(k-1) = V(Bottom) - 2$ and $V(Top') = 2(n-k+1) = V(Top) + 2$. Note that $Bottom'$ is 2-thin and is contained in $Bottom$, and Top' is 2-fat and contains Top:

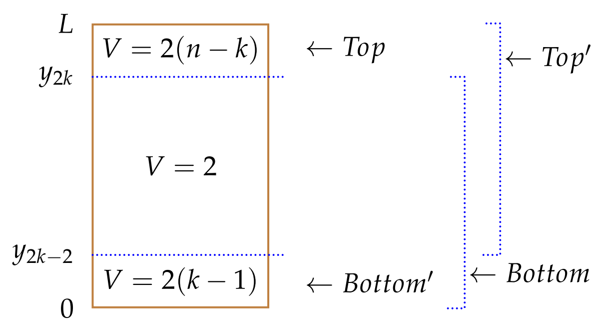

Because here $n \geq 2$, either $n - k \geq 1$ or $k - 1 \geq 1$ or both. Hence, at least one of the two 2-thin parts (Top, $Bottom'$) is non-empty and with value at least 2. Use the *thin procedure* to divide the non-empty thin part/s. In each part there are two possible outcomes: a smaller number of squares within 4 walls or a larger number of squares within 3 walls. There are several cases to consider.

— One easy case is that we get the 4-walls outcome in at least one of the parts — either in Top or in $Bottom'$ or in both. Suppose that we get the 4-walls outcome in $Bottom'$. So we have $k-1$ squares within the 4 walls of $Bottom'$. Ignore the outcome on Top and apply the fat procedure to Top'. This results in $n-k+1$ additional squares, so we have the required n squares. The situation is analogous if we get the 4-walls outcome in Top.

— Another easy case is that we get the 3-walls outcome in one part, and the other part is empty. Suppose that Top is empty (this implies $k = n$) and we get the 3-walls outcome in $Bottom'$. So we have $(k-1)+1 = n$ squares contained in $[0,1] \times [0, 1 - y_{2k-2}] \subseteq C$, as required. The situation is analogous if $Bottom'$ is empty and we get the 3-walls outcome in Top.

— The hard case is that both Top and $Bottom'$ are non-empty and the thin procedure on both of them returns the 3-walls outcome. Now we have k bottom squares and $n-k+1$ top squares, for a total of $n+1$ squares, e.g:

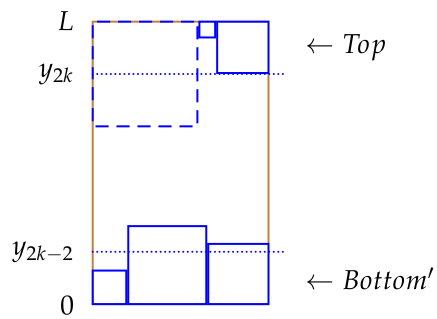

A potential problem in the last step is that some of the squares might overlap: some top squares might flow over the lower boundary of Top and overlap a bottom square, or some bottom squares might flow over the upper boundary of Bottom′ and overlap a top square. To prevent an overlap, we remove a single square — the largest of the $n+1$ squares (dashed square in the illustration above) — and return the remaining n squares.

It remains to prove that, indeed, after the largest square is removed, the remaining n squares do not overlap. The proof is purely geometric and it is delegated to Appendix 3.B. □

Thin procedure

At this point, the cake is a 2-thin rectangle with width 1 and height $L \in [2, \infty)$. Its total value is $2n-2$, and $n \geq 2$. The procedure is allowed to return one of two outcomes:

Outcome #1: $n-1$ squares bounded by the 4 walls of C, i.e, contained in $[0,1] \times [0,L]$, or —

Outcome #2: n squares bounded by the 3 walls of C, i.e, contained in $[0,\infty] \times [0,L]$. In this case, every square that flows over the rightmost border must have its leftmost edge adjacent to the leftmost edge of C (the edge $x=0$), and its side-length must be at most $L-1$ (the longer side of C minus its shorter side). This means that all n squares must be contained in $[0, L-1] \times [0, L]$.

We first handle the case $n=2$, in which $V=2$.

Select $y \in [0,L]$ such that $V([0,1] \times [0,y]) = V([0,1] \times [y,L]) = 1$. Proceed according to the value of y:

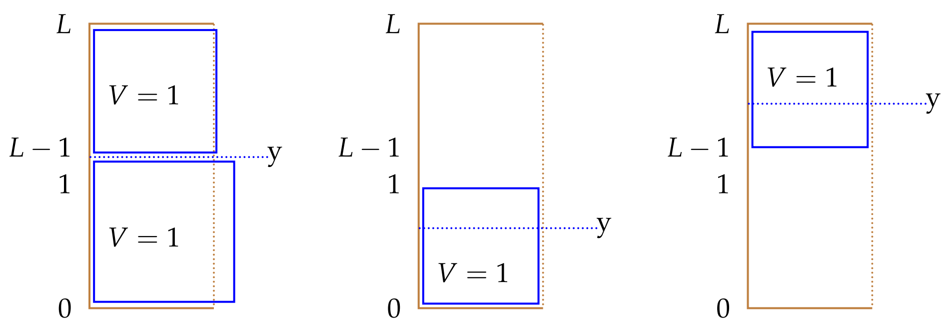

- If $y \in [1, L-1]$ (left) then return the two squares $[0,y] \times [0,y]$ and $[0, L-y] \times [y, L]$. Both squares are in $[0, L-1] \times [0, L]$ with their left side at $x=0$; this is an instance of outcome #2.
- If $y \in [0,1)$ (middle) then return $[0,1] \times [0,1]$; if $y \in (L-1, L]$ (right) then return $[0,1] \times [L-1, L]$. Both cases are instances of outcome #1.

From now on we assume that $n \geq 3$.

For every $u \in [0, 2n-2]$, define y_u as the value $y \in [0,L]$ such that the cake below y has value u: $V([0,1] \times [0, y_u]) = u$. By the positivity assumption, y_u is unique and $y_0 = 0$ and $y_{2n-2} = L$. Therefore, there exists a smallest $k \in [1, n-1]$ such that: $y_{2k} \geq \frac{1}{2}$. Mark the cake below y_{2k} ($[0,1] \times [0, y_{2k}]$) as *Bottom* and the part above it ($[0,1] \times [y_{2k}, L]$) as *Top*. We have $V(Bottom) = 2k$ and $V(Top) = 2(n-k-1)$.

Now there are two cases:

Case A: $L - y_{2k} \geq \frac{1}{2}$ (this implies $k < n-1$). Thus each of *Bottom* and *Top* is either 2-fat, or 2-thin with its longer side facing rightwards.

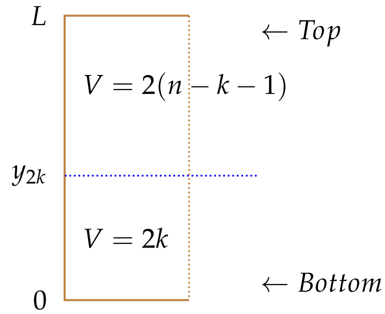

Apply the fat procedure or the thin procedure, whichever is appropriate, to *Bottom* and *Top*. In each part there are two possible outcomes: a smaller number of squares within 4 walls, or a larger number of squares within 3 walls.

— If we get the 4-walls outcome in both parts, then we have $k + (n-k-1) = n-1$ squares within the 4 walls of C, which is an instance of Outcome #1.

— If we get the 4-walls outcome in one part and the 3-walls outcome in the other part, then we have $k + (n-k) = n$ or $(k+1) + (n-k-1) = n$ squares within 3 walls. By the induction assumption, the thin procedure guarantees that all squares flowing over the rightmost border have their leftmost edge adjacent to the leftmost wall $x = 0$, and their side-length at most the longer side minus the shorter side. Here, the longer side of both Bottom and Top is less than L and their shorter side is 1, so all these squares are contained in $[0, L-1] \times [0, L]$, so we have an instance of Outcome #2.

— If we get the 3-walls outcome in both parts, then we have $k + (n-k) + 1 = n+1$ squares within 3 walls. We can discard one square arbitrarily and remain with n squares as in the above case, which is again an instance of Outcome #2.

Case B: $L - y_{2k} < \frac{1}{2}$, so *Bottom* is 2-fat or 2-thin facing rightwards, and *Top* is 2-thin facing downwards. Now consider y_{2k-2}. By definition of k, $y_{2k-2} < \frac{1}{2}$. let $Bottom' = [0,1] \times [0, y_{2k-2}]$ and $Top' = [0,1] \times [y_{2k-2}, L]$, so $V(Bottom') = 2(k-1) = V(Bottom) - 2$ and $V(Top') = 2(n-k) = V(Top) + 2$. Note that $Bottom'$ is 2-thin facing upwards and is contained in *Bottom*, and Top' is 2-fat or 2-thin facing rightwards and contains *Top*:

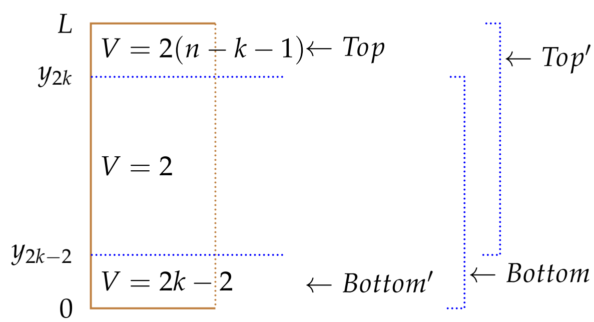

At this point $n \geq 3$, so either $n-k-1 \geq 1$ or $k-1 \geq 1$ or both. Hence, at least one of the two thin parts facing downwards/upwards (*Top*, $Bottom'$) is non-empty and with value at least 2. Use the *thin procedure* on the non-empty part/s facing downwards/upwards. In each part there are two possible outcomes: a smaller number of squares within 4 walls or a larger number of squares within 3 walls. There are several cases to consider.

— One easy case is that we get the 4-walls outcome in at least one of the parts — either in *Top* or in $Bottom'$ or in both. Suppose that we get the 4-walls outcome in $Bottom'$ (the situation is analogous if we get the 4-walls outcome in *Top*). So we have $k-1$ squares within the 4 walls of $Bottom'$. We ignore the outcome on *Top* and proceed to get additional squares from Top'. Apply to Top' either the fat procedure (if it is 2-fat) or the thin procedure (if it is 2-thin facing rightwards). One possibility is that we get $n-k$ additional squares contained in Top'; then we have a total of $n-1$ squares contained in C, which is an instance of Outcome #1. Another possibility is that we get $n-k+1$ additional squares bounded by only three walls of Top'; by the induction assumption and the guarantees of the Thin Procedure, the squares that flow over the rightmost border of Top' are adjacent to its leftmost wall, which coincides with the leftmost wall of C. Their side-length is at most the longer side-length of Top' minus its shorter side-length; the longer side-length of

Top' is less than L and its shorter side-length is 1, so the side-length of all the additional squares is at most $L-1$, and we have an instance of Outcome #2.

— Another easy case is that we get the 3-walls outcome in one part, and the other part is empty. Suppose that *Top* is empty (this implies $k = n-1$) and we get the 3-walls outcome in *Bottom'*. So we have $(k-1)+1 = n-1$ squares contained in $[0,1] \times [0, 1-y_{2k-2}] \subseteq C$, which is an instance of Outcome #1. The situation is analogous if *Bottom'* is empty and we get the 3-walls outcome in *Top*.

— The hard case is that both *Top* and *Bottom'* are non-empty and the thin procedure on both of them returns the 3-walls outcome. We now have the following squares:

- $k \geq 1$ bottom squares in $[0,1] \times [0, 1-y_{2k-2}]$;
- $n-k \geq 1$ top squares in $[0,1] \times [L-1+(L-y_{2k}), L]$.

Because $L \geq 2$, no squares overlap:

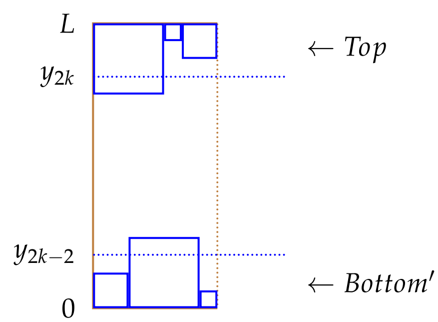

We now have n squares within the 4 walls of C, which is more than we need for Outcome #1. □

The guarantees of the Fat Procedure imply that, for all $n \geq 2$:

$$\text{PropSame}(\textit{Square with 4 walls},\ \textit{Squares},\ n) \geq \frac{1}{2n}$$

which exactly matches the upper bound of Claim 3.3.4.

Fat rectangle pieces

When the pieces are allowed to be R-fat rectangles, the above lower bound is of course still valid. But when $R \geq 2$, the Fat Procedure can give a slightly stronger guarantee - the required value is $2n-1$ instead of $2n$ (the Thin Procedure is unchanged). The required modifications in the Fat Procedure are briefly explained below:

- In the base case ($n = 1$), the cake value is 1 and it is 2-fat, so the procedure returns the entire cake as a single piece.
- In the main procedure ($n \geq 2$), we first try to cut the cake horizontally to two 2-fat rectangles and apply the Fat Procedure to each of them. For this, we need to find some $y \in [1/2, L-1/2]$ such that the value below y is at least $2k-1$ and the value above y is at least $2(n-k)-1$, for some integer $k \geq 1$. Then, both the part below y and the part above y are 2-fat. By the induction assumption, the Fat Procedure finds k 2-fat-rectangles in the bottom part and $n-k$ 2-fat-rectangles in the top part, so we are done.
- If we cannot find such y, this means that for all $y \in [1/2, L-1/2]$ and every integer k', either the value below y is less than $2k'-1$ or the value above y is less than $2(n-k')-1$. But the latter condition implies that the value below y is more than $2k'$, so the condition becomes: for all $y \in [1/2, L-1/2]$ and every integer k', the value below y is either less than $2k'-1$ or more than $2k'$. So for all $y \in [1/2, L-1/2]$, the value below y is in the open interval $(2k-2, 2k-1)$ for some integer $k \geq 1$. This means that the cake looks like this, for some integer k:

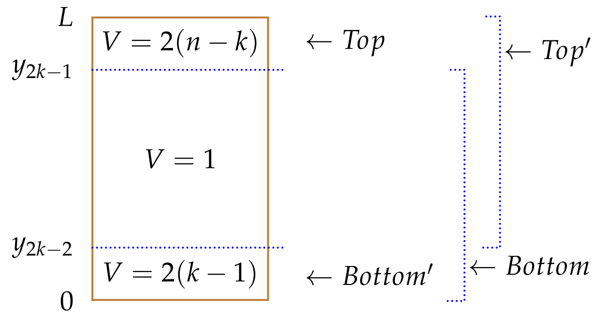

where $y_{2k-2} < 1/2$ and $y_{2k-1} > L - 1/2$. Hence, the parts $Top := [0,1] \times [y_{2k-1}, L]$ and $Bottom' := [0,1] \times [0, y_{2k-2}]$ are both 2-thin rectangles (one of these parts may be empty). $V(Top) = 2(n-k)$ and $V(Bottom') = 2(k-1)$. This is exactly the same situation as in Case B of the original procedure. We can now apply the Thin Procedure to Top and to $Bottom'$ and proceed according to the outcomes. □

Therefore, for all $n \geq 2$ and $R \geq 2$:

$$\text{PropSame}(\textit{Square with 4 walls},\ R\ \textit{fat rectangles},\ n) \geq \frac{1}{2n-1}$$

which exactly matches the upper bound of Claim 3.3.15.

Remark

The above procedures work only when the value measures are identical. The main reason is that the Thin procedure may return one of two outcomes. When there is a single value measure, the returned outcome is unique. But when there are different value measures, each value measure may induce a different outcome, and the different outcomes may be incompatible.

3.5.6 Compact cakes of any shape

As explained in Subsection 3.1.1, when the cake can be of an arbitrary shape, $\text{Prop}(C, S, n)$ may be arbitrarily small. Hence it makes sense to assess the fairness of an allocation for a particular agent relative to the total utility that this agent can get in an S-piece when given the entire cake. This intuition is captured by the following definition. It is an analogue of Definition 3.2.1, the only difference being that the normalization factor is the cake utility $V^S(C)$ instead of the cake value $V(C)$:

Definition 3.5.3. *(Relative proportionality)* For a cake C, a family of usable pieces S and an integer $n \geq 1$:

(a) The *relative proportionality level* of C, S and n, marked $\text{RelProp}(C, S, n)$, is the largest fraction $r \in [0,1]$ such that, for every set of n value measures $(V_i, ..., V_n)$, there exists an S-allocation $(X_1, ..., X_n)$ for which $\forall i:\ V_i(X_i)/V_i^S(C) \geq r$.

(b) The *same-value relative proportionality level* of C, S and n, marked $\text{RelPropSame}(C, S, n)$, is the largest fraction $r \in [0,1]$ such that, for every single value measure V, there exists an S-allocation $(X_1, ..., X_n)$ for which $\forall i:\ V(X_i)/V^S(C) \geq r$.

Our first result involves parallel squares.

Claim 3.5.3. For every cake C which is a compact subset of $\mathbb{R}^2$:

$$\text{RelProp}(C,\ \textit{Parallel squares},\ n) \geq \frac{1}{8n-6}$$

Proof. We normalize the valuations of all agents such that, for every agent i, $V_i^S(C) = 8n - 6$. We show a division procedure giving each agent a square with a value of at least 1.

(1) **Preparation:** Each agent i draws a "best square" in C — a square q_i that maximizes V_i. The existence of such a square can be proved based on the compactness of the set of squares in C; this is done in Appendix 3.C. By definition of the utility function V^S, for every i: $V_i(q_i) = V_i^S(C) = 8n - 6$.

(2) **Mark auction:** Let $N := 4n - 3$. Ask each agent i to mark, inside q_i, N pairwise-disjoint parallel squares with a value of 1 (the agent can do so by using the division procedure for identical value measures described in Subsection 3.5.5: this procedure finds N squares in q_i, each of which has a value of at least $V_i(q_i)/(2N) = 1$). Let Q_i be the collection of N squares marked by i.

An agent's bid is interpreted as saying "I am willing to give my entitlement to a piece of C in exchange for any square in Q_i". Our goal now is to allocate to each agent i a single piece from the collection Q_i such that the n allocated pieces are pairwise-disjoint.

(3) **Winner selection:** a smallest square in $\cup_i Q_i$ is selected as the winning bid (if there several smallest squares, one is selected arbitrarily). Denote the selected smallest square by q^* and suppose it belongs to agent i. Agent i now receives q^* and goes home.

(4) **Bid adjustment:** For each agent $j \neq i$, remove from Q_j all squares that overlap q^*. Since the squares in Q_j are all pairwise-disjoint and not smaller than q^*, the number of squares removed is at most 4. This is based on the following geometric fact: given a square q, there are at most 4 parallel squares that are larger than q, overlap q and do not overlap each other. This is because each square larger than q which overlaps q, must overlap one of its 4 corners, so there can be at most 4 such squares:

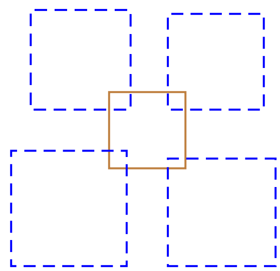

After the removal, each of the remaining $n - 1$ agents has a collection of at least $4(n - 1) - 3$ squares. If only a single agent remains, then his collection contains at least 1 square; allocate this square to the single agent and finish. Otherwise, go back to step (3) and select the next winner from the remaining $n - 1$ agents.

Finally, each agent $i \in \{1, \dots, n\}$ holds a square from the collection Q_i. This square has a value of at least 1, which proves the claim. □

The proof of Claim 3.5.3 can be generalized to other families of usable pieces:

Claim 3.5.4. For a family of pieces S, define:

- O_S = the largest number of pairwise-disjoint S-pieces that overlap an S-piece with a smaller diameter.
- $\text{PropSame}(S, S, n) = \inf_{C \in S} \text{PropSame}(C, S, n)$.

Then for every compact cake C and every $n \geq 1$:

$$\text{RelProp}(C,\ S,\ n)\ \geq\ \text{PropSame}(S,\ S,\ O_S \cdot (n-1) + 1)$$

The proof is exactly the same as that of Claim 3.5.3, with only the constant 4 replaced by O_S, 3 replaced by $O_S - 1$ and the function $1/(2N)$ replaced by $\text{PropSame}(S, S, N)$.

When S is the family of general (rotated) squares, $O_S = 8$:[18]

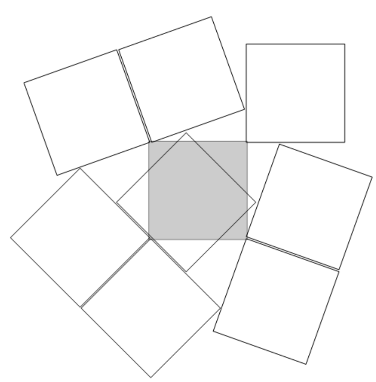

[18]We are grateful to Mark Bennet, Martigan, calculus, Red, Peter Woolfitt and Dejan Govc for their help in calculating this number in http://math.stackexchange.com/q/1085687/29780 . Image credit: Dejan Govc. Licensed under CC-BY-SA 3.0.

Corollary 3.5.4. *For every cake C which is a compact subset of* $\mathbb{R}^2$*:*

$$RelProp(C,\, Squares,\, n) \geq \frac{1}{16n-14}$$

When S is the family of parallel R-fat rectangles, $O_S = \lceil 2R+2 \rceil$:

Corollary 3.5.5. *For every cake C which is a compact subset of* $\mathbb{R}^2$*:*

$$RelProp(C,\, Parallel\ R\ fat\ rectangles,\, n) \geq \frac{1}{2\lceil 2R+2 \rceil (n-1)+2}$$

For completeness, we present the following trivial result regarding identical value measures:

Claim 3.5.5. For every cake C which is a compact subset of $\mathbb{R}^2$:

$$\text{RelPropSame}(C,\, Squares,\, n) = \frac{1}{2n}$$

Proof. Suppose the value measure of all n agents is V. Let q be a best square in C — a square that maximizes V. By definition of the utility function, $V(q) = V^S(C)$. Because q is a square, it is possible to allocate within it n disjoint squares with a value of at least $V(q)/(2n) = V^S(C)/(2n)$. □

Remarks

1. The constant O_S — the largest number of pairwise-disjoint S-pieces that overlap an S-piece with a smaller diameter — has been used for developing approximation procedures for the problem of finding a maximum non-overlapping set (Marathe et al., 1995). The approximation factors are not tight. For example, for $n = 2$, in step (b) we create $4n - 3 = 5$ axis-parallel squares for each agent, but it is possible to prove that 3 squares per agent suffice for guaranteeing that a pair of disjoint squares exists. Hence, RelProp(C, *Axis parallel squares*, $n = 2$) $\geq 1/6$. What is the smallest number of squares required to guarantee the existence of n disjoint squares? This open question is interesting because it affects both the proportionality coefficient in our fair cake-cutting procedure and the approximation coefficient in the maximum disjoint set algorithm of Marathe et al. (1995).

2. The Winner Selection procedure (step 3 in the proof) can be used even when the value functions of the agents are not additive or even not monotone (i.e. some parts of the land have negative utility to some agents). As long as every agent can draw N disjoint squares, the procedure guarantees that he receives one of these pieces.

3. Iyer and Huhns (2009) present a division procedure in which each agent marks n desired *rectangles*. Their goal is to allocate each agent a single desired rectangle. However, because the rectangles might be arbitrarily thin, it is possible that a single rectangle will intersect all other rectangles. In this case, the procedure fails and no allocations are returned. In contrast, our procedure requires the agents to draw fat pieces. This guarantees that it always succeeds.

3.6 Conclusions and Future Work

This chapter laid the foundations for fair cake-cutting with geometric constraints. This topic has a large potential for future research. Some possible directions are suggested below.

3.6.1 Open questions

We would like to close the gaps between the possibility and impossibility results in Tables 3.1 and 3.2. The most interesting gap, in our opinion, is related to an unbounded plane. Our impossibility result assumes that the squares are parallel to each other; if the squares are allowed to rotate arbitrarily, then we do not have an impossibility result, and we do not know whether a proportional division is possible.

Based on our current results, and some other results which we had to omit in order to keep the paper length at a reasonable level, we make the following conjecture:

Conjecture. *When a cake C is divided to n agents each of whom must receive a fat rectangle, the attainable proportionality is:*

$$\frac{1}{2n + Geom(C)}$$

Where $Geom(C)$ is a (positive or negative) constant that depends only on the geometric shape of the cake.

In other words: the move from a one-dimensional division to a two-dimensional division asymptotically decreases the fraction that can be guaranteed to every agent by a factor of 2.

Another direction is extending the results to cakes in three or more dimensions. We have some preliminary results in this direction.

It may be interesting to study cakes of different topologies, such as cylinders and spheres. We mention, in particular, the following potentially practical open question: *is it possible to divide Earth (a sphere) in a fair-and-square way?*

3.6.2 Different geometric constraints

The present chapter focused on constraints related to *geometric shape* — squareness or fatness. One could also consider constraints related to *size*, e.g. by defining the family S to be the family of all rectangles of length above 10 meters or area above 100 square meters. A problem with these constraints is that they are not scalable. For example, if the cake is 200-by-200 meters and there is either a length-minimum of 10 or an area-minimum of 100, then it is impossible to divide the land to more than 400 agents. Governments often cope with this problem by putting an upper bound on the number of people allowed to settle in a certain location. However, this limitation prevents people from taking advantage of new possibilities that become available as the number of people increases. For example, while in rural areas a land-plot of less than 10-by-10 meters may be considered useless because it cannot be efficiently cultivated, in densely populated cities even a land-plot as small as 2-by-2 meters can be used as a parking lot for rent or as a lemonade selling spot. Limiting the number of agents assures that each agent gets a land-plot that can be cultivated efficiently, but it may prevent more profitable ways of using the land-plots. In contrast, the squareness/fatness constraint is scalable because it does not depend on the absolute size of the land-cake. It is equally meaningful in both densely and sparsely populated areas.

The division problem can be extended by allowing each agent to have a different geometric constraint (a different family S of usable shapes) or even to have utility functions which combine different families of usable shapes (with an agent-specific weight for each family).

Finally, the two auction types used by our procedures (see Subsection 3.1.2) can possibly be generalized. For example, it may be interesting to see what can be attained if each agent receives two entitlements instead of one. This is common in some rural settlements, in which each settler receives two plots — one for housing and one for farming.

3.6.3 Web implementation

Some of the algorithms presented in this chapter have been implemented and can be tried online.[19] Currently the website is only used for demonstration purposes, but it can be made more practical, for example, by letting the users upload maps of lands that have to e divided, and by letting users submit their valuations remotely.

3.7 Acknowledgments

We are most grateful to the editor of the Journal of Mathematical Economics and to an anonymous referee for their detailed comments, which greatly improved the quality of the manuscript. We are also grateful for the comments and suggestions we received from the participants of the EuroCG 2016 conference, game-theory seminars at Bar-Ilan university and the Hebrew university and the Technion and the Corvinus

[19] http://tora.us.fm/geometry/fair-division.html

university at Budapest, computer science colloquium at the Technion, omputational geometry seminar at the Tel-Aviv university, as well as anonymous referees of ITCS 2015 and SOCG 2015 conferences.

Many members of the stackexchange.com communities kindly helped in coping with specific issues arising in this paper. In particular, we are grateful to the following members: Abel, Abishanka Saha, Acarchau, AcomPi, Adam P. Goucher, Adriano, aes, Alecos Papadopoulos, András Salamon, Anthony Carapetis, Aravind, Barry Cipra, Barry Johnson, BKay, Boris Bukh, Boris Novikov, BradHards, Brian M. Scott, calculus, Chris Culter, Christian Blatter, Claudio Fiandrino, coffeemath, Cowboy, D.W., dafinguzman, Dale M, David Eppstein, David Richerby, Debashish, Dejan Govc, Dömötör Pálvölgyi (domotorp), dtldarek, Emanuele Paolini, Flambino, Foo Barrigno, frafl, George Daccache, Hagen von Eitzen, Henno Brandsma, Henrik, Henry, Herbert Voss, Ilya Bogdanov, Ittay Weiss, J.R., Jan Kyncl, Jared Warner, Jeff Snider, John Gowers, John Hughes, JohnWO, Joonas Ilmavirta, Joseph O'Rourke, Kenny LJ, Kris Williams, Macavity, MappaGnosis, Marc van Leeuwen, Mari-Lou A, Mariusz Nowak, Mark Bennet, Martigan, Martin Van der Linden, Matt, Maulik V, Merk, Marcus Hum (mhum), Michael Albanese, Mikaël Mayer, Mike Jury, mjqxxxx, mvcouwen, Martin von Gagern (MvG), Nick Gill, njaja, Normal Human, NovaDenizen, Peter Franek, Peter Michor, Peter Woolfitt, pew, PhoemueX, Phoenix, Raphael Reitzig, R Salimi, Raymond Manzoni, Realz Slaw, Red, Robert Israel, Ross Millikan, sds, Stephen, Steven Landsburg, Steven Taschuk, StoneyB, tomasz, TonyK, Travis, vkurchatkin, Włodzimierz Holsztyński, Yury, Yuval Filmus and Zur Luria.

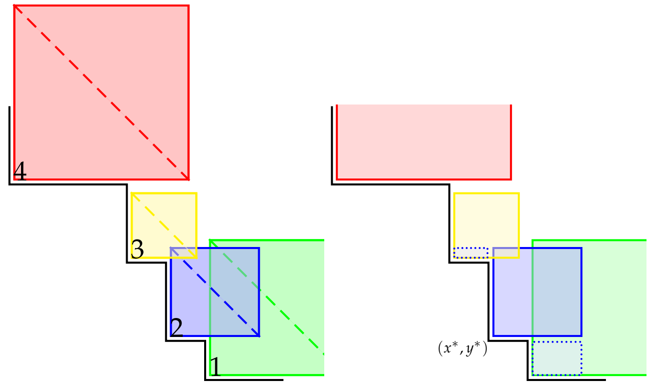

Figure 3.11: **a.** A staircase with $T = 3$ teeth and $T + 1 = 4$ corners and a square in each corner. The diagonal (dashed) represents t_j — the taxicab distance from the origin to the square center. The square at corner 2 is the winning square as its taxicab distance is minimal (the diagonal is closest to the origin).
b. The shadow of the winning square (dotted). Note that each rectangular component of the shadow is entirely contained in the square of the corresponding corner.

Chapter 3 Appendix

3.A Staircase Lemma

This appendix proves the following geometric lemma, which is used in Section 3.5.2:

Lemma 3.A.1. *(Staircase Lemma) Let C be a staircase-shaped polygonal domain with T teeth (and $T + 1$ corners). Suppose that in each inner corner $j \in \{1, \ldots, T + 1\}$, with coordinates (x_j, y_j), there is a square with side-length l_j (the square $[x_j, x_j + l_j] \times [y_j, y_j + l_j]$).*

Define the shadow *of square j as the intersection of C with the rectangle $[0, x_j + l_j] \times [0, y_j + l_j]$ (this is the area of C that is removed when cutting from the top-right corner of square j towards the bottom and left boundaries of C; see Figure 3.11/b).*

There exists a corner j such that the shadow of square j is contained in the union of the $T + 1$ squares.

Proof. For every $j \in \{1, \ldots, T + 1\}$, define:

$$t_j := x_j + y_j + l_j$$

t_j can be interpreted as the "taxicab distance" (ℓ_1 distance) from the origin to the center of the square at corner j, or equivalently to its bottom-right or top-left corner;

Define the *winning square* as the square j *for which t_j is minimized*. Denote its corner coordinates by (x^*, y^*) and its side-length by l^*. We now prove that the shadows of the winning square are contained in the other squares. We decompose the shadows of the winning square to pairwise-disjoint rectangular components in the following way.

- For each corner j to the top-left of the winning square, the component is a rectangle with coordinates: $[x_j, x^*] \times [y_j, y^* + l^*]$. Note that this component is empty if $y_j \geq y^* + l^*$, as in corner 4 in Figure 3.11.
- For each corner j to the bottom-right of the winning square, the component is a rectangle with coordinates: $[x_j, x^* + l^*] \times [y_j, y^*]$. This component is empty if $x_j \geq x^* + l^*$.

By definition of the winning square, for every $j \in \{1, \ldots, T + 1\}$:

$$x_j + y_j + l_j \geq x^* + y^* + l^* \tag{3.3}$$

Now:

- For each corner j to the top-left of the winning square, we have $x_j < x^*$. Combining this with (3.3) gives $y^* + l^* < y_j + l_j$. Moreover, if the component in that corner is not empty, then necessarily

$y_j < y^* + l^*$. Combining this with (3.3) gives $x^* < x_j + l_j$. Hence, the component $[x_j, x^*] \times [y_j, y^* + l^*]$ is contained in the square $[x_j, x_j + l_j] \times [y_j, y_j + l_j]$.

- For each corner j to the bottom-right of the winning square, we have $y_j < y^*$. Combining this with (3.3) gives $x^* + l^* < x_j + l_j$. Moreover, if the component in that corner is not empty, then necessarily $x_j < x^* + l^*$. Combining this with (3.3) gives $y^* < y_j + l_j$. Hence, the component $[x_j, x^* + l^*] \times [y_j, y^*]$ is contained in the square $[x_j, x_j + l_j] \times [y_j, y_j + l_j]$.

We proved that every component of the shadow of the winning square is contained in one of the $T + 1$ squares; hence, the winning square satisfies the requirement of lemma. □

3.B Non-intersection of Squares in Fat Procedure

This appendix proves that in the last step of the Fat Procedure (Subsection 3.5.5), the n returned squares do not overlap.

Recall that at this step, the cake has two distinguished regions: $Bottom' := [0,1] \times [0, y_b]$ and $Top := [0,1] \times [y_t, L]$, both of which are 2-thin rectangles, i.e, $0 < y_b < 1/2 \leq L - 1/2 < y_t < L$. In each region there is a family of squares: the *bottom squares* were returned by applying the Thin Procedure to Bottom', and the *top squares* were returned by applying the Thin Procedure to Top. The squares in each family are pairwise-disjoint, but squares from different families might overlap. Our goal is to prove that, after a single largest square is removed, the remaining squares do not overlap, as in the following illustration:

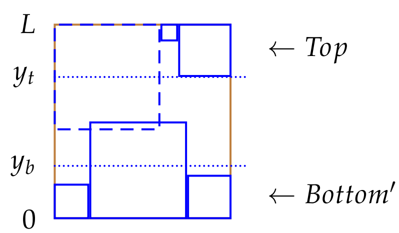

Recall that, by the specification of the Thin Procedure (Subsection 3.5.5), the squares in each family can be divided to two types, which we call "doves" and "hawks":

- *Doves* are squares generated by Outcome #1 of the Thin Procedure (or by recursive calls to the Fat Procedure). They are contained within the four walls of their rectangle: the bottom doves are contained in $[0,1] \times [0, y_b]$, and the top doves are contained in $[0,1] \times [y_t, L]$.

- *Hawks* are squares generated by Outcome #2 of the Thin Procedure. They are contained within only three walls of their rectangle, with one of their edges adjacent to the wall opposite the open side: the bottom edge of all bottom hawks is at $y = 0$, and the top edge of all top hawks is at $y = L$. Moreover, the side-length of each hawk is at most the longer side of its rectangle minus the shorter side of its rectangle; hence, the side-length of all bottom hawks is at most $1 - y_b$ and their top edge is in $y \in [y_b, 1 - y_b]$, and the side-length of all top hawks is at most $1 - (L - y_t)$ and their bottom edge is in $y \in [L - (1 - L + y_t), y_t]$.

Claim 3.B.1. In each family, the sum of the side-lengths of all hawks is at most 1.

Proof. The bottom hawks are all bounded in a rectangle of length 1: $[0,1] \times [0, 1 - y_b]$. Their bottom side is at $y = 0$. Since they do not overlap, the sum of their side-lengths must be at most 1. A similar argument holds for the top hawks. □

An immediate corollary of Claim 3.B.1 is that at most one hawk from each side has side-length more than 1/2. We call each of these two hawks (if it exists) the *dangerous hawk*.

We say that a square q *attacks* a square q' if q is larger than q' and q overlaps q'. This is possible only if q and q' are in two opposite families, since the squares in each family are pairwise-disjoint. The doves obviously do not attack each other because $y_b < y_t$. So the only possible attacks are: top hawks attacking bottom hawks/doves, or bottom hawks attacking top hawks/doves.

After removing the largest square, at most one dangerous hawk remains; it is only this hawk that might attack other squares in the opposite side. We now prove that even this dangerous hawk does not attack other squares.

Claim 3.B.2. No remaining hawk attacks any dove.

Proof. We prove that no remaining hawk even enters the rectangle of the opposite family (no remaining bottom-hawk enters Top and no remaining top-hawk enters $Bottom'$). Since all doves are contained in their rectangle, they are safe. There are two cases:

Case 1: $y_t \geq L - y_b$. Then also $y_t \geq 1 - y_b$. The side-length of all bottom hawks is at most $1 - y_b$, so no bottom hawk enters Top. If the top dangerous hawk enters $Bottom'$, then its side-length must be more than $L - y_b$, so it is larger than all bottom hawks. Hence, it is the largest square and it is removed.

Case 2: $y_t < L - y_b$. Then also $1 - (L - y_t) < 1 - y_b \leq L - y_b$. The side-length of all top hawks is at most $1 - (L - y_t)$, no top hawk enters $Bottom'$. If the bottom dangerous hawk enters Top, then its side-length must be more than y_t, so it is larger than all top hawks. Hence, it is the largest square and it is removed. □

Claim 3.B.3. No remaining hawk attacks any hawk.

Proof. There are two cases:

Case 1: There is only one hawk (either bottom *or* top) with side-length more than $1/2$. This is the largest square so it is removed. The remaining squares have side-length at most $1/2$ and thus do not attack each other.

Case 2: There are two hawks (bottom *and* top) with side-length more than $1/2$. W.l.o.g, assume the top hawk is the largest, with a side-length of $h_t \geq h_b$. By Claim 3.B.1, the sum of the side-lengths of all other top hawks is at most $1 - h_t$, hence the side-length of any single other top hawk is at most $1 - h_t$ which is at most $1 - h_b$ which is at most $L - h_b$. Hence, the bottom side of all remaining top hawks is above h_b. Hence the remaining bottom hawk cannot attack any of them. □

3.C Existence of Best Pieces

This appendix shows how to prove the existence of a usable piece with a maximum value (this is used in the proof of Claim 3.5.3). We start by defining a metric space of pieces (recall that a *piece* is a Borel subset of $\mathbb{R}^2$ and $Area$ is its Lebesgue measure).

