

Thesis Proposal:  
**Fair Division**

David Kurokawa

August 22, 2016

School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213

**Thesis Committee:**

Ariel D. Procaccia (Chair)

Manuel Blum

Ioannis Caragiannis (University of Patras)

Tuomas Sandholm

*Submitted in partial fulfillment of the requirements  
for the degree of Doctor of Philosophy.*



## **Abstract**

Fairly dividing goods among several (possibly strategic) parties is a common problem both in economics and computer science. For instance, a government may wish to divide tracts of land to various organizations or a cloud computing environment may wish to divide computing time among several users. Often, such situations are dealt with via the use of monetary transfers — such as in auctioning off the desired goods; however we are concerned with the case where such transfers are not allowed.

In this proposal we expound upon our work in both the divisible setting as well as the indivisible. We start with an examination of our work and future problems to tackle in classical envy-free cake-cutting, as well as the often ignored game-theoretic aspects of this area. We then move onto the indivisible setting and present our results on the MMS guarantee, and two new properties we coin as PMMS and EFX. The improvement of guaranteeable approximations is the main open problem regarding these properties. We end with a discussion on our applications of fair division techniques and research to real world problems: classroom allocation among schools, and the peer review process.

# Contents

- 1 Introduction** **1**
  
- 2 Divisible Goods** **3**
  - 2.1 Piecewise Uniform/Linear Valuation Functions . . . . . 3
  - 2.2 Gaming of Cake-Cutting Mechanisms . . . . . 4
  
- 3 Indivisible Goods** **7**
  - 3.1 Maximin Share (MMS) . . . . . 7
  - 3.2 Envy-Freeness Up to Any Good (EFX) and Pairwise Maximin Share Guarantee (PMMS) . . . . . 9
  
- 4 Applications of Fair Division** **11**
  - 4.1 Classroom Allocation to Charter Schools . . . . . 11
  - 4.2 Impartial Peer Review . . . . . 13
  
- 5 Open Problems** **15**
  
- Bibliography** **17**



# Chapter 1

## Introduction

Fairly dividing goods among several (possibly strategic) parties is a common problem both in economics and computer science. For instance, a government may wish to divide tracts of land to various organizations or a cloud computing environment may wish to divide computing time among several users. Often, such situations are dealt with via the use of monetary transfers — such as in auctioning off the desired goods; however, we are concerned with the case where such transfers are not allowed.

In its purest form, our setting is we have a set of agents  $\mathcal{N}$  and set of goods  $\mathcal{M}$  where each agent  $i \in \mathcal{N}$  has a function  $\text{val}_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}$  that maps sets of goods to  $i$ 's perceived value of said goods. We are then concerned with how to allocate the goods to the agents in some *fair* way. In this thesis proposal, we mainly restrict our attention to the common case where the valuations are additive. That is, for all  $i \in \mathcal{N}$  and  $A, B \subseteq \mathcal{M}$  where  $A \cap B = \emptyset$  we have  $\text{val}_i(A \cup B) = \text{val}_i(A) + \text{val}_i(B)$ .

In chapter 2, we first focus on the setting where the goods are divisible — commonly known as cake-cutting. We start by outlining the relationship of this general divisible setting in comparison to when the valuations are restricted to piecewise uniform or linear valuations and we are interested in *envy-free* allocations. We refer to our work in [15] that shows the surprising result that the seemingly strict restriction does not simplify this setting in any way, but with the introduction of a natural measure of the complexity of the valuation functions, can be solved quite easily in polynomial time. In a tangential line of our work (see [7]) we examine this setting when agents act strategically. We show in such settings that subgame perfect Nash equilibrium always do exist despite the uncountably infinite action space (thus not allowing for standard backward induction like arguments).

Chapter 3 then shifts the focus onto the setting where the goods are indivisible. Owing to the discrete nature of this setting, properties desired of an allocation are often not guaranteed. After all, how can we divide a single indivisible good among multiple agents? We examine one such property (the MMS guarantee) which was demonstrated to not always exist, but only for carefully constructed counterexamples (see [27]). We improve the results in this paper and further demonstrate that under a natural randomized model, the desired property exists with high probability (see [18]). We also examine two new properties briefly mentioned in our related paper [9]: PMMS and EFX.

Chapter 4 follows with a discussion on applications of fair division theory and its ideas to real world problems. Specifically, we summarize our work relating to classroom allocations in a

large Californian school district [17], and our work in the peer review process [16].

Chapter 5 ends our thesis with a curated list of open problems we will investigate in the coming year. These questions will be mentioned in relevant parts of this proposal prior, but we arrange them here for sake of clarity and transparency of future plans.

## Chapter 2

# Divisible Goods

In this section, we are solely focussed on the setting of *divisible* goods — commonly known as cake-cutting. That is, we have the set of agents  $\mathcal{N}$  and a cake  $\mathcal{M}$  which we, as an arbiter, wish to divide among them. Though the agents may have different valuations for different parts of the cake and they may disagree amongst each other, each agent has a valuation for every subset/slice of said cake. The important and defining assumption of the divisible good setting is that each agent must believe the cake can be cut in anyway and its value is divisible.

Formally, the setting is succinctly described by the following formulation:

1.  $\mathcal{M} = [0, 1]$  (referred to as the *cake*).
2.  $\forall i \in \mathcal{N}$  we have a function  $v_i : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$  with  $\int_0^1 v_i(z) dz = 1$ .
3.  $\forall i \in \mathcal{N}$  we have  $\text{val}_i([x, y]) = \int_x^y v_i(z) dz$ .

