On the tradeoff among efficiency, fairness and revenue in resource allocation

Hua, Xia

2012

Hua, X. (2012). On the tradeoff among efficiency, fairness and revenue in resource allocation. Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/50746>

<https://doi.org/10.32657/10356/50746>

NANYANG
TECHNOLOGICAL
UNIVERSITY

**ON THE TRADEOFF AMONG EFFICIENCY,
FAIRNESS AND REVENUE IN RESOURCE
ALLOCATION**

HUA XIA

School of Physical and Mathematical Sciences

2012

ON THE TRADEOFF AMONG EFFICIENCY, FAIRNESS AND
REVENUE IN RESOURCE ALLOCATION

HUA XIA

School of Physical and Mathematical Sciences

A thesis submitted to the Nanyang Technological University
in partial fulfillment of the requirement for the degree of
Doctor of Philosophy

2012

“We can only see a short distance ahead, but we can see plenty there that needs to be done.”

Alan Turing

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Assistant Professor Chen Ning, for his instructive guidance and helpful suggestions. His enthusiasm on research, creativity and excellent teaching skills have been a unfailing source of inspiration to my research undertakings. This thesis would not have been possible without his supervision.

I wish to give my special thanks to Associate Professor Wu Guohua. Since early in my undergraduate career in NTU, he gave me plenty useful advices and clear guidance, which led me into the field of computer science. He had a great impact on my development in academic studies.

I'm also pleased to thank Assistant Professor Edith Elkind, Associate Professor Sinai Robins, Associate Professor Wang Huaxiong and Professor Xing Chaoping for their excellent graduate lectures and seminars.

I would like to show my gratitude to John Augustine, Bei Xiaohui, Zhang Hongyang, Tao Biaoshuai and Yang Endong, for their insightful discussions and suggestions throughout the research projects. I'm also benefited from interaction with many of the faculty members and visitors here, including Radhika Arava, Angelo Fanelli, Fang Ning, Nikolay Gravin, Keshave Rao Kini, Wang Kejun, Li Mengling, Svetlana Obratzsova, Wu Weiwei, Le Truc Viet, Yu Lan, Shiryayev Dmitry Yurievitch, Zhao Kaifeng and Yair Zick. I am grateful for their friendship and help. I thank the Graduate Studies Office and the Office of Finance of NTU for providing me such a good chance to pursue my PHD degree in NTU with financial support.

Last but not least, I want to devote my sincere appreciation to my parents, for their consideration, trust and support. I also would like to thank my families and friends including He Yikai, Yin Xiang, Shelly Chen, Fang Wangxi, Huang Yuli, Li Chao, Liu Yi and Yang Fang.

To My Beloved Parents

Contents

ACKNOWLEDGEMENTS	iv
SUMMARY	ix
1 Introduction	1
1.1 Economic Efficiency	5
1.2 Fair division	9
1.3 Revenue	13
1.4 Problem Studied in This Thesis	15
2 Indivisible Resource Allocation with Payment	20
2.1 Introduction	20
2.2 Model and Solution Concepts	25
2.3 Unit Demand Consumers	27
2.3.1 Pareto Efficiency and Social Welfare	28
2.3.2 Determining Pareto Efficiency	34
2.3.3 Complexity of Computing REP	35
2.4 Single-Minded Consumers	38
2.5 Concluding Remarks	43
3 Indivisible Resource Allocation without Payment	44
3.1 Introduction	44
3.1.1 Related Work	48
3.2 Indivisible-Item Allocation	50

3.2.1	Existence	51
3.2.2	Impossibility	58
3.2.3	Polynomial Time Algorithms	59
3.2.4	Heterogeneous Agents	61
3.3	Multiple Fractional Knapsack Problem	62
3.3.1	Tight Ratios	63
3.3.2	Balanced Knapsacks	66
3.3.3	Polynomial Time Computation	70
3.4	Concluding Remarks	71
4	Divisible Resource Allocation without Payment	73
4.1	Introduction	73
4.2	Preliminaries	75
4.3	Linear Functions	76
4.3.1	Optimal Order	76
4.3.2	Polynomial Time Approximation Scheme	79
4.4	Piecewise Constant Functions	81
4.4.1	Fixed Order	81
4.4.2	NP-Hardness	83
4.4.3	Normalized Valuations	87
4.5	Concluding Remarks	90
5	Conclusions	91
	PUBLICATIONS RELATED TO THIS THESIS	94
	BIBLIOGRAPHY	95

Summary

Resource allocation is a fundamental problem studied extensively in economics, social sciences and computer science, due to its vast applications. Our focus is to consider different solution concepts that capture different perspectives in the market. These solution concepts generally have three factors: *fairness* from the consumers' perspective, *efficiency* from the centralized authority's global point of view, and *revenue* from the market maker's perspective (if we take payment into consideration).