Definition 3.C.1. The *symmetric difference (SD) pseudo-metric* is defined by:

$$d_{SD}(X, Y) = Area[(X \setminus Y) \cup (Y \setminus X)]$$

d_{SD} is not a metric because there may be different pieces whose symmetric difference has an area of 0, e.g, a square with an additional point and a square with a missing point. To make SD a metric, we consider only pieces X that are *regularly open*, i.e, the interior of the closure of themselves: $X = Int[Cl[X]]$.

Claim 3.C.1. SD is a metric on the set of all regularly-open pieces.

Proof. [20] Let X and Y be two regularly-open sets such that $d_{SD}(X, Y) = 0$. We prove that $X = Y$.

$d_{SD}(X, Y) = 0$ implies $Area[X \setminus Y] = Area[Y \setminus X] = 0$.

$Y \subseteq Cl[Y]$ so $X \setminus Y \supseteq X \setminus Cl[Y]$. Hence also $Area[X \setminus Cl[Y]] = 0$.

X is open and $Cl[Y]$ is closed; hence $X \setminus Cl[Y]$ is open (it is an intersection of two open sets).

The only open set with an area of 0 is the empty set (because any non-empty open set contains a ball with a positive measure). Hence: $X \setminus Cl[Y] = \emptyset$.

Equivalently: $X \subseteq Cl[Y]$.

By taking the Cl of both sides: $Cl[X] \subseteq Cl[Y]$

By a symmetric argument: $Cl[Y] \subseteq Cl[X]$

Hence: $Cl[Y] = Cl[X]$

By taking the Int of both sides and by the fact that they are regularly-open: $Y = X$. □

Thus when we allocate a square we actually allocate only its interior. This has no effect on the utility of the agents since the boundary has an area of 0 and so its value is 0 for all agents.

[20] We are thankful to Tony K., Phoemue X., Dafin Guzman, Henno Brandsma and Ittay Weiss for contributing to this proof via discussions in the math.stackexchange.com website (http://math.stackexchange.com/a/1099461/29780).

Claim 3.C.2. Let D be the metric space defined by d_{SD}. Let V be a measure absolutely continuous with respect to area. Then V is a uniformly continuous function from D to $\mathbb{R}$.

Proof. The fact that V is an absolutely continuous measure implies that, for every $\epsilon > 0$ there is a $\delta > 0$ such that every piece X with $Area(X) < \delta$ has $V(X) < \epsilon$ (Nielsen, 1997, Proposition 15.5 on page 251). Hence, for every two pieces X and Y, if $d_{SD}(X,Y) < \delta$ then $Area(X \setminus Y) < \delta$ and $Area(Y \setminus X) < \delta$, then $V(X \setminus Y) < \epsilon$ and $V(Y \setminus X) < \epsilon$, then $|V(X) - V(Y)| = |V(X \setminus Y) - V(Y \setminus X)| < \epsilon$. □

Claim 3.C.3. Let V be a measure absolutely continuous with respect to area and Q a set of pieces which is compact in the SD metric space. Then there exists a piece $q \in Q$ for which V is maximized.

Proof. By the previous claim, V is a uniformly continuous and hence a continuous real-valued function. By the extreme value theorem, it attains a maximum in every compact set. □

The value measures considered in this paper are always absolutely continuous with respect to area. Hence, to prove that a certain set of pieces Q contains a "best piece" it is sufficient to prove that Q is compact. We do this now for the special case in which Q is the set of open squares contained in a given cake (note that the same proof could be used for the set of closed squares):

Claim 3.C.4. Let C be a closed, bounded subset of $\mathbb{R}^2$. Let Q be the set of all open squares contained in C. Then Q is compact in the SD metric space.

Proof. It is sufficient to prove that Q is sequentially compact, i.e. every infinite sequence of open squares in C has a subsequence converging to an open square in C. Let $\{q_i\}_{i=1}^{\infty}$ be an infinite sequence of open squares in C. For every q_i, let (A_i, B_i) be a pair of opposite corners. Because C is compact, it contains $Cl[q]$ and hence contains the points A_i and B_i. Hence the infinite sequence of pairs of points, $\{(A_i, B_i)\}_{i=1}^{\infty}$, is an infinite sequence in $C \times C$. $C \times C$ is compact because it is a finite product of compact sets. Hence, the sequence has a subsequence converging to a limit point $(A^*, B^*) \in C$. From now on we assume that $\{(A_i, B_i)\}_{i=1}^{\infty}$ is that converging subsequence. Let q^* be the open square having A^* and B^* as two opposite corners. We show that: (a) q^* is an open square in C; (b) The subsequence $\{q_i\}_{i=1}^{\infty}$ converges to q^*.

(a) q^* is a obviously an open square by definition. We have to show that each point in q^* is also a point of C. To every square q_i, attach a local coordinate system in which corner A_i has coordinates $0,0$ and corner B_i has coordinates $1,1$ and every other point in $Cl[q_i]$ has coordinates in $[0,1] \times [0,1]$. For every coordinate $(x,y) \in [0,1] \times [0,1]$, let $q_i(x,y)$ be the unique point with these coordinates in $Cl[q_i]$ (e.g. $A_i = q_i(0,0)$ and $B_i = q_i(1,1)$).

For every (x,y), The sequence $\{q_i(x,y)\}_{i=1}^{\infty}$ is a sequence of points which are all in C, and they converge to $q^*(x,y)$. Since C is closed, $q^*(x,y) \in C$.

(b) For every i, the area of the symmetric difference between q^* and q_i is bounded and satisfies the following inequality:

$$d_{SD}(q^*, q_i) \leq 4 \cdot \max(d(A^*, A_i), d(B^*, B_i)) \cdot \max(d(A^*, B^*), d(A^*, B_i), d(A_i, B^*), d(A_i, B_i))$$

Since all distances are bounded and $d(A^*, A_i)$, $d(B^*, B_i)$ converge to 0, the same is true for $d_{SD}(q^*, q_i)$. Hence, the subsequence $\{q_i\}_{i=1}^{\infty}$ converges to q.

The previous paragraph proved that Q is sequentially compact. Hence it is compact. □

In a similar way it is possible to prove similar results for other families S, such as the family of R-fat rectangles or cubes.

3.D Non-Rectangular Pieces

In the main body of this chapter, the usable pieces were fat rectangles. Interestingly, we can get better results and simpler procedures by expanding the family of usable pieces to include other 2-fat polygons with angles that are multiples 45 degrees. We call such polygons *2-FFDPs* (2-fat Forty-Five Degree Polygons).[21] Our procedure is based on the following geometTric facts:

1. A right-angled isosceles triangle (RAIT) is a 2-FFDP.

[21] This idea was suggested by Galya Segal-Halevi.

2. Both a RAIT and a square can be partitioned into two congruent halves, each of which is a RAIT.

3. Each RAIT half in such a partition can be shrunk by translating the division line towards one of the corners, such that the smaller piece is a RAIT and the larger piece is a 2-FFDP.

We present a procedure for dividing a cake that can be either a RAIT or a square. The procedure requires that for every agent i: $V_i(C) \geq \max(1, 2n-2)$. It returns n disjoint 2FFDPs $\{X_i\}_{i=1}^n$ such that for every agent i: $V_i(X_i) \geq 1$.

The procedure is developed by induction on the number of agents. When there is a single agent ($n=1$), he can just be given the entire cake, which is a 2FFDP with value at least 1. We now assume that we can handle any number of agents less than n. Now there are n agents ($n \geq 2$), each of whom values C as at least $2n-2$. We proceed as follows.

(1) **Eval auction**. Cut C to two congruent RAITs: C' and C'':

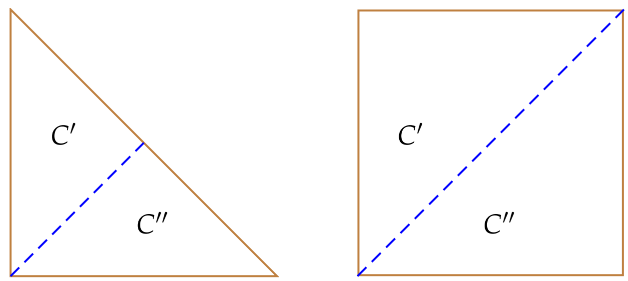

Do an eval auction on C'. Order the agents in a descending order of their bid, $V_1(C') \geq \cdots \geq V_n(C')$, and let n' be the largest integer with:

$$V_{n'}(C') \geq \max(2n'-2, 1)$$

If $n' = n$ then all agents value C' as the entire cake, so the other parts of the cake can be discarded and the division procedure can start again with C' as the cake. Hence, we assume that $n' < n$. There are two main cases to consider:

- *Easy case*: $1 \leq n' \leq n-2$. Make a diagonal guillotine cut between C' and C''. Divide C' recursively among the n' winners.

  The $n-n'$ losers value C' as less than $\max(2(n'+1)-2, 1) = 2n'$, so the value the remainder C'' as at least $(2n-2)-2n' = 2(n-n')-2$. Since $n-n' \geq 2$, this value is also larger than 1, so we can divide C'' recursively among the $n-n'$ losers.

- *Hard case*: $n' = 0$. This means that all agents value C' as less than 1, so they value C'' as more than $2n-1$.

  We have to shrink C'' towards the corner, until one of the agents decides that it is better to take a piece outside C'' and leave C'' to the remaining $n-1$ agents. This solution is implemented using a *mark auction*, which is described in step (2) below. But before proceeding there is one more case that must be handled:

- *Mixed case*: $n' = n-1$. This is handled according to the bid of the single losing agent (agent n): if $V_n(C') < 2n-1$, then the losing agent values C'' as more than 1, so we can proceed as in the Easy case (the winning agents receive C' and the losing agent receives C''). Otherwise, $V_n(C') \geq 2n-1$, so all agents value C' as at least $2n-1$ (because the agents are ordered in descending order of their bid). Switch the roles of C' and C'' and proceed as in the hard case to the next auction.

(2) **Mark auction**. Ask each agent to mark a 2FFDP with a value fo exactly 1, whose complement is a RAIT adjacent to the corner of C':

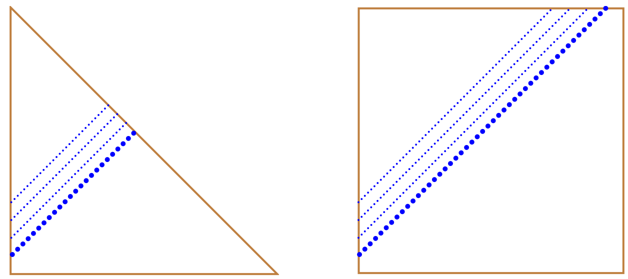

The winning bid (marked by thicker dots above) is a 2FFDP. It can be given to the winner, who values it as exactly 1 so i

The remaining cake is a RAIT and its value for the remaining $n-1$ agents is at least $V(C)-1 \geq 2n-1 \geq \max(2(n-1)-2, 1)$. Divide it recursively among the losers. $\square$

This procedure proves:

Claim 3.D.1. For every $n \geq 2$:

$$\text{Prop}(RAIT, 2\,FFDPs, n) \geq \frac{1}{2n-2}$$
$$\text{Prop}(Square, 2\,FFDPs, n) \geq \frac{1}{2n-2}$$

The procedure is clearly much simpler than when the pieces must be fat rectangles (as in Subsection 3.5.1) and the proportionality coefficient is better. In other words, it is easier to divide a cake fairly when 45-degree polygons are allowed. This might explain why practical land allocation maps usually contain more than just rectangles.

Chapter 4

Geometric Envy-Free Division

Contents

3.1 Introduction . . . **10**
3.1.1 Results . . . 11
3.1.2 Techniques . . . 14
3.1.3 Related work . . . 14
3.2 Model . . . **15**
3.3 Impossibility Results . . . **16**
3.3.1 Impossibility results for two, three and four walls . . . 17
3.3.2 Impossibility results for one wall . . . 18
3.3.3 Impossibility results for zero walls . . . 19
3.3.4 Impossibility results with fat rectangles . . . 22
3.4 Auctions and Covers . . . **24**
3.4.1 Mark auction . . . 24
3.4.2 Eval auction . . . 25
3.4.3 Cover numbers . . . 26
3.5 Division procedures . . . **27**
3.5.1 Four and three walls, guillotine cuts . . . 27
3.5.2 Two walls . . . 33
3.5.3 One and zero walls . . . 35
3.5.4 Three walls . . . 36
3.5.5 Four walls, guillotine cuts, identical valuations . . . 41
3.5.6 Compact cakes of any shape . . . 46
3.6 Conclusions and Future Work . . . **48**
3.6.1 Open questions . . . 48
3.6.2 Different geometric constraints . . . 49
3.6.3 Web implementation . . . 49
3.7 Acknowledgments . . . **49**
Chapter appendix . . . **51**
3.A Staircase Lemma . . . **51**
3.B Non-intersection of Squares in Fat Procedure . . . **52**
3.C Existence of Best Pieces . . . **53**
3.D Non-Rectangular Pieces . . . **54**

A preliminary version of this chapter appeared in the proceedings of AAAI 2015 (Segal-Halevi et al., 2015a).

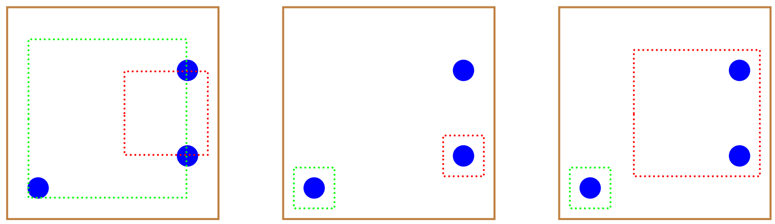

Figure 4.1: A square land-estate has to be divided between two people. The land-estate is mostly barren, except for three water-pools (discs). The agents have the same preferences: each agent wants a square land-plot with as much water as possible. The squares must not overlap. Hence:
(a) It is impossible to give both agents more than 1/3 of the water. Hence:
(b) An envy-free division must give each agent at most 1/3 of the water.
(c) But such a division cannot be Pareto-efficient since it is dominated by a division which gives one agent 1/3 and the other 2/3 of the water.
Hence, a Pareto-efficient envy-free allocation does not exist.

4.1 Introduction

In the previous chapter we focused on a single measure of fairness: proportionality. Our aim was to guarantee all agents a certain fraction of their total cake value, and we tried to make this fraction as large as possible.

In the present chapter we add a second measure of fairness: *envy-freeness*. Our aim now is to make sure that each agent believes that his/her allocated piece is at least as good as any other piece.

Envy-freeness on its own is trivially satisfied by the empty allocation. The task becomes more interesting when envy-freeness is combined with an *efficiency* criterion. The most common such criterion is Pareto efficiency. Indeed, Weller (1985) has proved that, when the agents' preferences are represented by non-atomic measures, there always exists a competitive-equilibrium with equal-incomes, and the equilibrium allocation is both Pareto-efficient and envy-free. However, Weller's equilibrium allocation gives no guarantees about the *geometric shape* of the allotted pieces. A "piece" in his allocation might even be a union of a countable number of disconnected cake-bits. So, Weller's positive result is valid only when the agents' preferences ignore the geometry of their allotted pieces. While such preferences may make sense when dividing an actual edible cake, they are not so sensible when dividing land.

Berliant and Dunz (2004) have studied a multi-dimensional cake model. Their results are mostly negative: when general value measures are combined with geometric preferences, a competitive-equilibrium might not exist. In fact, even regardless of competitive-equilibrium, a Pareto-efficient-envy-free allocation might not exist, as we show in Figure 4.1.

Thus, to get an envy-free allocation among agents with geometric preferences, we must replace Pareto-efficiency with a different efficiency criterion. A natural candidate is *proportionality* — every agent should receive at least $1/n$ of the total cake value. Since with geometric preferences, a proportional division does not always exist (see Figure 4.1), we relax the proportionality requirement and consider *partial proportionality*. Partial proportionality means that each agent receives a piece worth at least a fraction p of the total cake-value, where p is a positive constant, $0 < p \leq 1/n$ (see definition in section 2.5 in page 7). Obviously we would like p to be as large as possible.

In the previous chapter, we showed that partial-proportionality can be attained in various geometric settings. For example, when there is a square cake and two agents who want square pieces, each agent can be guaranteed at least a fraction 1/4 of the total cake-value, and this is the largest fraction that can be guaranteed. However, these results did not consider envy. This raises the following question, which is at the heart of the present chapter:

> When each agent wants a piece with a given geometric shape, what is the largest fraction of the cake-value that can be guaranteed to every agent in an envy-free allocation?

The following example shows that existing cake-cutting procedures are insufficient for answering this question.

Example 4.1.1. You and a partner are going to divide a square land-estate. It is 100-by-100 square meters and its western side is adjacent to the sea. Your desire is to build a house near the sea-shore. You decide to use the classic procedure for envy-free division: "You cut, I choose". You let your partner divide the land to two plots, knowing that you have the right to choose the plot that is more valuable according to your personal preferences. Your partner makes a cut parallel to the shoreline at a distance of only 1 meter from the sea. [1] Which of the two plots would you choose? The western plot contains a lot of sea shore, but it is so narrow that it has no room for building anything. On the other hand, the eastern plot is large but does not contain any shore land. Whichever plot you choose, the division is not proportional for you, because your utility is far less than half the utility of the original land estate.

Of course the cake *could* be cut in a more sensible way (e.g. by a line perpendicular to the sea), but the current division procedures say nothing about how exactly the cake should be cut in each situation in order to guarantee that the division is fair in a way that respects the geometric preferences. While the cut-and-choose procedure still guarantees envy-freeness, it does not guarantee partial-proportionality since it does not guarantee any positive utility to agents who want square pieces. □

This paper presents cake-cutting procedures that guarantee both envy-freeness and partial-proportionality. Our procedures focus on agents who want *fat pieces* — pieces with a bounded length/width ratio, such as squares (see definition in subsection 2.6 on page 8). The rationale is that a fat shape is more convenient to work with, build on, cultivate, etc.

4.1.1 Results

We prove that envy-freeness and partial-proportionality are *compatible* in progressively more general geometric scenarios. Our proofs are constructive: in every geometric scenario (geometric shape of the cake and preferred shape of the pieces), we present a procedure that divides the cake with the following guarantees:

- *Envy-freeness*: every agent weakly prefers his/her allotted piece over the piece given to any other agent.
- *Partial-proportionality*: every agent receives a piece worth for him at least a fraction p of his total cake-value, where p is a positive constant that depends on the geometric requirements.

In the following theorems, the partial-proportionality guarantee p is given in parentheses.

Theorem 4.1. *When dividing a cake to two agents, there is a procedure for finding an envy-free and partially-proportional allocation in the following cases:*
(a) The cake is square and the usable pieces are squares ($p \geq 1/4$).
(b) The cake is an R-fat rectangle and the usable pieces are R-fat rectangles, where $R \geq 2$ ($p \geq 1/3$).
(c) The cake is an arbitrary R-fat object and the pieces are $2R$-fat, where $R \geq 1$ ($p \geq 1/2$).

Value-shape trade-off: Theorem 4.1 illustrates a multiple-way trade-off between value and shape. Consider two agents who want to divide a square land-estate with no envy. They have the following options:

- By projecting a 1-dimensional division obtained by any classic cake-cutting procedure, they can achieve a proportional allocation (a value of at least 1/2) with rectangular pieces but with no bound on the aspect ratio — the pieces might be arbitrarily thin.
- By (a), they can achieve an allocation with square pieces but only partial proportionality — the proportionality might be as low as 1/4.
- By (b), they can achieve a proportionality of 1/3 with 2-fat rectangles, which is a compromise between the previous two options.
- By (c), they can achieve an allocation that is *both* proportional *and* with 2-fat pieces, but the pieces might be non rectangular.

[1] The reason why he decided to cut this way is irrelevant since a fair division procedure is expected to guarantee that the division is fair for every agent playing by the rules, regardless of what the other agents do.

The proportionality constants in Theorem 4.1 are tight in the following sense: it is not possible to guarantee an allocation with a larger proportionality, even if envy is allowed. This means that envy-freeness is compatible with the largest possible proportionality — we don't have to compromise on proportionality to prevent envy.

Our second theorem extends these results to any number of agents.

Theorem 4.2. *When dividing a cake to n agents, there is a procedure for finding an envy-free and partially-proportional allocation in the following cases:*
(a) The cake is square and the usable pieces are squares ($p \geq 1/(4n^2)$).
(b) The cake is an R-fat rectangle and the usable pieces are R-fat rectangles, where $R \geq 1$ ($p \geq 1/(4n^2)$).
(c) The cake is a d-dimensional R-fat object and the pieces are $\lceil n^{1/d} \rceil$ R-fat,[2] where $d \geq 2$ and $R \geq 1$ ($p \geq 1/n$).

Value-shape trade-off: Part (a) and part (c) are duals in the following sense:

- Part (a) guarantees an envy-free division with perfect pieces (squares) but compromises on the proportionality level;

- Part (c) guarantees an envy-free division with perfect proportionality ($1/n$) but compromises on the fatness of the pieces.

The "magnitude" of the first compromise is $4n$, since the proportionality drops from $1/n$ to $1/(4n^2)$. We do not know if this magnitude is tight: we know that it is possible to attain a division with square pieces and a proportionality of $1/O(n)$ which is not necessarily envy-free Segal-Halevi et al. (2017), but we do not know if a proportionality of $1/O(n)$ is compatible with envy-freeness.

The "magnitude" of the second compromise is $\lceil n^{1/d} \rceil$. This magnitude is asymptotically tight. We prove that, in order to guarantee a proportional division of an R-fat cake, with or without envy, we must allow the pieces to be $\Omega(n^{1/d})R$-fat.

4.1.2 Related Work

(See also the Related Work subsection in the previous chapter, page 14).

The main challenge in two-dimensional cake-cutting is that utility functions that depend on geometric shape are not *additive*. For example, consider an agent who wants to build a square house the utility of which is determined by its area. The utility of this agent from a 20×20 plot is 400, but if this plot is divided to two 20×10 plots, the utility from each plot is 100 and the sum of utilities is only 200. Most existing procedures for proportional cake-cutting assume that the valuations are additive, so they are not applicable in our case. While there are some previous works on cake-cutting with non-additive utilities, they too cannot handle geometric constraints:

- Berliant et al. (1992); Maccheroni and Marinacci (2003) focus on sub-additive, or concave, utility functions, in which the sum of the utilities of the parts is *more* than the utility of the whole. These utility functions are inapplicable in our scenario because, as illustrated in the previous paragraph, utility functions that consider geometry are not necessarily sub-additive — the sum of the utilities of the parts might be less than the utility of the whole.

- Dall'Aglio and Maccheroni (2009) do not explicitly require sub-additivity, but they require *preference for concentration*: if an agent is indifferent between two pieces X and Y, then he prefers 100% of X to 50% of X plus 50% of Y. This axiom may be incompatible with geometric constraints: the agent in the above example is indifferent between the two 20×10 rectangles, but he prefers 50% of their union (the 20×20 square) to 100% of a single rectangle.[3]

- Sagara and Vlach (2005); Hüsseinov and Sagara (2013) consider general non-additive utility functions but provide only non-constructive existence proofs.

- Su (1999); Caragiannis et al. (2011); Mirchandani (2013) provide practical division procedures for non-additive utilities, but they crucially assume that the cake is a 1-dimensional interval and cannot handle two-dimensional constraints.

[2] $\lceil x \rceil$ denotes the *ceiling* of x — the smallest integer which is larger than x.
[3] We are grateful to Marco Dall'Aglio for his help in clarifying this issue.

Cake	Pieces	Agents	Impossibility	Possibility
Square	Squares	2	1/4	1/4
R-fat rectangle	R-fat rectangles ($R \geq 2$)	2	1/3	1/3
R-fat object	$2R$-fat objects	2	1/2	1/2
Square	Squares	n	$1/(2n)$	$1/(4n^2)$
R-fat rectangle	R-fat rectangles	n	$1/(2n-1)$	$1/(4n^2)$
R-fat object	$\lceil n^{1/d} \rceil R$-fat objects	n	$1/n$	$1/n$

Table 4.1: Summary of results for geometric envy-free division: upper and lower bounds on the level of attainable proportionality.

When envy-free division protocols are applied to agents with non-additive utility functions, the division is still envy-free, but the utility per agent might be arbitrarily small. This is true for cut-and-choose (as shown in Example 4.1.1 above) and it is also true for all other procedures for envy-free division that we are aware of (Stromquist (1980); Brams and Taylor (1995); Reijnierse and Potters (1998); Su (1999); Barbanel and Brams (2004); Manabe and Okamoto (2010); Cohler et al. (2011); Deng et al. (2012); Kurokawa et al. (2013); Chen et al. (2013); Aziz and Mackenzie (2016)).

Our way to cope with this challenge is to explicitly handle the geometric constraints in the procedures. The main tool we use is the *geometric knife function*.

Moving-knife procedures have been used for envy-free cake-cutting since its earliest years (Dubins and Spanier, 1961; Stromquist, 1980; Brams et al., 1997; Saberi and Wang, 2009). For example, consider the following simple procedure for envy-free division among two agents. A referee moves a knife slowly over the cake, from left to right. Whenever an agent feels that the piece to the left of the knife is worth for him exactly half the total cake value, he shouts "stop!". Then, the cake is cut at the current knife location, the shouter receives the piece to its left and the non-shouter receives the piece to its right.

In this paper we formalize the notion of a knife and add geometric constraints guaranteeing that the final pieces have both the desired geometric shape and a sufficiently high value.

4.2 Model

We briefly recall some terminology from Chapter 2 (see there for formal definitions).

- C is the cake to be divided. In this chapter it will be a square or a fat object in $\mathbb{R}^d$.
- S is the family of pieces that are considered *usable*. An *S-piece* is an element of S. In this chapter it will be the family of squares or of fat objects.
- For each agent $i \in \{1, \ldots, n\}$, $V_i(X_i)$ is agent i's value-measure of the piece X_i.
- For each agent $i \in \{1, \ldots, n\}$, $V_i^S(X_i)$ is agent i's utility of the piece X_i. It is the value-measure of the most valuable S-piece contained in X_i.

When the utilities of all agents are determined by S-value functions, we can restrict our attention to allocations in which each agent receives an S-piece. An *S-allocation* is a vector of n S-pieces $X = (X_1, ..., X_n)$, one piece per agent, such that the X_i are pairwise-disjoint and their union is contained in C.

An S-allocation X is called *envy-free* if the utility of an agent from his allocated S-piece is at least as large as his utility from every piece allocated to another agent:

$$\forall i, j \in \{1, ..., n\} : \; V_i^S(X_i) \geq V_i^S(X_j)$$

In addition to envy-freeness, an allocation is assessed by the the fraction of the total cake value that is given to each agent. An allocation is called *proportional* if every agent receives a piece worth for him at least $1/n$ of the total cake value. Since a proportional S-allocation does not always exist (see e.g. Figure 4.1), we define:

Definition 4.2.1. For a cake C, a family of usable pieces S and an integer $n \geq 1$, the **envy-free proportionality** of C, S and n, marked $\text{PropEF}(C, S, n)$, is the largest fraction $p \in [0, 1]$ such that, for every set of n

value measures $(V_i, ..., V_n)$, there exists an envy-free S-allocation $(X_1, ..., X_n)$ for which: [4]

$$\forall i : \frac{V_i(X_i)}{V_i(C)} \geq p$$

This is very similar to the definition of $\text{Prop}(C, S, n)$ - Definition 3.2.1 on page 16. The only difference is that in $\text{Prop}(C, S, n)$, the supremum is taken over all allocations, and in $\text{PropEF}(C, S, n)$, the supremum is taken only on envy-free allocations. Obviously, because the supremum in $\text{PropEF}(C, S, n)$ is taken over a smaller set:

$$\forall C, n, S : \text{PropEF}(C, S, n) \leq \text{Prop}(C, S, n)$$

This means that, in theory, if we want to guarantee that there is no envy, we may have to "pay" in terms of proportionality. One of the goals of the present research is to study if and how much we may have to pay.

Classic cake-cutting results imply that for every cake C:

$$\text{Prop}(C,\ All,\ n) = \text{PropEF}(C,\ All,\ n) = 1/n$$

where *All* is the collection of all pieces. That is: when there are no geometric constraints, every cake can be divided among every group of n agents in an envy-free allocation in which the utility of each agent is at least $1/n$.

Our challenge in the rest of this paper will be to establish bounds on $\text{PropEF}(C, S, n)$ for various combinations of C and S. All our possibility results (lower bounds) are on $\text{PropEF}(C, S, n)$ and therefore are also valid for $\text{Prop}(C, S, n)$. Similarly, all the impossibility results (upper bounds) proved in section 3.3 on page 16 are for $\text{Prop}(C, S, n)$ and therefore are also valid for $\text{PropEF}(C, S, n)$.

4.3 Geometric Preliminaries

Example 4.1.1 illustrates that, in order to achieve a fair division that respects the geometric preferences, we should constrain the ways in which agents are allowed to cut the cake. This requires several definitions of geometric concepts, which are the topic of the present section.

4.3.1 Geometric loss

A key geometric concept in our analysis is the *geometric loss* — the maximum factor by which the utility of an agent can be reduced by his insistence on using pieces only from family S.

Definition 4.3.1. For a piece C and family of usable pieces S, the *geometric loss* factor of C relative to S is:

$$\text{Loss}(C, S) := \sup_V \frac{V(C)}{V^S(C)}$$

where the supremum is over all finite absolutely-continuous value measures V having $V^S(C) > 0$. If there is no supremum, then we write $\text{Loss}(C, S) = \infty$.

When $C \in S$ the loss is 1, which means is no loss, since in this case $V^S(C) = V(C)$. When $C \notin S$, the loss is generally larger than 1. For example, if C is a 30-by-20 rectangle. The largest square contained in C is 20-by-20. Hence, if the value density is uniform over C (as in Figure 4.2/a), then $\frac{V(C)}{V^S(C)} = \frac{600}{400} = \frac{3}{2}$, implying that $\text{Loss}(C,\ Squares) \geq 3/2$. But the loss may be larger: suppose V is uniform over the right and left sides of C (as in Figure 4.2/b). In this case $\frac{V(C)}{V^S(C)} = 2$, implying that $\text{Loss}(C,\ Squares) \geq 2$. As we will see in Subsection 4.3.3, the loss in this case is exactly 2, and in general the loss of a rectangle with a length/width ratio of L is $\lceil L \rceil$; a thinner rectangle has a larger loss.

For some combinations of C and S, the geometric loss factor might be infinite. For example, if C is a circle and that V is nonzero only in a very narrow strip near the perimeter (as in Figure 4.2/c), any square

[4] Shortly: $\text{PropEF}(C, S, n) = \inf_V \sup_X \min_i V_i(X_i)/V_i(C)$, where the infimum is on all combinations of n value measures $(V_1, ..., V_n)$, the supremum is on all envy-free S-allocations $(X_1, ..., X_n)$ and the minimum is on all agents $i \in \{1, ..., n\}$.

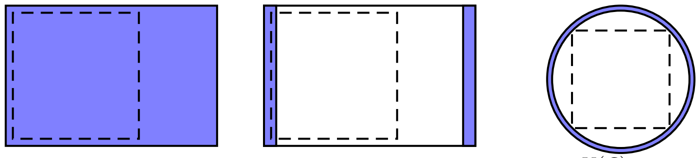

Figure 4.2: Geometric loss factors relative to the family of squares.

contained in C intersects the valuable strip only in the corners. and the intersection might be arbitrarily small. Hence, $V^S(C)$ might be arbitrarily small and Loss$(C, \mathit{Squares}) = \infty$.

4.3.2 Chooser Lemma

We now relate the geometric loss factor to cake partitions. Our goal is to prove that, if a cake is partitioned such that the sum of the geometric losses of its parts is sufficiently small, then an agent can choose at least one part with a large value. Formally:

Lemma 4.3.2. *For every cake C, integer m, partition $X_1 \sqcup \cdots \sqcup X_m = C$, family S and value measure V:*

$$\exists j : V^S(X_j) \geq \frac{V(C)}{\sum_{i=1}^{m} \mathit{Loss}(X_i, S)}$$

Proof. Denote the denominator in the right-hand side by:

$$\text{Loss}(X, S) := \sum_{i=1}^{m} \text{Loss}(X_i, S)$$

By additivity of V:

$$\sum_{i=1}^{m} V(X_i) = V(C) \tag{4.1}$$

Multiply both sides of (4.1) by the Loss$(X, S) = \sum_{i=1}^{m} \text{Loss}(X_i, S)$:

$$\sum_{i=1}^{m} V(X_i) \cdot \text{Loss}(X, S) = \sum_{i=1}^{m} \text{Loss}(X_i, S) \cdot V(C)$$

By the pigeonhole principle, at least one of the m summands in the left-hand side must be greater than or equal to the corresponding summand in the right-hand side. I.e., there exists j for which:

$$V(X_j) \cdot \text{Loss}(X, S) \geq \text{Loss}(X_j, S) \cdot V(C)$$

By Definition 4.3.1 and the definition of supremum, for every value measure V:

$$\text{Loss}(X_j, S) \geq \frac{V(X_j)}{V^S(X_j)}$$

Combining the above two inequalities yields:

$$V(X_j) \cdot \text{Loss}(X, S) \geq \frac{V(X_j) \cdot V(C)}{V^S(X_j)}$$

which is equivalent to:

$$V^S(X_j) \geq \frac{V(C)}{\text{Loss}(X,S)}$$

□

Motivated by the Chooser Lemma and its proof, we define the expression $\text{Loss}(X,S) := \sum_{i=1}^{m} \text{Loss}(X_i,S)$ as the *geometric loss of the partition* X. The Chooser Lemma implies that *smaller* geometric-loss is *better* for the chooser. This is easy to see in Example 4.1.1, where a 100-by-100 land-estate is divided using cut-and-choose:

- A partition to 100-by-1 and 100-by-99 rectangles has a geometric loss of 102 (the loss of the 100-by-1 sliver is 100 and the loss of the 100-by-99 rectangle is 2). Hence, the utility guarantee for a chooser who wants square pieces is only $1/102$.
- In contrast, a partition to two 100-by-50 rectangles has a geometric loss of 4 (2+2). By Lemma 4.3.2, the chooser can always get a square with a utility of at least $1/4$.

We will often use this simple implication of the Chooser Lemma:

Corollary 4.3.3 (Chooser Corollary)**.** *Suppose a cake-partition has a geometric loss of at most* M*. Each of two agents chooses a best piece, and the choices are different. Then the resulting allocation is envy-free, and each agent's value is at least* $1/M$ *of the total cake-value.*

4.3.3 Cover Numbers and Cover Lemma

Since smaller geometric loss is better, it is useful to have an upper bound on the geometric loss. Our upper bound uses the Cover Number — see Definition 3.4.3 on page 26.

Lemma 4.3.4. *For every cake* C *and family* S*:*

$$Loss(C,S) \leq CoverNum(C,S)$$

Proof. Let $m = \text{CoverNum}(C,S)$. By definition of CoverNum, there are m S-pieces $X_1, ..., X_m$, possibly overlapping, that cover the cake C:

$$X_1 \cup X_2 \cup ... \cup X_m = C$$

Let V be any value measure. By additivity:

$$V(X_1) + V(X_2) + ... + V(X_m) \geq V(C)$$

By the pigeonhole principle, there is at least one piece $X_i \in S$ with:

$$V(X_i) \geq V(C)/m$$

On the other hand, since X_i is an S-piece contained in C, its value is bounded by the supremum V^S:

$$V^S(C) \geq V(X_i)$$

Combining the above two inequalities yields:

$$V^S(C) \geq V(C)/m$$

Combining this into the definition $\text{Loss}(C,S) = \sup_V \frac{V(C)}{V^S(C)}$, yields:

$$\text{Loss}(C,S) \leq \sup_V \frac{V(C)}{V(C)/m} = \sup_V m = m$$

□

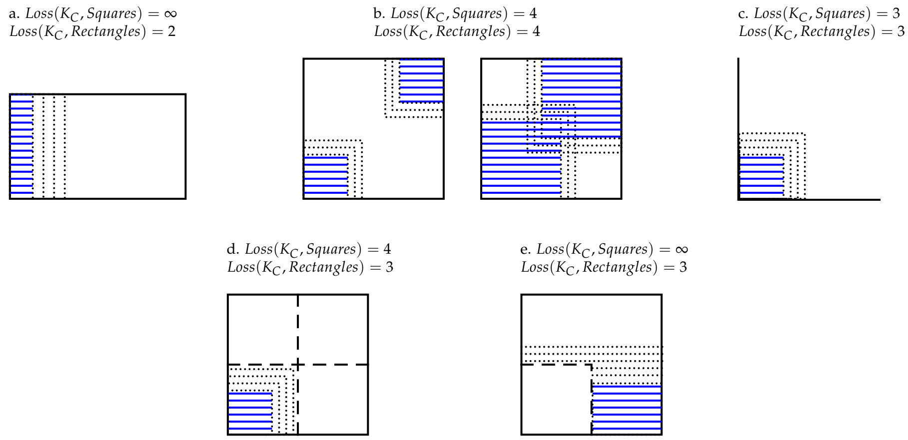

Figure 4.3: Several knife functions. The area filled with horizontal lines marks $K_C(t)$ in a certain intermediate time $t \in (0, 1)$. Dotted lines mark future knife locations.

By Definition 4.3.1, for every value measure V: $V^S(C) \geq \frac{V(C)}{\text{Loss}(C,S)}$. By Lemma 4.3.4, this implies $V^S(C) \geq \frac{V(C)}{\text{CoverNum}(C,S)}$. Thus, for example, in the 30×20 rectangle of Figure 4.2, $\text{CoverNum}(C, \mathit{Squares}) = 2$ so $\text{Loss}(C, \mathit{Squares}) \leq 2$ so $V^S(C) \geq V(C)/2$. This means that every agent, with any value measure, can get from C a utility of at least half its total value.

4.3.4 Knife functions

Moving knives have been used to cut cakes ever since the seminal paper of Dubins and Spanier (1961). We generalize the concept of a moving knife to handle geometric shape constraints.