We are then interested in the problem of dividing  $\mathcal{M} = [0, 1]$  into a *finite* number of intervals where each agent receives some collection of them. For example, if  $n = 2$ , we may give the intervals  $[0, 0.1]$  and  $[0.9, 1]$  to agent 1, and the remainder of  $(0.1, 0.9)$  to agent 2.

We begin in Section 2.1 with an exploration of constructing what are called *envy-free* (EF) allocations under restricted valuation functions. Section 2.2 follows with an analysis of agent behavior when they are allowed to lie (to increase their utility) in true game-theoretic fashion.

## 2.1 Piecewise Uniform/Linear Valuation Functions

Throughout this section, we will often concern ourselves with what are called envy-free (EF) allocations — which are known to always exist for any instance of our cake-cutting problem.

**Definition** An *allocation* of  $\mathcal{M}$  is a partition of  $\mathcal{M}$  into  $n$  sets  $A_1, \dots, A_n$ , such that for all  $i$ ,  $A_i$  is a finite set of intervals. An allocation is called *envy free* (EF) if for all  $i, j \in \mathcal{N}$  we have that  $\text{val}_i(A_i) \geq \text{val}_i(A_j)$ . That is, no agent envies another.

As such EF allocations always exist, our question is not of existence, but of efficiency. Troublingly, the uncountably infinite number of possibilities of the valuation functions dampers any hope on any sort of conventional complexity analysis.

To assuage these issues, we call upon the standard Robertson-Webb model of query complexity

first introduced in [29]<sup>1</sup>. In this paradigm, the arbiter can only communicate with the agents via the use of the following two types of queries (albeit adaptively):

1.  $\text{Evaluate}(i, a, b)$ : returns  $\text{val}_i([a, b])$
2.  $\text{Cut}(i, a, z)$ : returns the lowest  $b \geq a$  where  $\text{val}_i([a, b]) = z$ .

That is, the arbiter may ask an agent what an interval is worth, and he may also ask an agent to produce an interval (with a predetermined left endpoint) of a specific worth.

We are then concerned with protocols that produce EF allocations via communicating to the agents through the use of only these two queries. Moreover, we wish to have such a protocol that uses a polynomial (in  $n$ ) number of queries. [5] gives a protocol that always terminates, but unfortunately, the query complexity was not bounded by  $n$  (owing to the possible intricacies of the valuation functions). More recently, [2] showed that this can be done with a number of queries bounded by  $n$ , but the given bound is an extraordinary  $n^{n^{n^{n^n}}}$ .

Owing to the difficulty of this problem we restrict our attention to when the  $v_i$  functions are reined in to a more tractable realm. We then have the two somewhat conflicting results.

Our first result concerns that of piecewise uniform valuations. That is,  $\forall i \in \mathcal{N}$ , we have that  $v_i$  can only take on two values — one of which is zero. Intuitively, this describes a valuation function where all agents believe sections of the cake are either satisfactory or not and are indifferent among the satisfactory slices aside from the length. Although confining agent valuations to these piecewise uniform valuations may seem overly restrictive, our first result shows that this is not the case. In fact, EF cake-cutting for piecewise uniform valuations is just as hard as EF cake-cutting for general valuations, when seeking protocols that are bounded by a function of  $n$ .

**Theorem 2.1.1.** *Let  $\mathcal{A}$  be an algorithm that computes an EF allocation for  $n$  arbitrary piecewise uniform valuations in less than  $f(n)$  queries. Then  $\mathcal{A}$  can compute EF allocations in less than  $f(n)$  queries for general valuation functions.*

Now let us consider the case of piecewise linear valuations. That is,  $\forall i \in \mathcal{N}$ , we have that  $v_i$  is piecewise linear. Clearly, these valuation functions include all piecewise uniform valuations and therefore our previous theorem still applies. Surprisingly however, if we let  $k$  be the number of breakpoints<sup>2</sup> of all the piecewise linear valuations (for all the agents) then we can achieve an EF allocation with a polynomial number of queries w.r.t.  $n$  and  $k$ .

**Theorem 2.1.2.** *Algorithm 2 of [15] will produce an EF allocation and require at most  $\text{poly}(n, k)$  queries.*

An obvious question to be resolved is whether we can construct EF allocations in general with a polynomial number of queries in just  $n$ . [26] showed that this problem is in  $\Omega(n^2)$ , and as previously mentioned, [2] gave an upper bound of  $n^{n^{n^{n^n}}}$ , but this leaves much to be desired.

## 2.2 Gaming of Cake-Cutting Mechanisms

In general, cake-cutting mechanisms are often examined under an assumption of agent truthfulness. That is, agents will give true responses to queries. An important tangential line of research

<sup>1</sup>This is unrelated to the notion of query complexity in descriptive complexity theory.

<sup>2</sup>A breakpoint of a piecewise linear function is a point where the derivative of the function is undefined.

is understanding the behavior when we assume that agents may in fact lie to maximize their slice of the cake in true game-theoretic fashion.