In the thesis we will consider a number of problems raised in the market and try to capture fairness, efficiency and revenue in different circumstances. First, in the centralized combinatorial market, taking market maker's perspective into consideration, we study a solution concept called Envy-free Pareto-efficient pricing that strikes a balance among all the three factors. It captures fairness and balances the tradeoff between efficiency and revenue. Specifically, we consider Envy-free Pareto-efficient pricing in two domains, unit demand and single-minded consumers, and analyze its existence, computation, and economic properties.

Secondly, we address the following question to capture the tradeoff relation between fairness and efficiency: For *any* given constants r_f and r_w , does there exist an allocation that simultaneously achieves fairness approximation ratio r_f and welfare approximation ratio r_w ? We considered two classic resource allocation models under this question: the indivisible-item allocation with subadditive functions and the multiple fractional knapsack allocation with capacities and demands. For both models, we give (nearly) optimal constant bounds on fairness and welfare bifactor approximations regardless of computational constraints.

Furthermore, in the cake cutting model, we study the computational complexity of computing an efficiency optimal division given that the allocation satisfies proportional fairness

and assigns each agent a connected piece. For linear valuation functions, we provide a polynomial time approximation scheme to compute an efficiency optimal allocation. On the other hand, we show that the problem is NP-hard to approximate within a factor of $\Omega(\frac{1}{\sqrt{n}})$ for general piecewise constant valuation functions.

Chapter 1

Introduction

Because of its vast applications, resource allocation is a fundamental problem studied extensively in economics, social sciences and computer science. In resource allocation problems, we are given a fixed amount of resources, and a set of n self-interested consumers. We need to determine an allocation of these scarce resources to consumers such that a certain objective function is optimized. At the same time, we have only limited computational power in the system. This kind of problem exists in many aspects of the industrial process, such as bandwidth distribution, production planning, computer scheduling and portfolio selection. In order to show the ubiquity of resource allocation problems and to give readers a brief idea of the problems this thesis considers, we present a few examples below.

Example 1.1 (Bandwidth allocation: social welfare). *Imagine that you have a non-negatively weighted Internet network $G = (V, E, w)$, where the weight function $w(\cdot) : E \rightarrow \mathbb{R}^+ \cup \{0\}$ denotes the capacity of each edge. Assume the graph is connected. There are n consumers who want to transfer data between $(s_1, t_1), \dots, (s_n, t_n)$ along fixed paths $P_1, \dots, P_n$. The data cannot go in or go out at other vertices, and in each edge, the total data-transmitting ratio cannot exceed its capacity. Our target is to assign a data-transmitting ratio $X = \{x_1, x_2, \dots, x_n\}$ to each path so that the total data-transmitting ratio $\sum_i x_i$ in the network is maximized.*

This is the well-known bandwidth allocation problem. The problem itself as well as its variations have been widely studied in order to enable more efficient Internet communication [53, 83–87, 90, 98]. We can see that our target is to maximize the total data-transmitting ratio of all participants in the market. It is quite natural for a centralized authority that is in

charge of resource allocation to set such a target because it maximizes the total well-being (i.e. *social welfare* or *efficiency*) of the market. Now imagine that you have a simple network, as in the Figure 1.1. Suppose each of the edges has capacity 1 and we have $n + 1$

Figure 1.1: Bandwidth allocation example.

consumers: If our target is to maximize the total data-transmitting ratio in the network, that means that we must allocate a data-transmitting ratio 1 for consumers $1, \dots, n$, while consumer $n + 1$ needs to tolerate a zero-transmitting ratio. This may be *unfair* to her. This result means that the maximization of the total data-transmitting ratio may not be good for everybody in the market. Now we reset our target of optimization as the following:

Example 1.2 (Bandwidth allocation: fairness). *Imagine that you have a non-negatively weighted Internet network $G = (V, E, w)$, where the weight function $w(\cdot) : E \rightarrow \mathbb{R}^+ \cup \{0\}$ denotes the capacity of each edge. Assume the graph is connected. There are n consumers who want to transfer data between $(s_1, t_1), \dots, (s_n, t_n)$ along fixed paths $P_1, \dots, P_n$. We still have the constraints that the data cannot go in or go out at other vertices, and in each edge, the total data-transmitting ratio cannot exceed its capacity. Our target is to assign a data-transmitting ratio $X = \{x_1, x_2, \dots, x_n\}$ to each path so that the minimum transmitting ratio $\min_i x_i$ is maximized.*

In this case, the objective has been changed. Maximizing the minimum transmitting ratio may make everybody feel that she is treated *fairly* within the network. In this case, $x_i = 1/2$ for each of $1 \leq i \leq n + 1$. However, the total transmitting ratio inside the network will be significantly less than the optimal solution contained in Example 1.1 (about $1/2$ of the total transmitting ratio of Example 1.1).

Now we reset our objective again:

Example 1.3 (Bandwidth allocation: revenue). *Imagine that you have a non-negatively weighted Internet network $G = (V, E, w)$, where the weight function $w(\cdot) : E \rightarrow \mathbb{R}^+ \cup \{0\}$*

denotes the capacity of each edge. Assume the graph is connected. There are n consumers who want to transfer data between $(s_1, t_1), \dots, (s_n, t_n)$ along fixed paths $P_1, \dots, P_n$. We still have the constraints that the data cannot go in or go out at other vertices, and in each edge, the total data-transmitting ratio cannot exceed its capacity. A bandwidth allocation is an allocation of data-transmitting ratio $X = \{x_1, x_2, \dots, x_n\}$ to each consumer in each path. Suppose there is one market maker who will charge consumer i $P_i(\cdot) : X \rightarrow \mathbb{R}^+ \cup \{0\}$, and each consumer i has a valuation function $v_i(\cdot) : X \rightarrow \mathbb{R}^+ \cup \{0\}$ which denotes the price she is willing to pay for the allocation X . So the happiness (utility) of consumer i under allocation X is $v_i(X) - P_i(X)$. Our target is to maximize the market maker's revenue.

Again we take the network shown in Figure 1.1 for an example. Supposing consumer i receives a data-transmitting ratio of x_i , she will be charged x_i dollars for transmitting data along each unit length of the path (s_{n+1}, t_{n+1}) . Therefore, if we have a bandwidth allocation $X = \{x_1, x_2, \dots, x_{n+1}\}$, consumer $i, 1 \leq i \leq n$ will be charged x_i dollars and consumer $n+1$ will be charged $(2n+1)x_{n+1}$ dollars (path (s_{n+1}, t_{n+1}) has length $2n+1$). Suppose the valuation function of each consumer i is given by $v_i(X) = x_i \cdot (1 + \text{the length of path } P_i)$. In a revenue maximization solution, the market maker may assign a bandwidth 1 to consumer $n+1$ and 0 to others; she will get a revenue of $2n+1$. In this case, consumers $1, 2, \dots, n$ will have utility 0 and consumer $n+1$ will have utility 1. Hence the total social happiness (i.e., social welfare) will be $2n+1+1=2n+2$. In this case, the first n consumers may feel that they have been treated *unfairly* if they observe the outcome of player $n+1$. On the other hand, if we assign a bandwidth 1 to each of consumers $1, 2, \dots, n$, the market maker will have revenue n and the consumers will have a total utility $3n$. The total happiness of all participants in the market will be $4n$, which is much larger than the total happiness in the revenue maximization solution.

From the previous examples, we observe that:

- A *social welfare*-optimized solution (as we see in Example 1.1) may be *unfair* to individual participants.
- A *fair* solution may reduce the total social well-being (as we see in Example 1.2).
- Consumers and market makers will have conflict of interests (as we see in Example 1.3), since the market makers tend to increase their revenue but consumers want to pay less.

- A revenue-maximized solution may be neither *fair* nor *social welfare*-maximized (as we see in Example 1.3).