Definition 4.3.5. Given a cake C, a *knife function on* C is a function K_C from the real interval $[0, 1]$ to pieces of C with the following *monotonicity* property: for every $t' \geq t$, $K_C(t') \supseteq K_C(t)$.

If $K_C(0) = C_0$ and $K_C(1) = C_1$, where $C_0 \subseteq C_1 \subseteq C$, we say that K_C is *a knife function from* C_0 *to* C_1.
The *complement* of K_C, marked $\overline{K_C}$, is defined by:

$$\overline{K_C}(t) := C \setminus K_C(t)$$

Some examples are shown in Figure 4.3.
A knife function K_C on a cake C can be used to attain an envy-free division of C between two agents:

Generic Knife Procedure
Each agent $i \in \{A, B\}$ selects a time $t_i \in [0, 1]$ such that:

$$V_i^S(K_C(t_i)) = V_i^S(\overline{K_C}(t_i))$$

Rename the agents, if needed, such that $t_A \leq t_B$.
Select any time $t^* \in [t_A, t_B]$.
Give $K_C(t^*)$ to agent A and $\overline{K_C}(t^*)$ to agent B.

This procedure obviously generates an envy-free division, since it gives to each agent a piece worth for him at least as much as the other piece. The challenge is in the first step: we must be sure that each agent i can, indeed, select a time t_i such that the S-values on both sides of the knife are equal. This requires that both $V_i^S(K_C(t))$ and $V_i^S(\overline{K_C}(t))$ change continuously as a function of t. Hence, we define:

Definition 4.3.6. Given a family S of usable shapes, a knife-function K is called S*-good* if for every absolutely-continuous value-measure V, both $V^S(K(t))$ and $V^S(\overline{K}(t))$ are continuous functions of t.

How can we find S-good knife-functions? In Appendix 4.A, we define two different properties of knife-functions, each of which is a sufficient condition for S-goodness:

- *S-smoothness* means that the Lebesgue measure of $K(t)$ is a continuous function of t, and that both $K(t) \in S$ and $\overline{K}(t) \in S$. For example, the knife-function in Figure 4.3/a is rectangle-smooth (but not square-smooth).

- *S-continuity* means (informally) that all S-pieces in $K(t)$ grow continuously and all S-pieces in $\overline{K}(t)$ shrink continuously; no S-piece with a positive area is created abruptly in $K(t)$ and no S-piece with a positive area is destroyed abruptly in $\overline{K}(t)$. All knife-functions in Figure 4.3 are square-continuous (and also rectangle-continuous).

See Appendix 4.A for formal definitions, proofs and additional examples.

With an S-good knife, the Generic Knife Procedure can be executed:

Lemma 4.3.7. *Let C be a cake and C_0, C_1 pieces such that: $C_0 \subseteq C_1 \subseteq C$. Let K_C be an S-good knife-function from C_0 to C_1. Assume that an agent has a value function V such that:*

- $V^S(C_0) \leq V^S(C \setminus C_0)$
- $V^S(C_1) \geq V^S(C \setminus C_1)$

Then there exists a time $t_i \in [0,1]$ in which the utilities on both sides of the knife are equal:

$$V^S(K_C(t_i)) = V^S(\overline{K_C(t_i)})$$

Proof. When $t = 0$:

$$V^S(K_C(t)) = V^S(C_0) \leq V^S(C \setminus C_0) = V^S(\overline{K_C(t)})$$

and when $t = 1$:

$$V^S(K_C(t)) = V^S(C_1) \geq V^S(C \setminus C_1) = V^S(\overline{K_C(t)})$$

Since K_C is S-good, by Definition 4.3.6 both $V^S(K_C(t))$ and $V^S(\overline{K_C(t)})$ are continuous functions of t. Hence the lemma follows from the intermediate value theorem. $\square$

4.3.5 Geometric loss of knife functions

When a knife function K_C is "stopped" at a certain time $t \in [0,1]$, it induces a partition of the cake C to the part which was already covered by the knife, $K_C(t)$, and the part not covered, $\overline{K_C}(t)$. Based on this partition, the geometric loss of the knife can be defined:

Definition 4.3.8. Let C be a cake, K_C a knife function on C and S a family of pieces. Define the *geometric loss* of K_C as:

$$\text{Loss}(K_C,\, S) = \sup_{t\in[0,1]} \left(\text{Loss}(K_C(t),\, S) + \text{Loss}(\overline{K_C}(t),\, S) \right)$$

Whenever a knife is stopped, the resulting partition has a geometric loss of at most $Loss(K_C, S)$. Therefore, we can expect such a knife to be useful for fairly dividing a cake among agents who want S-pieces.

Recall that the smallest possible Loss of a single piece is 1 (which means "no loss"); hence the smallest possible loss of a knife function is 1+1=2. Some examples are illustrated in Figure 4.3, from left to right:

(a) Let $C = [0,L] \times [0,1]$ and $K_C(t) = [0,L] \times [0,t]$. Both $K_C(t)$ and its complement are rectangles so their geometric loss relative to the family of rectangles is 1. Hence $\text{Loss}(K_C,\, Rectangles) = 1 + 1 = 2$. In contrast, the geometric loss of these rectangles relative to the family of squares is unbounded, so: $\text{Loss}(K_C,\, Squares) = \infty$.

(b) Let $C = [0,1] \times [0,1]$ and $K_C(t) = [0,t] \times [0,t] \cup [1-t,1] \times [1-t,1]$. For every t, $K_C(t)$ is a union of two squares and its complement is also a union of two squares. By the Cover Lemma, each such union has a geometric loss of 2 (relative to the family of squares). Hence, $\text{Loss}(K_C,\, Squares) = 2 + 2 = 4$.

(c) Let C be the top-right quarter-plane and S the family of squares and quarter-planes (we consider a quarter-plane to be a square with infinite side-length). Define: $K_C(t) = [0, x/(1-x)] \times [0, x/(1-x)]$.

$K_C(t)$ is a square and its complement can be covered by two quarter-planes, so the geometric loss of K_C is 1+2=3.

(d) Let $C = [0,2] \times [0,2]$ and $K_C(t) = [0,t] \times [0,t]$. Note that $K_C(0) = \emptyset$ and $K_C(1) = [0,1] \times [0,1] =$ the bottom-left quarter of C. For every t, $K_C(t)$ is a square and its complement is an L-shape, similar to the L-shapes in Figure 3.6, which can be covered by 3 squares. Hence, $\text{Loss}(K_C, \, Squares) = 3 + 1 = 4$.

(e) Let $C = [0,2] \times [0,2]$, $C_1 = C \setminus [0,1] \times [0,1]$ (an L-shape), and $K_C(t) = C_1 \cap ([0,2] \times [0,t/2])$. This is a knife-function from $\emptyset$ to C_1; it covers C_1 continuously from bottom to top. The partition can be covered by at most 2+1=3 rectangles, but its square-loss is not bounded.

Lemma 4.3.9. *(Knife Lemma) Let C be a cake and C_0, C_1 pieces such that: $C_0 \subseteq C_1 \subseteq C$. Let K_C be an S-good knife-function from C_0 to C_1. If there are two agents and for every agent i:*

- $V_i^S(C_0) \leq V_i^S(C \setminus C_0)$ and

- $V_i^S(C_1) \geq V_i^S(C \setminus C_1)$,

then C can be divided using the Generic Knife Procedure (see Subsection 4.3.4) and every agent playing by the rules is guaranteed an envy-free share with a utility of at least:

$$\max\left(V_i^S(C_0),\, V_i^S(C \setminus C_1),\, \frac{1}{Loss(K_C,\, S)}\right)$$

Proof. Consider an agent, say Alice, who plays by the rules and declares a time t_A for which $V_A^S(K_C(t_A)) = V_A^S(\overline{K_C(t_A)})$. Denote this equal utility by U. There are two cases: if $t_A \leq t^* \leq t_B$, then Alice receives $K_C(t^*)$, which contains $K_C(t_A)$. Otherwise, $t_B \leq t^* \leq t_A$, and Alice receives $\overline{K_C(t^*)}$, which contains $\overline{K_C(t_A)}$ (because K_C is monotonically increasing). In both cases, Alice feels no envy and receives a utility of at least U. This utility is bounded from below in three ways:

(a) $U \geq V^S(C_0)$, because the piece $K_C(t_i)$ contains C_0.

(b) $U \geq V^S(C \setminus C_1)$, because the complement piece $\overline{K_C(t_i)}$ contains $C \setminus C_1$.

(c) $U \geq 1/\text{Loss}(K_C,\, S)$ by the Chooser Lemma, since the loss of the partition is at most $\text{Loss}(K_C,\, S)$. □

Note that the Generic Knife Procedure is discrete: it does not need to continuously move the knife until an agent shouts "stop"; the agents are asked in advance in what time they would like to "stop the knife".

The Chooser Lemma and the Knife Lemma are the main tools we use to construct division procedures.

4.4 Envy-Free Division for Two Agents

4.4.1 Squares and rectangles

Our first generic envy-free division procedure is based on a single knife function.

Lemma 4.4.1. (Single Knife Procedure). *Let C be a cake, S a family of pieces and $M \geq 2$ an integer. If there exists an S-good knife-function K_C from $\emptyset$ to C with*

$$Loss(K_C,\, S) \leq M,$$

then

$$PropEF(C, S, 2) \geq 1/M.$$

Proof. The cake can be divided using the Generic Knife Procedure, taking $C_0 = \emptyset$ and $C_1 = C$. The assumptions of the Knife Lemma (4.3.9) hold trivially because $C_0 = C \setminus C_1 = \emptyset$. Hence each agent playing by the rules receives an envy-free share worth at least $1/M$. □

The knife function in Figure 4.3/b is Square-good and its Square-loss is 4. Applying Lemma 4.4.1 to that knife function yields our first sub-theorem:

Theorem 4.1(a).. $\text{PropEF}(Square,\, Squares,\, 2) \geq 1/4$

The generality of Lemma 4.4.1 allows us to get more results with no additional effort. For example:

- By the knife function of Figure 4.3/b: PropEF($Square$, $Square\ pairs$, 2) $\geq 1/2$. I.e., if each agent has to receive a union of two squares (as is common when dividing land to settlers, e.g. one land-plot for building and another one for agriculture, etc.), then a proportional division is possible since the knife function in example (b) has a geometric loss of 2 relative to the family of square pairs.

- By Figure 4.3/c: PropEF($Quarter\ Plane$, $Generalized\ Squares$, 2) $\geq 1/3$.

All bounds presented above are tight in the strong sense stated in the introduction, i.e., it is not possible to guarantee both agents a larger utility *even if envy is allowed*. This is obvious for the $\geq 1/2$ results, since a proportionality of $1/n$ is the best that can be guaranteed to n agents. For the other results, the matching upper bounds are proved in Section 3.3 on page 16.

4.4.2 Cubes and archipelagos

In some cases it may be difficult to find a single knife function that covers the entire cake. This is so, for example, when the cakes are multi-dimensional cubes or unions of disjoint squares. To handle such cases, the following lemma suggests a generalized division procedure employing several knife functions.

Lemma 4.4.2. (Single Partition Procedure). *Let C be a cake, S a family of pieces and $M \geq 2$ an integer such that:*
(a) C has a partition with a geometric loss of at most M:

$$\bigsqcup_{j=1}^{m} C_j = C$$

$$\sum_{j=1}^{m} Loss(C_j, S) \leq M$$

(b) For every j, there are S-good knife functions from $\emptyset$ to C_j and from $\emptyset$ to $\overline{C_j}$ (where $\overline{C_j} := C \setminus C_j$).
(c) For every part C_j, the geometric loss of the knife-function on C_j is at most M:

$$\forall j: \ Loss(K_{C_j},\ S) \leq M$$

Then:

$$PropEF(C, S, 2) \geq 1/M$$

Proof. C can be divided using the following procedure.

(1) Each agent chooses the part C_j that gives him maximum utility. If the choices are different, then by the Chooser Corollary and condition (a), each agent receives an envy-free share worth at least $1/M$, so we are done

(2) If both agents chose the same part C_j, then ask each agent to choose either C_j or $\overline{C_j}$ (where $\overline{C_j} := C \setminus C_j$). If the choices are different, then by the Chooser Corollary and condition (a), each agent receives an envy-free share worth at least $1/M$, so we are done. If the choices are identical then there are two cases:

(3-a) Both agents chose C_j. By condition (c), there exists a knife function K_{C_j} from $\emptyset$ to C_j with a geometric loss of at most M. Apply the Generic Knife Procedure with that knife function. The requirements of the Knife Lemma (4.3.9) are satisfied since for both agents, $V_i^S(\emptyset) \leq V_i^S(C \setminus \emptyset)$ (trivially) and $V_i^S(C_j) \geq V_i^S(C \setminus C_j)$ (both agents preferred C_j to $\overline{C_j}$). Hence, the Knife Lemma guarantees each agent an envy-free share worth at least $1/\text{Loss}(K_{C_j},\ S) \geq 1/M$.

(3-b) Both agents chose $\overline{C_j}$. By condition (b), there exists a knife function $K_{\overline{C_j}}$ from $\emptyset$ to $\overline{C_j}$. There is no guarantee about the geometric loss of $K_{\overline{C_j}}$, but this is fine since we will not use its geometric loss below. Apply the Generic Knife Procedure. The requirements of the Knife Lemma (4.3.9) are met since for both agents, $V_i^S(\emptyset) \leq V_i^S(C \setminus \emptyset)$ (trivially) and $V_i^S(\overline{C_j}) \geq V_i^S(C \setminus \overline{C_j})$ (both agents preferred $\overline{C_j}$ over C_j). The Knife Lemma guarantees each agent an envy-free share with a utility of at least $V_i^S(C \setminus \overline{C_j}) = V_i^S(C_j)$. The fact that in step (2) both agents chose C_j implies, by the Chooser Lemma, that $\forall i: V_i^S(C_j) \geq 1/M$. □

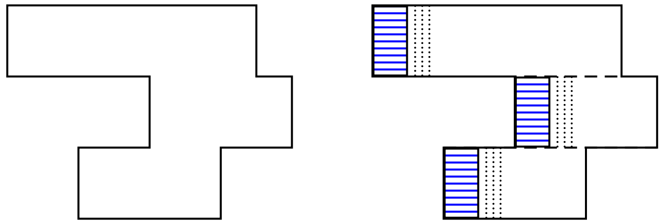

Figure 4.4: (a) A cake made of a union of 3 disjoint rectangles.
(b) Three knife functions, each having a geometric loss of at most 4, proving that PropEF(C, *rectangles*, 2) $\geq 1/4$.

Several applications of Lemma 4.4.2 are presented below.

(a) PropEF$(Square, Squares, 2) \geq 1/4$. *Proof*: A square cake can be partitioned to a 2-by-2 grid of squares. The loss of the partition relative to the family of squares is 4, satisfying condition (a). Each quarter C_j has a knife-function with a loss of 4 (see Figure 4.3/d), satisfying condition (c). For each complement $\overline{C_j}$, we can use e.g. a sweeping-line knife-function, as illustrated in Figure 4.3/e (see Lemma 4.A.7 in the appendix for a proof that such functions are S-good), satisfying condition (b).

The advantage of this result over the identical result presented in the previous subsection is that it can be easily generalized to higher dimensions:

(b) **Multi-dimensional cakes**: PropEF$(d\ dimensional\ cube, Squares, 2) \geq 1/2^d$. *Proof*: C can be partitioned to 2^d sub-cubes of equal side-length. For each sub-cube C_j there is a knife function analogous to Figure 4.3/d — a cube growing from the corner towards the center of C. Its geometric loss is 2^d. For each complement $\overline{C_j}$, there is a sweeping-plane knife-function (analogous to Figure 4.3/e, as described in Lemma 4.A.7).

(c) **Archipelagos**: Let C be an archipelago which is a union of m disjoint rectangular islands. Then PropEF(C, *Rectangles*, 2) $\geq \frac{1}{m+1}$. *Proof*: The geometric loss of the partition of C to m rectangles is obviously $m < m+1$, satisfying condition (a). For each part C_j, define a knife function K_{C_j} based on a line sweeping from one side of the rectangle to the other side, similar to Figure 4.3/a. $K_{C_j}(t)$ is always a rectangle. Its complement can be covered by m rectangles: one rectangle to cover $C_j \setminus K_{C_j}(t)$ and additional $m-1$ rectangles to cover $C \setminus C_j$. Hence the geometric loss of every K_{C_j} is $1+1+m-1 = m+1$, satisfying condition (c) (see Figure 4.4). A similar sweeping-line knife-function can be used for the complements, satisfying condition (b).

(d) Let C be an archipelago which is a union of m disjoint square islands. Then PropEF(C, *Squares*, 2) $\geq \frac{1}{m+3}$. The proof is the same as in (c), the only difference being that each of the knife functions on the C_j is a union of two squares, similar to Figure 4.3/b.

All bounds proved above are tight. The tightness of (a) is proved in Section 3.3. The tightness of (b) can be proved by an analogous d-dimensional cake, with a water-pool in each of its 2^d corners. (c) is tight in the following sense: for every m there is a cake C, which is a union of m disjoint rectangles, having Prop(C, *Rectangles*, 2) $\leq \frac{1}{m+1}$. (d) is tight in a similar sense by a similar proof.

4.4.3 Fat rectangles

More types of cakes can be handled by adding partition steps.

Lemma 4.4.3. (Multiple Partition Procedure). *Let C be a cake, S a family of pieces and $M \geq 2$ an integer such that:*

(a) C has a partition $C_1, \ldots, C_m$ with a geometric loss of at most M:

$$\bigsqcup_{j=1}^{m} C_j = C$$

$$\sum_{j=1}^{m} Loss(C_j, S) \leq M$$

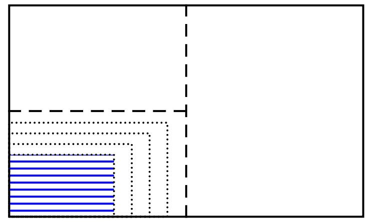

Figure 4.5: A knife function with a geometric loss of 3, proving that PropEF$(C, 2\, fat\, rectangles, 2) \geq 1/3$.

(b) Every part C_j can be further partitioned such that, if C_j is replaced with its partition, then the geometric loss of the resulting partition of C is at most M, i.e. for every j there exist $C_j^1, \ldots, C_j^{m_j}$ with:

$$\bigsqcup_{k=1}^{m_j} C_j^k = C_j$$

$$\sum_{j' \neq j} Loss(C_{j'}, S) + \sum_{k=1}^{m_j} Loss(C_j^k, S) \leq M$$

(c) For every j, k, there are S-good knife functions from $\emptyset$ to C_j and to $\overline{C_j}$ and to C_j^k and to $\overline{C_j^k}$.

(d) For every j, k, the geometric loss of the knife function from $\emptyset$ to C_j^k is at most M:

$$\forall j, k : \; Loss(K_{C_j^k},\, S) \leq M$$

Then:

$$PropEF(C, S, 2) \geq 1/M$$

Proof. The proof uses a refinement of the procedure used to prove Lemma 4.4.2. Steps (1) and (2) and (3-b) are exactly the same. We have to refine case (3-a), in which both agents prefer C_j over $\overline{C_j}$.

(3-a-1) Refine the partition of C by replacing C_j with its sub-partition:

$$(\bigsqcup_{j' \neq j} C_{j'}) \sqcup (\bigsqcup_{k=1}^{m_j} C_j^k) = C$$

Let each agent choose a best part from this refined partition. If the choices are different, then by condition (b) and the Chooser Corollary, each agent receives an envy-free share worth at least $1/M$.

(3-a-2) If both agents chose the same part from the main partition, e.g. $C_{j'}$ for some $j' \neq j$, then by condition (c) there exists a knife-function *from $\emptyset$ to C_j* (the part chosen by both agents at step 2). Apply the Generic Knife Procedure. The requirements of the Knife Lemma (4.3.9) are satisfied since for both agents, $V_i^S(\emptyset) \leq V_i^S(C \setminus \emptyset)$ (trivially) and $V_i^S(C_j) \geq V_i^S(C \setminus C_j)$ (both agents prefer C_j to $\overline{C_j}$). The Knife Lemma guarantees each agent an envy-free share with utility at least $V_i^S(\overline{C_j})$. This $\overline{C_j}$ contains all other parts of the main partition, including $C_{j'}$. The fact that both agents chose $C_{j'}$ in the refined partition proves, by the Chooser Lemma, that $V_i^S(C_{j'}) \geq 1/M$. Hence also $V_i^S(\overline{C_j}) \geq 1/M$.

(3-a-3) If both agents chose the same part from the sub-partition, e.g. C_j^k for some k, then ask each agent to choose either C_j^k or $\overline{C_j^k}$ (where $\overline{C_j^k} := C \setminus C_j^k$). If the choices are different, then by condition (b) and the Chooser Corollary, each agent receives an envy-free share worth at least $1/M$. If the choices are identical then there are two cases:

(3-a-4-a) Both agents chose C_j^k. By condition (d), there exists a knife function $K_{C_j^k}$ from $\emptyset$ to C_j^k with a geometric loss of at most M. Apply the Generic Knife Procedure with that knife function. The requirements of the Knife Lemma (4.3.9) are satisfied since for both agents, $V_i^S(\emptyset) \leq V_i^S(C \setminus \emptyset)$ (trivially) and $V_i^S(C_j^k) \geq V_i^S(C \setminus C_j^k)$ (both agents preferred C_j^k over $\overline{C_j^k}$). Hence, the Knife Lemma guarantees each agent an envy-free share worth at least $1/\text{Loss}(K_{C_j^k},\, S) \geq 1/M$.

(3-a-4-b) Both agents chose $\overline{C_j^k}$. By condition (c), there exists an S-good knife function from $\emptyset$ to $\overline{C_j^k}$. Apply the Generic Knife Procedure. The requirements of the Knife Lemma (4.3.9) are met since for both agents, $V_i^S(\emptyset) \leq V_i^S(C \setminus \emptyset)$ (trivially) and $V_i^S(\overline{C_j^k}) \geq V_i^S(C \setminus \overline{C_j^k})$ (both agents preferred $\overline{C_j^k}$ over C_j^k). The Knife Lemma guarantees each agent an envy-free share with utility at least $V_i^S(C \setminus \overline{C_j^k}) = V_i^S(C_j^k)$. The fact that in step (3-a-2) both agents chose C_j^k implies, by the Chooser Lemma, that $V_i^S(C_j^k) \geq 1/M$. □

Lemma 4.4.3 is used to get the second part of our Theorem 4.1:

Theorem 4.1(b). For every $R \geq 2$:

$$\text{PropEF}(R\,fat\,rectangle,\ R\,fat\,rectangles,\ 2) \geq 1/3$$

Proof. The proof relies on the following geometric fact: for every $R \geq 2$, an R-fat rectangle can be bisected to two R-fat rectangles using a straight line through the center of its longer sides (see Figure 4.5).

Apply Lemma 4.4.3 in the following way. Let C be an R-fat rectangle. Partition C in the middle of its longer side. The two halves are R-fat so the geometric loss of the partition is $1 + 1 < 3$, satisfying condition (a). Each half can be further partitioned along its longer side to two rectangles, which are also R-fat (each of these is exactly one quarter of C). When a part is replaced by its sub-partition, the geometric loss of the resulting partition is thus $2 + 1 = 3$, satisfying condition (b). Condition (c) is satisfied e.g. by knife-functions based on sweeping lines, as in Figure 4.3/e. For each quarter-rectangle, there is a knife function (growing from the corner towards the center, as in Figure 4.5) with a geometric loss of 3, satisfying condition (d). □

The bound of $1/3$ is tight; see Subsection 3.3.4 on page 22.

Lemma 4.4.3 can be further refined by adding more sub-partition steps. For example, by adding a third sub-partition step we can prove that if C is an archipelago of m disjoint R-fat rectangles (with $R \geq 2$) then:

$$\text{PropEF}(C,\ R\text{-fat rectangles},\ 2) \geq \frac{1}{m+2}$$

and this bound is tight. The proof is a analogous to examples (c) and (d) after Lemma 4.4.2.

Note that the upper bound of $1/3$ is valid when the pieces are R-fat rectangles for every finite R, while the upper bound of $1/4$ for square pieces is valid for every $R < 2$. This implies that 2-fat rectangles are a good practical compromise between fatness and fairness: if we require fatter pieces ($R < 2$) then the proportionality guarantee drops from $1/3$ to $1/4$, while if we allow thinner pieces ($R > 2$) the proportionality remains $1/3$ for all $R < \infty$.

4.4.4 Arbitrary fat objects

Our most general result involves cakes that are arbitrary Borel sets. The result is proved for cakes of any dimensionality; Figure 4.6 illustrates the proof for $d = 2$ dimensions.

Theorem 4.1(c). For every $R \geq 1$, If C is R-fat and S is the family of $2R$-fat pieces then:

$$\text{PropEF}(C, S, 2) = \text{Prop}(C, S, 2) = 1/2$$

Proof. The proof uses Lemma 4.4.2 (the Single Partition Procedure). We show a partition of C to two pieces and a knife-function on each piece. Scale, rotate and translate the cake C such that the largest cube contained in C is $B^- = [-1,1]^d$ (Figure 4.6/a). By definition of fatness (see Subsection 2.6), C is now contained in a cube B^+ of side-length at most $2R$.

Using the hyperplane $x = 0$, bisect the cube B^- to two 2-by-1 boxes $B_1 = [-1,0] \times [-1,1]^{d-1}$ and $B_2 = [0,1] \times [-1,1]^{d-1}$. This hyperplane also bisects C to two parts, C_1 and C_2 (Figure 4.6/b). Every C_j contains B_j which contains a cube with a side-length of 1. Every C_j is of course still contained in B^+ which is cube with a side-length of $2R$. Hence every C_j is $2R$-fat. Hence the geometric loss of the partition $C = C_1 \sqcup C_2$, relative to the family or $2R$-fat objects, is 2, satisfying condition (a) of Lemma 4.4.2.

For every $j \in \{1,2\}$, define the following knife function K_j on C_j (see Figure 4.6/c,d):

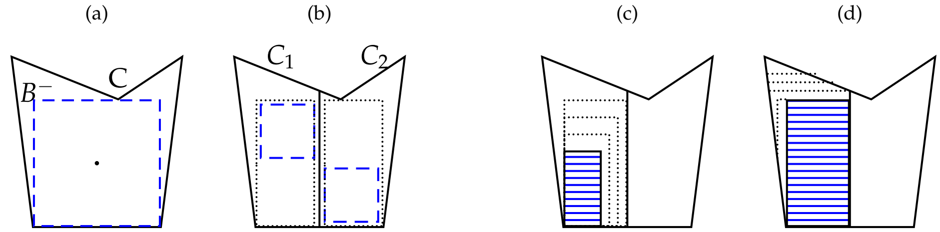

Figure 4.6: Dividing a general R-fat cake to two people.
(a) The R-fat cake C and its largest contained square B^- (the smallest containing square B^+ is not shown).
(b) The sub-cakes C_1 and C_2 (solid), the two rectangles B_1 and B_2 (dotted) and their largest contained squares (dashed).
(c) The knife function on C_1 in $t \in [0, \frac{1}{2}]$.
(d) The knife function on C_1 in $t \in [\frac{1}{2}, 1]$.

- For $t \in [0, \frac{1}{2}]$, $K_j(t) = (B_j)^{2t}$, i.e., the box B_i dilated by a factor of $2t$. Hence $K_j(0) = \emptyset$ and $K_j(\frac{1}{2}) = B_j$.
- For $t \in [\frac{1}{2}, 1]$, $K_j(t)$ is any knife-function from B_j to C_j with continuous Lebesgue-measure (see Subsection 4.A.1 for a proof that such a function exists).

$K_j(t)$ is always $2R$-fat, since in $[0, \frac{1}{2}]$ it is a scaled-down version of the box B_j (which is 2-fat) and in $[\frac{1}{2}, 1]$ it contains B_j and is contained in the cube B^+. $C \setminus K_j(t)$ is also $2R$-fat, since it contains B_{3-j} and is contained in B^+. Moreover, the Lebesgue measure of $K_j(t)$ is a continuous function of t. Hence, by Subsection 4.A.1, K_j is an S-good knife function, satisfying condition (b) of Lemma 4.4.2.

Since both K_j and $\overline{K_j}$ are $2R$-fat, the geometric loss of K_j relative to the family of $2R$-fat shapes is $1 + 1 = 2$, satisfying condition (c) of Lemma 4.4.2.

All conditions of Lemma 4.4.2 are satisfied, and its conclusion is exactly the claimed theorem. □

Theorem 4.1(c) implies that we can satisfy the two main fairness requirements: proportionality and envy-freeness, while keeping the allocated pieces sufficiently fat. The fatness guarantee means that each allotted piece: (a) contains a sufficiently *large* square, (b) is contained in a sufficiently *small* square. In the context of land division, these guarantees can be interpreted as follows: (a) Each land-plot has sufficient room for building a large house in a convenient shape (square); (b) The parts of the land that are valuable to the agent are close together, since they are bounded in a sufficiently small square.

Finally we note that a different technique leads to a version of Theorem 4.1(c) which guarantee that the pieces are not only $2R$-fat but also *convex* (if the original cake is convex); hence an agent can walk in a straight line from his square house to his valuable spots without having to enter or circumvent the neighbor's fields. See Appendix 4.B for details.

4.4.5 Between envy-freeness end proportionality

For all cakes C and families of usable pieces S studied in this section, we proved that there exists a positive constant p such that $\text{PropEF}(C, S, 2) \geq p$. Moreover, for the cases in which $p < 1/2$, we proved in Section 3.3 that $\text{Prop}(C, S, 2) \leq p$ (for the cases in which $p = 1/2$ the latter inequality is obvious). Since $\text{PropEF}(C, S, 2) \leq \text{PropEF}(C, S, 2)$ always, we get that for all settings studied here:

$$\text{PropEF}(C, S, 2) = \text{Prop}(C, S, 2)$$

In other words, in these cases, envy-freeness is compatible with the best possible partial-proportionality.

It is an open question whether this equality holds for *every* combination of cakes C and families S.

What *can* we say about the relation between proportionality and envy-freeness for arbitrary C and S? In addition to the trivial upper bound $\text{PropEF}(C, S, 2) \leq \text{Prop}(C, S, 2)$, we have the following lower bound:

Lemma 4.4.4. *For every cake C and family S:*

$$PropEF(C, S, 2) \geq Prop(C, S, 2) \cdot \inf_{s \in S} PropEF(s, S, 2)$$

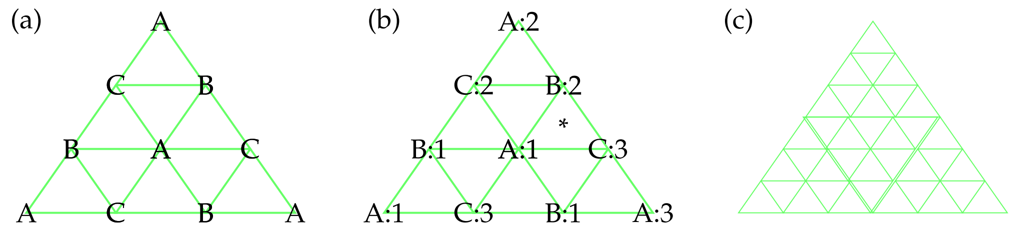

Figure 4.7: An illustration of the Simmons-Su procedure for $n = 3$ agents, A B and C.
(a) A triangulation of the simplex of partitions in which each vertex is assigned to an agent.
(b) Each vertex is labeled with the index of the piece preferred by its assigned agent. The fully-labeled triangle is starred.
(c) The process is repeated with a finer triangulation of the original simplex.

Proof. Let $p = \text{Prop}(C, S, 2)$ and $e = \inf_{s \in S} \text{PropEF}(s, S, 2)$. The following meta-procedure yields an envy-free partition of C in which the utility of each agent is at least $p \cdot e$.

By the definition of $\text{Prop}(C, S, 2)$, there exists an S-allocation $X = (X_1, X_2)$ with a proportionality of at least p, i.e, each agent i receives an S-piece X_i with $V_i(X_i) \geq p$.

Ask each agent whether he envies the other agent and proceed accordingly:

(a) If no agent envies the other agent, then the partition is already envy-free. The utility of each agent is at least p, which is at least $p \cdot e$ (since $e \leq 1$).

(b) If both agents envy each other, then let them switch the pieces. The resulting partition is envy-free and the utility of each agent is more than $p \geq p \cdot e$.

(c) The remaining case is that only one agent envies the other agent. W.l.o.g, assume it is agent 1 who envies agent 2. This means that the S-piece X_2 has a utility of at least p to both agents. By the assumptions of the lemma, since $X_2 \in S$, $\text{PropEF}(X_2, S, 2) \geq e$. Therefore, there exists an envy-free S-allocation of X_2 in which the utility of each agent i is at least $e \cdot V_i(X_2) \geq e \cdot p$. $\square$

So by previous results we have the following partial-compatibility results for *every* cake C:

- $\text{Prop}(C,\ Squares,\ 2) \geq \text{PropEF}(C,\ Squares,\ 2) \geq \frac{1}{4}\text{Prop}(C,\ Squares,\ 2)$
- $\text{Prop}(C,\ R\,fat\,rects,\ 2) \geq \text{PropEF}(C,\ R\,fat\,rects,\ 2) \geq \frac{1}{3}\text{Prop}(C,\ R\,fat\,rects,\ 2)$ (for $R \geq 2$)
- $\text{Prop}(C,\ Rectangles,\ 2) \geq \text{PropEF}(C,\ Rectangles,\ 2) \geq \frac{1}{2}\text{Prop}(C,\ Rectangles,\ 2)$

4.5 Envy-Free Division For n agents

4.5.1 The one-dimensional procedure

Existence of envy-free allocations in one dimension was first proved by Stromquist (1980). A procedure for finding such allocations was developed by Simmons and first described by Su (1999). Our procedure for n agents is a generalization of that procedure. We briefly describe the 1-dimensional procedure below.

The cake is the 1-dimensional interval $[0, 1]$ and S is the family of intervals. A partition of the cake to n intervals can be described by a vector of length n whose elements are the lengths of the intervals. The sum of all lengths in a partition is 1, so the set of all partitions is an $(n-1)$-dimensional simplex in $\mathbb{R}^n$. The procedure proceeds as follows (see Figure 4.7):

(a) *Preparation*. Triangulate the simplex of partitions to a collection of $(n-1)$-dimensional sub-simplexes. Assign each vertex of the triangulation to one of the n agents, such that in each sub-simplex, all n agents are represented. Su shows that there always exists such a triangulation.

(b) *Evaluation*. Recall that each vertex of the triangulation corresponds to a partition of the cake to n intervals. For each vertex, ask its assigned agent: "if the cake is partitioned according to this vertex, which piece would you prefer?". The answer is an integer between 1 and n; label that vertex with that integer.

The labeling created in step (b) has a special structure. First, each of the n main vertexes of the large simplex corresponds to a partition in which a single piece $i \in \{1, \ldots, n\}$ encompasses the entire cake and all other pieces are empty. Any agent prefers the entire cake over an empty piece, so this vertex will surely be labeled by i (see Figure 4.7/b, where the three vertexes of the large triangle are labeled by 1, 2 and 3). Moreover, each point on the segment between vertex i_1 and vertex i_2 corresponds to a partition in which the cake is divided between pieces i_1 and i_2, and all other pieces are empty. Therefore, each such point will be labeled by either i_1 or i_2. The same is true in any number of dimensions: in each face of the simplex, all interior points are labeled by one of the labels of the main vertexes that span that face. A labeling that has such a structure is called a *Sperner labeling*. By Sperner's lemma, any triangulation with a Sperner labeling has a *fully-labeled sub-simplex*, in which all vertexes are labeled differently.

(c) *Refinement*. Steps (a) and (b) can be repeated again and again, each time with a finer triangulation. This yields an infinite sequence of fully-labeled simplexes. By compactness of the simplex, there is a subsequence that converges to a single point. By the continuity of the agents' valuations, this point corresponds to a partition in which each of the n agents prefers a different piece. By definition, this partition is envy-free.

Note that the above procedure is infinite — the envy-free partition is found only at the limit of an infinite sequence. In fact, Stromquist (2008) proved that when $n \geq 3$, an envy-free partition to n agents with connected pieces cannot be found by a finite procedure. Therefore, Simmons' infinite procedure is the best that can be hoped for. Deng and Qi and Saberi (Deng et al., 2012) show that an approximately-envy-free division can be found in bounded time. For example, suppose that an interval is divided among several agents and they all agree that a 1 centimeter movement of the border between their plots is irrelevant. Then the simplex of partitions can be divided to sub-simplices of side-length 1 cm. If the total length of the cake is L centimeters, then a fully-labeled simplex can be found using $O(L^{n-2})$ queries (Deng et al., 2012, Theorem 5). All points in that simplex correspond to a division that is approximately-envy-free up to the agents' tolerance.

4.5.2 Knife tuples

Both Stromquist's existence proof and the Simmons–Su and the Deng–Qi–Saberi algorithms do not work directly on the cake — they work on the unit simplex, each point of which represents a cake-partition. Therefore, we can extend these algorithms to two dimensions if we find an appropriate way to map each point of the unit simplex to a two-dimensional cake-partition.

Our main tool is a *knife-tuple* — an extension of the knife-function defined in Definition 4.3.5.

Definition 4.5.1. Given a cake C, an *n-knife-tuple on C* is a vector of n functions $(K_1, \dots, K_n)$, which is a function from Δ^n (the $(n-1)$-dimensional unit-simplex in $\mathbb{R}^n$) to the partitions of C, such that for every nonempty subset of indexes $I \subseteq \{1, \dots, n\}$, if:

$$\sum_{i \in I} t_i = 1 \qquad \text{and} \qquad \forall i \notin I : t_i = 0,$$

then the pieces whose indexes are in I form a partition of the cake and the other pieces are empty:

$$\bigsqcup_{i \in I} K_i(t_1, \dots, t_n) = C \qquad \text{and} \qquad \forall i \notin I : K_i(t_1, \dots, t_n) = \emptyset.$$

In particular, at endpoint #i of the simplex, piece #i comprises the entire cake. I.e, if $t_i = 1$ and $t_{i' \neq i} = 0$, then $K_i = C$ and $K_{i \neq i} = \emptyset$.

Knife-tuples can be constructed from knife-functions.