To this end, we arrive at an immediate snag. Cake-cutting protocols are quite clearly sequential games when viewed under the lens of game-theoretic agents. However, due to the infinite nature of the action spaces the standard solution concept in such settings, the subgame perfect Nash equilibrium (SPNE), is not even guaranteed to exist. In [7] we define a class of cake-cutting protocols which we call GCC protocols that not only encompass nearly all classic cake-cutting protocols<sup>3</sup>, but are guaranteed to have approximate SPNEs.

**Theorem 2.2.1.** *For any  $n$ -agent GCC protocol  $\mathcal{P}$  with a bounded number of steps, any  $n$  valuation functions, and any  $\varepsilon > 0$ , the game induced by  $\mathcal{P}$  and the valuation functions has an  $\varepsilon$ -SPNE.*

We would hope that in fact we could ensure the existence of exact (and not approximate) SPNE, but unfortunately, there are simple and natural examples where SPNE *do not* exist. Consider the following mechanism.

1. Agent 1 makes a mark on the cake (i.e.  $[0, 1]$ ) say at  $x$ .
2. Agent 2 makes a mark on the cake say at  $y$ .
3. If  $x \leq y$ , agent 1 gets  $[0, x]$ , and agent 2 gets  $(x, 1]$ . Otherwise, agent 1 gets  $(y, 1]$  and agent 2 gets  $[0, y]$ .

This unfortunately has no exact SPNE when (for example)  $v_2(z) = 1$ . To see this, consider when  $x = 1$ . Agent 2 will wish to set  $y$  s.t.  $y < x = 1$ , but as close to  $x$  as possible. Agent 2 will therefore not have a best response strategy as whatever  $y$  he chooses will be dominated by the choice  $(1 + y)/2$ .

In this example the key issue is that when  $x = y$  there is a “tie” which causes a one-sided discontinuity. Fortunately, it turns out that the crux of the nonexistence in this example, this issue of tie-breaking, is the only true problem in this infinite action space GCC paradigm. We therefore arrive at the following theorem that skirts this issue by allowing the arbiter to break ties in an informed manner.

**Theorem 2.2.2.** *For any  $n$ -agent GCC protocol  $\mathcal{P}$  with a bounded number of steps and any  $n$  valuation functions, there exists an informed<sup>4</sup> GCC protocol  $\mathcal{P}'$  that is equivalent to  $\mathcal{P}$  up to tie-breaking, such that the game induced by  $\mathcal{P}'$  and the valuation functions has an exact SPNE.*

This work largely conclusively answers questions on the existence of cake-cutting protocols, but the generalization of these results to other infinite action space and infinite horizon settings is one area of further exploration.

<sup>3</sup>The class does not include the Brams-Taylor protocol given in [5], due to its unbounded nature. As [5] gave an analysis of its equilibrium properties, [7] does not examine it further.

<sup>4</sup>Intuitively, an informed GCC protocol is one where the arbiter in the case of “ties” may choose to break said ties in any way he chooses. The exact definition of this concept is rather involved and we therefore refer the reader to [7] for the precise definition.



## Chapter 3

# Indivisible Goods

In this section, we are interested in the fair allocation of *indivisible* goods. The divisible good setting (i.e. cake-cutting) is a nice metaphor for real-world problems such as land division; the study of cake-cutting distills insights about fairness that are useful in related settings, such as the allocation of computational resources [11, 24, 14, 25]. However, typical real-world situations where fairness is a chief concern, notably divorce settlements and the division of an estate between heirs, involve *indivisible goods* (e.g., houses, cars, and works of art) — which in general preclude envy-free, or even proportional, allocations. As a simple example, if there are several agents and only one indivisible item to be allocated, the allocation cannot possibly be proportional or envy free. Even more vexingly, we note that no allocation can be even *approximately* (in a multiplicative sense) fair according to these notions, because some agents receive an empty allocation of zero value.

Whereas the divisible setting is largely focussed on the property of envy-freeness, this problematic nature of the indivisible setting forces a more fundamental question to the forefront: what is fair? In Section 3.1 we examine a popular axiomatic notion of fairness: maximin share (MMS) guarantee. Section 3.2 follows with two alternates: PMMS and EFX.

### 3.1 Maximin Share (MMS)

Due to the discreteness of the indivisible good setting, the elegant and encompassing results of the divisible good (i.e. cake-cutting) literature do not transfer well. Notably, in the indivisible setting, an envy-free, let alone proportional allocation may not even exist. Take for instance the simple case where there is a single good to be divided among multiple agents. As all but one agent will be left with zero utility, there is no hope of even assuring an approximation to the properties of envy-freeness or its weaker cousin of proportionality.

Introduced by Budish in [8], the idea of the maximin share was in some sense introduced to remedy this gap of axiomatic guarantees. An agent  $i \in \mathcal{N}$ 's MMS value is defined as:

$$\text{MMS}_i(\mathcal{N}) := \max_{A_1, \dots, A_n} \min_j \text{val}_i(A_j).$$

Intuitively, if  $i$  is asked to partition the goods into  $n$  bundles and given the least valuable bundle then  $i$ 's MMS value is the value he can guarantee himself by producing an optimal partition. The

MMS guarantee is the fairness property that every agent achieves at *least* his MMS value.

Somewhat motivated by its natural generalization of proportionality to the indivisible setting, the concept is one of few properties discussed in the cake-cutting literature. Unfortunately, although it does always exist when  $n = 2$ <sup>1</sup>, we now know that the MMS guarantee need not exist when  $n \geq 3$ , even for a small number of goods.