One basic assumption behind these models is that each participants in the market is a utility maximizer. We see three parties in the market. The first party is that of the potential consumers. Consumers cannot decide the allocation or the pricing of the resources in question. In general, consumers prefer to receive as much utility as possible. If no payments are needed, consumers actually desire as many resources as possible. In the mean time, individual consumers may compete with each other, and it is possible for them to quit the market if they feel they are *unfairly* treated. The second party is the centralized authority. The centralized authority can price the resources and decide how to allocate resources to potential consumers, but it doesn't directly participate in the market. The target of centralized authority is to ensure the performance of the market from a global point of view, i.e., to ensure total social well-being, or in another words, to ensure efficiency. The third party is the market maker (seller). This party may not exist in some cases. If the market maker exists, consumers will pay resources according to pricing. The market maker desires high revenue through high pricing, which decreases the welfare that potential consumers can enjoy.

A central question raised by the nature of resource allocation problems is *how* to allocate a given set of resources to a set of potential consumers. From the above examples and discussion, we can see that different parties have different interests. For different application domains, different requirements need to be achieved, or different factors need to be balanced. In particular, consumers require a *fair* allocation; a centralized authority may have global objectives; and if a market maker (seller) exists, she will demand as much revenue as possible. Therefore, a new series of questions arise: What kind of output should the market produce? In other words, what objective or solution concept should we select in different applications with different focuses? Are there any objectives or solution concepts that can simultaneously capture the different interests of different parties? Furthermore, are these objective or solution concepts achievable with limited computational power? Before addressing the problems considered in this thesis, we will examine a few critical criteria that are considered to be *efficient* from the centralized authority's perspective or to be *fair* from individual consumer's perspective, so that we can achieve the objective to make the different parties happy.

1.1 Economic Efficiency

In economy, the *welfare* (or *utility*) is a measure of the happiness or well-being received by a consumer or a market maker when consuming or selling goods or services, possessing wealth and spending leisure time. As we have seen in the previous examples, when a consumer receives a given set of resources, there is a valuation function that denotes the price she is willing to pay for these resources. This function also quantifies the happiness (i.e. *utility*) of the consumer if no payments are needed. If payments are needed, the utility of the consumer is figured as valuation less payments. When a market maker receives payments, her utility is the total amount of money that she receives. Based on these utility measures, economic efficiency evaluates the performance of a market in terms of the welfare from a global point of view, which is important to the centralized authority's perspective. The following are important criteria for economic efficiency.

Pareto Efficiency Pareto efficiency was firstly introduced by the Italian economist Vilfredo Pareto [110]. It is a condition that captures the overall performance of the system. Given a Pareto-efficient solution, there is no way to improve any participant's welfare without that of hurting others.

For a simple illustration, take the *stable marriage problem* instance, which includes n boys and n girls. Each person has ranked all members of the opposite sex within his or her own list (i.e. *preference list*). We need to marry the boys and girls such that there are no two persons of opposite sexes who would both rather marry each other than their current partners. We call the valid solution a *stable matching*. The one of the beauties of this problem is that stable matchings are actually Pareto-efficient, by which we mean that we want to improve the welfare of any given person (i.e., marry she/he to a person ranked higher on her/his preference list), there must be someone who loses welfare (i.e., someone who is married to a person ranked lower in her/his preference list).

We normally call a change of an allocation or price within a given outcome a *Pareto improvement* if at least one participant strictly becomes “better off” after the change and none of the other participants becomes “worse off.” Hence, an allocation or price vector is Pareto-efficient only when no Pareto improvement can be made.

Pareto efficiency is accepted by nearly all economists [120], and it is used as a bottom line for economic efficiency in many literatures.

About mechanism design, Crawford(1979) described a simple, operational procedure that always generates Pareto-efficient egalitarian-equivalent allocations [47]. In Abdulkadiroglu and Sonmez’s several works [1–3] regarding the house allocation problem and the student assignment problem, they use Pareto efficiency as their guarantee of social efficiency. In 1994, MacKie-Mason and H. Varian analyze a generalized version of Vickrey auctions, that guarantees Pareto efficiency [94]. Bagnoli and McKee considered a mechanism for the voluntary provision of public good that is proven to be Pareto optimal [10]. Baliga and Maskin designed a mechanism for environmental scientists that, is Pareto optimal when externalities such as pollution are non-excludable [13]. Pazner and Schmeidler considered Pareto efficiency as a basic efficiency requirement for the *cake cutting* problem [111]. Reijnierse and Potters used Pareto optimality as an efficiency guarantee for a finite number of homogeneous infinitely divisible good problem under linear utility functions [116]. Pareto optimality has also been considered when fairly allocating indivisible items to consumers [59].