Lemma 4.5.2. *Let C be a cake and K a knife-function from $\emptyset$ to C. Define functions K_1, K_2:*

$$K_1(t_1, t_2) := K(t_1)$$
$$K_2(t_1, t_2) := C \setminus K(1 - t_2)$$

Then, (K_1, K_2) is a 2-knife-tuple on C.

Proof. We verify the knife-tuple property for all nonempty subset of indexes $I \subseteq \{1, 2\}$:

- $I = \{1, 2\}$: since we are on the unit simplex, $t_1 + t_2 = 1$, $K_1 = K(t_1)$ and $K_2 = C \setminus K(t_1)$ so indeed $K_1 \sqcup K_2 = C$.

- $I = \{1\}$: When $t_1 = 1$ and $t_2 = 0$, $K_1 = K(1) = C$ and $K_2 = C \setminus K(1) = \emptyset$.
- $I = \{2\}$: When $t_2 = 1$ and $t_1 = 0$, $K_1 = K(0) = \emptyset$ and $K_2 = C \setminus K(0) = C$.

□

Longer knife-tuples can be constructed recursively, by replacing an element of an existing knife-tuple with two elements separated by a knife-function. We exemplify this construction with a 3-knife-tuple.

Lemma 4.5.3. *Let C be a cake and (K_1, K_2) a 2-knife-tuple on C. Suppose that, for every t_1 and every $t_2 > 0$, we have a knife-function K^{t_1,t_2} from $\emptyset$ to $K_2(t_1, t_2)$. Then, replacing the function K_2 with two complementary functions K'_2 and K'_3 gives a 3-knife-tuple (K'_1, K'_2, K'_3):*

$$
\begin{aligned}
K'_1(t_1, t_2, t_3) &:= && K_1(t_1, t_2 + t_3) && \\
K'_2(t_1, t_2, t_3) &:= && K^{t_1,t_2+t_3}\left(\frac{t_2}{t_2 + t_3}\right) && [t_2 + t_3 > 0] \\
& && \emptyset && [t_2 + t_3 = 0] \\
K'_3(t_1, t_2, t_3) &:= && K_2(t_1, t_2 + t_3) \setminus K^{t_1,t_2+t_3}\left(\frac{t_2}{t_2 + t_3}\right) && [t_2 + t_3 > 0] \\
& && \emptyset && [t_2 + t_3 = 0]
\end{aligned}
$$

Proof. We verify that (K'_1, K'_2, K'_3) satisfies the knife-tuple property for all nonempty subsets of indexes. Recall that the knife-tuple property of the original (K_1, K_2) implies that: $K_1(1, 0) = K_2(0, 1) = C$ and $K_1(0, 1) = K_2(1, 0) = \emptyset$ and $K_1(t_1, 1 - t_1) \sqcup K_2(t_1, 1 - t_1) = C$.

- When $t_1 = 1$ and $t_2 = t_3 = 0$, $K'_1 = K_1(1, 0) = C$ and $K'_2 = K'_3 = \emptyset$ by definition.
- When $t_2 = 1$ and $t_1 = t_3 = 0$, $K'_1 = K_1(0, 1) = \emptyset$ and $K'_2 = K^{0,1}(1) = K_2(0, 1) = C$ and $K'_3 = K_2(0, 1) \setminus K^{0,1}(1) = \emptyset$.
- When $t_3 = 1$ and $t_1 = t_2 = 0$, $K'_1 = K_1(0, 1) = \emptyset$ and $K'_2 = K^{0,1}(0) = \emptyset$ and $K'_3 = K_2(0, 1) \setminus K^{0,1}(0) = K_2(0, 1) = C$.
- When $t_1 + t_2 = 1$ and $t_3 = 0$, $K'_1 = K_1(t_1, t_2)$ and $K'_2 = K^{t_1,t_2}(1) = K_2(t_1, t_2)$ so $K'_1 \sqcup K'_2 = K_1 \sqcup K_2 = C$, and $K'_3 = 0$.
- When $t_1 + t_3 = 1$ and $t_2 = 0$, $K'_1 = K_1(t_1, t_3)$ and $K'_3 = K_2(t_1, t_3) \setminus K^{t_1,t_3}(0) = K_2(t_1, t_3)$ so $K'_1 \sqcup K'_3 = K_1 \sqcup K_3 = C$, and $K'_2 = 0$.
- When $t_2 + t_3 = 1$ and $t_1 = 0$, $K'_1 = K_1(0, 1) = \emptyset$ and $K'_2 = K^{0,1}(\frac{t_2}{t_2+t_3})$ and $K'_3 = K_2(0, 1) \setminus K^{0,1}(\frac{t_2}{t_2+t_3})$, so $K'_2 \sqcup K'_3 = K_2(0, 1) = C$.
- When $t_1 + t_2 + t_3 = 1$, $K'_2 \sqcup K'_3 = K_2(t_1, t_2 + t_3)$, so $K'_1 \sqcup K'_2 \sqcup K'_3 = K_1(t_1, t_2 + t_3) \sqcup K_2(t_1, t_2 + t_3) = C$.

□

So to build a 3-knife-tuple, we start with a single knife-function on C that cuts it to $K_1 \sqcup K_2$. Then, for every point in time, we use another knife-function on K_2 that cuts it to $K'_2 \sqcup K'_3$. Alternatively, we can use a knife-function on K_1 that cuts it to $K'_1 \sqcup K'_3$; the proof is entirely analogous.

An example of a 3-knife-tuple is shown in Figure 4.8. There, the first knife-function ($K_1 \equiv K'_1$) is a growing pair-of-squares, identical to the knife-function in Figure 4.3/b. K_2 is its complement (which is also a pair-of-squares). For every point in time, the second knife-function (K'_2) is a growing union-of-four-squares. It starts at an empty set and grows until it covers all of K_2. K'_3 is the remainder, which is also a union of four squares.

The previous lemma can be generalized to create knife-tuples of arbitrary length.

Lemma 4.5.4. *Let C be a cake and $(K_1, \ldots, K_n)$ an n-knife-tuple on C. Suppose that for some $i \in \{1, \ldots, n\}$, for every $t_1, \ldots, t_n$ where $t_i > 0$, we have a knife-function $K^{t_1,\ldots,t_i,\ldots,t_n}$ from $\emptyset$ to $K_i(t_1, \ldots, t_i, \ldots, t_n)$. Then, replacing* *the index i with two indexes i1 and i2 and replacing the function K_i with two complementary functions K'_{i1} and K'_{i2} gives an $(n+1)$-knife-tuple $(K'_1, \ldots, K'_{i1}, K'_{i2}, \ldots, K'_n)$:*

$$
\begin{aligned}
K'_{i1}(t_1, \ldots, t_{i1}, t_{i2}, \ldots, t_n) \quad & := K^{t_1, \ldots, t_{i1}+t_{i2}, \ldots, t_n}\left(\frac{t_{i1}}{t_{i1}+t_{i2}}\right) && [t_{i1}+t_{i2} > 0] \\
& := \emptyset && [t_{i1}+t_{i2} = 0] \\
K'_{i2}(t_1, \ldots, t_{i1}, t_{i2}, \ldots, t_n) \quad & := K_i(t_1, \ldots, t_{i1}+t_{i2}, \ldots, t_n) \\
& \quad \setminus K^{t_1, \ldots, t_{i1}+t_{i2}, \ldots, t_n}\left(\frac{t_{i1}}{t_{i1}+t_{i2}}\right) && [t_{i1}+t_{i2} > 0] \\
& := \emptyset && [t_{i1}+t_{i2} = 0] \\
\forall j \neq i : K'_j(t_1, \ldots, , t_{i1}, t_{i2}, \ldots, t_n) \quad & := K_j(t_1, \ldots, t_{i1}+t_{i2}, \ldots, t_n)
\end{aligned}
$$

Proof. We verify that $(K'_1, \ldots, K'_{i1}, K'_{i2}, \ldots, K'_n)$ satisfies the knife-tuple property for every nonempty subset of indexes, I'. There are four cases, depending on whether I' contains $i1$ or $i2$ or both.

- $i1 \notin I'$ and $i2 \notin I'$. Then, $\sqcup_{j \in I'} K'_j = \sqcup_{j \in I'} K_j = C$ by the knife-tuple property of $(K_1, \ldots, K_n)$, and $K'_{i1} = K'_{i2} = \emptyset$ by definition.
- $i1 \in I'$ and $i2 \notin I'$. When $t_{i2} = 0$, $t_{i1} + t_{i2} = t_{i1}$, so $K'_{i1} = K^{t_1, \ldots, t_{i1}, \ldots, t_n}(1) = K_i(t_1, \ldots, t_{i1}, \ldots, t_n)$. Define an alternative subset of indexes: $I := I' \setminus \{i1\} \cup \{i\}$. Then, $\sqcup_{j \in I'} K'_j = \sqcup_{j \in I} K_j = C$ by the knife-tuple property of $(K_1, \ldots, K_n)$, and $K'_{i2} = \emptyset$ by definition.
- $i2 \in I'$ and $i1 \notin I'$. When $t_{i1} = 0$, $t_{i1} + t_{i2} = t_{i2}$, so $K'_{i2} = K_i(t_1, \ldots, t_{i2}, \ldots, t_n) \setminus K^{t_1, \ldots, t_{i2}, \ldots, t_n}(0) = K_i(t_1, \ldots, t_{i2}, \ldots, t_n)$. Define an alternative subset of indexes: $I := I' \setminus \{i2\} \cup \{i\}$. Then, $\sqcup_{j \in I'} K'_j = \sqcup_{j \in I} K_j = C$ by the knife-tuple property of $(K_1, \ldots, K_n)$, and $K'_{i1} = \emptyset$ by definition.
- $i2 \in I'$ and $i1 \in I'$. Note that when $t_i := t_{i1} + t_{i2}$, $K'_{i1} \sqcup K'_{i2} = K_i$. Define a subset of indexes $I := I' \setminus \{i1, i2\} \cup \{i\}$. Then, $\sqcup_{j \in I'} K'_j = \sqcup_{j \in I} K_j = C$ by the knife-tuple property of $(K_1, \ldots, K_n)$.

$\square$

Definition 4.3.6 on page 65 and Definition 4.3.8 on page 66 can be naturally generalized from a knife-function to a knife-tuple:

Definition 4.5.5. The *geometric-loss* of a knife-tuple $(K_1, \ldots, K_n)$ is the supremum geometric loss of the resulting partitions:

$$\text{Loss}((K_1, \ldots, K_n), S) := \sup_{(t_1, \ldots, t_n) \in \Delta^n} \left(\sum_{j=1}^{n} \text{Loss}(K_j(t_1, \ldots, t_n), S) \right)$$

In Figure 4.8, K'_1 can be covered by two squares and K'_2 and K'_3 can be covered by four squares each, so the square-geometric-loss of this 3-knife-tuple is 10.

Definition 4.5.6. A knife-tuple $(K_1, \ldots, K_n)$ is called *S-good* if for every $i \in \{1, \ldots, n\}$ and every absolutely-continuous value-measure V, the function $V^S(K_i(t_1, \ldots, t_n))$ is a continuous function of $t_1, \ldots, t_n$.

In the knife-tuple of Figure 4.8, the squares meet only at their corners, no square is created or destroyed abruptly, so the knife-tuple is square-good. This can be proved formally as in Appendix 4.A.2; the details are omitted for the sake of brevity.

Now we can generalize Lemma 4.4.1 from 2 to n agents:

Lemma 4.5.7. *Let C be a cake and S a family of pieces. If there is an S-good n-knife-tuple on C with a geometric loss of at most M, then:*

$$PropEF(C, S, n) \geq 1/M$$

Proof. Use the Simmons–Su procedure described in Subsection 4.5.1. The Preparation step (a) is exactly the same. In the Evaluation step (b), for each vertex $(t_1, \ldots, t_n)$ of the triangulation, use the n-knife-tuple to

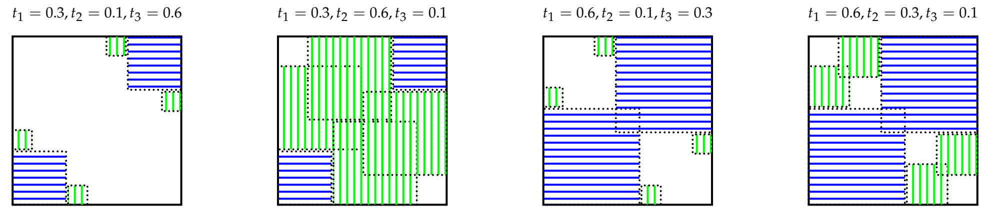

Figure 4.8: Four partitions induced by the 3-knife-tuple (K_1', K_2', K_3') in different points (t_1, t_2, t_3) of the unit-simplex. $K_1'(\cdot)$ is filled with horizontal blue lines, $K_2'(\cdot)$ is filled with vertical green lines and $K_3'(\cdot)$ is blank.

create the partition: $K_1(t_1, \ldots, t_n), \ldots, K_n(t_1, \ldots, t_n)$. Ask the owner of that vertex (e.g. agent i) to indicate its favorite piece in this partition, namely:

$$\arg \max_{j \in \{1,\ldots,n\}} V_i^S(K_j(t_1, \ldots, t_n))$$

and label that vertex with the agent's reply. By the properties of a knife-tuple, whenever $t_j = 0$, $K_j = \emptyset$, so $V_i^S(K_j) = 0$, so the agent will never reply j. Therefore, the resulting labeling is a Sperner labeling, so a fully-labeled sub-simplex exists.

By repeating steps (a) and (b) infinitely many times with finer and finer triangulations, we get a sub-sequence of fully-labeled triangles that converges to a single point. Because the knife-tuple is S-good, all agents' S-value functions are continuous, so the limit point corresponds to an envy-free partition. The loss of the knife-tuple is at most M, so the proportionality of the limit partition is at least $1/M$. [5] □

We now apply Lemma 4.5.7 to prove our Theorem 4.2.

4.5.3 Squares and fat rectangles

Theorem 4.2(a). For every $n \geq 1$:

$$\text{PropEF}(Square, Squares, n) \geq \frac{1}{2^{2\lceil log_2 n \rceil}} > \frac{1}{4n^2}$$

Proof. For every n which is a power of 2, we construct an n-knife-tuple $(K_1, \ldots, K_n)$, in which for every $(t_1, \ldots, t_n) \in \Delta^n$, and for every $j \in \{1, \ldots, n\}$ for which $t_j > 0$, $K_j(t_1, \ldots, t_n)$ is a union of at most n squares. Hence, the partition induced by $(K_1, \ldots, K_n)$ has a geometric loss of $n \cdot n = n^2$.

The construction is recursive. The base is $n = 2$. Take the knife-function in Figure 4.3/b (a union of two corner-squares growing towards the center). By Lemma 4.5.2, it defines a 2-knife-tuple which we denote by: (K_1, K_2). For each t_1 and t_2, $K_1(t_1, t_2)$ and $K_2(t_1, t_2)$ are square-pairs (unions of two squares).

Consider next the case $n = 4$. In every square-pair in the above 2-knife-tuple, define a knife-function as shown in Figure 4.8 — a union of four corner-squares growing from opposite corners towards the center. By Lemma 4.5.4, we can replace K_1 by K_1', K_2' and K_2 by K_3', K_4'. For each $i \in \{1, 2, 3, 4\}, t_1, t_2, t_3, t_4$, $K_i(t_1, t_2, t_3, t_4)$ is a union of four squares.

After l steps, we have a 2^l-knife-tuple in which each component is a union of 2^l squares. We split each component using a knife-function made of a union of 2^{l+1} squares growing from opposite corners. This gives a new, 2^{l+1}-knife-tuple in which each component is a union of 2^{l+1} squares. After $\log_2 n$ steps, we get the desired n-knife-tuple.

This knife-tuple is square-good since no squares are created or destroyed abruptly; this is apparent in the illustration, since the squares from opposite sides meet only at their corners. We suppress a formal proof of this geometric fact.

When n is not a power of two, it can be rounded to the next power of two — $2^{\lceil log_2 n \rceil}$. The geometric loss is then at most $2^{2\lceil log_2 n \rceil}$, which is always less than $4n^2$. □

[5] When $n = 3$, the three-knives procedure of Stromquist (1980) can be used instead of Simmons' procedure. See the conference version (Segal-Halevi et al., 2015a).

For completeness we prove the following very simple theorem:

Theorem 4.2(b). If C is an R-fat rectangle and S the family of R-fat rectangles then:

$$\text{PropEF}(C,S,n) \geq \frac{1}{2^{2\lceil log_2 n\rceil}} > \frac{1}{4n^2}$$

Proof. Scale the coordinate system such that C becomes a square. Use Theorem 4.2(a) and get a division with square pieces. Scale the coordinate system back. Now the pieces are R-fat rectangles. □

We do not know if the $1/(4n^2)$ lower bound is asymptotically tight. The upper bound from Claim 3.3.4 on page 18 is $\text{Prop}(Square, squares, n) \leq 1/(2n)$. Moreover, the procedure of Subsection 3.5.1 on page 27 proves that $\text{Prop}(Square, squares, n) \geq 1/(4n-4)$, but ignores envy considerations. We do not know if it is possible to attain an envy-free division with a proportionality of $1/O(n)$.

In the following subsection we show that it is possible to attain an envy-free *and proportional* division for every n, in return to a compromise on the family of usable pieces.

4.5.4 Arbitrary fat objects

Theorem 4.2(c). Let C be a d-dimensional R-fat cake and $n \geq 2$ an integer. Let S be the family of mR-fat pieces, where m be the smallest integer such that $n \leq m^d$ (i.e. $m = \lceil n^{1/d} \rceil$). Then:

$$\text{PropEF}(C,S,n) = 1/n$$

Proof. The proof is illustrated in Figure 4.9 for the case of $d = 2$ dimensions. Let C be an R-fat d-dimensional cake. By definition of fatness it contains a cube B^- of side-length x and it is contained in a parallel cube B^+ of side-length $R \cdot x$, for some $x > 0$.

Partition the cube B^- to a grid of m^d sub-cubes, $B_1, ..., B_{m^d}$, each of side-length $\frac{x}{m}$. For every i, denote by B_{-i} the union of all $m^d - 1$ squares different than B_i, i.e:

$$B_{-i} := \bigcup_{j \neq i} B_j = B^- \setminus B_i$$

Denote by $\overline{B^-}$ the cake outside the enclosed cube, i.e:

$$\overline{B^-} := C \setminus B^-$$

Define the following knife function K on C (see Figure 4.9):

- For $t \in [0, \frac{1}{3}]$: $K(t) = (B_1)^{3t}$, i.e., the cube B_1 dilated by a factor of $3t$. Hence $K(0) = \emptyset$ and $K(\frac{1}{3}) = B_1$.
- For $t \in [\frac{1}{3}, \frac{2}{3}]$: $K(t)$ is any knife function from B_1 to $C \setminus B_{-1}$ with continuous Lebesgue measure. See Subsection 4.A.1 for a proof that such a function exists.
- For $t \in [\frac{2}{3}, 1]$: $K(t)$ is $C \setminus [(B_{-1})^{3(1-t)}]$, i.e., the cake not yet covered by the knife is B_{-1} dilated by a factor proportional to the remaining time. Hence $K(1) = C$.

By Lemma 4.5.2, K induces a 2-knife-tuple (K_1, K_2) where $K_1 := K$ and $K_2 := C \setminus K_1$. For every t_1, t_2 with $t_1 + t_2 = 1$, $K_1(t_1, t_2)$ is mR-fat:

- When $t_1 \in [0, \frac{1}{3}]$, K_1 it is a cube, which is 1-fat.
- When $t_1 \in [\frac{1}{3}, 1]$, K_1 contains the cube B_1, whose side-length is x/m, and is contained in the cube B^+, whose side-length is $x \cdot R$.

and $K_2(t_1, t_2)$ is also mR-fat:

- When $t_1 \in [0, \frac{2}{3}]$, K_2 contains e.g. the cube B_n, whose side-length is x/m, and is contained in the larger cube B^+, whose side-length is $x \cdot R$.

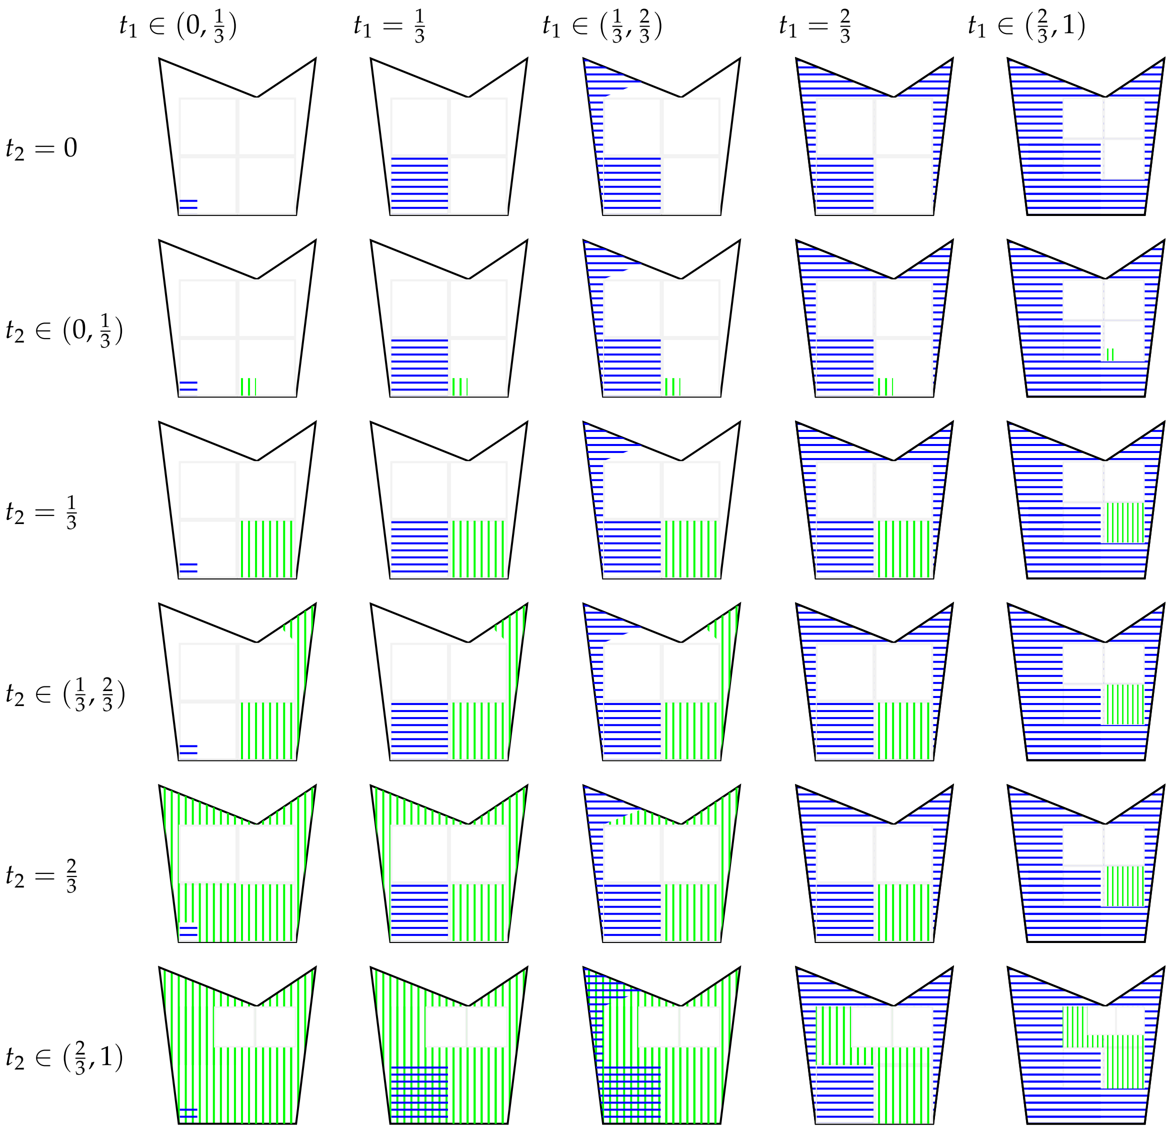

Figure 4.9: Dividing a general R-fat cake to $n = 3$ people. K_1 is filled with horizontal lines, K_2 is filled with vertical lines and K_3 is white. Note that each of these three pieces is $2R$-fat, where R is the fatness of the original cake.

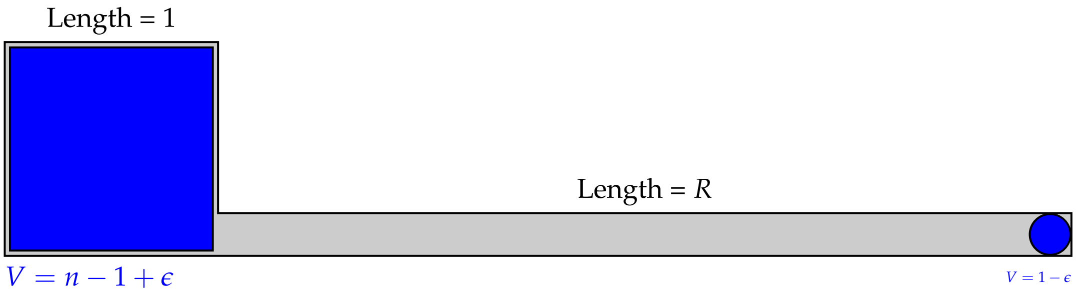

Figure 4.10: A fat cake in which every proportional division must use slim pieces. See Lemma 4.5.8.

- When $t_1 \in [\frac{2}{3}, 1]$, K_2 contains a dilated B_n and it is contained in a dilated B^-; since they are dilated by the same factor, the ratio between their side-lengths is always m.

For every t_1, t_2 with $t_1 + t_2 = 1$, we now define a knife-function K^{t_1,t_2} from $\emptyset$ to $K_2(t_1, t_2)$. K^{t_1,t_2} is analogous to K but uses the sub-cube B_2. This is possible because:

- When $t_1 \in [0, \frac{2}{3}]$, K_2 contains the cube B_2 itself;
- When $t_1 \in (\frac{2}{3}, 1]$, K_2 contains a dilated B_2, which is contained in a dilated B^+.

The function K^{t_1,t_2} is defined as follows:

- For $t \in [0, \frac{1}{3}]$: $K^{t_1,t_2}(t) = (B_2)^{3t}$.
- For $t \in [\frac{1}{3}, \frac{2}{3}]$: $K^{t_1,t_2}(t)$ is any knife-function from B_2 to $K_2 \setminus B_{-2}$ with continuous Lebesgue measure.
- For $t \in [\frac{2}{3}, 1]$: $K^{t_1,t_2}(t)$ is $K_2 \setminus [(B_{-2})^{3(1-t)}]$.

By Lemma 4.5.3, this induces a 3-knife-tuple (K'_1, K'_2, K'_3).

To define an n-knife-tuple, proceed in a similar way for the pieces $B_1, \ldots, B_n$. All components in the knife-tuple are mR-fat, and their Lebesgue measure changes continuously. Therefore, by the proofs in Subsection 4.A.1, the knife-tuple is S-good, as required by Lemma 4.5.7. $\square$

Figure 4.9 shows an example of the construction for $d = 2$ dimensions and $n = 3$ agents. Here $m = \lceil \sqrt{3} \rceil = 2$ so each agent receives an envy-free $2R$-fat land-plot with a utility of at least $1/3$.

Theorem 4.2(c) implies that we can guarantee proportionality by compromising on the fatness of the pieces — allowing the pieces to be thinner than the cake by a factor of $\lceil n^{1/d} \rceil$. This factor is asymptotically optimal even when envy is allowed:

Lemma 4.5.8. *For every $R \geq 1$, there is an $(R+1)$-fat cake C for which, for every $m' \leq (n-1)^{1/d}$:*

$$PropEF(C, m'R\,fat\,objects, n) \leq Prop(C, m'R\,fat\,objects, n) < 1/n$$

Proof. Let δ, ϵ be small positive constants. Let C be a cake with the following two components:

- The left component is a cube with all sides of length 1;
- The right component is a box with one side of length R and the other sides of length δ.

See Figure 4.10 for an illustration for $d = 2$. C is contained in a cube of side-length $R + 1$ and it contains a cube of side-length 1, so it is $(R + 1)$-fat.

C represents a desert with the following water sources:

- The left cube contains $n - 1 + \epsilon$ water units;

- A small disc at the end of the right box contains $1-\epsilon$ water units.

C has to be divided among n agents whose value functions are proportional to the amount of water. To get a proportional division, each agent must receive exactly 1 unit of water. This means that at least one piece, e.g. X_i, must overlap both the right pool and the left pool.

The smallest cube containing X_i has a side-length of at least R. For the largest cube contained in X_i, there are two options:

- If the largest contained cube is in the left side, then its side-length must be at most $\left(\frac{1}{n-1+\epsilon}\right)^{1/d}$, since it must contain at most 1 unit of water.

- If the largest contained cube is in the right side, then its side-length must be at most δ.

If δ is sufficiently small (in particular, $\delta < \left(\frac{1}{n-1}\right)^{1/d}$), then the piece X_i is not $m'R$-fat for every $m' \leq (n-1)^{1/d}$. This means that, if all pieces must be $m'R$-fat, a proportional division is impossible. $\square$

4.6 Conclusions and Future Work

This chapter presented the problem of dividing a cake to agents whose utility functions depend on geometric shape, where the division should be both partially-proportional and envy-free. The main contributions in this chapter are several generic division procedures for envy-free division. For two agents, these procedures have the best possible partial-proportionality guarantees in various geometric scenarios. For n agents, the procedures guarantee a positive partial proportionality.

The tools developed in this chapter are generic and can work for cakes and pieces of other geometric shapes. In fact, our tools reduce the envy-free division problem to a geometric problem — the problem of finding appropriate knife functions.

Some topics not covered in the present chapter are:

- Utility functions that takes into account both the value contained in the best usable piece and the total value of the piece, e.g.: $U(X) = w \cdot V^S(X) + (1-w) \cdot V(X)$, where w is some constant.

- Absolute size constraints on the usable pieces instead of the relative fatness constraints studied here, e.g. let S be the family of all rectangles with length and width of at least 10 meters.

- Personal geometric preferences — letting each agent i specify a different family S_i of usable pieces.

4.7 Acknowledgments

We are grateful to the anonymous referees of the AAAI 2015 conference for their helpful comments. We are also grateful to participants of the AAAI 2015 conference and the Dagstuhl seminar on fair division for their comments.

Nerly Biton's comments on a preliminary version of this manuscript greatly improved its readability.

Many members of the math.stackexchange.com, mathoverflow.com and cs.stackexchange.com communities kindly helped in coping with specific issues arising in this paper. In particular, we are grateful to: Julián Aguirre, Christian Blatter, Ilya Bogdanov, Beni Bogosel, Boris Bukh, Anthony Carapetis, Scott Carnahan, Christopher Culter, Gabriel C. Drummond-Cole, David Eppstein, Yuval Filmus, Christopher Fish, Peter Franek, Nick Gill, John Gowers, Michael Greinecker, Dafin Guzman, Włodzimierz Holsztyński, Marcus Hum, Robert Israel, Barry Johnson, Joonas Ilmavirta, V. Kurchatkin, Jonathan Lubin, Raymond Manzoni, Ross Millikan, Lee Mosher, Rahul Narain, Sebastian Negraszus, Mariusz Nowak, Boris Novikov, Joseph O'Rourke, Emanuele Paolini, Raphael Reitzig, David Richerby, András Salamon, Realz Slaw, B. Stoney, Eric Stucky, Steven Taschuk, Marc van Leeuwen, Martin van der Linden, Hagen von Eitzen, Martin von Gagern, Jared Warner, Frank W., Ittay Weiss, Willie Wong, Phoemue X and Amitai Yuval.

Chapter 4 Appendix

4.A Geometric conditions for S-good knife functions

Recall Definition 4.3.6:

> Given a cake C and a family S, a knife function K_C is called *S-good* if for every absolutely-continuous value-measure V, both $V^S(K_C(t))$ and $V^S(\overline{K_C}(t))$ are continuous functions of t.

This section presents two different geometric properties of a knife function K_C, each of which guarantees that it is S-good.

4.A.1 S-smoothness

The first property is simple: both the region covered and the region not covered by the knife function should always return S-pieces whose Lebesgue measure changes continuously.

Definition 4.A.1. Let S be a family of pieces. A knife function $K(t)$ is called *S-smooth* if:
(a) The Lebesgue measure of $K(t)$ (and hence of $\overline{K}(t)$) is a continuous function of t, and:
(b) for all t, both $K(t) \in S$ and $\overline{K}(t) \in S$.

Lemma 4.A.2. *If V is a measure absolutely-continuous with respect to Lebesgue measure, and K is an S-smooth knife-function, then the real functions $V^S \circ K$ and $V^S \circ \overline{K}$ are continuous.*

Proof. The measure V is absolutely continuous with respect to Lebesgue measure, and $Lebesgue(K(t))$ is a continuous function of t by condition (a). Hence, $V(K(t))$ is also a continuous function of t. Condition (b), namely $K(t) \in S$, implies that $\forall t \in [0,1] : V^S(K(t)) = V(K(t))$, so $V^S(K(t))$ is also a continuous function of t. An analogous proof applies to $V^S(\overline{K}(t))$. □

The knife function in Figure 4.3/a is Rectangle-smooth but not Square-smooth. The other knife functions in that figure are neither Rectangle-smooth nor Square-smooth (e.g in Figure 4.3/c, $\overline{K}(t)$ is not a rectangle).

We now prove a useful lemma that will help us find S-smooth functions. Recall that S-smoothness has two conditions: $Lebesgue(K(t))$ should be continuous, and $K(t)$ should be in S. We now focus on the first condition — continuity of $Lebesgue(K(t))$.

Given two bounded Borel subsets of $\mathbb{R}^d$, A and B, does there always exist a knife function K from A to B such that $Lebesgue(K(t))$ is continuous? By the monotonicity of a knife-function, a necessary condition is that $A \subset B$. By the following lemma, this condition is also sufficient.

Lemma 4.A.3. *Let A and B be two bounded Borel subsets of $\mathbb{R}^d$ with $A \subseteq B$. There exists a knife function K from A to B, such that $Lebesgue(K(t))$ is a continuous function of t.*

Proof (based on Fish (2014)). Pick a point $O \in B$. For every $r \geq 0$ let $D(r)$ be the open d-ball of radius r around O. Since B is bounded, there is a certain radius r_{max} such that $B \subseteq D(r_{max})$. For every $t \in [0,1]$, define $D^*(t) = D(t \cdot r_{max})$, so $D^*(t)$ is an open ball whose radius grows continuously from 0 to r_{max}. Define: $K(t) := [A \cup D^*(t)] \cap B$. Clearly, $K(0) = A$, $K(1) = B$ and K is (weakly) monotonically increasing. Hence, K is a knife-function from A to B. The continuity of $Lebesgue(K(t))$ follows from the fact that $Lebesgue(D^*(t))$ is continuous and for every Δt: $Lebesgue(K(t + \Delta t)) - Lebesgue(K(\Delta t)) \leq Lebesgue(D^*(t + \Delta t)) - Lebesgue(D^*(\Delta t))$. □

We call any function satisfying the requirements of Lemma 4.A.3 a *knife-function with continuous Lebesgue-measure*. Any S-smooth knife-function has continuous Lebesgue-measure. Any knife-function with continuous Lebesgue-measure in which $K(t) \in S$ and $\overline{K}(t) \in S$ is S-smooth.

4.A.2 S-continuity

The second property is more involved. The knife function may return pieces that are not from S. However, it must change in a way that S-pieces are not created or destroyed abruptly, but rather grow or shrink in a continuous manner.

Definition 4.A.4. A piece s is called a *ϵ-predecessor* of a piece s' if $s \subseteq s'$ and $Lebesgue(s' \setminus s) < \epsilon$.

Definition 4.A.5. Let S be a family of pieces. A knife function $K(t)$ is called *S-continuous* if for every $\epsilon > 0$ there exists $\delta > 0$ such that, for all t and t' having $|t' - t| < \delta$:
(a) Every S-piece $s_{t'} \subseteq K(t')$ has an ϵ-predecessor S-piece $s_t \subseteq K(t)$.
(b) Every S-piece $s_{t'} \subseteq \overline{K}(t')$ has an ϵ-predecessor S-piece $s_t \subseteq \overline{K}(t)$.

Lemma 4.A.6. *If V is a measure absolutely-continuous with respect to Lebesgue measure, and K is an S-continuous knife function, then the real functions $V^S \circ K$ and $V^S \circ \overline{K}$ are uniformly-continuous.*

Proof. Given $\epsilon' > 0$, we show the existence of $\delta > 0$ such that, for every t, t', if $|t' - t| < \delta$ then $|V^S(K(t')) - V^S(K(t))| < \epsilon'$.

Given ϵ', by the continuity of V, there is an $\epsilon > 0$ such that:

$$Lebesgue(s) < \epsilon \qquad \Longrightarrow \qquad V(s) < \epsilon' \tag{4.2}$$

Given that ϵ, by the S-continuity of K there is a $\delta > 0$ such that, if $|t' - t| < \delta$, then every S-piece $s_{t'} \subseteq K(t')$ has an ϵ-predecessor S-piece $s_t \subseteq K(t)$. This means that $s_t \subseteq s_{t'}$ and:

$$Lebesgue(s_{t'} \setminus s_t) < \epsilon$$

which by (4.2) implies

$$V(s_{t'} \setminus s_t) < \epsilon'$$

which by additivity of V implies

$$V(s_t) > V(s_{t'}) - \epsilon'$$

The latter inequality is true for every S-piece $s_{t'} \subseteq K(t')$, so it is also true for the supremum:

$$\sup_{s_t \in S, s_t \subseteq K(t)} V(s_t) \geq V(s_t) > \sup_{s_{t'} \in S, s_{t'} \subseteq K(t')} V(s_{t'}) - \epsilon'$$

By definition, the S-value is the supremum, so:

$$V^S(K(t)) > V^S(K(t')) - \epsilon'$$

By symmetric arguments (replacing the roles of t and t'), $V^S(K(t')) > V^S(K(t)) - \epsilon'$. Hence $|V^S(K(t')) - V^S(K(t))| < \epsilon'$ as we wanted to prove.

An analogous proof applies to the function $V^S \circ \overline{K}$. □

The following lemma demonstrates how the existence of S-continuous functions can be proved.