**Theorem 3.1.1.** *For any set of agents  $\mathcal{N}$  such that  $n \geq 3$  there exist a set of items  $\mathcal{M}$  of size  $m \leq 3n + 4$ , and (additive) valuation functions, that do not admit an MMS allocation.*

Although this lack of guarantee may seem quite problematic, we also know that the property does admit a multiplicative approximation. Let  $\gamma_n := \frac{2\lfloor n \rfloor_{\text{odd}}}{3\lfloor n \rfloor_{\text{odd}} - 1}$ , where  $\lfloor n \rfloor_{\text{odd}}$  is the largest odd number that is smaller or equal to  $n$ . Then we have the following.

**Theorem 3.1.2.** *There always exists an allocation  $A_1, \dots, A_n$  such that for all  $i \in \mathcal{N}$ ,  $\text{val}_i(A_i) \geq \gamma_n \text{MMS}_i(n, \mathcal{M})$ . Moreover, for every  $\varepsilon > 0$ , an allocation  $A_1, \dots, A_n$  such that for all  $i \in \mathcal{N}$ ,  $\text{val}_i(A_i) \geq (1 - \varepsilon)\gamma_n \text{MMS}_i(n, \mathcal{M})$  can be computed in polynomial time in  $n$  and  $m$ .*

Note that  $\gamma_n$  is always greater than  $2/3$  (and it is equal to  $3/4$  for the important cases of three and four agents — i.e.  $n \in \{3, 4\}$ ). Theorem 3.1.2 therefore asserts that a  $2/3$  approximation always exists.

This would largely complete the picture of MMS guarantee existence, if it were not for the fact that all proofs of Theorem 3.1.1 explicitly construct instances where the MMS guarantee fails, but only by a very small fraction (approximately 0.99999). In addition, randomly generated instances did not contain any counterexamples of MMS guarantee existence [4]. We formalize these observations by considering the regime where for each  $i \in \mathcal{N}$  there is a distribution  $\mathcal{D}_i$  such that the values  $\text{val}_i(g)$  are drawn independently from  $\mathcal{D}_i$ .

**Theorem 3.1.3.** *Assume that for all  $i \in \mathcal{N}$ ,  $\mathbb{V}[\mathcal{D}_i] \geq c$  for a constant  $c > 0$ . Then for all  $\varepsilon > 0$  there exists  $K = K(c, \varepsilon)$  such that if  $\max(n, m) \geq K$ , then the probability that an MMS allocation exists is at least  $1 - \varepsilon$ .*

In words, an MMS allocation exists with high probability as the number of agents *or* the number of goods goes to infinity. It was previously known that an envy-free allocation (and, hence, an MMS allocation) exists with high probability when  $m \in \Omega(n \ln n)$  [10]. The analysis in [18] therefore focuses on the case of  $m \in O(n \ln n)$ . In this case, an envy-free allocation is unlikely to exist (such an allocation certainly does not exist when  $m < n$ ), but (as we show in [18]) the existence of an MMS allocation is still likely. Specifically, we develop an allocation algorithm and show that it finds an MMS allocation with high probability. The algorithm's design and analysis leverage techniques for matching in random bipartite graphs.

Natural questions and lines of research thus emerge out of the confusing woodwork. For example, what is the highest guaranteeable approximation? After all, if the instances where the MMS guarantee could not be assured still admit extremely high approximations, practically, the non-existence results are negligible. If higher MMS guarantee approximations do always exist, do there exist polynomial algorithms computing such allocations?

<sup>1</sup>A simple mechanism that achieves the MMS guarantee for  $n = 2$  is for the first agent to produce a partition of  $\mathcal{M}$  into two sets s.t. both sets are worth at least his MMS value (he can do this by the definition of MMS value). The second agent then receives the set that is of higher value to him while the first receives the other set. Now note that the set the second agent receives is worth at least  $1/2 \cdot \text{val}_i(\mathcal{M}) \geq$  his MMS value.

## 3.2 Envy-Freeness Up to Any Good (EFX) and Pairwise Maximin Share Guarantee (PMMS)

With the problematic unguaranteeability of the MMS guarantee, a natural avenue of exploration is the search for other sensible properties. In this section we introduce two: pairwise maximin share guarantee (PMMS) and envy-freeness up to any good (EFX). They are defined as follows.

**Definition** An allocation  $A_1, \dots, A_n$  is PMMS if for every  $i, j \in \mathcal{N}$  we have that

$$\text{val}_i(A_i) \geq \max_{B,C \text{ is a partition of } A_i \cup A_j} \min(\text{val}_i(B), \text{val}_i(C)).$$

In other words, for every pair of agents, each agent believes that he achieves his MMS value when restricted only to those two agents and the goods given to them.

**Definition** An allocation  $A_1, \dots, A_n$  is EFX if for every  $i, j \in \mathcal{N}$  we have at least one of:

- $A_j = \emptyset$ .
- $\text{val}_i(A_i) \geq \text{val}_i(A_j) - \min_{g \in A_j} \text{val}_i(\{g\})$ .

That is, no agent envies any strict subset of another's allocation.