Besides mechanisms, the Pareto efficiencies of many kinds of economic systems have been studied, and some of them have actually made an impact on our every-day life. Arrow(1969) analyzed public goods and their externalities using Pareto efficiency as a basic condition [5]. Pareto-efficient taxation has been a big topic and is studied in lots of literature [80, 121, 126]. Ross(1973) used Pareto efficiency as a primary criterion when he considered the relationship between agencies [119]. The remarkable BitTorrent file distribution system was also designed to be Pareto-efficient [42].

Kaldor-Hicks Efficiency Kaldor-Hicks efficiency [73, 79], which is named after Nicholas Kaldor and John Hicks, is also one of the most important efficiency criteria in microeconomics. *Kaldor-Hicks improvement* signifies that a change in an allocation can still be a Pareto improvement after participants who are better off “compensate” for participants who are worse off. Hence, an allocation or price vector is said to be Kaldor-Hicks efficient when we cannot make a Kaldor-Hicks improvement on it. Kaldor-Hicks improvement may not necessarily qualify as a Pareto improvement, but Pareto improvement must always qualify as a Kaldor-Hicks improvement, which means that Kaldor-Hicks efficiency is stronger than Pareto efficiency.

Now take a simple (actually stupid) gambling game as an example. Suppose there are

	Player 2 bids A	Player 2 bids B
Player 1 bids A	$(0, 0)$	$(-5, 10)$
Player 1 bids B	$(10, -5)$	$(0, 0)$

Table 1.1: Pay off for the gambling bids.

two players participating in the gambling game, Player 1 and Player 2. There are two kinds of bids, A and B . If both players make the same bid, they will get nothing. If Player 1 bids B and Player 2 bids A , player 1 will receive 10 dollars and player 2 will pay 5 dollars. If Player 2 bids B and Player 1 bids A , Player 2 will receive 10 dollars and Player 1 will pay 5 dollars. The payoff of the game is shown in Table 1.1. If we take Pareto's criterion as the measurement of economic efficiency, there will be no difference among these four outcomes, since none of these outcomes can result in a Pareto improvement of another one. Therefore, all four outcomes are Pareto-efficient. From here we can see a limitation of Pareto's criterion: if the two players cooperate, they receive some economic profit instead of receiving nothing. Kaldor-Hicks's criterion overcomes this limitation of Pareto efficiency. $(10, -5)$ and $(-5, 10)$ will be Kaldor-Hicks improvements of $(0, 0)$; this means that the outcomes $(10, -5)$ and $(-5, 10)$ are Kaldor-Hicks efficient, which is also welfare-maximized. The shortcoming of Kaldor-Hicks's criterion is that it needs to quantify all the given preferences. For example, the stable matching example in the previous paragraph cannot be examined using Kaldor-Hicks's criterion.

Kaldor-Hicks's criterion is an important notion in economics, and there has been a lot of literature about it. For example, Buchanan discussed the debate over the criteria of welfare efficiency [31]. In Posner's highly cited work [113], he used Kaldor-Hicks's criterion as a "consensual base" of efficiency. Kronman actually believed that Kaldor-Hicks's criterion is more practical than Pareto's criterion for wealth maximization [89]. Balkin took Kaldor-Hicks efficiency as the guarantee of social efficiency in his work [14]. Stavins et al. related Kaldor-Hicks's criterion to the sustainability of the economy [124]. Parisi listed Kaldor-Hicks's criterion as one of the important efficiency criteria in the institutional design of lawmaking [107].

Welfare maximization *Social Welfare* is defined as the sum of the utility of all participants in the market. When allocating resources, social welfare maximization involves the

maximization of the total social welfare in the market, which means, to allocate the given resources to the people who most desire them.

Social welfare maximization is a natural aim of market mechanism design. Sociologists and economists, such as Jeremy Bentham and John Stuart Mill, believed that society should aim to maximize the total welfare of its individuals, or aim for “the greatest happiness for the greatest number of people”. Intuitively, whenever the market does not achieve social welfare maximization, it has the potential to generate economic surplus. Generating additional economic surplus means that a given society has the ability to gain more resources that it needs to do the things it wishes to. That is to say, whenever the market does not achieve social welfare maximization, there is a waste, which is undesirable to the society.