Lemma 4.A.7. *Let S be the family of d-dimensional cubes. For every bounded cake C in $\mathbb{R}^d$, there exists an S-continuous knife-function from $\emptyset$ to C.*

Proof. Since C is bounded, it can be moved and scaled such that it is contained in the unit cube $[0,1]^d$. For every $t \in [0,1]$, Let $H(t)$ be the half-space defined by: $x < t$. Define: $K_C(t) := H(t) \cap C$. Clearly, $K_C(0) = \emptyset$, $K_C(1) = C$ and K_C is (weakly) monotonically increasing. Hence, K_C is a knife-function from $\emptyset$ to C.

The proof that K_C is S-continuous is based on the following geometric fact: for every cube $s_{t'}$ contained in the half-space $H(t + \delta)$, there exists a cube $s_t \subseteq s_{t'}$ contained in the half-space $H(t)$, such that the side-length of s_t is smaller than that of $s_{t'}$ by at most δ (it is smaller by exactly δ when $s_{t'}$ is adjacent to the rightmost side of $H(t + \delta)$ and parallel to the axes; see Figure 4.11 for an illustration of the two-dimensional

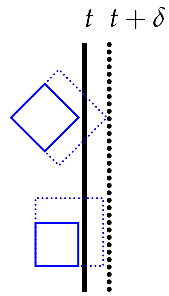

Figure 4.11: Square-continuity of the knife-function defined in Lemma 4.A.7.
The solid line describes the knife location at time t; the dotted line describes its location at time $t+\delta$.
The dotted squares are squares contained in $H(t+\delta)$; the solid squares are their predecessors in $H(t)$.
At the bottom, the side-length of the solid square is smaller than the dotted square by exactly δ.
At the top, the side-length of the solid square is smaller than the dotted square by less than δ.

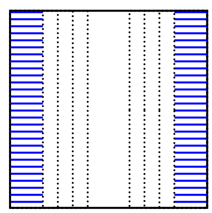

Figure 4.12: A knife-function that is not S-continuous.

case). Suppose $s_{t'}$ is also contained in C. Since C is contained in the unit cube, the side-length of $s_{t'}$ is at most 1. Therefore, the area of s_t is smaller than that of $s_{t'}$ by at most $1-(1-\delta)^d \leq d\cdot\delta$.

Consider now the definition of S-continuity. For every $\epsilon > 0$, take $\delta := \epsilon/d$, let $t' = t+\delta$ and let $s_{t'}$ be an S-piece contained in $K_C(t')$. By definition of K_C, $s_{t'}$ is contained in both C and $H(t')$. By the geometric fact, $s_{t'}$ has an ϵ-predecessor s_t that is contained in $H(t)$. Since $s_t \subseteq s_{t'}$, it is also contained in C. Hence, it is contained in $K_C(t)$. □

Using similar arguments, it is possible to prove that the function K_C described above is S-continuous also when S is the family of boxes or fat boxes. Full characterization of the the families S for which K_C is S-continuous is an interesting question that is beyond the scope of the present paper.

4.A.3 Examples

The knife-function in Figure 4.3/a, $K_C(t) = [0,L]\times[0,t]$, is a special case of the 'sweeping plane' function of Lemma 4.A.7. Hence it is square-continuous (and also rectangle-continuous).

As a negative example, consider the knife function $K_C(t) = [0,t]\times[0,1]\cup[1-t,1]\times[0,1]$ defined on the cake $C = [0,1]\times[0,1]$. This function describes two rectangles that approach each other from two opposite sides of the cake (see Figure 4.12). It is not square-continuous. Intuitively, a square of side-length 1 is created at time $t = 0.5$, when the two components of $K_C(t)$ meet. Formally, let $\epsilon = 0.75$. For every $\delta > 0$, select $t = 0.5 - \frac{\delta}{3}$ and $t' = 0.5 + \frac{\delta}{3}$. Then $K_C(t')$ contains the square $s' = [0,1]\times[0,1]$, but all squares $s \subseteq K_C(t)$ have a side-length of less than 0.5, hence $Lebesgue(s'\setminus s) > 0.75 = \epsilon$.

The knife-functions in Figure 4.3/b,c,d,e are S-continuous but not S-smooth. Thus one may think that S-continuity is more permissive than S-smoothness. But this is not the case: S-continuity and S-smoothness are two independent properties. To see this, let S' be the family of *rectangle-pairs* (defined as unions of two rectangles). The function K_C defined in the previous paragraph (and Figure 4.12) is S'-smooth, because both $K_C(t)$ and $\overline{K_C}(t)$ are rectangle-pairs. However, K_C is not S'-continuous because some rectangle-pairs (e.g. $[0,1]\times[0,0.2]\cup[0,1]\times[0.8,1]$) are created abruptly at time $t=0.5$.

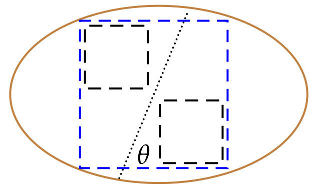

Figure 4.13: Dividing a convex R-fat cake to two people.
The cake (the ellipse) is divided by a rotating knife (dotted line) to two $2R$-fat convex pieces. This is a convex variant of Figure 4.6.

4.A.4 Conclusion

We proved two independent sufficient conditions for S-goodness. Combining Lemmas 4.A.2 and 4.A.6 gives:

Corollary 4.A.8. *If a knife-function is either S-smooth or S-continuous (or both), then it is S-good.*

Each of the two conditions, S-smoothness and S-continuity, is sufficient but not necessary for S-goodness.

4.B Dividing a convex fat cake to convex fat pieces

The following theorem is a variant of Theorem 4.1(c) in Subsection 4.4.4, in which the cake must be convex and the pieces are guaranteed to be convex.

The convexity requirement, while seemingly simple, implies that we cannot use the usual knife functions anymore. For example, if C is a circle then every knife function (which must be a straight line to keep the pieces convex) must start with an infinitely slim piece. Hence we must use another technique which can be called a *rotating-knife*.

Theorem. *For every $R \geq 1$, If C is an R-fat 2-dimensional convex figure and S is the family of convex $2R$-fat pieces then:*

$$PropEF(C, S, 2) = 1/2$$

Proof. Scale, rotate and translate the cake C such that the largest square contained in C is $B^- = [-1, 1] \times [-1, 1]$. By definition of fatness, C is now contained in a square B^+ of side-length at most $2R$.

Consider a line passing through the origin at angle $\theta \in [0°, 360°]$ from the x axis (see Figure 4.13). This line cuts the contained square B^- into two quadrangles, each of which contains a square with side-length 1. Because C is convex, this line also cuts the boundary of C at exactly two points, splitting C to two convex pieces. Each of these two pieces is $2R$-fat since it contains a square with side-length 1 and it is contained in B^+ whose side-length is $2R$.

Let $W(\theta)$ be the value of the piece for agent #1 at the left-hand side of the line when facing at angle θ. Because the value measure is continuous, W is continuous. When θ rotates by 180°, the piece that was at the left-hand side is now at the right-hand side and vice versa (e.g. when $\theta = 0°$ the left-hand side is above the line and when $\theta = 180°$ the right-hand side is above the line). Hence if $W(\theta) > 1/2$ then $W(180° + \theta) = 1 - W(\theta) < 1/2$ and vice versa. Hence by the continuity of W there must be a θ for which $W(\theta) = 1/2$. Cut the cake at the line in angle θ. Let agent #2 choose a piece and give the other piece to agent #1. Now both agents have a piece which is convex and $2R$-fat and their value is at least $1/2$. □

So far we have not managed to generalize the rotating-knife technique to more than two agents.

Remark 4.B.1. The rotating-knife technique was introduced by Robertson and Webb (1998) as an algorithm to produce an envy-free cake-allocation among three agents. The algorithm works only when the cake has two or more dimensions, and it does not provide geometric shape guarantees.

Rotating knives are also used by Barbanel and Brams (2011) for finding an approximate envy-free, equitable and Pareto-optimal division of a *pie* (a one-dimensional circle) between two players.

Chapter 5

Re-division and Price-of-Fairness

Contents

4.1 Introduction . . . **58**
4.1.1 Results . . . 59
4.1.2 Related Work . . . 60
4.2 Model . . . **61**
4.3 Geometric Preliminaries . . . **62**
4.3.1 Geometric loss . . . 62
4.3.2 Chooser Lemma . . . 63
4.3.3 Cover Numbers and Cover Lemma . . . 64
4.3.4 Knife functions . . . 65
4.3.5 Geometric loss of knife functions . . . 66
4.4 Envy-Free Division for Two Agents . . . **67**
4.4.1 Squares and rectangles . . . 67
4.4.2 Cubes and archipelagos . . . 68
4.4.3 Fat rectangles . . . 69
4.4.4 Arbitrary fat objects . . . 71
4.4.5 Between envy-freeness end proportionality . . . 72
4.5 Envy-Free Division For n agents . . . **73**
4.5.1 The one-dimensional procedure . . . 73
4.5.2 Knife tuples . . . 74
4.5.3 Squares and fat rectangles . . . 77
4.5.4 Arbitrary fat objects . . . 78
4.6 Conclusions and Future Work . . . **81**
4.7 Acknowledgments . . . **81**
Chapter appendix . . . **82**
4.A Geometric conditions for S-good knife functions . . . **82**
4.A.1 S-smoothness . . . 82
4.A.2 S-continuity . . . 83
4.A.3 Examples . . . 84
4.A.4 Conclusion . . . 85
4.B Dividing a convex fat cake to convex fat pieces . . . **85**

5.1 Introduction

The classic cake-cutting setting assumes a one-shot division: the resource is divided once and for all, like a cake that is divided and eaten soon after it comes out of the oven. But in practice, it is often required to re-divide an already-divided resource. One example is a cloud-computing environment, where new agents come and require resources held by other agents. A second example is fair allocation of radio spectrum among several broadcasting agencies: it may be required to re-divide the frequencies to accommodate new broadcasters. A third example is land-reform: large land-estates are held by a small number of landlords, and the government may want to re-divide them to landless citizens. A fourth example arises when new shares of a company are issued, and one desires that the old share-holders not feel unduly hurt by the dilution of their shares.[1]

In the classic one-shot division setting, there are n agents with equal rights. The goal is to give each agent a fair share of the cake. Ideally we would like to give each agent a piece worth at least $1/n$ of the total cake value — a requirement called "proportionality" (see Chapter 2 on page 6). If this is not possible, we would like to give each agent at least a fraction r/n of the total cake value, where $r \in (0,1)$ is constant independent of n. We call this requirement *r-proportionality*.

In contrast, in the re-division setting, there is an existing division of the cake among the n agents. This division is not necessarily fair; in particular, there may be some agents whose allocation is empty. If the cake is re-divided, it may be required to give extra rights to the existing landlords. In particular, it may be required to give each landlord the opportunity to keep a substantial fraction of its current value. This may be due either to efficiency reasons (in the cloud computing scenario) or economic reasons (in the radio spectrum scenario) or political reasons (in the land-reform scenario). We call this requirement *ownership*. Given a constant $w \in (0,1)$, *w-ownership* means that each agent receives at least w times its old value. The main question in this chapter is:

> Can positive levels of proportionality and ownership be attained simultaneously?

5.1.1 Results

Our first result answers this question positively.

Theorem 5.1. *For every constants $r, w \in [0,1]$ where $r + w \leq 1$, and for every existing division of the cake, there exists a division that simultaneously satisfies r-proportionality and w-ownership. Moreover, when r, w are constant rational numbers, such a division can be found in time $O(n^2)$.*

The parameters r, w represent the level of balance between two principles: large r means more emphasis on fairness while large w means more emphasis on ownership rights. As an example, taking $r = w = 1/2$, Theorem 5.1 implies that it is possible to re-divide the cake, giving each agent at least half its previous value, while simultaneously giving each agent at least $1/(2n)$ of the total cake value.

The balance parameters can also be given probabilistic interpretation. Suppose the government wants to do a land reform and needs the agreement of the current landowners. Naturally, the current landowners do not want to give away their lands. However, they may fear that, without land-reform, the landless citizens might revolt and they might lose all their lands. If the landowners believe that the probability of a successful revolt is $1 - w$, then they will agree to a land-reform that guarantees w-ownership. Theorem 5.1 implies that, in this case, it is possible to carry out a land-reform that guarantees $(1 - w)$-proportionality.

The following proposition shows that the balance given by Theorem 5.1 is tight:

Proposition 5.1. *For every constants $r, w \in [0,1]$ where $r + w > 1$, it may be impossible to simultaneously guarantee r-proportionality and w-ownership.*

Geometric constraints

While Theorem 5.1 is encouraging, it ignores an important aspect of practical division problems: geometry. The division it guarantees may be highly fractional, giving each agent a large number of disconnected pieces. But in land division (as well as many other practical division problems), the agents may want to receive a single connected piece. Can partial-proportionality and partial-ownership be attained simultaneously with a connectivity constraint? The following proposition answers this question negatively.

Proposition 5.2. *When the cake is a 1-dimensional interval and each piece must be an interval, for every positive constants $r, w \in (0,1)$, it may be impossible to simultaneously satisfy r-proportionality and w-ownership. Moreover, for every $r > 0$ and every integer $k \in \{1, \ldots, n\}$, there might be k agents who, in any r-proportional division, receive at most a fraction $1/\lfloor \frac{n}{k} \rfloor$ of their old value.*

[1] I am grateful to a referee for this example.

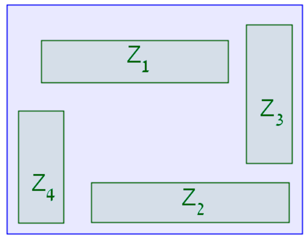

Figure 5.1: With geometric constraints, a Pareto-efficient division might paradoxically have to discard some of the cake.

The latter part of the proposition involves a property much weaker than proportionality: all we want is to guarantee each agent a positive value. With the connectivity constraint, even this weak "positivity" requirement is incompatible with w-ownership for every constant $w > 0$: a positive division might require us to give one agent at most $1/n$ of its previous value, give two agents at most $2/n$ of their previous value, give $n/3$ agents at most $1/3$ of their previous value, etc.

Proposition 5.2 motivates the following weaker ownership requirement: for every k, at least $n - k$ agents receive at least a fraction $1/\lfloor \frac{n}{k} \rfloor$ of their old value. For example (taking $k = n/3$ and assuming all quotients are integers), at least $2n/3$ agents should receive at least $1/3$ of their old value. This criterion is inspired by the "90th percentile" criterion common in Service-Level-Agreements and Quality-of-Service analysis, e.g. Zhang et al. (2014); Delimitrou and Kozyrakis (2014). It can also be justified by political reasoning: in a democratic country, it may be sufficient to win the support of a sufficiently large majority.

Our following results almost match this relaxed ownership criterion. Formally, the **democratic ownership** property means that, for every integer $k \in \{1, \ldots, n\}$, at least $n - k$ agents receive at least a fraction $1/\lceil \frac{n}{k} \rceil$ of their previous value. Democratic-ownership is almost the same as the upper bound implied by Proposition 5.2; the only difference is that in the upper bound the fraction is rounded downwards ($1/\lfloor \frac{n}{k} \rfloor$) while in democratic-ownership the fraction is rounded upwards.

Theorem 5.2. *When the cake is a 1-dimensional interval and each piece must be an interval, it is possible to find in time $O(n^2 \log n)$ a division simultaneously satisfying democratic-ownership and $1/3$-proportionality.*

It is an open question whether democratic-ownership is compatible with r-proportionality for some $r > 1/3$.

Theorem 5.2, like most cake-cutting papers, assumes that the cake is 1-dimensional. In realistic division scenarios, the cake is often 2-dimensional and the pieces should have a pre-specified geometric shape, such as a rectangle or a convex polygon. Rectangularity and convexity requirements are sensible when dividing land, exhibition space in museums, advertisement space in newspapers and even virtual space in web-pages. Moreover, in the frequency-range allocation problem, it is possible to allocate frequency ranges for a limited time-period; the frequency-time space is two-dimensional and it makes sense to require that the "pieces" are rectangles in this space (Iyer and Huhns, 2009).

2-dimensional cake-cutting introduces new challenges over the traditional 1 dimensional setting. As an example, in one dimension, it can be assumed that the initial allocation is a partition of the entire cake; this is without loss of generality, since any "blank" (unallocated part) can just be attached to a neighboring allocated interval without harming its shape or value. However, in two dimensions, the initial allocation might contain blanks that cannot be attached to any allocated piece due to the rectangularity or convexity constraints. For example, suppose the cake is as the rectangle in Figure 5.1. There are 4 agents and each agent i has positive value-density only inside the rectangle Z_i. The most reasonable division (e.g. the only Pareto-efficient division) is to give each Z_i entirely to agent i. But, this allocation leaves a blank in the center of the cake, and this blank cannot be attached to any allocated piece due to the rectangularity constraint.

This counter-intuitive scenario cannot happen in a one-dimensional cake. Handling such cases requires new geometry-based tools. Using such tools we can prove analogues of Theorem 5.2 to two common 2-dimensional settings.

Theorem 5.3. *When the cake is a rectangle and each piece must be a parallel rectangle, it is possible to find in time $O(n^2 \log n)$ a division simultaneously satisfying democratic-ownership and $1/4$-proportionality.*

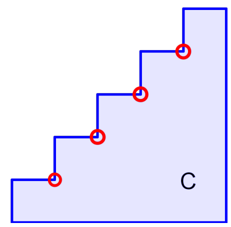

Figure 5.2: A rectilinear polygon with 4 reflex vertexes (circled).

Cake	Pieces	Value guarantee	Ownership
Arbitrary	Arbitrary	r/n for any $r \in [0,1]$	$1-r$
Interval	Intervals	$1/(3n)$	Democratic
Rectangle	Rectangles	$1/(4n)$	Democratic
Convex 2-d	Convex 2-d	$1/(5n)$	Democratic
Rectilinear with T ref.vert.	Rectangles	$1/(4n+T)$	Democratic

Table 5.1: Summary of results for cake redivision: ownership and proportionality guarantees.

Theorem 5.4. *When the cake is a 2-dimensional convex polygon and each piece must be convex, it is possible to find in time $O(n^2 \log n)$ a division simultaneously satisfying democratic-ownership and $1/5$-proportionality.*

Remark. In the interval, rectangle and convex settings, the geometric constraints are mostly harmless without the ownership requirement: when the cake is an interval or a rectangle or a convex object, classic algorithms for proportional cake-cutting, such as Even–Paz (Even and Paz, 1984), can be easily made to return interval/rectangle/convex pieces by ensuring that the cuts are parallel. Similarly, the ownership requirement is easy to satisfy without the geometric constraints, as shown by Theorem 5.1. It is the combination of these two requirements that leads to interesting challenges.

Our next result generalizes Theorem 5.3 to a cake that is a *rectilinear polygon* — a polygon all whose angles are 90° or 270°. Rectilinearity is a common assumption in polygon partition problems (Keil, 2000). The "complexity" of a rectilinear polygon is characterized by the number of its *reflex vertexes* — vertexes with a 270° angle. We denote the cake complexity by T. A rectangle — the simplest rectilinear polygon — has $T = 0$. The cake in Figure 5.2 has $T = 4$ reflex vertexes.

Theorem 5.5. *When the cake is a rectilinear polygon with T reflex vertexes, and each piece must be a rectangle, it is possible to find in time $O(n^2 \log n + poly(T))$ a division satisfying democratic-ownership, in which each agent receives at least $1/(4n+T)$ of the total cake value.*[2]

Price-of-fairness

Redivision protocols can be used not only to compromise between old and new agents, but also to compromise between fairness and efficiency. Often, the most economically-efficient allocation is not fair, while a fair allocation is not economically-efficient. The trade-off between fairness and efficiency is quantified by the *price-of-fairness* (Bertsimas et al., 2011, 2012; Caragiannis et al., 2012; Aumann and Dombb, 2010). It is defined as the worst-case ratio of the maximum attainable social-welfare to the maximum attainable social-welfare of a fair allocation. The social welfare is usually defined as the arithmetic mean of the agents' values (also called *utilitarian welfare*) or their geometric mean (also called *Nash welfare*, see Moulin (2004) and Caragiannis et al. (2016)).

[2] The guarantee of $1/(4n+T)$ is calculated as a fraction of the total cake value. However, with a rectilinear cake and a rectangular piece, even a single agent cannot always get the entire cake value to itself. Therefore, one could think of an alternative guarantee where the benchmark for each agent is the largest value that this agent can attain in a rectangle. For example, we could guarantee each agent a fraction $1/(4n)$ of the value of its most valuable rectangle. However, such guarantee might be much worse than the guarantee of Theorem 5.5. The proof in Appendix 5.A implies that the value of the most valuable rectangle might be as small as $1/(T+1)$ of the total cake value. Therefore, the alternative guarantee of $1/(4n)$ this value translates to a guarantee of $1/(O(n \cdot T))$ — much worse than the $1/(O(n+T))$ guaranteed by Theorem 5.5.

Cake	Pieces	Value guarantee	Utilitarian price	Nash price
Arbitrary	Arbitrary	r/n for any $r \in [0,1]$	$1/(1-r)$	1
Interval	Intervals	$1/(3n)$	$O(\sqrt{n})$	8.4
Rectangle	Rectangles	$1/(4n)$	$O(\sqrt{n})$	11.2
Convex 2-d	Convex 2-d	$1/(5n)$	$O(\sqrt{n})$	14

Table 5.2: Summary of price-of-fairness upper bounds, Note that the price is a ratio. This means that a price of 1 means "no price". Indeed, the Nash price of proportionality is 1, since the Nash-optimal division is always envy-free, hence also proportional.

A redivision protocol can be used to calculate an upper bound on the price of fairness in the following way. Take a welfare-maximizing allocation as the initial allocation; use a redivision protocol to produce a partially-proportional allocation in which the utility of each agent is close to its initial utility; conclude that the new welfare is close to the initial (maximal) welfare.

Without geometric constraints, we have the following upper bound:

Theorem 5.6. *For every constant $r \in [0,1]$, the utilitarian-price of r-proportionality is at most $1/(1-r)$.*

Note that when $r = 1$, the bound is infinity. Indeed, Caragiannis et al. (2012) proved that the price of 1-proportionality in this setting is $\Theta(\sqrt{n})$, which is not bounded by any constant. Our results show that by making a small compromise on the level of proportionality we can get a constant (independent of n) bound on the utilitarian-price. The parameter r sets the level of trade-off between fairness and efficiency.

With geometric constraints, we have the following upper bounds:

Theorem 5.7. *When the cake is an interval and each piece must be an interval, for every $B \geq 3$:*

- *The utilitarian-price of $(1/B)$-proportionality is $O(\sqrt{n})$;*
- *The Nash-price of $(1/B)$-proportionality is at most 8.4.*

Theorem 5.8. *When the cake is a rectangle and each piece must be a rectangle, for every $B \geq 4$:*

- *The utilitarian-price of $(1/B)$-proportionality is $O(\sqrt{n})$;*
- *The Nash-price of $(1/B)$-proportionality is at most 11.2.*

Theorem 5.9. *When the cake is convex polygon and each piece must be convex, $\forall B \geq 5$:*

- *The utilitarian-price of $(1/B)$-proportionality is $O(\sqrt{n})$;*
- *The Nash-price of $(1/B)$-proportionality is at most 14.*

Note that the first claim in Theorem 5.7 is subsumed by Aumann and Dombb (2010), who prove that the utilitarian-price of 1-proportionality in this setting is $\Theta(\sqrt{n})$. We bring this claim only for completeness. The second claim in this theorem, as well as the following theorems which deal with two-dimensional constraints, are not implied by previous results.

5.1.2 Related Work

Dynamic fair division

Our cake redivision problem differs from several division problems studied recently.

1. Dynamic resource allocation (Kash et al., 2013; Friedman et al., 2015) is a common problem in cloud-computing environments. The server has several resources, such as memory and disk-space. Agents (processes) come and depart. The server has to allocate the resources fairly among agents. When new agents come, the server may have to take some resources from existing agents. The goal is to do the re-allocation with minimal disruption to existing agents (Friedman et al., 2015). A different but related problem is the *food-bank problem*, where a charity organization receives food donations and must decide on-line to whom each donation should be allocated (Aleksandrov et al., 2015). In these problems, the resources are *homogeneous*, which means that the only thing that matters is what quantity of each resource is given to each agent. In contrast, our cake is *heterogeneous* and different agents may have different valuations on it, so our protocol must decide which parts of the cake should be given to which agent.

2. Population monotonicity (Thomson, 1983; Moulin, 1990a, 2004; Thomson, 2011) is an axiom that describes a desired property of allocation rules. When new agents arrive and the same division rule is re-activated, the value of all old agents should be weakly smaller than before. This axiom represents the virtue of solidarity: if sacrifices have to be made to support an additional agent, then everybody should contribute (Thomson, 1983). We, too, assume that old agents are taking part in supporting the new agents. However, we add the *ownership* requirement, which means that old agents should be allowed to keep at least some of their previous value. In addition, while their approach is axiomatic and mainly interested in existence results, our approach is constructive and our goal is to provide an actual re-division protocol.

Recently, we have started to study monotonicity axioms, such as population-monotonicity and resource-monotonicity, in the context of cake-cutting; see Sziklai and Segal-Halevi (2015).

3. Private endowment in economics resource allocation problems means that each agent is endowed with an initial bundle of resources. Then, agents exchange resources using a market mechanism. The classic problem in economics involves homogeneous resources, but it has also been studied in the cake-cutting framework (Berliant and Dunz, 2004; Aziz and Ye, 2014). A basic requirement in these works is *individual rationality*, which means that the final value allocated to each agent must be weakly *larger* than the value of the initial endowment (note the contrast with the population monotonicity axiom). In our problem we do not make this assumption as it is incompatible with fairness: since some agents may initially own no land, individual rationality would mean that they might not receive anything in the exchange.

4. Online cake-cutting (Walsh, 2011) is characteristic of a birthday party in an office, in which some agents come or leave early while others come or leave late. It is required to give some cake to agents who come early while keeping a fair share to those who come late. In contrast to our model, there it is impossible to re-divide allocated pieces, since they are eaten by their receivers. The fairness guarantees are inevitably weaker.

5. Land reform is the re-division of land among citizens. It has been attempted in numerous countries around the globe and in many periods throughout history. Some books on land reform are Powelson (1988); Bernstein (2002); Rosset et al. (2006); Lipton (2009). The earliest recorded land-reform was done in ancient Egypt in the times of King Bakenranef, 8th century BC. The most recent land-reform act has been legislated in Scotland in 2016 AD. The balance between fairness and ownership rights is a major concern in such reforms (Sellar, 2006; Hoffman, 2013; Wightman, 2015; MacInnes and Shields, 2015).

6. The cake-cutting procedures of Fink and Austin (Brams and Taylor, 1996, pages 40-44) handle a situation in which all agents — old and new alike — are entitled to a proportional share; however, the agents come sequentially. Initially two agents come and divide the land using cut-and-choose such that each agent has a value of at least $1/2$; then, the third agent comes and he should be given a part of each existing piece such that each of the three agents (old and new alike) will have a value of at least $1/3$; and so on. This is different than our setting since the only requirement is proportionality — there is no "ownership" requirement.

Price of fairness

The price-of-fairness has been studied in various contexts, such as routing and load-balancing (Bertsimas et al., 2011, 2012) and kidney exchange (Dickerson et al., 2014). The price-of-fairness in cake-cutting has been studied in two settings:

- The cake is a *one-dimensional interval* and the pieces must be *intervals* (Aumann and Dombb, 2010). The utilitarian-price-of-proportionality in this case is $\Theta(\sqrt{n})$.
- The cake is *arbitrary* and the pieces may be *arbitrary* (Caragiannis et al., 2012). The utilitarian-price-of-proportionality in this case is $\Theta(\sqrt{n})$ too.

Both papers study the price of other fairness criteria such as envy-freeness and equitability, but do not study the price in Nash-welfare. Additionally, they do not handle two-dimensional geometric constraints such as rectangularity or convexity.

Several authors study the algorithmic problem of finding a welfare-maximizing cake-allocation allocation in various settings:

1. The cake is an interval and the pieces must be connected (Aumann et al., 2013);

2. The cake is an interval and the pieces must be connected, and additionally, the division must be proportional (Bei et al., 2012);

3. The cake and pieces are arbitrary, and the division must be envy-free (Cohler et al., 2011).

4. The cake and pieces are arbitrary, and the division must be equitable (Brams et al., 2012).

5.2 Model

We briefly recall some terminology from Chapter 2 (see there for formal definitions).

- C is the cake to be divided. In this chapter it will be an interval or a polygon in $\mathbb{R}^2$.
- S is the family of pieces that are considered *usable*. In this chapter it will be the family of intervals, rectangles or convex objects. An S-allocation is an allocation in which all pieces are elements of S.
- For each agent $i \in \{1, \dots, n\}$, $V_i(X_i)$ is agent i's value-measure of the piece X_i.

In this chapter, for every constant $r \in (0,1)$, an allocation X is called *r-proportional* if every agent receives at least r/n of the total cake value:

$$\forall i \in \{1, \dots, n\} : V_i(X_i) \geq \frac{r}{n} \cdot V_i(C)$$

(note that this definition is slightly different than in the previous two chapters). A 1-proportional division is usually called in the literature "proportional".

5.2.1 Cake redivision

There is an existing S-allocation of the cake: $Z_1, \dots, Z_n$. It is assumed that the old pieces Z_j are pairwise-disjoint and $\forall j : Z_j \in S$, but nothing else is assumed on the division. In particular, the initial division is not necessarily proportional, and some of C may be undivided.

It is required to create a new S-allocation of C to all agents: $X_1, \dots, X_n$. For every constant $w \in (0,1)$, the re-allocation satisfies the *w-ownership* property if every agent receives at least a fraction w of its old value:

$$\forall j \in \{1, \dots, n\} : \qquad V_j(X_j) \geq w \cdot V_j(Z_j)$$

Since w-ownership is not always compatible with r-proportionality for any $r > 0$, we define the following weaker property. A re-allocation satisfies the *democratic-ownership* property if, for every $k \in \{1, \dots, n\}$, there are at least $n - k$ indexes $j \in \{1, \dots, n\}$ for which:

$$V_j(X_j) \geq \frac{1}{\lceil n/k \rceil} \cdot V_j(Z_j).$$

5.2.2 Social-welfare and Price-of-fairness

In addition to fairness, it is often required that a division has a high *social welfare*. The social welfare of an allocation is a certain aggregate function of the normalized values of the agents (the normalized value is the piece value divided by the total cake value). Common social welfare functions are sum (utilitarian) and product (Nash); see Moulin (2004). We normalize them such that the maximum welfare is 1:

- *Utilitarian welfare* — the arithmetic mean of the agents' normalized values:

$$W_{util}(X) = \frac{1}{n} \sum_{i \in \{1,\dots,n\}} \frac{V_i(X_i)}{V_i(C)}$$

- *Nash welfare* — the geometric mean of the agents' normalized values:

$$W_{Nash}(X) = \left(\prod_{i \in \{1,\dots,n\}} \frac{V_i(X_i)}{V_i(C)} \right)^{1/n}$$

The goal of maximizing the social welfare is not always compatible with the goal of guaranteeing a fair share to every agent. For example, Caragiannis et al. (2012) describe a simple example in which the maximum utilitarian welfare of a proportional allocation is $O(1/n)$ while the maximum utilitarian welfare of an arbitrary (unfair) allocation is $O(1/\sqrt{n})$. This means that society has to pay a price, in terms of social-welfare, for insisting on fairness. This is called the *price of fairness*. Formally, given a social welfare function W and a fairness criterion F, the price-of-fairness relative to W and F (also called: "the W-price-of-F") is the ratio:

$$\frac{\sup_X W(X)}{\sup_{Y \in F} W(Y)} \tag{*}$$

where the supremum at the nominator is over all allocations X and the supremum at the denominator is over all allocations Y that also satisfy the fairness criterion F. The cited example shows that the utilitarian-price-of-proportionality might be $\Omega(\sqrt{n})$.

When there are geometric constraints, they affect both the numerator and the denominator of (*), i.e, the suprema are taken only on S-allocations. Therefore, it is not a-priori clear whether the price-of-fairness with constraints is higher or lower than without constraints.

5.3 Arbitrary Cake and Arbitrary Pieces

This section proves Theorem 5.1, which assumes no geometric constraints on the cake or pieces. The main lemma is:

Lemma 5.3.1. *Given cake-allocations Z and Y and a constant $r \in [0,1]$, there exists an allocation X such that, for every agent i:*

$$V_i(X_i) \geq rV_i(Y_i) + (1-r)V_i(Z_i)$$

Moreover, when r is a constant rational number, X can be found using $O(n^2)$ mark/eval queries.

Proof. We first give an existential proof. Consider the set of all possible cake-partitions. For each cake-partition, consider the $n \times 1$ vector of utilities of the agents. The Dubins–Spanier theorem (Dubins and Spanier, 1961) says that the set of all such vectors is convex. Therefore, there exists an allocation X satisfying the requirement as an equality: $\forall i : V_i(X_i) = rV_i(Y_i) + (1-r)V_i(Z_i)$.

Since the Dubins–Spanier theorem (Dubins and Spanier, 1961) is not constructive, we give here a constructive protocol for creating the allocation Z when r is a rational number, $r = p/q$ with $p < q$ some positive integers. For every pair of agents i, j (including when $i = j$), the protocol works as follows:

Step 1. Agent i divides the piece $Z_i \cap Y_j$ to q pieces that are equal in its eyes.
Step 2. Agent j takes the p pieces that are best in its eyes.
Step 3. Agent i takes the remaining $q - p$ pieces.
(Note that when $i = j$, agent i receives the entire piece $Z_i \cap Y_i$).

Each agent i is allocated a piece X_i which is a union of nq pieces: np pieces that agent i took from other agents (including itself) in piece Y_i and $n(q-p)$ pieces that were left for agent i from other agents in piece Z_i.

From every piece $Y_i \cap Z_j$ (for $j \in \{1, \ldots, n\}$), agent i picks the best p out of q pieces, which give it a value of at least $\frac{p}{q} V_i(Y_i \cap Z_j)$. Its total value of these np pieces is thus at least $rV_i(Y_i)$.

In addition, from every piece $Z_i \cap Y_j$ (for $j \in \{1, \ldots, n\}$), agent i receives $q - p$ out of q equal pieces, which give it a value of exactly $\frac{q-p}{q} V_i(Z_i \cap Y_j)$. Its total value of these $n(q-p)$ pieces is thus exactly $(1-r)V_i(Z_i)$. □

Proof of Theorem 5.1. We are given a pair r, w where $r + w \leq 1$. Apply Lemma 5.3.1, with:

Y — any proportional allocation, which can be found by classic protocols such as Steinhaus (1948); Even and Paz (1984).

Z — the initial allocation.

The new division satisfies r-proportionality and $(1-r)$-ownership. By assumption $1 - r \geq w$. □

Note that the redivision protocol gives to each agent a piece that is not only worth at least $(1-r)V_i(Z_i)$, but is also a subset of Z_i (in addition to a subset of Y_i). This may be desirable in some cases, e.g. in land division, the old landlords may care not only about their value but also about getting a subset of their old plot.

Remark. The $O(n^2)$ complexity assumes the integers p, q are constant (not part of the input). If they are considered part of the input, then the complexity becomes linear in q which is exponential in the number of input bits. The number of queries can be reduced using concepts from number theory, but this is beyond the scope of this paper. See McAvaney et al. (1992); Robertson and Webb (1998) for details.

Finally we show that the balance guaranteed by Theorem 5.1 is asymptotically tight.

Proof of Proposition 5.1. We are given a pair r, w where $r + w > 1$. Consider the following scenario. In the initial allocation, a single agent owns the entire cake. All n agents have the same value-density and they value the entire cake as 1. In any r-proportional division, the $n - 1$ landless citizens must receive a total value of $(n-1)r/n = r - r/n$. Therefore the old landlord receives at most $1 - r + r/n$. By assumption, $1 - r < w$. Therefore, if n is sufficiently large, the old landlord receives less than w of his previous value, so the division does not satisfy w-ownership. □

5.4 Interval Cake and Interval Pieces

In this section, the cake is an interval in $\mathbb{R}$. Each piece in the initial division is an interval in C and each piece in the new division must be an interval in C. We begin by proving the impossibility result (Proposition 5.2), using a lemma opposite to Lemma 5.3.1.

Lemma 5.4.1. *Let Z be a connected allocation, $r \in (0,1)$ a positive constant and $k \leq n$ an integer. Then there exist valuations such that, in every connected r-proportional allocation X, for every agent $j \in \{1, \ldots, k\}$: $V_j(X_j) \leq V_j(Z_i)/\lfloor \frac{n}{k} \rfloor$.*

Proof. Assume that the valuations are as follows. Each agent $j \in \{1, \ldots, k\}$ values the piece Z_j as $\lfloor \frac{n}{k} \rfloor$ and the rest of the cake as 0. The value-density of j in Z_j is piecewise-uniform: It has $\lfloor \frac{n}{k} \rfloor$ regions with a value of 1 and $\lfloor \frac{n}{k} \rfloor - 1$ "gaps" — regions with a value of 0. The other $n - k$ agents are divided to k groups of roughly equal size: the size of each group is either $\lfloor \frac{n-k}{k} \rfloor = \lfloor \frac{n}{k} \rfloor - 1$ or $\lceil \frac{n-k}{k} \rceil = \lceil \frac{n}{k} \rceil - 1$. Each agent in group j assigns a positive value only to a unique gap in the piece Z_j (so when the group size is $\lfloor \frac{n}{k} \rfloor - 1$, each gap is wanted by exactly one agent; otherwise, there is one gap wanted by two agents). The following figure illustrates the value-densities that are positive in piece Z_1. The solid boxes represent the value-density of agent #1; each dotted box represents the value-density of a single agent from group #1.

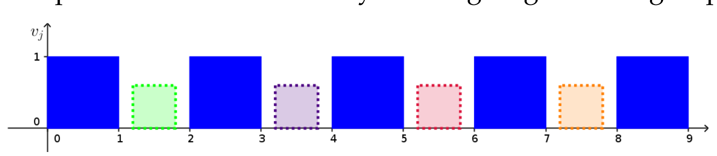

In any positive division, each gap in Z_j must be at least partially allocated to an agent in group j. Hence, the interval allocated to agent j must contain at most a single positive region in Z_j — it is not allowed to overlap any gap. Therefore the value of agent j is at most $V_j(Z_j)/\lfloor \frac{n}{k} \rfloor$. □

Proof of Proposition 5.2. Apply Lemma 5.4.1 with Z = the initial allocation. □

To prove the matching positive result (Theorem 5.2), we introduce a protocol for fair division of an "archipelago" — a cake made of one or more interval "islands".