Both are generalizations of the concept of envy-freeness in the divisible setting in that they imply envy-freeness in the cake-cutting setting — analogously to how the MMS guarantee is a generalization of proportionality. Though PMMS is a strictly stronger property than EFX, the latter has its merits in that it is a natural property in itself. For one, the concept itself strengthens the popular notion of envy-freeness up to one good (EF1):

**Definition** An allocation  $A_1, \dots, A_n$  is EF1 if for every  $i, j \in \mathcal{N}$  we have at least one of:

- $A_j = \emptyset$ .
- $\text{val}_i(A_i) \geq \text{val}_i(A_j) - \max_{g \in A_j} \text{val}_i(\{g\})$ .

The only difference between EFX and EF1 being that EFX considers the least valued element, whereas EF1 considers the highest valued one. The hierarchy of property strength is therefore:  $\text{PMMS} \Rightarrow \text{EFX} \Rightarrow \text{EF1}$ .

Another compelling strength of EFX as a fairness property is that it is a natural envy-freeness like property that can be verified in polynomial time (and in fact linear time). This is notable because to check whether an agent has achieved his MMS value in an allocation or is in a PMMS allocation is an  $\mathcal{NP}$ -hard problem. Arguments for the unrealistic nature of MMS and PMMS depending on validation complexity issues therefore hold no weight in the criticism of EFX.

A pressing question with both properties is the one of: can it be guaranteed? This is an open question we are currently exploring. Similarly to the MMS guarantee, no constructions where the stronger property of PMMS cannot be guaranteed have been found neither in extensive random simulations or Spliddit data. This is to be expected, since with appropriate finesse, the randomized Theorem 3.1.3 can be adapted to give a similar result for PMMS. We do however, have the following  $\frac{4}{3+\sqrt{5}} \approx 0.76$  approximation.

**Theorem 3.2.1.** *There always exists an allocation  $A_1, \dots, A_n$  s.t. for every  $i, j \in \mathcal{N}$  we have*

$$\text{val}_i(A_i) \geq \gamma \max_{B,C \text{ is a partition of } A_i \cup A_j} \min(\text{val}_i(B), \text{val}_i(C)).$$

where  $\gamma = \frac{4}{3+\sqrt{5}} \approx 0.76$ . Moreover, this can be found in polynomial time.



## Chapter 4

# Applications of Fair Division

Over the course of the last seven decades, the study of fair division has given rise to a slew of elegant solutions to a variety of problems [6, 29, 21], which span the practicability spectrum from abstract (e.g., cake-cutting [25]) to everyday (e.g., rent division [1]). Building on its rich history, the field of fair division — and *computational* fair division, in particular — is poised to make a significant impact on society through applications that are beginning to emerge. For example, Budish’s fair division approach [8] — which leads to challenging computational questions [23, 22] — is now regularly used by the Wharton School of the University of Pennsylvania to allocate seats in MBA courses. And the not-for-profit website *Spliddit* ([www.spliddit.org](http://www.spliddit.org)) — which offers provably fair solutions for the division of goods, rent, and credit — has already been used by tens of thousands of people [12].

In this chapter we deal with two of our applications of fair division techniques and its ideas to practical scenarios. First, in Section 4.1 we introduce our work in dividing unused public school resources to charter schools in a major Californian school district. Section 4.2 follows with an examination of the peer review process under the restriction that we wish to have impartiality.

### 4.1 Classroom Allocation to Charter Schools

Our work [17] presents a solution to a real world fair division problem posed by a representative of one of the largest school districts in California. Since the details are confidential, we will refer to the school district as the Pentos Unified School District (PUSD), and to the representative as Illyrio Mopatis. Mr. Mopatis contacted us in May 2014 after learning about Spliddit (and fair division, more generally) from an article in the New York Times.<sup>1</sup> He was tasked with the allocation of unused space (most importantly, classrooms) in PUSD’s public schools to the district’s charter schools, according to California’s Proposition 39, which states that “public school facilities should be shared fairly among all public school pupils, including those in charter schools”.<sup>2</sup> While the law does not elaborate on what “fairly” means, Mr. Mopatis was motivated by the belief that a provably fair solution would certainly fit the bill. He asked us to design

<sup>1</sup><http://goo.gl/Xp3omV>

<sup>2</sup><http://goo.gl/bGH6dT>

an automated allocation method that would be evaluated by PUSD, and potentially replace the existing manual system.

To be a bit more specific, the setting consists of charter schools and facilities (public schools). Each facility has a given number of unused classrooms — its *capacity*, and each charter school has a number of required classrooms — its *demand*. In principle the classrooms required by a charter school could be split across multiple facilities, but such offers have always been declined in the past, so we assume that an agent’s demand must be satisfied in a single facility (if it is satisfied at all). Other details are less important and can be abstracted away. For example, classroom size turns out to be a nonissue, and the division of time in shared space (such as the school gym or cafeteria) can be handled *ad hoc*.

Of course, to talk of fairness we must also take into account the preferences of charter schools, but preference representation is a modeling choice, intimately related to the design and guarantees of the allocation mechanism. Moreover, fairness is not our only concern: to be used in practice, the mechanism must be relatively intuitive (so it can be explained to decision makers) and computationally feasible. The challenge we address is therefore to

... *design and implement a classroom allocation mechanism that is provably fair as well as practicable.*

In [17] we model the preferences of charter schools as being *dichotomous*: charter schools think of each facility as either acceptable or unacceptable. This choice is motivated by current practice: Under the 2015/2016 request form issued by PUSD, charter schools are essentially asked to indicate acceptable facilities (specifically, they are asked to “provide a description of the district school site and/or general geographic area in which the charter school wishes to locate” using free-form text). In other words, formally eliciting dichotomous preferences — by having charter schools select acceptable facilities from the list of all facilities — is similar to the *status quo*, a fact that increases the practicability of the approach.