In general, Pareto efficiency and Kaldor-Hicks’s efficiency are weaker than social welfare maximization since they only give locally optimal solutions for the maximization of social welfare. Therefore, we can characterize the three efficiency criteria into the following structure:

$$\text{Pareto efficiency} \leq \text{Kaldor-Hicks Efficiency} \leq \text{Welfare maximization}$$

In the remarkable papers [41, 67, 129], Vickrey (1961), Clarke (1971), and Groves (1973) introduced a mechanism whereby multiple goods are allocated to bidders after bidders sealed reported their valuations for the items. The mechanism charges each bidder for the harm she cause to other bidders, which ensures social welfare maximization.

There is much literature that targets the question of social welfare in resource allocation problems: Many efforts have been made to design heuristics for maximizing social welfare given consumers’ *submodular* valuations, for example: A highly cited Lehmann, Lehmann and Nisan’s work [91] provided a simple greedy approach; several linear programming approaches appeared in [54, 56, 105]; and randomized algorithm approaches can be found in [56, 130]. Besides these instances, *subadditive* and *fractionally subadditive* valuations are considered in [56, 61], respectively.

There are also a lot of work done with the aim of achieving good social welfare in other applications such as mechanism design for combinatorial auctions with various valuation functions [55, 88, 101], unsplittable flow routing and multicast routing [30], etc.

Furthermore, there is the first and simplest general market equilibrium model [131], which was introduced in 1874 by the great French economist Léon Walras as well as an

extended and refined model by Arrow. Debreu [6] and McKenzie [97] are guaranteed to produce solutions that maximize social welfare.

Above all, social welfare is one of the most important factor for evaluation of the outcome of a system from a global point of view. Most seminal designs and benchmarks aim to maximize social welfare.

1.2 Fair division

Fairness, which balances individual participants' competitive interests, is also one of the most important solution conditions considered in practice. Fair allocation divides resources in a way that each participant in the market believes that she receives a relatively fair amount of resources, which is crucial from consumers' individual perspectives. Two critical features of this solution condition are that people may disagree about the value of a given portion of resources and that people may disagree with the given definition of "fairness." Imagine, for instance, that a mother who wants to divide a piece of cake for two hungry children. What may look like fair halves to her may not look like halves at all to the children, and merely weighing two equal pieces and assigning them to the children may produce dissatisfaction because of these features. To establish a common ground, it is necessary to consider some generally accepted conditions for fairness. The following are some of the fairness criteria that economists often use to ensure fairness. Some of them can be achieved simultaneously while some of them conflict:

Proportionality An allocation is proportional if each consumer receives at least an average share based on his own measurement, i.e., if there are n consumers dividing the resources, then in the allocation, each consumer receives at least $1/n$ of her valuation of all the resources. Proportional division is also called *simple fare division*.

Take the example of two children cutting a cake. If either child divides the cake into two halves she thinks are equal, and the other child chooses the half she prefers, then this will result a proportional division since the two children will receive at least $1/2$ of their total valuation.

The origin of proportionality can be traced back to Aristotle, who argued in *Ethics* that goods should be divided in proportion to each claimant's contribution [22]. In recent years

proportionality has become a basic and classic fairness criterion of the cake-cutting problem and it has received a lot of attention in the literature [22]. Even and Paz gave a protocol that produce proportional division with continuous pieces [60]. Brams, Jones, and Klamler gave a procedure for cutting a pie (which is different from cutting a cake) proportionally between two people [25]. Caragiannis et al. measured the loss of efficiency if what is desired is a proportional cake cutting [33]. Aumann and Dombb extended Caragiannis's results by restricting the division of the cake to connected pieces [7]. Y. Chen et al. considered the consumer behavior within the cake-cutting model and obtained a truthful mechanism for dividing a cake proportionally [39].

Envy-Freeness We are assuming all consumers are utility maximizers, i.e., for a given allocation and price vector (if any), these consumers would like to compete with each other and be allocated only the subsets of resources that give them maximum utility. Envy-freeness captures such self-interested considerations of consumers and requires that no consumer would envy any other's allocation for the given market allocation and price vector (if any). In other words, every consumer's utility is maximized at his current allocation, and therefore, each consumer is happy with the outcome of the market.

A resource allocation outcome is envy-free if no player feels that another has a strictly larger piece of the allocation.