Lemma 5.4.2. *Let C be a cake made of $m \geq 1$ pairwise-disjoint intervals: $C = Z_1 \cup \cdots \cup Z_m$. There exists a division X of C among n agents, in which (a) Each agent i receives an interval entirely contained in one of the islands: $\forall i : \exists j : X_i \subseteq Z_j$, and (b) Each agent receives a value of at least $V_i(C)/(n+m-1)$. Moreover, X can be found using $O(mn \log n)$ mark/eval queries.*

Proof. We normalize the value measures of all agents such that the total value of C is $n+m-1$. The following recursive protocol allocates each agent an interval with a value of at least 1.

Base: $m = 1$. The cake is a single interval and its total value is n. Use the Even–Paz protocol (Even and Paz, 1984) to allocate each agent an interval with a value of at least 1.
Step: $m > 1$.

1. Ask each agent $i \in \{1, \ldots, n\}$ to evaluate the island Z_m.
2. Order the agents in descending order of their evaluation: $V_1(Z_m) \geq \cdots \geq V_n(Z_m)$.
3. Let q be the largest integer such that $V_q(Z_m) \geq q$ (or 0 is already $V_1(Z_m) < 1$).
4. If $q = 0$, discard the island Z_m. Otherwise ($q \geq 1$),
   divide Z_m proportionally among the agents $\{1, \ldots, q\}$ using Even–Paz protocol (Even and Paz, 1984).
5. Divide the remaining $m-1$ islands recursively among the remaining $n-q$ agents.

The descending order of the agents guarantees that: $V_1(Z_m) \geq \cdots \geq V_q(Z_m) \geq q$. So in step #4, the interval Z_m is divided proportionally among q agents that value it as at least q, and each of these agents receives an interval with a value of at least 1.

By definition of q, $V_{q+1}(Z_m) < q+1$ (this is true even when $q = 0$). By the descending order of the agents, the same is true for all remaining agents $\{q+1, \ldots, n\}$. Therefore, all remaining agents value the remaining cake as more than $(m+n-1)-(q+1) = (n-q)+(m-1)-1$. Since there are $n-q$ agents and $m-1$ islands, the recursive algorithm gives each agent an interval with value at least 1. The Even–Paz protocol requires $O(n \log n)$ queries, and it is done at most m times, so the total number of queries is $O(mn \log n)$. □

Remark. The fraction of $1/(n+m-1)$, guaranteed by Lemma 5.4.2, is the largest that can be guaranteed. To see this, assume that all agents $i \in \{1, \ldots, n\}$ have the same value-measures — they value the islands $Z_1, \ldots, Z_{m-1}$ as 1 and the island Z_m as n (so their total cake value is $n+m-1$). The piece of every agent must be entirely contained in a single island. If any agent receives a piece in islands $Z_1, \ldots, Z_{m-1}$, then that agent receives a value of at most 1. Otherwise, if all n agents receive a piece in Z_m, then the value of at least one agent is at most 1. In both cases, at least one agent receives a fraction of at most $1/(n+m-1)$ the cake value. □

Proof of Theorem 5.2. Our protocol for re-division of an interval has three steps.

Step 1. Given the original partial allocation $Z_1 \cup \cdots Z_n \subseteq C$, extend it to a complete allocation $Z'_1 \cup \cdots Z'_n = C$, by attaching each "blank" (unallocated interval in C) arbitrarily to one of the two adjacent allocated intervals. This, of course, does not harm the old values: $\forall j \in \{1, \ldots, n\} : V_j(Z'_j) \geq V_j(Z_j)$.

Step 2. For each agent $j \in \{1, \ldots, n\}$, add a “helper agent” j^* and assign it a value-density function v^*_j:

$$
\begin{aligned}
v^*_j(x) &= v_j(x) && \text{if } x \in Z'_j \\
v^*_j(x) &= 0 && \text{if } x \notin Z'_j
\end{aligned}
$$

Use the protocol of Lemma 5.4.2 with $n+n$ agents, regarding the cake C as an archipelago and the pieces $Z'_1, \ldots, Z'_n$ as the islands.

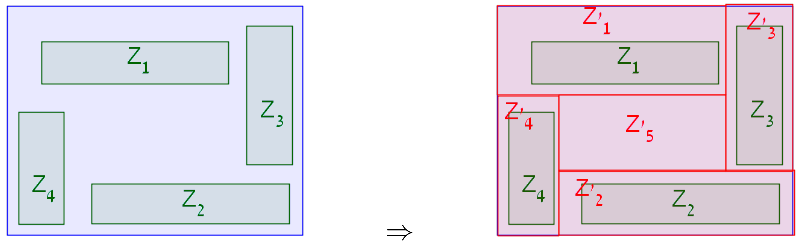

Figure 5.3: The allocation-completion step: input and output.

> **Step 3.** Give each agent $j \in \{1, \ldots, n\}$ either the interval allocated to its normal agent j or the interval allocated to its helper agent j^*, whichever is more valuable for it.

We now prove that the resulting allocation is $1/3$-proportional and satisfies the democratic-ownership property.

(a) *Proof of 1/3-proportionality.* We apply Lemma 5.4.2 with $2n$ agents and $m = n$ islands. Each of the $2n$ agents receives an interval contained in one of the pieces $Z'_1, \ldots, Z'_n$, with a value of at least $1/((2n) + n - 1)$ its total cake value. This value is larger than $1/(3n)$.

(b) *Proof of democratic-ownership.* We focus on the n helper agents. First, by Lemma 5.4.2, every helper agent j^* must receive an interval contained in Z'_j, since its value is positive only in the island Z_j. Moreover, by the pigeonhole principle, for every integer $k \leq n$, at most k islands are populated by at least $\lceil \frac{n}{k} \rceil$ normal agents. Hence, at least $n - k$ islands are populated by at most $\lceil \frac{n}{k} \rceil - 1$ normal agents. Adding the helper agent, these islands are populated by at most $\lceil \frac{n}{k} \rceil$ agents. Hence, the proportional allocation of step #4 in the protocol of Lemma 5.4.2 gives these helper agents an interval subset of Z'_j, which is worth for agent j at least $V_j(Z'_j)/\lceil \frac{n}{k} \rceil$. □

5.5 Polygonal Cake and Polygonal Pieces

Rectangle cake and pieces. Initially, we assume that the cake is a rectangle in $\mathbb{R}^2$. Each piece in the initial division is a rectangle parallel to C and each piece in the new division must be a rectangle parallel to C.

We would like to use the re-division protocol of Theorem 5.2. Steps #2 and #3 are easily adapted: the Even–Paz protocol (Even and Paz, 1984) can operate on a rectangular cake, requiring the agents to make cuts parallel to the cake sides. This guarantees that the pieces are rectangles.

However, there is one obstacle. Step #1, the allocation-completion step, is no longer trivial. We cannot just attach each unallocated part of C to an allocated rectangle, since the result will not necessarily be a rectangle. We still need to extend the initial partial allocation $Z_1 \cup \cdots Z_n \subseteq C$ to a complete allocation, but the number of rectangles in the complete allocation might be larger than n, since we might have unattached blanks.

Our goal, then, is to find a partition of C to rectangles, $Z'_1 \cup \cdots Z'_{n+b} = C$, with $b \geq 0$, such that every input rectangle is contained in a unique output rectangle: $\forall j \in \{1, \ldots, n\} : Z_j \subseteq Z'_j$. The additional b rectangles are called *blanks*. In Step 3, we will have $m = n + b$ islands and $2n$ agents, so the value guarantee per agent will be $1/((2n) + (n + b) - 1) = 1/(3n + b - 1)$; therefore, we would like the number of blanks b to be as small as possible.

An example of the input and output of the allocation-completion step is shown in Figure 5.3. Here, $b = 1$ since there is one blank — Z'_5. In this case $b = 1$ is minimal.

Convex cake and pieces. The situation is similar when C is convex and the pieces should be convex. The Even–Paz protocol can operate on a convex cake, requiring the agents to make cuts parallel to the each other. This guarantees that the pieces will be convex. In Step #1, a similar challenge arises. We have an initial partial allocation $Z_1 \cup \cdots Z_n \subseteq C$, where each Z_j is convex. We need a complete allocation $Z'_1 \cup \cdots Z'_{n+b} = C$, where each Z'_j is convex, every input piece is contained in a unique output piece, and the number of blanks b is minimal.

Rectilinear cake and rectangular pieces. There are efficient algorithms for partitioning a rectilinear polygon to a minimal number of rectangles. A rectilinear polygon with T reflex vertexes can be partitioned in time $O(poly(T))$ to at most $T+1$ rectangles (Keil, 2000; Eppstein, 2010), and this bound is tight when the vertexes of C are in general position. Our goal is to bound b — the number of blank *rectangles*. Therefore, it is expected that the bound should depend on T, in addition to m.

The allocation-completion step for all two-dimensional settings is handled by the Akopyan and Segal-Halevi (2016), who prove the following lemma:

Lemma 5.5.1 (**(?)**). *There is an $O(m)$-time algorithm that extends a partial allocation $Z_1 \cup \cdots Z_m \subseteq C$ to a complete allocation $Z'_1 \cup \cdots Z'_{m+b} = C$, such that there are:*
(a) at most $m - 2\sqrt{m} - O(1)$ rectangular blanks when the cake & pieces are parallel rectangles:
(b) at most $2m-5$ convex blanks when the cake and pieces are convex polygons;
(c) at most $m + T - 2\sqrt{m} - O(1)$ rectangular blanks when the cake is rectilinear with T reflex vertexes and the pieces are rectangles. In this case the run-time is $O(m + poly(T))$.
The numbers of blanks in all cases are tight.

Proof of Theorems 5.3, 5.4 ,5.5 . Use the protocol of Theorem 5.2, plugging into Step #1 the algorithm of Lemma 5.5.1 with $m = n$. The value per agent is at least $1/(3n+b-1)$, which is:
(a) at least $1/(4n - 2\sqrt{n}) > 1/(4n)$ in the rectangle case — satisfying 1/4-proportionality;
(b) at least $1/(5n-6) \geq 1/(5n)$ in the convex case — satisfying 1/5-proportionality;
(c) at least $1/(4n + T - 2\sqrt{n}) > 1/(4n+T)$ in the rectilinear case. □

5.6 Price-of-Fairness Bounds

In this section, our redivision protocols are used to prove upper bounds on the price of partial-proportionality.

Theorem 5.6 follows directly from Theorem 5.1 by taking the original division to be a utilitarian-optimal division.

The proofs of Theorems 5.7, 5.8 and 5.9 are similar; only the constants are different. We present below only the proof of Theorem 5.8; to get the proofs of the other theorems, replace the constant "4" with "3" or "5" respectively.

The first part of Theorem 5.8 — regarding the **utilitarian** price — is proved by the following:

Lemma 5.6.1. *Let Z be a utilitarian-optimal rectangular division of a cake C among n agents who value the entire cake C as 1. Let U be the utilitarian welfare of Z:*

$$U := \frac{1}{n} \sum_{j=1}^{n} V_j(Z_j)$$

Then, there exists a $(1/4)$-proportional rectangular allocation of C to these same n agents with utilitarian welfare W, such that $U/W \in O(n^{1/2})$.

Proof. Apply the redivision protocol of Section 5.5 to the existing division by setting $m = n$ and treating all n agents as "old". The partial-proportionality guarantee of that protocol ensures that the new division is 1/4-proportional. The partial-ownership of that protocol ensures that for every integer $k \in \{0, \ldots, n\}$, there is a set S_k containing at least $n-k$ agents whose value is more than $\max(\frac{kV_j(Z_j)}{2n}, \frac{1}{4n})$. Renumber the agents in the following way. Pick an agent from S_{n-1} (which contains at least one agent) and number it $n-1$. Pick an agent from S_{n-2} (which contains at least one other agent) and number it $n-2$. Continue this way to number the agents by $k = n-1, \ldots, 0$. Now, the utilitarian welfare of the new division is lower-bounded by:

$$W > \frac{1}{n} \sum_{k=0}^{n-1} \max(\frac{kV_k(Z_k)}{2n}, \frac{1}{4n}) \geq \frac{1}{n} \cdot \frac{1}{4n} \cdot \sum_{k=0}^{n-1} \max(kV_k(Z_k), 1)$$

and the utilitarian welfare ratio is at most:

$$\frac{U}{W} < 4n \cdot \frac{\sum_{k=0}^{n-1} V_k(Z_k)}{\sum_{k=0}^{n-1} \max(k \cdot V_k(Z_k), 1)}$$

Denote the ratio in the right-hand side by $\frac{\textbf{NUM}}{\textbf{DEN}}$. Let $a_k = V_k(Z_k)$, so that $\textbf{NUM} = \sum_{k=0}^{n-1} a_k$ and $\textbf{DEN} = \sum_{k=0}^{n-1} \max(k \cdot a_k, 1)$. To get an upper bound on U/W, we find a sequence $a_0, \dots, a_{n-1}$ that maximizes $\frac{\textbf{NUM}}{\textbf{DEN}}$ subject to $\forall k : 0 \le a_k \le 1$.

Observation 1. in a maximizing sequence, $a_0 = 1$ and there is no $k > 0$ such that $a_k < 1/k$. *Proof*: Setting such a_k to $1/k$ increases **NUM** and does not change **DEN**.

Observation 2. A maximizing sequence must be weakly-decreasing (for all $k < k'$, $a_{k'} \ge a_k$). *Proof*: if there exists $k < k'$ such that $a_k < a_{k'}$, then we can swap a_k with $a_{k'}$. This does not change **NUM** but strictly decreases **DEN**.

Observation 3. In a maximizing sequence, there is no $k > 0$ such that $1/k < a_k < 1$. *Proof:*[3] If $1/k < a_k < 1$ then for some sufficiently small $\epsilon > 0$, both $a_k + \epsilon$ and $a_k - \epsilon$ are in $(1/k, 1)$ and replacing a_k with $a_k \pm \epsilon$ makes the ratio strictly smaller than the maximum. Replacing a_k with $a_k + \epsilon$ makes the ratio $\frac{\textbf{NUM}+\epsilon}{\textbf{DEN}+k\epsilon}$; this new ratio is smaller than $\frac{\textbf{NUM}}{\textbf{DEN}}$ so $\epsilon \cdot \textbf{DEN} < k\epsilon \cdot \textbf{NUM}$. Replacing a_k with $a_k - \epsilon$ makes the ratio $\frac{\textbf{NUM}-\epsilon}{\textbf{DEN}-k\epsilon}$; that new ratio is smaller than $\frac{\textbf{NUM}}{\textbf{DEN}}$ so $-\epsilon \cdot \textbf{DEN} < -k\epsilon \cdot \textbf{NUM}$. But the two latter inequalities $\epsilon \cdot \textbf{DEN} < k\epsilon \cdot \textbf{NUM}$ and $-\epsilon \cdot \textbf{DEN} < -k\epsilon \cdot \textbf{NUM}$ are contradictory. Hence, the assumption $1/k < a_k < 1$ is false.

Observations 1-3 imply that a maximizing sequence has a very specific format. It is characterized by an integer $l \in \{0, \dots, n-1\}$ such that, for all $k \le l$, $a_k = 1$ and for all $k \ge l+1$, $a_k = 1/k$. So:

$$\frac{\textbf{NUM}}{\textbf{DEN}} = \frac{\sum_{k=0}^{n-1} a_k}{\sum_{k=0}^{n-1} \max(k \cdot a_k, 1)} = \frac{(l+1) + (H_{n-1} - H_l)}{\frac{1}{2}l(l+1) + (n-l-1)} < \frac{2(l + H_n + 1)}{l^2 - l + 2(n-1)}$$

where $H_n = \sum_{k=1}^{n}(1/k)$ is the n-th harmonic number.

The number l is integer, but the expression is bounded by the maximum attained when l is allowed to be real. By standard calculus we get that the real value of l which maximizes the above expression is $l = \sqrt{2(n-1) + (H_n+1)(H_n+2)} - (H_n+1) = \Theta(\sqrt{n})$. Substituting into the above inequality gives:

$$\frac{\textbf{NUM}}{\textbf{DEN}} \le \frac{\Theta(n^{1/2})}{\Theta(n)} = \Theta(n^{-1/2}) \qquad \Longrightarrow \qquad \frac{U}{W} < 4n \cdot \frac{\textbf{NUM}}{\textbf{DEN}} = O(n^{1/2})$$

as claimed. ☐

The second part of Theorem 5.8 — regarding the **Nash** price — is proved by the following:

Lemma 5.6.2. *Let Z be a Nash-optimal rectangular division of a cake C among n agents who value the entire cake C as 1. Let U be the Nash welfare of Z (the geometric mean of the values):*

$$U^n = \prod_{j=1}^{n} V_j(Z_j)$$

Then, there exists a $(1/4)$-proportional rectangular allocation of C to these same n agents with Nash welfare W, and $U/W < 11.2$.

Proof. Apply the redivision protocol of Section 5.5 to redivide the existing n pieces among the n agents. Renumber the agents as in Lemma 5.6.1. The Nash welfare of the new division, raised to the n-th power, can be bounded as:

$$W^n > \prod_{k=0}^{n-1} \max\left(\frac{k \cdot V_k(Z_k)}{2n}, \frac{1}{4n}\right) \ge \left(\frac{1}{4n}\right)^n \prod_{k=0}^{n-1} \max(k \cdot V_k(Z_k), 1)$$

and the ratio of the new welfare to the previous welfare can be bounded as:

$$\frac{U^n}{W^n} < (4n)^n \cdot \frac{\prod_{k=0}^{n-1} V_k(Z_k)}{\prod_{k=0}^{n-1} \max(kV_k(Z_k), 1)} = \frac{(4n)^n}{\prod_{k=0}^{n-1} \max(k, 1/V_k(Z_k))}$$

The nominator does not depend on the valuations, so the ratio is maximized when the denominator is minimized. This happens when each factor in the product is minimized. The minimal value of the 0-th

[3]We are grateful to Varun Dubey for suggesting this proof in: http://math.stackexchange.com/q/1609071/29780

factor is 1 and the minimal value of the other factors is k. Hence:

$$\frac{U^n}{W^n} < \frac{(4n)^n}{\prod_{k=1}^{n-1} k} = \frac{(4n)^n}{(n-1)!} = \frac{n(4n)^n}{n!} \approx \frac{n(4n)^n}{\sqrt{2\pi n}(n/e)^n} = \sqrt{\frac{n}{2\pi}} \cdot (4e)^n$$

where e is the base of the natural logarithm. Taking the n-th root gives $U/W < (4e) \cdot \sqrt{n/2\pi}^{1/n}$. A calculation in Wolfram Alpha shows that the rightmost term $\sqrt{n/2\pi}^{1/n}$ is bounded globally by 1.03, so all in all $U/W < 1.03 \cdot 4 \cdot e < 11.2$ as claimed. □

5.7 Conclusions and Future Work

Two-dimensional division , the price-of-fairness and the re-division problem are relatively new topics, and there is a lot of room for future research in each of them.

5.7.1 Handling other geometric constraints

Two steps in our redivision algorithm are sensitive to the geometric constraint: the *allocation-completion algorithm* (Step #1 in Theorem 5.2), and the *Even–Paz protocol* (Step #4 in Lemma 5.4.2). We describe how these steps are affected by several alternative constraints.

1. Convexity in three or more dimensions. The Even–Paz protocol can easily operate on a multi-dimensional convex object, requiring the agents to cut using hyper-planes parallel to each other. However, we currently do not have an allocation-completion algorithm for convex objects (or even for boxes) in three or more dimensions.

2. Path-connectivity in two dimensions. If the pieces have to be path-connected but not necessarily convex, then the allocation-completion step is much easier and no blanks are created (Akopyan and Segal-Halevi, 2016). However, it is not clear how to use the Even–Paz protocol in this case: when the cake is connected but not convex, making parallel cuts might create disconnected pieces.

3. Fatness. A fat object is an object with a bounded length/width ratio, such as a square. Fatness makes sense in land division: if you are entitled to a 900 square meters of land, you will probably prefer to get them as a 30×30 square or a 45×20 rectangle rather than 9000×0.1 sliver. A division problem with fatness requirement cannot be reduced to one-dimensional division. There exist specialized division protocols that support fatness constraints Segal-Halevi et al. (2015a); **?** and they can be used instead of the Even–Paz protocol. However, we do not have an allocation-completion algorithm with fatness constraints.

4. Two pieces per agent. Theorem 5.1 allows an unlimited number of pieces per agent, while the other theorems allow only a single piece per agent. We do not know what happens between these two extremes. For example, if the cake is a one-dimensional interval and each agent can get at most two intervals, what ownership-proportionality combinations are attainable?

5.7.2 Handling other fairness requirements

1. Envy-freeness. In this paper we took proportionality as a benchmark of fairness. An alternative benchmark is *envy-freeness*. Envy-freeness means that each agent values its piece at least as much as each of the other pieces. Similarly, r-envy-freeness means that each agent values its piece as at least r times the value of each of the other pieces. For what pairs r, w is r-envy-freeness compatible with w-ownership? With democratic-ownership?

2. Pareto-efficiency. From an existential point of view, Pareto-efficiency does not add much difficulty. Both r-proportionality and w-ownership are preserved by Pareto-improvements. therefore, if there exists a division satisfying r-proportionality and w-ownership (or democratic-ownership), then there also exists a Pareto-optimal division satisfying these properties. However, the algorithmic task of finding such a division is not yet solved.

Note that Pareto-efficiency is "automatically" satisfied when the division is connected, envy-free, and all value-densities are strictly positive (Brams and Taylor, 1996, page 150).

5.7.3 Improving the constants

Our redivision protocol is 1/3 or 1/4 or 1/5-proportional (depending on the geometric constraint). We see two potential ways to improve these numbers.

1. In Step #2 of our redivision protocol, we add n helper agents, so the total number of agents is $2n$. However, in the Step #3, each agent chooses either its helper or its normal agent, while the other agent is "wasted". If we could know the n choices of the agents in advance, we could employ only n agents overall and this would subtract 1 from the constant (the constants would become 1/2 or 1/3 or 1/4). One way to analyze this scenario is to define a *strategic game* in which each agent has two possible strategies: "normal" vs. "helper". A pure-strategy Nash equilibrium in this game corresponds to an allocation satisfying the partial-proportionality and the democratic-ownership requirements. We conjecture that a pure-strategy Nash equilibrium indeed exists. While finding a Nash equilibrium is usually a computationally-hard problem, it may be useful as an existential result.

2. In Lemma 5.4.2, we treat each existing piece Z_j as an "island" and insist that each new piece be entirely contained in an existing piece, i.e, we do not cross the existing division lines. This may be desirable in the context of land division, since it respects the *Uti Possidetis* principle (Lalonde, 2002). However, it implies that the resulting division can only be partially-proportional and never fully proportional (as shown by the remark following Lemma 5.4.2). It may be possible to improve the proportionality guarantees by devising a different redivision procedure that does cross the existing division lines. This may require some new geometric techniques.

These possibilities invoke the following open question: what is the highest level of proportionality that is compatible with democratic-ownership?

5.7.4 Price-of-fairness

It is not clear whether the upper bounds of our Theorems 5.6-5.9 are tight.

In particular, for the case of interval cake and interval pieces, there is a lower bound of $\Omega(\sqrt{n})$ on the utilitarian price of proportionality. However, we could not generalize it to the price of *partial* proportionality, and it is interesting to know which of the following two options is correct: (a) there is a lower bound of $\Omega(\sqrt{n})$ matching our Theorem 5.7, or (b) the actual price of partial-proportionality is $o(\sqrt{n})$. The latter option would imply that partial-proportionality is asymptotically "cheaper" than full proportionality, in social welfare terms.

Regarding the Nash price-of-fairness, it is known (Sziklai and Segal-Halevi, 2015) that with arbitrary pieces, every Nash-optimal allocation is envy-free (hence also proportional), so the Nash price of envy-freeness (hence, of proportionality) is 1. However, this is not true when the pieces must be connected. We do not have a lower bound for this case.

5.8 Acknowledgments

This paper benefited a lot from discussions and comments by Galya Segal-Halevi, Reuven Cohen, Panagiotis Kanellopoulos, Laszlo Csato, Herve Moulin, Yehuda Levy, Emel Ozturk, the participants in Bar-Ilan University game-theory seminar, Glasgow university micro-economics seminar, Corvinus university game-theory seminar, and anonymous reviewers of the AAMAS 2016, EC 2016 and SODA 2017 conferences.

I am also grateful to Math Stack Exchange members, particularly, Arseiny Akopyan, Chris Culter, Varun Dubey (vnd), Alex Ravsky, Swami Sarvattomananda (Shreesh), David K and JustPassingBy.

All figures were created with GeoGebra 5 (Hohenwarter et al., 2013).

Chapter 5 Appendix

5.A Fair Division of a Rectilinear Polygon

This appendix shows what proportionality guarantees are possible when the cake is a rectilinear polygon, the pieces have to be rectangles (parallel to the sides of the cake), and there are no ownership requirements. It can be seen as a baseline for Theorem 5.5.

Lemma 5.A.1. *Let C be a rectilinear polygon with T reflex vertexes. It is possible to divide C among n agents such that the value of each agent is at least $1/(n+T)$ of the total cake value:*

$$\forall i \in \{1, \ldots, n\} : V_i(X_i) \geq \frac{V_i(C)}{n+T}$$

The fraction $1/(n+T)$ is the largest that can be guaranteed.

Proof. A rectilinear polygon with T reflex vertexes can be partitioned in time $O(poly(T))$ to at most $T+1$ rectangles (Keil, 2000; Eppstein, 2010). Denote these rectangles by Z_j, so that:

$$C = Z_1 \cup \cdots \cup Z_{T+1}$$

Apply the archipelago-division protocol of Lemma 5.4.2 with $m = T+1$. The value-guarantee per agent is at least $1/(n+m-1)$ which is at least $1/(n+T)$, as claimed.

For the upper bound, consider a staircase-shaped cake with $T+1$ stairs. as illustrated below (for $T=4$):

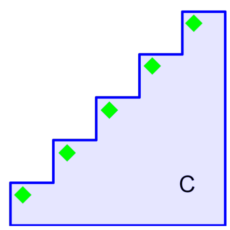

All agents have the same value-measure, which is concentrated in the diamond-shapes: the top diamond is worth n and each of the other diamonds is worth 1 (so for all agents, the total cake value is $n+T$).

Any rectangle in C can touch at most a single diamond. There are two cases:

(a) All n agents touch the top diamond. Then, their total value is n and at least one of them must receive a value of at most 1.

(b) At least one agent touches one of the T bottom diamonds. Then, the value of that agent is at most 1.

In any case, at least one agent receives at most a fraction $1/(n+T)$ of the total cake value, as claimed. □

Chapter 6

Family Ownership

Contents

5.1 Introduction **86**
5.1.1 Results 87
5.1.2 Related Work 90
5.2 Model **92**
5.2.1 Cake redivision 92
5.2.2 Social-welfare and Price-of-fairness 92
5.3 Arbitrary Cake and Arbitrary Pieces **93**
5.4 Interval Cake and Interval Pieces **94**
5.5 Polygonal Cake and Polygonal Pieces **96**
5.6 Price-of-Fairness Bounds **97**
5.7 Conclusions and Future Work **99**
5.7.1 Handling other geometric constraints 99
5.7.2 Handling other fairness requirements 99
5.7.3 Improving the constants 100
5.7.4 Price-of-fairness 100
5.8 Acknowledgments **100**
Chapter appendix **101**
5.A Fair Division of a Rectilinear Polygon **101**

6.1 Introduction

In most fair division problems, the resource is divided among individual agents, and the fairness of a division is assessed based on the valuation of each agent. However, in real life, goods are often owned and used by groups. As an example, consider a land-estate inherited by k families, or a nature reserve that has to be divided among k states. The land should be divided to k pieces, one piece per group. Each group's share is then used by *all* members of the group simultaneously. The land-plot allotted to a family is inhabited by the entire family. The share of the nature-reserve alloted to a state becomes a national park open to all citizens of that state. In economic terms, the alloted piece becomes a "club good" (Buchanan, 1965). The happiness of each group member depends on his/her valuation of the entire share of the group. But, in each group there are different people with different valuations. The same division can be considered proportional by some family members and not proportional by other members of the same family. The main question in this chapter is:

> How should we assess the fairness of a division among families?

6.1.1 Results

One option that comes to mind is to aggregate the valuations in each family to a single *family valuation* (also known as: *collective welfare function*). Following the utilitarian tradition (Bentham, 1789), the family-valuation can be defined as the sum or (equivalently) the arithmetic average of the valuations of all family members. We call a division **average-proportional** if every family receives a share with an average value (averaged over all family members) of at least $1/k$ of its average value of the entire cake. This definition makes the family-division problem easy, since each family can be regarded as a single agent, so the problem reduces to fair division among k agents. Classic results (Steinhaus, 1948) imply that an average-proportional division always exists (Section 6.3).

The problem with average-proportionality is that it makes sense only when the numeric values of the agents' valuations are meaningful and they are all measured in the same units, e.g. in dollars (see chapter 3 of Moulin (2004) for some real-life examples of such situations). However, if the valuations represent individual happiness measures that cannot be put on a common scale, then their sum is meaningless, and other fairness criteria should be used.

A second option is to require that all members of every family agree that the division is fair. We call a division **unanimous-proportional** if every agent values his/her family's share as at least $1/k$ of the total value. The advantage of this definition is that it does not need to assume that all valuations share a common scale. A unanimous-proportional division always exists (Section 6.4).

A disadvantage of unanimous-proportionality, compared to average-proportionality, is that unanimous-proportional divisions might be highly fractional. As an illustration, if the cake is an interval, then there always exists an average-proportional division in which each family receives an interval. However, a unanimous-proportional division in which each family receives an interval might not exist, and moreover, in all unanimous-proportional divisions, the total number of intervals might have to be at least n - the number of agents (Section 6.4). When the number of agents is large, as in the case of dividing land between states, such divisions might be impractical.

In democratic societies, decisions are almost never accepted unanimously. In fact, when the number of citizens is large, it may be impossible to attain unanimity on even the most trivial issue. The simplest decision rule in such societies is the majority rule. Inspired by this rule, we suggest a third fairness criterion. We call a division **democratic-proportional** if at least half the citizens in each family value their family's share as at least $1/k$. This definition can be justified according to the following process. After a division is proposed, each group conducts a referendum in which each citizen approves the division if he/she feels that the division is proportional. The division is implemented only if, in every group, at least half of its members approve it. The democratic-proportionality criterion combines some advantages of the other two criteria. It is similar to unanimous-proportionality in that it does not need to assume that all valuations share a common scale. When there are $k = 2$ families with equal rights, it is similar to average-proportionality in that it can be satisfied with connected pieces - there always exists a democratic-proportional division in which each family receives a single connected piece. An additional advantage of democratic-proportionality in this case is that it can be found by an efficient division protocol (Section 6.5). [1]

The present paper compares the three fairness criteria in different settings: the number of families can be two or more than two, and the entitlements of the families can be equal or different. In the common case when there are two families with equal entitlements, democratic-proportionality is apparently the most practical criterion, since it guarantees the existence of connected divisions without assuming a common utility scale. Although democratic-fairness might leave up to half the citizens unhappy, this may be unavoidable in real-life situations. This adds an aspect to Winston Churchil's dictum: "democracy is the worst form of government, except all the others that have been tried".

6.1.2 Related Work

Group-envy-freeness and on-the-fly coalitions

Berliant et al. (1992); Hüsseinov (2011) study the concept of *group-envy-free* cake-cutting. Their model is the standard cake-cutting model in which the cake is divided among *individuals* (and not among families as in our model). They define a group-envy-free division as a division in which no coalition of individuals can

[1] In contrast, average-proportional and unanimous-proportional allocations cannot be found by *any* finite protocol. We omit the details here since the present chapter focuses on existence rather than computational efficiency. More details can be found in Segal-Halevi and Nitzan (2016).

take the pieces allocated to another coalition with the same number of individuals and re-divide the pieces among its members such that all members are weakly better-off. Coalitions are also studied by Dall'Aglio et al. (2009); Dall'Aglio and Di Luca (2014).

In our setting, the families are pre-determined and the agents do not form coalitions on-the-fly. In an alternative model, in which agents *are* allowed to form coalitions based on their preferences, the family-cake-cutting problem becomes easier. For instance, it is easy to achieve a unanimous-proportional division with connected pieces between two coalitions: ask each agent to mark its median line, find the median of all medians, then divide the agents to two coalitions according to whether their median line is to the left or to the right of the median-of-medians.

Fair division with public goods

In our setting, the piece given to each family is considered a "public good" in this specific family. The existence of fair allocations of homogeneous goods when some of the goods are public has been studied e.g. by Diamantaras (1992); Diamantaras and Wilkie (1994, 1996); Guth and Kliemt (2002). In these studies, each good is either private (consumed by a single agent) or public (consumed by all agents). In the present paper, each piece of land is consumed by all agents in a single family - a situation not captured by existing public-good models.

Family preferences in matching markets

Besides land division, family preferences are important in matching markets, too. For example, when matching doctors to hospitals, usually a husband and a wife who are both doctors want to be matched to the same hospital. This issue poses a substantial challenge to stable-matching mechanisms (Klaus and Klijn, 2005, 2007; Kojima et al., 2013; Ashlagi et al., 2014).

Fairness in group decisions

The notion of fairness between groups has been studied empirically in the context of the well-known *ultimatum game*. In the standard version of this game, an individual agent (the *proposer*) suggests a division of a sum of money to another individual (the *responder*), which can either approve or reject it. In the group version, either the proposer or the responder or both are groups of agents. The groups have to decide together what division to propose and whether to accept a proposed division.

Experiments by Robert and Carnevale (1997); Bornstein and Yaniv (1998) show that, in general, groups tend to act more rationally by proposing and accepting divisions which are less fair. Messick et al. (1997) studies the effect of different group-decision rules while Santos et al. (2015) uses a threshold decision rule which is a generalized version of our majority rule (an allocation is accepted if at least M agents in the responder group vote to accept it).

These studies are only tangentially relevant to the present paper, since they deal with a much simpler division problem in which the divided good is *homogeneous* (money) rather than heterogeneous (cake/land).

Non-additive utilities

As explained in Sections 6.4 and 6.5, the difficulty with unanimous-proportionality and democratic-proportionality is that the associated family-valuation functions are not additive. In the previous chapters we encountered value-functions that are not additive because of the geometry; here, the value-functions are not additive because of the family constraints. See subsection 4.1.2 on page 60 for related work.

6.2 Model

We briefly recall some terminology from Chapter 2 (see there for formal definitions).

- C is the cake to be divided. In this chapter we return to the one-dimensional model and assume that C is an interval in $\mathbb{R}$.
- n is the number of agents participating in the division. In this chapter, n is the total number of agents in all families together.

- For each agent $i \in \{1, \ldots, n\}$, $V_i(X_i)$ is agent i's value-measure of the piece X_i. In this chapter we adapt the normalization assumption common in the cake-cutting literature, and assume that $\forall i$: $V_i(\emptyset) = 0$, $V_i(C) = 1$.

6.2.1 Families and entitlements

There are k families, denoted by F_j, $j \in \{1, ..., k\}$.

The number of agents in F_j is n_j. Each agent is a member of exactly one family, so $n = \sum_{j=1}^{k} n_j$.

For each family j, there is a positive weight w_j representing the entitlement of this family. The sum of all weights is one: $\sum_{j=1}^{k} w_j = 1$.

In the simplest setting, the families have equal entitlements, i.e, for each $j \in \{1, \ldots, k\}$: $w_j = 1/k$. Equal entitlements make sense, for example, when k married siblings inherit their parents' estate. While an heir will probably like to take his family's preferences into account when selecting a share, each heir is entitled to $1/k$ of the estate regardless of the size of his/her family.

In general, each family may have a different entitlement. The entitlement of a family may depend on its size but may also depend on other factors. For example, when two states jointly discover a new island, they will probably want to divide the island between them in proportion to their investment and not in proportion their population.

6.2.2 Allocations and components

An *allocation* is a vector of k pieces, $X = (X_1, \ldots, X_k)$, one piece per family, such that the X_j are pairwise-disjoint and $\cup_j X_j = C$.

Each piece is a finite union of intervals. We denote by $Comp(X_j)$ the number of connected components (intervals) in the piece X_j, and by $Comp(X)$ the total number of components in the allocation X, i.e:

$$Comp(X) = \sum_{j=1}^{k} Comp(X_j)$$

Ideally, we would like that each piece be connected, i.e, $\forall i$: $Comp(X_i) = 1$ and $Comp(X) = k$. This requirement is especially meaningful when the divided resource is land, since a contiguous piece of land is much easier to use than a collection of disconnected patches.

However, a division with connected pieces is not always possible. Several countries have a disconnected territory. A striking example is the India-Bangladesh border. According to Wikipedia,[2] "Within the main body of Bangladesh were 102 enclaves of Indian territory, which in turn contained 21 Bangladeshi counter-enclaves, one of which contained an Indian counter-counter-enclave... within the Indian mainland were 71 Bangladeshi enclaves, containing 3 Indian counter-enclaves". Another example is Baarle-Hertog - a Belgian municipality made of 24 separate parcels of land, most of which are exclaves in the Netherlands.[3]

In case a division with connected pieces is not possible, it is still desirable that the number of connectivity components - $Comp(X)$ - be as small as possible. This is a common requirement in the cake-cutting literature. When the cake is an interval, the components are sub-intervals and their number is one plus the number of *cuts*. Hence, the number of components is minimized by minimizing the number of cuts (Robertson and Webb, 1995; Webb, 1997; Shishido and Zeng, 1999; Barbanel and Brams, 2004, 2014). In a realistic, 3-dimensional world, the additional dimensions can be used to connect the components, e.g, by bridges or tunnels. Still, it is desirable to minimize the number of components in the original division in order to reduce the number of required bridges/tunnels. The goal of minimizing the number of components is also pursued in real-life politics. Going back to India and Bangladesh, after many years of negotiations they finally started to exchange most of their enclaves during the years 2015–2016. This is expected to reduce the number of components from 200 to a more reasonable number.