A natural starting point, therefore, is the seminal paper of [3], who study the special case of our setting with unit demands and capacities, under dichotomous preferences. They propose the *leximin mechanism*, which returns a random allocation with the following intuitive property: it maximizes the lowest probability of any charter school having its demand satisfied in an acceptable facility; subject to this constraint, it maximizes the second lowest probability; and so on.

In our work we show that the leximin mechanism remains compelling in the classroom allocation setting. Specifically, we prove that it satisfies the following properties: (i) proportionality — each charter school receives its proportional share of available classrooms; (ii) envy-freeness — each charter school prefers its own allocation to the allocation of any other school; (iii) Pareto optimality (a.k.a. *ex-ante efficiency*) — no other randomized allocation is at least as good for all charter schools, and strictly better for at least one; and (iv) group strategyproofness — even coalitions of charter schools cannot benefit by misreporting their preferences. The beauty of these properties, as well as the leximin mechanism itself, is that they are intuitive and can easily be explained to a layperson. This feature, once again, significantly contributes to the practicability of the approach. As an interesting aside, we show that the leximin mechanism still satisfies the foregoing properties in a much more general setting, thereby generalizing results from a variety of other papers in fair division and mechanism design.

A critical result of [17] is that the expected number of classrooms allocated by the leximin mechanism is always at least  $1/4$  of the maximum number of classrooms that can possibly be allocated. We do not view this result as enhancing the practicability of our approach, but rather as significantly contributing to its intellectual merit.

## 4.2 Impartial Peer Review

In this section, we consider an entirely different application of fair division work: peer review. The Sensors and Sensing Systems (SSS) program of the National Science Foundation (NSF) recently experimented with a drastically different peer review method. Traditionally, grant proposals submitted to a specific program are evaluated by a panel of reviewers. Potential conflicts of interest play a crucial role in composing the panel; most importantly, principal investigators (PIs) whose proposals are being evaluated by the panel cannot serve on the panel. In stark contrast, the new peer review method — originally designed by Merrifield and Saari [19] for the review of proposals for telescope time — requires the PIs themselves to review each other’s proposals!

Under the Saari-Merrifield mechanism, each PI must review  $m$  proposals submitted by other PIs; in the NSF pilot,  $m = 7$ . The PI then ranks the  $m$  proposals according to their quality. These reviews are aggregated using the Borda count voting rule, so each PI awards  $m - i$  points to the proposal she ranks in position  $i$ . A proposal’s overall rating is the average over the points awarded by the  $m$  PIs who reviewed it. Additionally, a PI’s own proposal receives a small bonus based on the similarity between the PI’s submitted ranking and the aggregate ranking of the proposals she reviewed; this is meant to encourage PIs to make an effort to produce accurate reviews.

The NSF pilot sparked a lively debate amongst mechanism design and social choice researchers in the blogosphere [28, 30, 20]. While most researchers seem to agree that the NSF should be commended for trying out an ambitious peer review method, serious concerns were raised regarding the pilot mechanism itself. Perhaps most strikingly, while the NSF announcement [13] states that the “theoretical basis for the proposed review process lies in an area of mathematics referred to as mechanism design”, the pilot mechanism provides no theoretical guarantees. In particular, the mechanism is susceptible to strategic manipulation: PIs will often be able to advance their own proposals by giving low scores to competitive proposals (even though they may forfeit some of the small bonus for similarity to others’ reviews). Furthermore, while most researchers who sit on NSF panels are well-respected, the pilot mechanism cannot control the quality (or morality) of PIs who submit proposals (and review proposals)—leaving open the very real possibility of game-theoretic mayhem.

In [16], we alleviate these concerns by proposing a peer review mechanism which is not susceptible to such manipulations. Each PI who submits a proposal or paper will review some other PIs’ proposals or papers. Our mechanism is *impartial*: reviewers will not be able to affect the chances of their own proposals being selected. The research challenge is therefore to *design provably impartial peer review mechanisms that provide formal quality guarantees*.

In our setting there are  $n$  PIs, each associated with a proposal. Each PI  $i$  has a hypothetical (honest) evaluation of the quality of the proposal  $j$ , which is the rating  $i$  would give  $j$  if she were asked to review that proposal (and could not affect her own chances of selection). The (honest) *score* of a proposal is the average (honest) rating given to it by other PIs. As NSF program

directors, if our budget is sufficient to fund  $k$  proposals, we would ideally want to select a set of  $k$  proposals with maximum honest score.<sup>3</sup> There are two obstacles we must overcome: we cannot possibly ask each PI to review all other proposals, and the reviews may be dishonest.

To address the first problem, we consider only mechanisms which request  $m$  reviews per PI (much like the NSF pilot). We define an  $(m, k)$ -*selection mechanism* as follows. First, the mechanism asks each PI to review  $m$  proposals, in a way that each proposal is reviewed by exactly  $m$  PIs; for every such pair  $(i, j)$ , PI  $i$ 's evaluation for proposal  $j$  is revealed. Based on these elicited reviews, the mechanism selects  $k$  vertices. The most natural  $(m, k)$ -selection mechanism is an abstract version of the NSF pilot mechanism, which we fondly refer to as the VANILLA mechanism; it chooses  $m$  reviews per PI uniformly at random (subject to the constraint that each proposal is reviewed by  $m$  PIs), and then selects the  $k$  vertices with highest average rating, based only on the sampled reviews.