Still, take the two-child cake-cutting example. The cut-and-choose protocol in paragraph above will still provide an envy-free allocation since each child will not prefer the other's piece more than her own.

Envy-freeness is generally stronger than proportionality. If there are only two consumers, it is understood that envy-freeness and proportionality are equivalent.

Envy-freeness is a fairness criterion that is studied extensively in both indivisible and divisible resource allocation problems.

For divisible resource allocation problems (for instance, cake cutting), a very elegant envy-free cake-cutting protocol was introduced by Brams and Taylor in 1995[27]. Brams et al. considered cake-cutting problems with money among several consumers, who required envy-freeness as a fairness guarantee [24]. Robertson, Webb and Pikhurko separately gave a construction of envy-free cake cutting solution after Brams and Taylor's results [112, 117]. Caragiannis et al. and Aumann et al. measured the loss of efficiency within a envy-free cake

division in a general setting and a continuous-piece setting, respectively [7, 33]. Cohler et al. provided algorithms to discover welfare-optimized envy-free allocation under various preference settings [43].

For indivisible-item allocation problems, Brams and Fishburn considered an indivisible-item allocation that ensures envy-freeness between two persons [23]. Envy-freeness has also been considered when efficiently allocating indivisible items to consumers [59]. Herreiner and Puppe provided a simple procedure for discovering envy-free allocations of indivisible goods [72]. Lipton et al. considered an approximation algorithm in order to minimize the number of envies and also provided an bound for that [92]. Bouveret and Lang considered the computational issue in order to find an efficient, envy-free allocation in a general setting [21].

For indivisible-item allocation with pricing, Guruswami et al. provided an approximation of a *profit-maximized envy-free* solution given unit demand and single-minded consumers in [69]. This is a frequently cited paper. Chen et al. considered the envy-free pricing problem when there is metric substitutability among the items [37]. Briest gave an inapproximability upper bound which showed that Guruswami et al.'s result in [69] is nearly optimal [28].

Exact Division Exact division as an allocation ensures that each consumer thinks that all of the other consumers receive the same amount of allocation.

Take a divisible resource allocation problem as an example. If R is the set of resources and there are n consumers among whom to divide the resources, suppose $v_i(\cdot) : T \rightarrow \mathbb{R}^+ \cup \{0\}$, $T \subseteq R$ is valuation function for consumer i , $0 \leq i \leq n$ and $X = (X_i)$ is a valid allocation of resources. An exact division is an allocation X such that $v_i(X_i) = v_i(X_j)$ for all $j \neq i$.

This actually implies that each consumer will get exactly $1/n$ of her valuation of all the resources, no more or less. Exact division is an even stronger condition than proportionality and envy-freeness. Therefore, in the cake cutting problem, the above three fairness criteria can be characterized into the following structure:

$$\text{Proportionality} \leq \text{Envy-freeness} \leq \text{Exact Division}$$

Exact division has been considered in some of the literature. Roberson and Webb for

example gave a *nearly exact* division algorithm for cake cutting. “Nearly exact” means that for any given small ϵ , one consumer believes that other consumers will receive an allocation that differs by at most ϵ [117].

Equitable Division An allocation is equitable if each consumer receives exactly the same utility by her own valuation.

In the divisible resource allocation problem in the paragraph above, an equitable division is an allocation X such that $v_i(X_i) = v_j(X_j)$ for all $j \neq i$.

Equitable Division is widely investigated in the literature since it actually balances the welfare of each consumer. Note that exact division is also equitable division, but an equitable division may not either be envy-free or proportional. Ichiishi and Idzik gave a non-constructive proof of the existence of equitable allocation in the cake-cutting problem [75]. Jones gave a protocol for a case involving two consumers that resulted in a equitable and efficient division [77]. In [24], equitable division was considered as a guarantee of fairness in the cake-cutting problems with money. Barbanel et al. constructed a procedure for cutting a pie equitably among several consumers [16]. Caragiannis et al. and Aumann et al. measured the loss of efficiency for equitable divisions in a general setting and a continuous-piece setting, respectively [7, 33].

Max-min Fairness Max-min fairness, is the allocation that least maximizes allocation among all consumers. It is a very natural benchmark for fairness since the target of this approach is to maximize the welfare of the most “unhappy” person. Unlike other fairness criteria, max-min fairness is a concrete amount that can be approximated. This is similar to maximum social welfare.