[2] Wikipedia page "India–Bangladesh enclaves".

[3] Wikipedia page "Baarle-Hertog". Many other examples are listed in Wikipedia page "List of enclaves and exclaves". We are grateful to Ian Turton for the references.

6.2.3 Three fairness criteria

To define the criterion of average-proportionality, consider the following family-valuation functions:

$$W_j^{avg}(X_j) = \frac{\sum_{i \in F_j} V_i(X_j)}{n_j} \qquad \text{for } j \in \{1, ..., k\}.$$

An allocation X is called **average-proportional** if

$$\forall j \in \{1, \dots, k\} : W_j^{avg}(X_j) \geq w_j$$

An allocation X is called **unanimous-proportional** if:

$$\forall j \in \{1, \dots, k\} : \forall i \in F_j : V_i(X_j) \geq w_j$$

An allocation X is called **democratic-proportional** if for all $j \in \{1, \dots, k\}$, for at least half the members $i \in F_j$:

$$V_i(X_j) \geq w_j$$

where w_j is the entitlement of family j.

Of these three fairness criteria, unanimous-proportionality is clearly the strongest: it implies both average-proportionality and democratic-proportionality. The other two definitions do not imply each other, as shown in the following example.

Consider a land-estate consisting of four districts. It has to be divided between two families: (1) {Alice,Bob,Charlie} and (2) {David,Eva,Frankie}. The families have equal entitlements, i.e, $w_1 = w_2 = 1/2$. Each member's valuation of each district is shown in the table below:

Alice	60	30	3	3
Bob	50	40	3	3
Charlie	10	80	3	3
David	3	3	60	30
Eva	3	3	60	30
Frankie	3	3	0	90

Note that the value of the entire land is 96 according to all agents, so proportionality implies that each family should get at least 48.

If the two leftmost districts are given to family 1 and the two rightmost districts are given to family 2, then the division is *unanimous-proportional*, since each member of each family feels that his family's share is worth 90. This division is also, of course, average-proportional and democratic-proportional.

If only the single leftmost district is given to family 1 and the other three districts are given to family 2, then the division is still *democratic-proportional*, since Alice and Bob feel that their family received more than 48. However, Charlie feels that his family received only 10, so the division is not unanimous-proportional. Moreover, the division is not average-proportional since the average valuation of family 1 is only (60+50+10)/3=40.

If the three leftmost districts are given to family 1 and only the rightmost district is given to family 2, then the division is *average-proportional*, since family 2's average valuation of its share is (30+30+90)/3=50. However, it is not unanimous-proportional and not even democratic-proportional, since David and Eva feel that their share is worth only 30. □

A property of cake partitions is called *feasible* if for every k families and n agents there exists an allocation satisfying this property. Otherwise, the property is called *infeasible*. In the following sections we study the feasibility of the three fairness criteria in turn.

6.3 Average fairness

Given any n additive value functions V_i, the k family-valuations W_j^{avg} defined above are also additive. Therefore, the family cake-cutting problem can be reduced to the classic problem of cake-cutting among individuals: there are k individual agents, indexed by $j \in \{1, \ldots, k\}$, and the valuation of agent j is the additive value measure W_j^{avg}. This implies the following easy positive result:

Theorem 6.3.1. *When families have equal entitlements, average-proportionality with connected pieces is feasible.*

Proof. This follows from classic results proving the existence of connected proportional allocations for individual agents (Steinhaus, 1948; Even and Paz, 1984). □ □

The situation is more difficult with different entitlements, as shown by the following negative result.

Theorem 6.3.2. *When families have different entitlements, average-proportionality with connected pieces may be infeasible. Moreover, at least $2k-1$ components may be required to attain an average-proportional allocation.*

Proof. Suppose there are k families, the entitlement of family 1 is $\frac{k^2}{k^2+k-1}$ and the entitlement of each of the the other families is $\frac{1}{k^2+k-1}$. The cake consists of $2k-1$ districts and the average family valuations in these districts are:

Family 1	1	0	1	0	1	0	1	...	1	0	1
Family 2	0	1	0	0	0	0	0	...	0	0	0
Family 3	0	0	0	1	0	0	0	...	0	0	0
Family 4	0	0	0	0	0	1	0	...	0	0	0
...								...			
Family k	0	0	0	0	0	0	0	...	0	1	0

Family 1 must receive more than $(k-1)/k$ of the cake, so it must receive a positive slice of each of its k positive districts. But, it cannot receive a single interval that touches two of its positive districts, since such an interval will leave one of the other families with zero value. Therefore, family 1 must receive at least k components. Each of the other families must receive one component, so the total number of components is at least $2k-1$. □ □

We do not know if the lower bound of $2k-1$ is tight even for individual agents.[4] Interestingly, our results on unanimous-proportional division with different entitlements can be used to attain a non-trivial upper bound on the number of cuts required for dividing a cake among k individuals with different entitlements.

Lemma 6.3.3. *Given k agents with different entitlements, a proportional division with $\lceil \log_2 k \rceil \cdot (2k-2)+1$ components is feasible.*

Proof. In Theorem 6.4.7 we will prove that, given n agents in k families with different entitlements, a unanimous-proportional division with $\lceil \log_2 k \rceil \cdot (2n-2)+1$ components is feasible. Now, suppose each family has a single member and let $n = k$. □ □

This immediately implies the same upper bound for average-proportionality:

Theorem 6.3.4. *Given k families with different entitlements, an average-proportional division with $\lceil \log_2 k \rceil \cdot (2k-2)+1$ components is feasible.*

This matches the lower bound of $2k-1$ for $k=2$ families, but leaves a gap for $k \geq 3$ families.

6.4 Unanimous fairness

Before presenting our results, we note that unanimous-proportionality, like average-proportionality, can also be defined using family-valuation functions. Define:

$$W_j^{\min}(X_j) := \min_{i \in F_j} V_i(X_j) \qquad \text{for } j \in \{1, ..., k\}.$$

[4] McAvaney et al. (1992); Robertson and Webb (1997, 1998) discuss the computational aspect of this question - how many intermediate "cut" marks are required (mainly for two agents). But they do not discuss the existential question of how many cuts are needed in the final division.

Then, a division is unanimous-proportional if-and-only-if:

$$\forall j : W_j^{\min}(X_j) \geq w_j$$

However, in contrast to the functions W^{avg} defined in Section 6.3, the functions $W^{\min}$ are in general not additive. For example, consider a cake with three districts and a family with the following valuations:

	C_1	C_2	C_3	$C_1 \cup C_2 \cup C_3$
Alice	1	1	1	$3 = 1 + 1 + 1$
Bob	0	2	1	$3 = 0 + 2 + 1$
Charlie	0	1	2	$3 = 0 + 1 + 2$
$W^{\min}$	0	1	1	$3 > 0 + 1 + 1$

While the individual valuations are additive, $W^{\min}$ is not additive (it is not even subadditive). Therefore, the classic cake-cutting results on proportional cake-cutting cannot be used, and different techniques are needed.

6.4.1 Exact division

Initially, we assume that the entitlements are equal, i.e: $w_j = 1/k$ for all j. We relate unanimous-proportionality to a classic cake-cutting problem of finding an *exact division*:

Definition 6.4.1. Exact(N, K) is the following problem. Given N agents and an integer K, find a division of the cake to K pieces, such that each of the N agents assigns exactly the same value to all pieces:

$$\forall j = 1, ..., K : \ \forall i = 1, ..., N : \ V_i(X_j) = 1/K.$$

Exact division is a difficult problem, since it requires all agents to agree on the values of all pieces, not only their own piece. In this section we prove that finding a unanimous-proportional division is similarly difficult: we show a two-way reduction between the problem of unanimous-proportional division and the problem of exact division.

Denote by UnanimousProp(n, k) the problem of finding a unanimous-proportional division when there are n agents grouped in k families with equal entitlements.

6.4.2 UnanimousProp $\implies$ Exact

Lemma 6.4.2. *For every pair of integers $N \geq 1, K \geq 1$, a solution to UnanimousProp $(N(K-1)+1,\ K)$ implies a solution to Exact (N, K).*

Proof. Given an instance of Exact(N, K) (N agents and a number K of required pieces), create K families. Each of the first $K - 1$ families contains N agents with the same valuations as the given N agents. The K-th family contains a single agent whose valuation is the average of the N given valuations:

$$V^* = \frac{1}{N} \sum_{i=1}^{N} V_i.$$

The total number of agents in all K families is $N(K-1)+1$. Use UnanimousProp $(N(K-1)+1,\ K)$ to find a unanimous-proportional division, X. For each agent i in family j: $V_i(X_j) \geq 1/K$.

By construction, each of the first $K - 1$ families has an agent with valuation V_i. Hence, all N agents value each of the first $K - 1$ pieces as at least $1/K$ and:

$$\forall i = 1, ..., N : \quad \sum_{j=1}^{K-1} V_i(X_j) \geq \frac{K-1}{K}.$$

Hence, by additivity, every agent values the K-th piece as at most $1/K$:

$$\forall i = 1, ..., N : \quad V_i(X_K) \leq 1/K.$$

The piece X_K is given to the agent with value measure V^*, so by proportionality: $V^*(X_K) \geq 1/K$. By construction, $V^*(X_K)$ is the average of the $V_i(X_K)$. Hence:

$$\forall i = 1, ..., N : \quad V_i(X_K) = 1/K.$$

Again by additivity:

$$\forall i = 1, ..., N : \quad \sum_{j=1}^{K-1} V_i(X_j) = \frac{K-1}{K}.$$

Hence, necessarily:

$$\forall i = 1, ..., N, \quad \forall j = 1, ..., K-1 : \quad V_i(X_j) = 1/K.$$

So we have found an exact division and solved Exact(N, K) as required. □ □

Alon (1987) proved that for every N and K, an Exact(N, K) division might require at least $N(K-1)+1$ components. Combining this result with the above lemma implies the following negative result:

Theorem 6.4.3. *For every N, K, let $n = N(K-1)+1$. A unanimous-proportional division for n agents in K families might require at least n components.*

This implies that, in particular, unanimous-proportionality with connected pieces is infeasible.

6.4.3 Exact $\implies$ UnanimousProp

Lemma 6.4.4. *For each n, k, a solution to Exact $(n-1, k)$ implies a solution to UnanimousProp (n, k) for any grouping of the n agents to k families.*

Proof. Suppose we are given an instance of UnanimousProp(n, k), i.e, n agents in k families. Select $n-1$ agents arbitrarily. Use Exact$(n-1, k)$ to find a partition of the cake to k pieces, such that each of the $n-1$ agents values each of these pieces as exactly $1/k$. Ask the n-th agent to choose a favorite piece; by the pigeonhole principle, this value is worth at least $1/k$ for that agent. Give that piece to the family of the n-th agent. Give the other $k-1$ pieces arbitrarily to the remaining $k-1$ families. The resulting division is unanimous-proportional. □ □

Alon (1987) proved that for every N and K, Exact(N, K) has a solution with at most $N(K-1)+1$ components (at most $N(K-1)$ cuts). Combining this result with the above lemma implies the following positive result:

Theorem 6.4.5. *Given n agents in k families with equal entitlements, a unanimous-proportional division with $(n-1) \cdot (k-1) + 1$ components is feasible.*

For $k = 2$ families, the positive result of Theorem 6.4.5 is n, which matches the lower bound of Theorem 6.4.3.

For $k > 2$ families, the number of components can be made smaller, as explained in the following subsections.

6.4.4 Less components: equal entitlements

We start with an example. Assume there are $k = 4$ families. By Theorem 6.4.5, using $3(n-1)$ cuts, the cake can be divided to 4 subsets which are considered equal by all n members. But for a unanimous-proportional division, it is not required that all members think that all pieces are equal, it is only required that all members believe that their family's share is worth at least $1/4$. This can be achieved as follows:

- Divide the cake to two subsets which all n agents value as exactly $1/2$. This is equivalent to solving Exact$(n, 2)$, which by Alon (1987), can be done with at most n cuts. Call the two resulting subsets West and East.
- Assign arbitrary two families to West and the other two families to East. Mark by n_W the total number of members in the families assigned to West and by n_E the total number of members assigned to East.

- Divide the West to two pieces which all n_W agents value as exactly $1/4$; this can be done with n_W cuts. Give a piece to each family. Divide the East similarly using n_E cuts.

The first step requires n cuts and the second step requires $n_W + n_E = n$ cuts too. Hence the total number of cuts required is only $2n$, rather than $3n - 1$.

In fact, two cuts can be saved in each step by excluding two members (from two different families) from the exact division. These members will not think that the division is equal, but they will be allowed to choose the favorite piece for their family. Thus only $2(n-2)$ cuts are required. A simple inductive argument shows that whenever k is a power of 2, $(\log_2 k) \cdot (n - k/2)$ cuts are required.

When k is not a power of 2, a result by Stromquist and Woodall (1985) can be used. They prove that, for every fraction $r \in [0,1]$, it is possible to cut a piece of cake such that all n agents agree that its value is exactly r using at most $2n-2$ cuts.[5] This can be used as follows:

- Select integers $l_1, l_2 \in \{1, ..., k-1\}$ such that $l_1 + l_2 = k$.
- Apply Stromquist and Woodall (1985) with $r = l_1/k$: using $2n-4$ cuts, cut a piece X_1 that $n-1$ agents value as exactly l_1/k. This means that these $n-1$ agents value the other piece, X_2, as exactly l_2/k.
- Let the n-th agent choose a piece for his family; assign the other families arbitrarily such that l_1 families are assigned to piece X_1 and the other l_2 families to piece X_2.
- Recursively divide piece X_1 to its l_1 families and piece X_2 to its l_2 families.

After a finite number of recursion steps, the number of families assigned to each piece becomes 1 and the procedure ends. The number of cuts in each level of the recursion is at most $(2n-4)$. The depth of recursion can be bounded by $\lceil \log_2 k \rceil$ by dividing k to halves (if it is even) or to almost-halves (if it is odd; i.e. take $l_1 = (k-1)/2$ and $l_2 = (k+1)/2$). Hence:

Theorem 6.4.6. *Given n agents in k families with equal entitlements, a unanimous-proportional division with* $\lceil \log_2 k \rceil \cdot (2n-4) + 1$ *components is feasible.*

Note that Theorem 6.4.5 and Theorem 6.4.6 both give upper bounds on the number of components required for unanimous-proportionality. The bound of Theorem 6.4.5 is stronger when k is small and the bound of Theorem 6.4.6 is stronger when k is large.

6.4.5 Less components: different entitlements

When the families have different entitlements, the procedure of the previous subsection cannot be used. We cannot let the n-th agent select a piece for his family, since the pieces are different. For example, suppose there are two families with entitlements $w_1 = 1/3, w_2 = 2/3$. We can divide the cake to two pieces X_1, X_2 such that $n-1$ agents value X_1 as $1/3$ and X_2 as $2/3$. So all of them agree that X_1 should be given to family 1 and X_2 should be given to family 2. But, the n-th agent might select the wrong piece for his family. Therefore, the procedure should be modified as follows.

- Select an integer $l \in \{1, ..., k\}$.
- Divide the families to two subsets: $F_1, \ldots, F_l$ and $F_{l+1}, \ldots, F_k$.
- Apply Stromquist and Woodall (1985) with $r = \sum_{j=1}^{l} w_j$: using $2n-2$ cuts, cut a piece X_1 which all n agents value as exactly $\sum_{j=1}^{l} w_j$. This means that all n agents value the other piece, X_2, as exactly $\sum_{j=l+1}^{k} w_j$.
- Recursively divide piece X_1 to $F_1, \ldots, F_l$ and piece X_2 to $F_{l+1}, \ldots, F_k$.

Here, the number of cuts in each level of the recursion is at most $(2n-2)$. The depth of recursion can be bounded by $\lceil \log_2 k \rceil$ by choosing $l = k/2$ (if k is even) or $l = (k-1)/2$ (if k is odd). Hence:

[5] They prove that, if the cake is a circle, the number of connected components is $n-1$. Hence, the number of cuts is $2n-2$. This is also true when the cake is an interval, although the number of connected components in this case is n.

Algorithm 1 Finding a democratic-envy-free division for two families

INPUT:
- A cake, which is assumed to be the unit interval $[0, 1]$.
- n additive agents, all of whom value the cake as 1.
- A grouping of the agents to 2 families, F_1, F_2.

OUTPUT:
A democratic-envy-free division of the cake to 2 pieces.

ALGORITHM:
- Each agent $i = 1, ..., n$ marks an $x_i \in [0, 1]$ such that $V_i([0, x_i]) = V_i([x_i, 1]) = 1/2$.
- For each family $j = 1, 2$, find the median of its members' marks: $M_j = \text{median}_{i \in F_j} x_i$. Find the median of the family medians: $M^* = (M_1 + M_2)/2$.
- If $M_1 < M_2$ then give $[0, M^*]$ to F_1 and $[M^*, 1]$ to F_2.
Otherwise give $[0, M^*]$ to F_2 and $[M^*, 1]$ to F_1.

Theorem 6.4.7. *Given n agents in k families with different entitlements, a unanimous-proportional division with* $\lceil \log_2 k \rceil \cdot (2n - 2) + 1$ *components is feasible.*

In concluding the analysis of unanimous-proportionality, recall that, even for $k = 2$ families, unanimous-proportionality is as difficult as exact division and might require the same number of components - n. In the worst case, we might need to give a disjoint component to each member, which negates the concept of division to families. Therefore we now turn to the analysis of an alternative fairness criterion that yields more useful results.

6.5 Democratic fairness

Like unanimous-proportionality (Section 6.4), democratic-proportionality can also be defined using family-valuation functions. Define:

$$W_j^{med}(X_j) := \frac{\text{median}_{i \in F_j} V_i(X_j)}{n_j} \qquad \text{for } j \in \{1, ..., k\}.$$

A division is democratic-proportional if-and-only-if:

$$\forall j : W_j^{med}(X_j) \geq w_j$$

However, the W^{med} functions are not additive,[6] so classic cake-cutting results cannot be used.

6.5.1 Two families: a division procedure

We start with a positive result for two families with equal entitlements, which shows that democratic-proportionality is substantially easier than unanimous-proportionality.

Theorem 6.5.1. *When there are* $k = 2$ *families with equal entitlements, democratic-proportionality with connected pieces is feasible.*

Proof. Algorithm 1 finds a democratic-proportional division between two families. For each family, a location M_j is calculated such that, if the cake is cut at M_j, half the members value the interval $[0, M_j]$ as at least $1/2$ and the other half value the interval $[M_j, 1]$ as at least $1/2$. Then, the cake is cut between the two family medians, and each family receives the piece containing its own median. By construction, at least half the members in each family value their family's share as at least $1/2$, so the division is democratic-envy-free. In contrast to the impossibility results of the previous sections, here each family receives a single connected piece. □

[6] See the example in the beginning of Section 6.4. In that example W^{med} is identical to $W^{\min}$.

Unfortunately, this positive result is not applicable when there are more than two families, as shown in the following subsection.

6.5.2 Three or more families: an impossibility result

Given a specific allocation of cake to families, define a *zero* agent as an agent who values his family's share as 0 and a *positive* agent as an agent who believes his family received a share with a positive value. Note that positivity is a much weaker requirement than proportionality.

Lemma 6.5.2. *Assume there are $n = mk$ agents, divided into k families with m members in each family. To guarantee that at least q members in each family are positive, the total number of components may need to be at least:*

$$k \cdot \frac{kq - m}{k - 1}$$

Proof. Number the families by $j = 0, ..., k-1$ and the members in each family by $i = 0, ..., m-1$. Assume that the cake is the interval $[0, mk]$. In each family j, each member i wants only the following interval: $(ik + j,\ ik + j + 1)$. Thus there is no overlap between desired pieces of different members. The table below illustrates the construction for $k = 2$, $m = 3$. The families are {Alice,Bob,Charlie} and {David,Eva,Frankie}:

Alice	1	0	0	0	0	0
Bob	0	0	1	0	0	0
Charlie	0	0	0	0	1	0
David	0	1	0	0	0	0
Eva	0	0	0	1	0	0
Frankie	0	0	0	0	0	1

Suppose the piece X_j (the piece given to family j) is made of $l \geq 1$ components. We can make l members of F_j positive using l intervals of positive length inside their desired areas. However, if $q > l$, we also have to make the remaining $q - l$ members positive. For this, we have to extend $q - l$ intervals to length k. Each such extension totally covers the desired area of one member in each of the other families. Overall, each family creates $q - l$ zero members in each of the other families. The number of zero members in each family is thus $(k-1)(q-l)$. Adding the q members which must be positive in each family, we get the following necessary condition: $(k-1)(q-l) + q \leq m$. This is equivalent to:

$$l \geq \frac{kq - m}{k - 1}.$$

The total number of components is $k \cdot l$, which is at least equal to the expression stated in the Lemma. □
□

In a unanimous-proportional division, all members in each family must be positive. Taking $q = m$ gives $l \geq m$ and the number of components is at least $km = n$, which coincides with the bound of Theorem 6.4.3. In a democratic-proportional division, at least half the members in each family must be positive. Taking $q = m/2$ yields the following negative result:

Theorem 6.5.3. *In a democratic-proportional division with n agents grouped into k families, the number of components may need to be at least*

$$n \cdot \frac{k/2 - 1}{k - 1}$$

Note that for $k = 2$ the lower bound is 0, and indeed we already saw that in this case a connected allocation is feasible.

6.5.3 Three or more families: positive results

Suppose we do want a democratic-proportional division for three or more families. How many components are sufficient?

Algorithm 2 Finding a democratic-proportional division for $k \geq 2$ families.

INPUT:

- A cake, which is assumed to be the unit interval $[0, 1]$.
- n additive agents, all of whom value the cake as 1.
- A grouping of the agents to k families, $F_1, ..., F_k$.

OUTPUT:

A democratic-proportional division of the cake to k pieces.

ALGORITHM:

- Each agent $i = 1, ..., n$ selects an $x_i \in [0, 1]$ such that $V_i([0, x_i]) = \frac{\lceil k/2 \rceil}{k}$ (this means $\frac{1}{2}$ if k is even and $\frac{k+1}{2k}$ if k is odd). Note: $V_i([x_i, 1]) = \frac{\lfloor k/2 \rfloor}{k}$.
- For each family $j = 1, ..., k$, find the median of its members' selections: $M_j = \text{median}_{i \in F_j} x_i$.
- Order the families in increasing order of their medians. Find the median of the family-medians: $M^* = M_{\lceil k/2 \rceil}$. Cut the cake at $x = M^*$.
- Define the *western families* as the F_j with $j = 1, ..., \lceil k/2 \rceil$. Let n_W be the total number of members in these families. Divide the interval $[0, M^*]$ among the western families using UnanimousProp($n_W/2$, $\lceil k/2 \rceil$).
- Similarly, define the *eastern families* as the F_j with $j = \lceil k/2 \rceil + 1, ..., k$. There are $\lfloor k/2 \rfloor$ such families. Let n_E be their total number of members. Divide the interval $(M^*, 1]$ among the eastern families using UnanimousProp($n_E/2$, $\lfloor k/2 \rfloor$).

As a first positive result, we can use Theorem 6.4.7, substituting $n/2$ instead of n: select half of the members in each family arbitrarily, then find a division which is unanimous-proportional for them while ignoring all other members. This leads to:

Theorem 6.5.4. *Given n agents in k families with different entitlements, democratic-proportionality with* $\lceil \log_2 k \rceil \cdot (n-2) + 1$ *components is feasible.*

However, for families with equal entitlements we can do much better. Algorithm 2 generalizes Algorithm 1: for any number of families.

The algorithm works in two steps.

Step 1: Halving. For each family, a location M_j is calculated such that, if the cake is cut at M_j, half the family members value the interval $[0, M_j]$ as at least $\frac{\lceil k/2 \rceil}{k}$ and the other half value the interval $[M_j, 1]$ as at least $\frac{\lfloor k/2 \rfloor}{k}$. Then, the cake is cut in M^* - the median of the family medians. The $\lceil k/2 \rceil$ "western families" - for which $M_j \leq M^*$ - are assigned to the western interval of the cake - $[0, M^*]$. By construction, at least half the members in each of the western families value $[0, M^*]$ as at least $\frac{\lceil k/2 \rceil}{k}$. We say that these members are "happy". Similarly, the $\lfloor k/2 \rfloor$ eastern families - for which $M_j \geq M^*$ - are assigned to the eastern interval $(M^*, 1]$; at least half the members in each of these families are "happy", i.e, value the interval $(M^*, 1]$ as at least $\frac{\lfloor k/2 \rfloor}{k}$.

If there are only two families ($k = 2$), then we are done: there is exactly one western family and one eastern family ($\lceil k/2 \rceil = \lfloor k/2 \rfloor = 1$). For each family $j \in \{1, 2\}$, at least half the members of each family value their family's share as at least $1/2$. Hence, the allocation of X_j to family j is democratic-proportional.

If there are more than two families ($k > 2$), an additional step is required.

Step 2: Sub-division. Each of the two sub-intervals should be further divided to the families assigned to it. In each family F_j, at least $n_j/2$ members are happy. So for each F_j, select exactly $n_j/2$ members who are happy. Our goal now is to make sure that these agents remain happy. This can be done using a unanimous-proportional allocation, where only $n_j/2$ happy members in each family (hence $n/2$ members overall) are counted. The unanimous-proportional allocation guarantees that every western-happy-member believes that his family's share is worth at least $\frac{\lceil k/2 \rceil}{k} \cdot \frac{1}{\lceil k/2 \rceil} = \frac{1}{k}$. Similarly, every eastern-happy-member believes that his family's share is worth at least $\frac{\lfloor k/2 \rfloor}{k} \cdot \frac{1}{\lfloor k/2 \rfloor} = \frac{1}{k}$. Hence, the resulting division is democratic-proportional.

We now calculate the number of components in the resulting division. One cut is required for the halving step. For the unanimous-proportional division of the western interval, the number of required cuts is at most $(\lceil k/2 \rceil - 1) \cdot (n_W/2 - 1)$ by Theorem 6.4.5, and at most $\lceil \log_2 \lceil k/2 \rceil \rceil \cdot (n_W - 4)$ by Theorem

Proportionality	#Families (k)	#Connectivity Components	
		Lower	Upper
Average (Sec. 6.3)	k	k	k (connected)
Unanimous	2	n	n
(Sec. 6.4)	k	n	$\min(1 + \lceil \log_2 k \rceil \cdot (2n-4),$ $(k-1)\cdot(n-1)+1)$
Democratic	2	2	2 (connected)
(Sec. 6.5)	k	$n \cdot \frac{k/2-1}{k-1}$	$\min(2 + \lceil \log_2 \lceil k/2 \rceil \rceil \cdot (n-8),$ $2 + (\lceil k/2 \rceil - 1) \cdot (n/2-2))$

Table 6.1: Summary of results for dividing a cake among families: upper and lower bounds on number of cuts

6.4.6. Similarly, for the eastern interval the number of required cuts is at most the minimum of $(\lfloor k/2 \rfloor - 1) \cdot (n_E/2 - 1)$ and $\lceil \log_2 \lfloor k/2 \rfloor \rceil \cdot (n_E - 4)$. The total number of cuts is thus at most $1 + (\lceil k/2 \rceil - 1) \cdot (n/2 - 2)$ and at most $1 + \lceil \log_2 \lceil k/2 \rceil \rceil \cdot (n-8)$. The total number of components is larger by one. We obtain:

Theorem 6.5.5. *Given n agents in k families with equal entitlements, democratic-proportionality is feasible with at most*

$$\min\left(2 + (\lceil k/2 \rceil - 1) \cdot (n/2 - 2) \quad , \quad 2 + \lceil \log_2 \lceil k/2 \rceil \rceil \cdot (n-8)\right)$$

components.

6.6 Conclusions and Future Work

Table 6.1 compares the three fairness criteria studied in the present paper, for families with equal entitlements. Recall that n is the total number of agents in all families.

The case of $k = 2$ families is well-understood. The results for all fairness criteria are tight: by all fairness definitions, we know that a fair division exists with the smallest possible number of connectivity components.

6.6.1 Open questions

The case of $k > 2$ families opens some questions:

- Is unanimous-proportionality with n components feasible for all k? (particularly, with $k = 3$ families, is the number of required components n as in the lower bound, or $2n - 1$ as in the upper bound?).
- Is democratic-proportionality with $n \cdot \frac{k/2-1}{k-1}$ components feasible for all k? (particularly, with $k = 3$ families, is the number of required components $n/4$ as in the lower bound, or $n/2$ as in the upper bound?).

The case of different entitlements is much less understood even for individual agents. As far as we know, it is an open question whether cake-cutting among k individuals with $2k - 1$ components is feasible for $k > 2$. This has direct implications on the number of required components for average-proportionality.

6.6.2 Alternative fairness criteria

One could consider the following alternative fairness criterion: an allocation is *individually-proportional* if the allocation $X = (X_1, \ldots, X_k)$ admits a refinement $Y = (Y_1, \ldots, Y_n)$, where for each family F_j, $\cup_{i \in F_j} Y_i = X_j$, such that for each agent i, $V_i(Y_i) \geq 1/n$. Individually-proportional allocations always exist and can be found by using any classic proportional cake-cutting procedure on the individual agents, disregarding their families. The number of components is at most n. Individual-proportionality makes sense if, after the division of the land among the families, each family intends to further divide its share among its members. However, often this is not the case. When an inherited land-estate is divided between two families, the members of each family intend to live and use their entire share together, rather than dividing it among them. Therefore, the happiness of each family member depends on the entire value of his family's share, rather than on the value of a potential private share he would get in a hypothetic sub-division.

Instead of proportionality, it is possible to use *envy-freeness* as the basic fairness criterion. Envy-freeness means that the valuation of each family in its share should be at least as large as the valuation of the family in another share. Then, **average-envy-freeness** means that the average value of each family in its allocated share (averaged over all family members) is at least as large as its average value in each of the other shares; **unanimous-envy-freeness** means that every agent values his family's share at least as much as any other share; **democratic-envy-freeness** means that at least half the members in each family believe that their family received the best share. Note that this definition inherently assumes that the families have equal entitlements. Section 6.3 (the equal-entitlements case) holds as-is for average-envy-freeness. In Theorems 2 and 3, the recursive-halving procedure cannot be used, and the number of components in the positive results is $O(nk)$ instead of $O(n \log k)$. More details are available in Segal-Halevi and Nitzan (2016).

Finally, the combination of *envy-freeness and Pareto-efficiency* is very interesting, regardless of geometric or computational constraints. Among individuals, an envy-free and Pareto-efficient cake-division always exists (Weller, 1985). Does there always exist an unanimous-envy-free and Pareto-efficient division among families?

The latter question is open not only in the cake-cutting setting, but also in the classic economic setting of dividing homogeneous resources.

6.7 Acknowledgments

We are grateful to Galya Segal-Halevi, Noga Alon, Christian Klamler, Ulle Endriss, Neill Clift and anonymous referees in the Econonic Theory journal and the SAGT 2016 conference for their helpful comments.

This paper started with a discussion in the MathOverflow website.[7] We are grateful to the members who participated in the discussion: Pietro Majer, Tony Huynh and Manfred Weis. Other members of the StackExchange network contributed useful answers and ideas: Alex Ravsky, Andrew D. Hwang, BKay, Christian Elsholtz, Daniel Fischer, David K, D.W, Hurkyl, Ittay Weiss, Kittsil, Michael Albanese, Raphael Reitzig, Real, babou, Domotor Palvolgyi (domotorp), Ian Turton (iant) and ivancho.

[7] http://MathOverflow.net/questions/203060/fair-cake-cutting-between-groups

Bibliography

Abe, K. (2012). A geometric approach to temptation. *Journal of Mathematical Economics*, 48(2):92–97.

Agarwal, P. K., Katz, M. J., and Sharir, M. (1995). Computing depth orders for fat objects and related problems. *Computational Geometry*, 5(4):187–206.

Akopyan, A. and Segal-Halevi, E. (2016). Cutting a Cake without Harming the Toppings. arXiv preprint 1604.00960.

Aleksandrov, M., Aziz, H., Gaspers, S., and Walsh, T. (2015). Online Fair Division: analysing a Food Bank problem.

Alon, N. (1987). Splitting necklaces. *Advances in Mathematics*, 63(3):247–253.

Alvarez-Valdés, R., Parajón, A., and Tamarit, J. M. (2002). A tabu search algorithm for large-scale guillotine (un)constrained two-dimensional cutting problems. *Computers & Operations Research*, 29(7):925–947.

Arzi, O. (2012). Cake Cutting: Achieving Efficiency While Maintaining Fairness. Master's thesis, Bar-Ilan University. Under the supervision of Prof. Yonatan Aumann.

Ashlagi, I., Braverman, M., and Hassidim, A. (2014). Stability in large matching markets with complementarities. *Operations Research*, 62(4):713–732.

Aumann, Y. and Dombb, Y. (2010). The Efficiency of Fair Division with Connected Pieces. *Web, Internet and Network Economics*, 6484:26–37.

Aumann, Y. and Dombb, Y. (2015). The Efficiency of Fair Division with Connected Pieces. *ACM Trans. Econ. Comput.*, 3(4).

Aumann, Y., Dombb, Y., and Hassidim, A. (2013). Computing Socially-Efficient Cake Divisions. In *The International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 343–350.

Aziz, H. and Mackenzie, S. (2016). A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. In *FOCS 2016*.

Aziz, H. and Ye, C. (2014). Cake Cutting Algorithms for Piecewise Constant and Piecewise Uniform Valuations. In Liu, T.-Y., Qi, Q., and Ye, Y., editors, *Web and Internet Economics*, volume 8877 of *Lecture Notes in Computer Science*, pages 1–14. Springer International Publishing.

Azrieli, Y. and Shmaya, E. (2014). Rental harmony with roommates. *Journal of Economic Theory*, 153:128–137.

Bar-Yehuda, R. and Ben-Hanoch, E. (1996). A linear-time algorithm for covering simple polygons with similar rectangles. *International Journal of Computational Geometry & Applications*, 6.01:79–102.

Barbanel, J. B. (2005). *The Geometry of Efficient Fair Division*. Cambridge University Press.

Barbanel, J. B. and Brams, S. J. (2004). Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. *Mathematical Social Sciences*, 48(3):251–269.

Barbanel, J. B. and Brams, S. J. (2011). Two-person pie-cutting: The fairest cuts. *The College Mathematics Journal*, 42(1):25–32.

Barbanel, J. B. and Brams, S. J. (2014). Two-Person Cake Cutting: The Optimal Number of Cuts. *The Mathematical Intelligencer*, 36(3):23–35.

Barbanel, J. B., Brams, S. J., and Stromquist, W. (2009). Cutting a pie is not a piece of cake. *American Mathematical Monthly*, 116(6):496–514.

Beck, A. (1987). Constructing a Fair Border. *The American Mathematical Monthly*, 94(2):157–162.

Bei, X., Chen, N., Hua, X., Tao, B., and Yang, E. (2012). Optimal Proportional Cake Cutting with Connected Pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 1263–1269.

Bentham, J. (1789). *An Introduction to the Principles of Morals and Legislation (Dover Philosophical Classics)*. Dover Publications.

Berliant, M. and Dunz, K. (2004). A foundation of location theory: existence of equilibrium, the welfare theorems, and core. *Journal of Mathematical Economics*, 40(5):593–618.

Berliant, M. and Raa, T. (1988). A foundation of location theory: Consumer preferences and demand. *Journal of Economic Theory*, 44(2):336–353.

Berliant, M., Thomson, W., and Dunz, K. (1992). On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216.

Bernstein, H. (2002). Land Reform: Taking a Long(er) View. *Journal of Agrarian Change*, 2(4):433–463.

Bertsimas, D., Farias, V. F., and Trichakis, N. (2011). The Price of Fairness. *Operations Research*, 59(1):17–31.

Bertsimas, D., Farias, V. F., and Trichakis, N. (2012). On the Efficiency-Fairness Trade-off. *Manage. Sci.*, 58(12):2234–2250.

Bornstein, G. and Yaniv, I. (1998). Individual and Group Behavior in the Ultimatum Game: Are Groups More Rational Players? *Experimental Economics*, 1(1):101–108.

Bouveret, S., Chevaleyre, Y., and Maudet, N. (2016). Fair Allocation of Indivisible Items. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Procaccia, A. D., editors, *Handbook of Computational Social Choice*, chapter 13, pages 284–310. Cambridge University Press.

Bouveret, S. and Lemaître, M. (2015). Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, pages 1–32.

Brams, S. J. (2007). *Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures*. Princeton University Press, first edition edition.

Brams, S. J., Feldman, M., Lai, J. K., Morgenstern, J., and Procaccia, A. D. (2012). On Maxsum Fair Cake Divisions. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI-12)*, pages 1285–1291.

Brams, S. J., Jones, M., and Klamler, C. (2008). Proportional pie-cutting. *International Journal of Game Theory*, 36(3-4):353–367.

Brams, S. J., Jones, M. A., and Klamler, C. (2011). Divide-and-Conquer: A Proportional, Minimal-Envy Cake-Cutting Algorithm. *SIAM Rev.*, 53(2):291–307.

Brams, S. J. and Taylor, A. D. (1995). An Envy-Free Cake Division Protocol. *The American Mathematical Monthly*, 102(1):9–18.

Brams, S. J. and Taylor, A. D. (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

Brams, S. J., Taylor, A. D., and Zwicker, W. S. (1997). A Moving-Knife Solution to the Four-Person Envy-Free Cake Division Problem. *Proceedings of the American Mathematical Society*, 125(2):547–554.

Brânzei, S. (2015). *Computational Fair Division*. PhD thesis, Faculty of Science and Technology in Aarhus university.

Brânzei, S., Caragiannis, I., Kurokawa, D., and Procaccia, A. D. (2016). An Algorithmic Framework for Strategic Fair Division. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, pages 411–417. AAAI Press.

Brânzei, S. and Miltersen, P. B. (2015). A Dictatorship Theorem for Cake Cutting. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 482–488. AAAI Press.

Brânzei, S., Procaccia, A. D., and Zhang, J. (2013). Externalities in Cake Cutting. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, pages 55–61. AAAI Press.

Buchanan, J. M. (1965). An economic theory of clubs. *Economica*, pages 1–14.

Budish, E. (2011). The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103.

Cantillon, E. and Rangel, A. (2002). A graphical analysis of some basic results in social choice. *Social Choice and Welfare*, 19(3):587–611.

Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., and Kyropoulou, M. (2012). The Efficiency of Fair Division. *Theory of Computing Systems*, 50(4):589–610.

Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016). The Unreasonable Fairness of Maximum Nash Welfare. In *17th ACM conference on Economic and Computation (EC 16)*.

Caragiannis, I., Lai, J. K., and Procaccia, A. D. (2011). Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 127–132. AAAI Press.

Chambers, C. P. (2005). Allocation rules for land division. *Journal of Economic Theory*, 121(2):236–258.

Chatterjee, S., Peters, H., and Storcken, T. (2016). Locating a public good on a sphere. *Economics Letters*, 139:46–48.

Chen, Y., Lai, J. K., Parkes, D. C., and Procaccia, A. D. (2013). Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297.

Cohler, Y. J., Lai, J. K., Parkes, D. C., and Procaccia, A. D. (2011). Optimal Envy-Free Cake Cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI-11)*, pages 626–631.

Cui, Y., Yang, Y., Cheng, X., and Song, P. (2008). A recursive branch-and-bound algorithm for the rectangular guillotine strip packing problem. *Computers & Operations Research*, 35(4):1281–1291.

Dall'Aglio, M., Branzei, R., and Tijs, S. (2009). Cooperation in dividing the cake. *TOP - the official journal of the Spanish Society of Statistics and Operations Research*, 17(2):417–432.

Dall'Aglio, M. and Di Luca, C. (2014). Finding maxmin allocations in cooperative and competitive fair division. *Annals of Operations Research*, 223(1):121–136.

Dall'Aglio, M. and Maccheroni, F. (2005). Fair Division without Additivity. *American Mathematical Monthly*, pages 363–365.

Dall'Aglio, M. and Maccheroni, F. (2009). Disputed lands. *Games and Economic Behavior*, 66(1):57–77.

Delimitrou, C. and Kozyrakis, C. (2014). Quasar: Resource-efficient and QoS-aware Cluster Management. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '14, pages 127–144, New York, NY, USA. ACM.

Deng, X., Qi, Q., and Saberi, A. (2012). Algorithmic Solutions for Envy-Free Cake Cutting. *Operations Research*, 60(6):1461–1476.

Devulapalli, R. (2014). *Geometric partitioning algorithms for fair division of geographic resources*. PhD thesis, University of Minnesota.

Diamantaras, D. (1992). On equity with public goods. *Social Choice and Welfare*, 9(2):141–157.

Diamantaras, D. and Wilkie, S. (1994). A Generalization of Kaneko's Ratio Equilibrium for Economies with Private and Public Goods. *Journal of Economic Theory*, 62(2):499–512.

Diamantaras, D. and Wilkie, S. (1996). On the set of Pareto efficient allocations in economies with public goods. *Economic Theory*, 7(2):371–379.

Dickerson, J. P., Procaccia, A. D., and Sandholm, T. (2014). Price of Fairness in Kidney Exchange. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 1013–1020, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Dubins, L. E. and Spanier, E. H. (1961). How to Cut A Cake Fairly. *The American Mathematical Monthly*, 68(1):1+.

Edmonds, J. and Pruhs, K. (2006a). Balanced Allocations of Cake. In *FOCS*, volume 47, pages 623–634. IEEE Computer Society.

Edmonds, J. and Pruhs, K. (2006b). Cake cutting really is not a piece of cake. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, SODA '06, pages 271–278, New York, NY, USA. ACM.

Edmonds, J., Pruhs, K., and Solanki, J. (2008). Confidently Cutting a Cake into Approximately Fair Pieces. *Algorithmic Aspects in Information and Management*, pages 155–164.

Eppstein, D. (2010). Graph-Theoretic Solutions to Computational Geometry Problems. In Paul, C. and Habib, M., editors, *Graph-Theoretic Concepts in Computer Science*, volume 5911 of *Lecture Notes in Computer Science*, pages 1–16. Springer Berlin Heidelberg.

Even, S. and Paz, A. (1984). A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–296.

Fish, C. (2014). A continuous path between shapes. Mathematics Stack Exchange. URL:http://math.stackexchange.com/a/1015267/29780.

Friedman, E., Psomas, C. A., and Vardi, S. (2015). Dynamic Fair Division with Minimal Disruptions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 697–713, New York, NY, USA. ACM.

Gasarch, W. (2015). Which Unbounded Protocol for Envy Free Cake Cutting is Better?

Goeree, J. K. and Kushnir, A. I. (2011). A Geometric Approach to Mechanism Design. *Social Science Research Network Working Paper Series*.

Goldman, J. and Procaccia, A. D. (2015). Spliddit: Unleashing Fair Division Algorithms. *SIGecom Exch.*, 13(2):41–46.

Guth, W. and Kliemt, H. (2002). Non-Discriminatory, Envy Free Provision of a Collective Good: A Note. *Public Choice*, 111(1-2):179–184.

Hifi, M., Saadi, T., and Haddadou, N. (2011). High Performance Peer-to-Peer Distributed Computing with Application to Constrained Two-Dimensional Guillotine Cutting Problem. In *2011 19th International Euromicro Conference on Parallel, Distributed and Network-Based Processing*, pages 552–559. IEEE.

Hill, T. P. (1983). Determining a Fair Border. *The American Mathematical Monthly*, 90(7):438–442.

Hines, J. R., Hlinko, J. C., and Lubke, T. J. F. (1995). From each according to his surplus: Equi-proportionate sharing of commodity tax burdens. *Journal of Public Economics*, 58(3):417–428.

Hoffman, M. (2013). Why community ownership? Understanding land reform in Scotland. *Land Use Policy*, 31:289–297.

Hohenwarter, M. (2002). GeoGebra: Ein Softwaresystem für dynamische Geometrie und Algebra der Ebene. Master's thesis, Paris Lodron University, Salzburg, Austria. (In German.).

Hohenwarter, M., Borcherds, M., Ancsin, G., Bencze, B., Blossier, M., Delobelle, A., Denizet, C., Éliás, J., Fekete, A., Gál, L., Konečný, Z., Kovács, Z., Lizelfelner, S., Parisse, B., and Sturr, G. (2013). GeoGebra 5.0.
urlhttp://www.geogebra.org.

Hüsseinov, F. (2011). A theory of a heterogeneous divisible commodity exchange economy. *Journal of Mathematical Economics*, 47(1):54–59.

Hüsseinov, F. and Sagara, N. (2013). Existence of efficient envy-free allocations of a heterogeneous divisible commodity with nonadditive utilities. *Social Choice and Welfare*, pages 1–18.

Ichiishi, T. and Idzik, A. (1999). Equitable allocation of divisible goods. *Journal of Mathematical Economics*, 32(4):389–400.

Iyer, K. and Huhns, M. N. (2009). A Procedure for the Allocation of Two-Dimensional Resources in a Multiagent System. *International Journal of Cooperative Information Systems*, 18:1–34.

Johnson, H. G. (1971). Trade and growth : A geometrical exposition. *Journal of International Economics*, 1(1):83–101.

Kash, I., Procaccia, A. D., and Shah, N. (2013). No agent left behind: dynamic fair division of multiple resources. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, AAMAS '13, pages 351–358. IFAAMAS.

Katz, M. J. (1997). 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects. *Computational Geometry*, 8(6):299–316.

Keil, J. M. (2000). Polygon Decomposition. In *Handbook of Computational Geometry*, pages 491–518. University of Saskatchewan Saskatoon, Sask., Canada.

Klaus, B. and Klijn, F. (2005). Stable matchings and preferences of couples. *Journal of Economic Theory*, 121(1):75–106.

Klaus, B. and Klijn, F. (2007). Paths to stability for matching markets with couples. *Games and Economic Behavior*, 58(1):154–171.

Kojima, F., Pathak, P. A., and Roth, A. E. (2013). Matching with Couples: Stability and Incentives in Large Markets*. *The Quarterly Journal of Economics*, 128(4):1585–1632.

Kurokawa, D., Lai, J. K., and Procaccia, A. D. (2013). How to Cut a Cake Before the Party Ends. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 555–561.

Lalonde, S. N. (2002). *Determining boundaries in a conflicted world: the role of Uti Possidetis*. McGill-Queen's Press-MQUP.

Legut, J., Potters, J. A. M., and Tijs, S. H. (1994). Economies with Land - A Game Theoretical Approach. *Games and Economic Behavior*, 6(3):416–430.

Lipton, M. (2009). *Land Reform in Developing Countries: Property Rights and Property Wrongs*. Routledge.

Maccheroni, F. and Marinacci, M. (2003). How to cut a pizza fairly: Fair division with decreasing marginal evaluations. *Social Choice and Welfare*, 20(3):457–465.

MacInnes, M. and Shields, K. (2015). The Land Reform (Scotland) Bill and Human Rights: Key Points and Recommendations. *SSRN Electronic Journal*.

Manabe, Y. and Okamoto, T. (2010). Meta-Envy-Free Cake-Cutting Protocols. In Hliněný, P. and Kučera, A., editors, *Mathematical Foundations of Computer Science 2010*, volume 6281 of *Lecture Notes in Computer Science*, chapter 44, pages 501–512. Springer Berlin Heidelberg, Berlin, Heidelberg.

Marathe, M. V., Breu, H., Hunt, H. B., Ravi, S. S., and Rosenkrantz, D. J. (1995). Simple heuristics for unit disk graphs. *Networks*, 25(2):59–68.

McAvaney, K., Robertson, J. M., and Webb, W. A. (1992). Ramsey partitions of integers and fair divisions. *Combinatorica*, 12(2):193–201.

Messick, D. M., Moore, D. A., and Bazerman, M. H. (1997). Ultimatum Bargaining with a Group: Underestimating the Importance of the Decision Rule. *Organizational Behavior and Human Decision Processes*, 69(2):87–101.

Michaelides, S. (2006). *Euclidean Economics: A Mathematical Method for Construction and Application of Economics Models*. Mesogios Publications - Ellinika Grammata.

Mill, J. S. (1863). *Utilitarianism*. Hackett Pub Co, 2 edition.

Mirchandani, R. S. (2013). Superadditivity and Subadditivity in Fair Division. *Journal of Mathematics Research*, 5(3):78–91.

Moulin, H. (1990a). Fair division under joint ownership: Recent results and open problems. *Social Choice and Welfare*, 7(2):149–170.

Moulin, H. (1990b). Uniform externalities: Two axioms for fair allocation. *Journal of Public Economics*, 43(3):305–326.

Moulin, H. (2004). *Fair Division and Collective Welfare*. The MIT Press.

Nicolò, A., Perea, and Roberti, P. (2012). Equal opportunity equivalence in land division. *SERIEs - Journal of the Spanish Economic Association*, 3(1-2):133–142.

Nicolò, A. and Yu, Y. (2008). Strategic divide and choose. *Games and Economic Behavior*, 64(1):268–289.

Nielsen, O. A. (1997). *An introduction to integration and measure theory*. Wiley.

Öztürk, M., Peters, H., and Storcken, T. (2013). Strategy-proof location of a public bad on a disc. *Economics Letters*, 119(1):14–16.

Öztürk, M., Peters, H., and Storcken, T. (2014). On the location of public bads: strategy-proofness under two-dimensional single-dipped preferences. *Economic Theory*, 56(1):83–108.

Pikhurko, O. (2000). On Envy-Free Cake Division. *American Mathematical Monthly*, 107:736–738.

Plott, C. R. (1967). A Notion of Equilibrium and Its Possibility Under Majority Rule. *The American Economic Review*, 57(4):787–806.

Powelson, J. P. (1988). *The Story of Land: A World History of Land Tenure and Agrarian Reform*. Lincoln Institute of Land Policy.

Procaccia, A. D. (2009). Thou Shalt Covet Thy Neighbor's Cake. In *IJCAI'09 Proceedings of the 21st international joint conference on Artificial intelligence*, pages 239–244.

Procaccia, A. D. and Wang, J. (2014). Fair Enough: Guaranteeing Approximate Maximin Shares. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 675–692, New York, NY, USA. ACM.

Reijnierse, J. H. and Potters, J. A. M. (1998). On finding an envy-free Pareto-optimal division. *Mathematical Programming*, 83(1-3):291–311.

Robert, C. and Carnevale, P. J. (1997). Group Choice in Ultimatum Bargaining. *Organizational Behavior and Human Decision Processes*, 72(2):256–279.

Robertson, J. M. and Webb, W. A. (1995). Approximating fair division with a limited number of cuts. *Journal of Combinatorial Theory, Series A*, 72(2):340–344.

Robertson, J. M. and Webb, W. A. (1997). Extensions of Cut-and-Choose Fair Division. *Elemente der Mathematik*, 52(1):23–30.

Robertson, J. M. and Webb, W. A. (1998). *Cake-Cutting Algorithms: Be Fair if You Can*. A K Peters/CRC Press, first edition.

Rosset, P., Patel, R., and Courville, M., editors (2006). *Promised Land: Competing Visions of Agrarian Reform*. Food First Books, first trade paper edition edition.

Saberi, A. and Wang, Y. (2009). Cutting a Cake for Five People. In Goldberg, A. V. and Zhou, Y., editors, *Algorithmic Aspects in Information and Management*, volume 5564 of *Lecture Notes in Computer Science*, pages 292–300. Springer Berlin Heidelberg.

Sagara, N. and Vlach, M. (2005). Equity and Efficiency in a Measure Space with Nonadditive Preferences: The Problem of Cake Division. *Journal of Economic Literature*, 90:1.

Santos, F. P., Santos, F. C., Paiva, A., and Pacheco, J. M. (2015). Evolutionary dynamics of group fairness. *Journal of Theoretical Biology*, 378:96–102.

Segal-Halevi, E., Hassidim, A., and Aumann, Y. (2015a). Envy-Free Cake-Cutting in Two Dimensions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, pages 1021–1028. arXiv preprint 1609.03938.

Segal-Halevi, E., Hassidim, A., and Aumann, Y. (2015b). Waste Makes Haste: Bounded Time Protocols for Envy-Free Cake Cutting with Free Disposal. In *The 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 901–908.

Segal-Halevi, E. and Nitzan, S. (2016). Envy-Free Cake-Cutting among Families. arXiv 1607.01517.

Segal-Halevi, E., Nitzan, S., Hassidim, A., and Aumann, Y. (2017). Fair and square: Cake-cutting in two dimensions. *Journal of Mathematical Economics*.

Sellar (2006). The great land debate and the Land Reform (Scotland) Act 2003. *Norsk Geografisk Tidsskrift - Norwegian Journal of Geography*, 60(1):100–109.

Shishido, H. and Zeng, D.-Z. (1999). Mark-Choose-Cut Algorithms For Fair And Strongly Fair Division. *Group Decision and Negotiation*, 8(2):125–137.

Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16(1):101–104.

Stromquist, W. (1980). How to Cut a Cake Fairly. *The American Mathematical Monthly*, 87(8):640–644.

Stromquist, W. (2007). A pie that can't be cut fairly. In *Dagstuhl Seminar Proceedings 07261*.

Stromquist, W. (2008). Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics*.

Stromquist, W. and Woodall, D. R. (1985). Sets on which several measures agree. *Journal of Mathematical Analysis and Applications*, 108(1):241–248.

Su, F. E. (1999). Rental Harmony: Sperner's Lemma in Fair Division. *The American Mathematical Monthly*, 106(10):930–942.

Sziklai, B. and Segal-Halevi, E. (2015). Resource-monotonicity and Population-monotonicity in Cake-cutting.

Thomson, W. (1983). The Fair Division of a Fixed Supply Among a Growing Population. *Mathematics of Operations Research*, 8(3):319–326.

Thomson, W. (2007). Children Crying at Birthday Parties. Why? *Economic Theory*, 31(3):501–521.

Thomson, W. (2011). *Fair Allocation Rules*, volume 2, pages 393–506. Elsevier.

Tsuruta, S., Oka, M., Todo, T., Sakurai, Y., and Yokoo, M. (2015). Fairness and False-Name Manipulations in Randomized Cake Cutting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 909–917.

Varian, H. R. (1974). Equity, Envy, and Efficiency. *Journal of Economic Theory*.

Walsh, T. (2011). Online Cake Cutting. *Algorithmic Decision Theory*, 6992:292–305.

Webb, W. A. (1990). A Combinatorial Algorithm to Establish a Fair Border. *European Journal of Combinatorics*, 11(3):301–304.

Webb, W. A. (1997). How to cut a cake fairly using a minimal number of cuts. *Discrete Applied Mathematics*, 74(2):183–190.

Weller, D. (1985). Fair division of a measurable space. *Journal of Mathematical Economics*, 14(1):5–17.

Wightman, A. (2015). *The Poor Had No Lawyers: Who Owns Scotland (and How They Got it)*. Birlinn Ltd.

Woeginger, G. J. and Sgall, J. (2007). On the complexity of cake cutting. *Discrete Optimization*, 4(2):213–220.

Woodall, D. R. (1980). Dividing a cake fairly. *Journal of Mathematical Analysis and Applications*, 78(1):233–247.

Zhang, Y., Laurenzano, M. A., Mars, J., and Tang, L. (2014). SMiTe: Precise QoS Prediction on Real-System SMT Processors to Improve Utilization in Warehouse Scale Computers. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO-47, pages 406–418, Washington, DC, USA. IEEE Computer Society.

Zivan, R. (2011). Can trust increase the efficiency of cake cutting algorithms. In *The International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1145–1146.

1. חלוקה גיאומטרית. כשמחלקים עוגה, הצורה הגיאומטרית של הפרוסות לא כל כך חשובה (פרט אולי לכמה אניני טעם), ממילא אנחנו מתכוונים לאכול אותן. ואכן התהליכים הקיימים לחלוקת עוגה אינם מתייחסים לצורה גיאומטרית. במקרים רבים, הם מניחים שהעוגה היא קטע חד-ממדי, וכל פרוסה היא תת-קטע או איחוד של תת-קטעים.

אבל כשמחלקים קרקע, הצורה הגיאומטרית חשובה. קרקע מלבנית, למשל, שימושית יותר לבנייה ולעיבוד מקרקע בצורת ספירלה. גם קרקע ריבועית היא שימושית יותר מקרקע בצורת מלבן צר וארוך כמו מקרוני.

בעבודה זו אנחנו מתייחסים בפירוש לעוגה רב-ממדית. בדרך-כלל, העוגה תהיה מצולע במישור, אבל חלק מהאלגוריתמים מתאימים גם לגופים רב-ממדיים כלליים יותר. אנחנו מאפשרים להציב דרישות מפורשות על הצורה של פרוסות-העוגה. למשל, אפשר לדרוש שכל שחקן יקבל פרוסה בצורת ריבוע, או מלבן שמן (מלבן עם יחס גובה/רוחב מאוזן), או מצולע שמן כלשהו. אנחנו מציגים אלגוריתמים שמכבדים דרישות אלו, יחד עם הדרישות הקלאסיות של פרופורציונליות-חלקית ואי-קיום קנאה.

2. חלוקה מחדש. חלוקת עוגה היא תהליך חד-פעמי. העוגה יוצאת מהתנור טריה וחדשה, ומחולקת מייד. אבל חלוקת קרקע היא תהליך מתמשך. במקרים רבים, הקרקע כבר מחולקת, וצריך לבצע חלוקה מחדש --- רפורמה אגרארית. רפורמות אגראריות בוצעו במקומות רבים בעולם בתקופות שונות בהסטוריה. ידועים ניסיונות לבצע רפורמה כזאת החל מהמאה השמינית לפני הספירה במצרים העתיקה (ע"י המלך המצרי בקנרנף), ועד שנת 2016 (ע"י ממשלת סקוטלנד). חלוקה מחדש של קרקע מעוררת שאלה לא פשוטה על האיזון הראוי בין זכויות-הקניין של בעלי-הקרקע הנוכחיים, לבין הזכות לחלק הוגן של חסרי-הקרקע. בעבודה זו אנו מניחים שהעוגה כבר מחולקת , והחלוקה הנוכחית אינה הוגנת. אנו מציעים אלגוריתמים המאפשרים למצוא חלוקה הוגנת-חלקית, שבה זכויות-הקניין של בעלי-הקרקע הנוכחיים נשמרים באופן חלקי. גורם האיזון בין ההגינות לבין זכות-הקניין ניתן לויסות על-ידי הממשלה. אנו מתייחסים גם לשילוב בין דרישות החלוקה-מחדש לבין דרישות גיאומטריות, כגון: הדרישה שהפרוסות יהיו קשירות, או מלבניות, או קמורות.

האלגוריתמים לחלוקה מחדש מאפשרים לנו גם להשיג חסם עליון על מחיר ההגינות. מחיר ההגינות הוא היחס בין החלוקה היעילה ביותר (החלוקה שבה סכום התועלות של השחקנים הוא גבוה ביותר), לבין החלוקה ההוגנת היעילה ביותר. זהו המחיר שהחברה צריכה לשלם, במונחים של תועלת כלכלית, תמורת הדרישה להגינות. המשמעות של "חסם עליון על מחיר ההגינות" היא, שהחברה לא צריכה לשלם הרבה. כלומר, אנחנו לא מפסידים הרבה מהדרישה להגינות. בעבודות קודמות נמצאו חסמים למחיר ההגינות בעוגה חד-ממדית; בעבודה זו אנחנו משתמשים באלגוריתמים לחלוקה-מחדש כדי למצוא חסמים עליונים למחיר ההגינות בעוגה דו-ממדית, כמו קרקע.

3. חלוקה למשפחות. כשמחלקים עוגה, כל פרוסה נאכלת על-ידי אדם אחד, ולכן ההגינות של החלוקה נקבעת רק על-ידי הטעם האישי של כל שחקן. אולם כשמחלקים קרקע, במקרים רבים כל חלקת-אדמה שייכת לקבוצה של אנשים, כגון: משפחה או מדינה. לכל אדם בקבוצה ישנן העדפות שונות. בעבודה זו אנו מציגים תהליכים לחלוקת עוגה בין קבוצות, המתייחסים להעדפות השונות של חברי הקבוצה.

ההרחבות השונות למודל חלוקת העוגה עשויות להיות שימושיות גם במצבים נוספים, פרט לחלוקת קרקע. לדוגמה: מגבלות גיאומטריות עשויות להיות שימושיות בחלוקת שטחי-פרסום דו-ממדיים בין מפרסמים, או שטחים בתערוכה בין מציגים. חלוקה מחדש מתרחשת גם במערכות מיחשוב-ענן, שבהם יש לחלק משאבי-מיחשוב לתהליכים חדשים תוך מיזעור הפגיעה בתהליכים הקיימים. בעלות משפחתית יש לא רק בקרקע אלא גם במשאבים אחרים השייכים למשפחה. לכן, העבודה הנוכחית עשויה לתרום למחקר בחלוקה הוגנת באופן כללי, מעבר לתרומה הישירה בנושא חלוקת קרקעות.

תקציר

עבודה זו מציגה תהליכים ואלגוריתמים לחלוקה הוגנת של קרקעות בין אנשים עם העדפות שונות.

חלוקה הוגנת של קרקעות יכולה להיות שימושית במספר מצבים: חלוקת נחלה משפחתית בין יורשים, חלוקת רכוש בין שותפים, חלוקת קרקעות חקלאיות במושב בין החברים, חלוקת קרקעות ציבוריות בין אזרחים ועוד. חלוקה הוגנת של קרקעות נזכרת כבר בתורה. בני ישראל מצווים לחלק את ארץ ישראל ביניהם בהגינות (במדבר כו), ומצווה זו חוזרת גם בחזונות הנבואיים לאחרית הימים (יחזקאל מז). חלוקה הוגנת של קרקעות רלבנטית גם בימינו כל אדם צריך קרקע כדי להקים עליה את ביתו. ככל שהאוכלוסיה גדלה, הקרקע נעשית יקרה יותר, מחירי הדיור עולים, ולצעירים (ואני ביניהם) קשה יותר להגיע לבית משלהם.

איך אפשר לחלק קרקע בצורה הוגנת? קרקע היא הרי משאב הטרוגני: יש בה הרים ועמקים, עצים וחולות, בארות נפט וחופי ים. לכל אחד מהאנשים המשתתפים בחלוקה עשויה להיות דעה אחרת לגבי השווי של כל איזור בקרקע. האם אפשר לחלק את הקרקע באופן שכל המשתתפים יסכימו שהם קיבלו חלק הוגן?

מפתיע לגלות שהתשובה היא כן! כדי להבין איך זה ייתכן, נניח שיש שני אנשים הרוצים לחלק ביניהם חלקת-אדמה. אנחנו מציעים להם להשתמש בתהליך הבא:

אחד מחלק. השני בוחר.

התהליך הזה מוכר לכל מי שניסה לחלק עוגה בין שני ילדים עם טעמים שונים. הוא גם נרמז בתנך, כשאברם ולוט מחלקים ביניהם את אזורי המרעה בארץ כנען (בראשית יג). למרות פשטותו הרבה, התהליך הזה הוא "הוגן" בשני מובנים:

א. כל שחקן יכול להבטיח לעצמו שיקבל לפחות $1/2$ משווי הקרקע, לפי ההערכה האישית שלו. המחלק פשוט צריך לחתוך לשני חלקים; הבוחר צריך רק לבחור את החלק עם השווי הגדול יותר. חלוקה המקיימת תכונה זו נקראת "חלוקה פרופורציונלית".

ב. כל שחקן יכול להבטיח לעצמו חלק טוב לפחות כמו החלק של השחקן השני. חלוקה המקיימת תכונה זו נקראת "חלוקה ללא קנאה".[1]

ההצלחה של התהליך "אחד מחלק והשני בוחר" מעוררת מייד את השאלה: מה קורה כשיש יותר משני אנשים? תשובה ראשונה לשאלה זו ניתנה בשנת 1948 במאמר של המתמטיקאי היהודי-פולני הוגו שטיינהוז. הוא פירסם אל-גוריתם, שפותח על-ידי שני תלמידיו בנך וקנסטר, המייצר חלוקה פרופורציונלית לכל מספר של אנשים. האלגוריתם מבטיח, שכל אדם שישחק לפי ההוראות, יקבל חלקה ששווה עבורו לפחות $1/n$ מהשווי הכולל, כאשר n הוא מספר האנשים.

המאמר של שטיינהוז פתח תחום-מחקר חדש שנקרא: חלוקה הוגנת. הוא פתר שאלה אחת, ועורר הרבה שאלות חדשות: איך אפשר למצוא חלוקה לכל מספר של אנשים, שהיא גם ללא קנאה? כמה פעולות צריך לבצע כדי למצוא חלוקה כזאת - כלומר, מה הסיבוכיות החישובית של האלגוריתם? מה קורה אם בקרקע ישנם חפצים שאינם ניתנים לחלוקה, כגון בניינים? מה קורה אם לקרקע יש ערך שלילי, למשל, כשהקרקע היא פיסת דשא שצריך לכסח, וכל אחד רוצה לכסח שטח כמה שיותר קטן? מה קורה כשהשחקנים נוהגים בצורה אסטרטגית ומנסים לעשות מניפולציות על המשחק? האם אפשר להשיג חלוקה הוגנת שהיא גם יעילה כלכלית (יעילה פארטו)? מה עושים אם לשחקנים שונים יש זכויות שונות על הקרקע? ועוד. המחקר בחלוקה הוגנת פעיל מאד בשנים האחרונות, והוא יוצר שיתופי-פעולה מעניינים בין מתמטיקאים, מדעני-מחשב, כלכלנים ואנשי מדעי-המדינה.[2]

המחקר העשיר בנושא חלוקה הוגנת עדיין לא ענה על כמה שאלות שהן חשובות במיוחד כשמחלקים קרקע. כשמדברים על חלוקת משאב הטרוגני, בדרך-כלל ממשילים אותו לעוגה - בעיית החלוקה ההוגנת של משאב הטרוגני נקראת בספרות "בעיית חלוקת העוגה". אבל, קרקע היא לא עוגה. יש כמה הבדלים חשובים בין קרקע לבין עוגה, המקשים עלינו להשתמש בתהליכי חלוקת-עוגה לצורך חלוקת קרקעות. כדי לחלק קרקעות, צריך לשפר את מודל ה"עוגה" כך שיתייחס לנושאים אלה.

[1] שתי התכונות הללו שקולות זו לזו כאשר ישנם רק שני אנשים וכל השטח מחולק ביניהם, אבל בדרך כלל אלו שתי תכונות עצמאיות.

[2] לפני כשנתיים עלה לאויר האתר *spliddit.org*. זהו אתר ללא כוונת רווח, שמטרתו להנגיש אלגוריתמי חלוקה הוגנת לציבור הרחב. כיום האתר מציע אלגוריתמים לחמש בעיות-חלוקה שונות. לפי עדות מנהלי האתר, עשרות אלפי אנשים כבר השתמשו בו לפתרון בעיות מהחיים.

6.5 הגינות דמוקרטית 204
6.5.1 שתי משפחות: תהליך חלוקה 204
6.5.2 שלוש משפחות או יותר: תוצאת אי-אפשרות 205
6.5.3 שלוש משפחות או יותר: תוצאה חיובית 207
6.6 מסקנות ורעיונות להמשך מחקר 209
6.6.1 שאלות פתוחות 209
6.6.2 תנאי-הגינות חלופיים 210
6.7 תודות 211
ביבליוגרפיה 212

תקציר עברי א

4 חלוקה גיאומטרית ללא קנאה 101
4.1 מבוא 102
4.1.1 תוצאות 105
4.1.2 סקר ספרות 107
4.2 הגדרות 109
4.3 מושגים גיאומטריים 111
4.3.1 הפסד גיאומטרי 111
4.3.2 למת הבוחר 112
4.3.3 מספרי כיסוי ולמת הכיסוי 114
4.3.4 פונקציות סכין 115
4.3.5 הפסד גיאומטרי של פונקציית סכין 118
4.4 חלוקה ללא קנאה לשני אנשים 120
4.4.1 ריבועים ומלבנים 120
4.4.2 קוביות וארכיפלגים 121
4.4.3 מלבנים שמנים 124
4.4.4 צורות שמנות כלליות 127
4.4.5 אי-קיום-קנאה לעומת פרופורציונליות 129
4.5 חלוקה ללא קנאה ליותר משני שחקנים 131
4.5.1 האלגוריתם הקיים לעוגה חד-ממדית 131
4.5.2 סט סכינים 133
4.5.3 ריבועים ומלבנים שמנים 138
4.5.4 צורות שמנות כלליות 140
4.6 מסקנות ורעיונות להמשך מחקר 144
4.7 תודות 145
נספח לפרק 146

5 חלוקה מחדש ומחיר ההגינות 155
5.1 מבוא 156
5.1.1 תוצאות 157
5.1.2 סקר ספרות 163
5.2 הגדרות 166
5.2.1 חלוקת-עוגה מחדש 167
5.2.2 רווחה חברתית ומחיר ההגינות 168
5.3 עוגות כלליות ופרוסות כלליות 169
5.4 עוגה שהיא קטע ופרוסות שהן קטע 171
5.5 עוגה שהיא מצולע ופרוסות שהן מצולעים 174
5.6 חסמים על מחיר ההגינות 176
5.7 מסקנות ורעיונות להמשך מחקר 180
5.8 תודות 183
נספח לפרק 184

6 חלוקה בין משפחות 187
6.1 מבוא 188
6.1.1 תוצאות 188
6.1.2 סקר ספרות 190
6.2 הגדרות 192
6.2.1 משפחות וזכויות 192
6.2.2 הקצאות ורכיבי-קשירות 193
6.2.3 שלושה קריטריונים להגינות 194
6.3 הגינות ממוצעת 196
6.4 הגינות פה-אחד 198
6.4.1 חלוקה מדוייקת 198
6.4.2 מחלוקה פה-אחד לחלוקה מדוייקת 199
6.4.3 מחלוקה מדוייקת לחלוקה פה-אחד 200
6.4.4 זכויות שוות 201
6.4.5 זכויות שונות 203

תוכן העניינים

תקציר i

1 מבוא 1
1.1 קרקע 1
1.2 חלוקה הוגנת 2
1.3 חלוקה הוגנת של קרקע 4
1.4 סקר ספרות 6

2 הגדרות 10
2.1 עוגה 10
2.2 שחקנים 10
2.3 שאילתות 11
2.4 חלוקות 12
2.5 הגינות 13
2.6 גיאומטריה 14

3 חלוקה גיאומטרית פרופורציונלית 17
3.1 מבוא 17
3.1.1 תוצאות 20
3.1.2 טכניקות 24
3.1.3 סקר ספרות 25
3.2 הגדרות 27
3.3 תוצאות אי-אפשרות 29
3.3.1 תוצאות אי-אפשרות לשניים, שלושה וארבעה קירות 30
3.3.2 תוצאות אי-אפשרות לקיר אחד 33
3.3.3 תוצאות אי-אפשרות למישור בלי קירות 34
3.3.4 תוצאות אי-אפשרות עם מלבנים שמנים 40
3.4 מכרזים וכיסויים 43
3.4.1 מכרז סימון 43
3.4.2 מכרז הערכה 44
3.4.3 מספרי כיסוי 46
3.5 תהליכי חלוקה 49
3.5.1 ארבעה ושלושה קירות, חתכי גיליוטינה 49
3.5.2 שני קירות 59
3.5.3 קיר אחד ואפס קירות 63
3.5.4 שלושה קירות 64
3.5.5 ארבעה קירות, חתכי גיליוטינה, העדפות זהות 73
3.5.6 עוגות קומפקטיות כלליות 82
3.6 מסקנות ורעיונות להמשך מחקר 86
3.6.1 שאלות פתוחות 86
3.6.2 אילוצים גיאומטריים מסוגים אחרים 87
3.6.3 אפליקציית רשת 88
3.7 תודות 88
נספח לפרק 90

תודות

המחקר הנוכחי הוא חלק ממסע ארוך, וזו ההזדמנות להודות למי שליווה אותי בחלקים מהמסע הזה.

קודם כל, תודה לה' המדריך אותנו תמיד, שנתן לנו תורה ומצוות. בפרט, המצווה לחלק את ארץ ישראל בצדק בין בני ישראל , הנזכרת בספר במדבר כ"ו ובספר יחזקאל מ"ז, היא ההשראה העיקרית לעבודת המחקר הנוכחית.

תודה רבה להוריי, דב ז"ל ודינה תבדל"א, שתמכו בי וחינכו אותי מילדותי לאהוב תורה ומדע. החינוך והאהבה שקיבלתי מהם מלווים אותי עד עכשיו.

הפרק הנוכחי בחיים האקדמיים שלי התחיל כששרית קראוס ועידו דגן קיבלו אותי כעוזר-מחקר במחלקה למדעי-המחשב באוניברסיטת בר-אילן. הם עזרו לי מאד להיכנס לעולם המחקר, ועודדו אותי לעבור לשלב הבא ולהתחיל במחקר לדוקטורט. למרות שעברתי לנושא אחר, הם המשיכו לעודד אותי ולתמוך בי לאורך כל הדרך. נהניתי מאד לעבוד עם חברים ושותפים במעבדות שלהם, בפרט: ינון צוקרמן, אבי רוזנפלד, אריאל רוט ואסנת דריין מהמעבדה של שרית, מני אדלר, עדן שלום ארז, עופר ברונשטיין, יונתן ברנט, אייל שנרך, אמנון לוטן ואורן מלמוד מהמעבדה של עידו.

המחקר הנוכחי התחיל כעבודה סמינריונית בקורס ניצוצות שחר - קורס רב-תחומי לתורה ומדע. מנהלי הקורס, הרב שבתאי רפפורט ורולי בלפר, עודדו אותי להמשיך את המחקר הלאה וקישרו אותי לשמואל ניצן.

שמואל ניצן, מהמחלקה לכלכלה בבר-אילן, הציג לי לראשונה את המחקר הכלכלי על חלוקה הוגנת. הוא עזר לי מאד להבין את עולם התיאוריה הכלכלית ולהשתלב בקהילת הכלכלנים, ועדיין עוזר לי רבות.

תודה מיוחדת כמובן ליונתן אומן ואבינתן חסידים, שהסכימו להיות המנחים שלי בתהליך לא שגרתי. הם עודדו אותי ועזרו לי מאד לפתח את הרעיונות הראשוניים שלי ולהפוך אותם למאמרים מחקריים, הדריכו אותי במחקרים נוספים בכלכלה חישובית ותורת המשחקים, וקישרו אותי לחוקרים מובילים בתחום. קיבלתי ואני עדיין מקבל מהם תמיכה רבה וחיונית בצעדיי הראשונים כחוקר.

המזכירות באוניברסיטת בר-אילן - הילה עטיה ודפנה גד וסילבי ברוך וסיגל שגב-זרחיה ממדעי המחשב, אורית ניסים ואפרת אוזן ומיכל גלזרמן מכלכלה - עזרו לי מאד בכל העניינים המנהליים. נחום וייל ודורון קורצברג ופינחס וייסברג וישראל ידידיה עזרו המון בענייני מיחשוב.

נהניתי להשתתף בסמינר תורת המשחקים בבר-אילן , הצגתי שם חלק מהמחקר הנוכחי וקיבלתי הערות מועילות מהמשתתפים, בפרט: יגאל מילכטייך, אייל בהרד, רון פרץ, גלעד בבלי, אמנון שרייבר, זיו הלמן, שירי אלון-עירון וגבי גייר. קיבלתי עצות מועילות גם מחברי המחלקה למתמטיקה, בפרט: רון עדין, טל נוביק, מיכאל כץ, נתן קלר, ראובן כהן ושמחה הבר.

אחרונה ואהובה ביותר - אשתי גליה. קיבלתי ממנה המון אהבה תמיכה עידוד והשראה מאז שהכרנו ועד עכשיו. היא נתנה לי רעיונות מקוריים, הקשיבה לכל מצגות-החזרה שלי והעירה הערות מועילות, עמדה בכל אתגרי האפייה שהצגתי לה והכינה עוגות נהדרות ☺

--- אראל סגל-הלוי

--- הוד השרון, סיון ה'תשע"ז

עבודה זו נעשתה בהדרכתם של

פרופ' יונתן אומן

ופרופ' אבינתן חסידים

מן המחלקה למדעי המחשב

של אוניברסיטת בר אילן.

חלוקה הוגנת של קרקע

חיבור לשם קבלת התואר "דוקטור לפילוסופיה"

מאת:

אראל סגל-הלוי

המחלקה למדעי המחשב

הוגש לסנט של אוניברסיטת בר-אילן

רמת-גן כסלו ה'תשע"ז