Returning to the second problem — dishonest reviewing — we consider only mechanisms where reviewers cannot affect their chances of being selected by misreporting their reviews. A selection mechanism is *impartial* if the probability of proposal  $i$  being selected is independent of the ratings given by PI  $i$ . The motivation for our work stems from the observation that the VANILLA mechanism is not impartial: we seek mechanisms that are.

In [16] we present an impartial  $(m, k)$ -selection mechanism, CREDIBLE SUBSET, which (usually) selects  $k$  proposals at random from a slightly larger pool (of size  $k + m$ ) of eligible proposals. We prove that CREDIBLE SUBSET gives an approximation ratio of  $\frac{k}{k+m}$  to the status quo mechanism VANILLA. We think of  $m$ , the number of reviews per PI, as being a small constant, and we would like to think of  $k$ , the number of proposals to be selected, as significantly larger. In particular, when  $m = o(k)$ , the approximation ratio goes to 1 as  $k$  goes to infinity (in an ideal world, where growth in funding outpaces growth in the quantity of work for reviewers). We also show that CREDIBLE SUBSET is the *optimal* impartial mechanism, in the sense that its approximation ratio of  $\frac{k}{k+m}$  is asymptotically tight (when  $k = m^2$  is a constant and the number of PIs  $n$  grows).

One immediate and practical extension to this work is to consider the more general case of dividing a divisible fund instead of simply  $k$  “acceptances”. As not all grants need be equal in value this more accurately models the grant application problem and introduces a complexity into the setting not captured by our techniques.

<sup>3</sup> We distill the strategic aspects of the NSF reviewing setting and abstract away some other practical aspects, such as the fact that PIs may submit multiple proposals to the same program. However, our model and results easily extend.

# Chapter 5

## Open Problems

1. *The query complexity of EF cake-cutting is known to be in  $\Omega(n^2)$  [26] and  $O\left(n^{n^{n^n}}\right)$  [2].*

*Can we close this gap to better understand the difficulty of EF cake-cutting?*

The [2] is a quite recent advancement and with its introduction of new ideas is likely to spur more advancements in the field. We propose to examine this work in depth in the hopes of improving the results.

2. *Can we extend our results in [7] on approximate and exact SPNE to more general settings with infinite action spaces and/or infinite horizons?*

The approach used in [7] to prove the exact SPNE existence results were far more general than the setting entailed. Finding interesting and useful generalizations or applications of the approaches could certainly answer similar questions for other related settings.

3. *We know that a 2/3 approximation to the MMS guarantee is guaranteed, but can we increase this approximation factor?*

This was originally proven in [27] but we recently have found an entirely different proof. Comparing these proofs may help us understand how to increase the approximation factor (by leveraging their respective advantages) or find examples where we cannot do much better than this 2/3 bound.

4. *Can we prove/disprove that PMMS and/or EFX allocations are guaranteed to exist or increase the approximation bound of  $\frac{4}{3+\sqrt{5}} \approx 0.76$ ?*

Experimental results as well as probabilistic theorems indicate that PMMS and EFX allocations almost certainly exist — suggesting that this bound can be improved (admittedly, this was true of MMS guarantee allocations as well). Moreover, as the approximation result is quite recent, further investigation is likely to reap improvements. The extreme naturalness of EFX along with a dearth of fairness properties in the indivisible good setting also raise the importance of this question greatly.

5. *Budish's introduction of the MMS guarantee property in [8] referred to a slightly different property than explored in Chapter 3. He considered the MMS value of an agent  $i$  to be:*

$$\text{MMS}_i(\mathcal{N}) := \max_{A_1, \dots, A_n, A_{n+1}} \min_j \text{val}_i(A_j).$$

*That is,  $i$ 's MMS value is the optimal value of the least valuable bundle when  $i$  partitions the*

*goods into  $n + 1$  goods (instead of simply  $n$ ). Can this property be guaranteed?*

As mentioned, instances that demonstrate the unguaranteeability of our definition of MMS rely on extraordinarily pathological and delicate constructions. This suggests that the  $n + 1$  nature of Budish's original definition of MMS will be guaranteeable. Similarly to the PMMS and EFX questions, this question is of raised importance due to fairness property scarcity.

6. *Can we extend our impartial peer review work in [16] to settings with varying levels of reward? For example, the value which a grant application gets is not a set value but differs per application.*

This is a natural extension to our work and an important one practically. Furthermore, the area offers many interesting and important questions if we consider different peer review settings and thus has a potential rich bank of research directions.

7. *Can we find further instances of fair division problems in the real world to tackle?*

The application of our work, and the burgeoning fair division research in general, is a key societal contribution researchers are poised to make in the coming years. Ariel Procaccia, has thrown massive effort into this front (such as in the website Spliddit.org) and this is an avenue that can be further reaped for important problems to investigate and settings for application.

# Bibliography

- [1] A. Abdulkadiroğlu, T. Sönmez, and M. U. Ünver. “Room assignment-rent division: A market approach”. In: *Social Choice and Welfare* 22.3 (2004), pages 515–538 (page 11).
- [2] H. Aziz and S. Mackenzie. *A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents*. CoRR abs/1604.03655. 2016 (pages 4, 15).
- [3] A. Bogomolnaia and H. Moulin. “Random Matching under Dichotomous preferences”. In: *Econometrica* 72 (2004), pages 257–279 (page 12).
- [4] S. Bouveret and M. Lemaître. “Characterizing conflicts in fair division of indivisible goods using a scale of criteria”. In: *13th*. 2014, pages 1321–1328 (page 8).
- [5] S. J. Brams and A. D. Taylor. “An Envy-Free Cake Division Protocol”. In: *The American Mathematical Monthly* 102.1 (1995), pages 9–18 (pages 4, 5).
- [6] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996 (page 11).
- [7] S. Brânzei, I. Caragiannis, and D. Kurokawa A. D. Procaccia. *An Algorithmic Framework for Strategic Fair Division*. 2016 (pages 1, 5, 15).
- [8] E. Budish. “The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes”. In: *Journal of Political Economy* 119.6 (2011), pages 1061–1103 (pages 7, 11, 15).
- [9] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. “The Unreasonable Fairness of Maximum Nash Welfare”. In: *17th*. 2016, pages 305–322 (page 1).
- [10] J. P. Dickerson, J. Goldman, J. Karp, A. D. Procaccia, and T. Sandholm. “The Computational Rise and Fall of Fairness”. In: *28th*. 2014, pages 1405–1411 (page 8).
- [11] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. “Dominant Resource Fairness: Fair Allocation of Multiple Resource Types”. In: *8th*. 2011, pages 24–37 (page 7).
- [12] J. Goldman and A. D. Procaccia. “Spliddit: Unleashing Fair Division Algorithms”. In: *SIGecom Exchanges* 13.2 (2014), pages 41–46 (page 11).
- [13] G. A. Hazelrigg. *Dear Colleague Letter: Information to Principal Investigators (PIs) Planning to Submit Proposals to the Sensors and Sensing Systems (SSS) Program October 1, 2013, Deadline*. Retrieved on June 17, 2014. 2013. URL: [http://www.nsf.gov/pubs/2013/nsf13096/nsf13096.jsp?WT.mc\\_id=USNSF\\_25#reference1](http://www.nsf.gov/pubs/2013/nsf13096/nsf13096.jsp?WT.mc_id=USNSF_25#reference1) (page 13).
- [14] I. Kash, A. D. Procaccia, and N. Shah. “No Agent Left Behind: Dynamic Fair Division of Multiple Resources”. In: *Journal of Artificial Intelligence Research* 51 (2014), pages 579–603 (page 7).
- [15] D. Kurokawa, J. K. Lai, and A. D. Procaccia. “How to cut a cake before the party ends”. In: *27th*. 2013, pages 555–561 (pages 1, 4).
- [16] D. Kurokawa, O. Lev, J. Morgenstern, and A. D. Procaccia. “Impartial Peer Review”. In: *24th*. 2015, pages 582–588 (pages 2, 13, 14, 16).

- [17] D. Kurokawa, A. D. Procaccia, and N. Shah. “Leximin Allocations in the Real World”. In: *16th*. 2015, pages 345–362 (pages 2, 11–13).
- [18] D. Kurokawa, A. D. Procaccia, and J. Wang. “When can the maximin share guarantee be guaranteed?” In: *30th*. 2016, pages 523–529 (pages 1, 8).
- [19] M. Merrifield and D. Saari. “Telescope Time Without Tears: a Distributed Approach to Peer Review”. In: *Astronomy and Geophysics* 50.4 (2009), pages 2–6 (page 13).
- [20] M. Mitzenmacher. *NSF Reviewing Trial Run*. Retrieved on June 17, 2014. 2013. URL: <http://mybiasedcoin.blogspot.com/2013/06/nsf-reviewing-trial-run.html> (page 13).
- [21] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003 (page 11).
- [22] A. Othman, C. H. Papadimitriou, and A. Rubinstein. “The complexity of fairness through equilibrium”. In: *15th*. 2014, pages 209–226 (page 11).
- [23] A. Othman, T. Sandholm, and E. Budish. “Finding approximate competitive equilibria: Efficient and fair course allocation”. In: *9th*. 2010, pages 873–880 (page 11).
- [24] D. C. Parkes, A. D. Procaccia, and N. Shah. “Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities”. In: *13th*. 2012, pages 808–825 (page 7).
- [25] A. D. Procaccia. “Cake Cutting: Not Just Child’s Play”. In: *Communications of the ACM* 56.7 (2013), pages 78–87 (pages 7, 11).
- [26] A. D. Procaccia. “Thou Shalt Covet Thy Neighbor’s Cake”. In: *21st*. 2009, pages 239–244 (pages 4, 15).
- [27] A. D. Procaccia and J. Wang. “Fair Enough: Guaranteeing Approximate Maximin Shares”. In: *14th*. 2014, pages 675–692 (pages 1, 15).
- [28] Ariel D. Procaccia. *NSF (actually) reviewing via social choice*. Retrieved on June 17, 2014. 2013. URL: <http://agtb.wordpress.com/2013/06/10/nsf-actually-reviewing-via-social-choice/> (page 13).
- [29] J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters, 1998 (pages 4, 11).
- [30] R. V. Vohra. *A Mechanism Design Approach to Peer Review*. Retrieved on June 17, 2014. 2013. URL: <http://theoryclass.wordpress.com/2013/06/06/a-mechanism-design-approach-to-peer-review/> (page 13).