In the divisible resource allocation problem above, a max-min optimal division is an allocation $X^* = (X_i^*)$ such that $\min_i v_i(X_i^*) \geq \min_i v_i(X_i)$ for all valid allocation X .

Max-min fairness has been widely used as a measurement of fairness in resource allocation. Herreiner and Puppe provided a simple procedure for finding max-min efficient allocations of indivisible goods [72]. Dall’Aglio and Hill derived a method for dividing a cake in order to reach max-min shares [49]. Brams et al. compared envy-free and max-min solutions in instances of efficient indivisible resource allocation [26]. There are several results which gave approximation algorithms of max-min fairness for indivisible resource

allocation with additive valuations [9, 15, 32]. Goemans et al. and Knot et al. provided approximation algorithms of max-min fairness for indivisible resource allocation with sub-modular and subadditive functions, respectively [66, 82].

In electric network design and other problems, Megiddo (1974) gave an allocation with optimal fairness and welfare for single-source networks with splittable fractional flows [98]. Hahne studied a strategy for max-min fair allocation of link capacity in a point-to-point packet network [70]. Ma et al. used routing algorithms for high-bandwidth traffic that improved resource allocation while maintaining max-min fair share semantics [93]. Max-min fairness-based algorithms for congestion control in broadband packet networks have also been considered [78, 127]. The results listed here are just the tip of the iceberg; guaranteeing max-min fairness is a very common methodology to ensure fairness in the domain of engineering.

1.3 Revenue

If payment exists, the market maker's utility, which is the total revenue received from consumers, will be taken into account. Under normal conditions, a reasonable market maker would like to obtain as much revenue as possible.

In order to generate revenue for the market maker, the centralized authority will price the resources in question. Price has basically two impacts on the market. The first is that market price will change the utility of consumers for receiving some resources, since if a consumer pays for something, the utility she receives will be the valuation less its price. The higher the price, the less utility consumers enjoy. On the other hand, from the market makers or providers' point of view, price guides production. Whenever the current market price of goods or resources exceeds its cost of production, producing more of the goods is profitable. Market makers or sellers will increase the output at the same time that new providers will have incentive to enter the market. More goods will be provided. This is called the principle of the "invisible hand" [123]. For market makers, revenue maximization is a natural goal. Therefore, while consumers prefer to pay less in order to increase their utilities, the market maker's interest is revenue maximization. Thus consumers and market makers will display contrary interests in the terms of payments.

Revenue maximization is different from social welfare maximization. For example, in a

combinatorial market, there are m items to be allocated to n consumers, and each consumer has a valuation on each subset of items. Normally, we assume the pricing in this case is under “budget balance condition”[108], where the price of each allocated subsets are less or equal to its valuation by its targeting consumer. If the price of each subset of items equals to the total price of all the items in the subset, then the revenue of the market maker equals to total price of all the allocated items. On the other hand, the social welfare equals to the total valuation of all the allocated consumers, which is different from the revenue.

Resource allocation model with the existence of a market maker has significant real-life applications. In fact, the combinatorial market above describes the Internet advertising market. Web search engines like Google and Yahoo! auction off advertising spaces next to their standard searching results. Figure 1.2 shows an example of an Internet advertisement in Google when the query “apple” is entered. Generally, there are a number of advertising slots provided by web search engines for certain query words. Advertisers bid for the placement on the search pages, and they are charged by web search engines accordingly once they are allocated the advertising slots. In this case, the advertising slots are the resources to be allocated. The advertisers are the consumers and they have their own valuations on each advertising slot. By 2007, Google and Yahoo! boasted a combined market capitalization

Ad - Why this ad?

Official **Apple®** Store | store.apple.com
store.apple.com/sg
 Buy the new iPod touch & the latest **Apple** products.

Official Apple Store	MacBook Air
The new iPad with Wi-Fi	MacBook Pro
The new iPad with Wi-Fi + 4G	iPod touch

[Singapore - Apple](#)
www.apple.com/sg/
Apple designs and creates iPod and iTunes, Mac laptop and desktop computers, the OS X operating system, and the revolutionary iPhone and iPad.

Singapore iPhone 4S - iPad - MacBook Pro - iPod touch - MacBook Air - iMac	How to Contact Us Visit the online Apple Store to purchase Apple hardware ...
iPad iPad is a magical window where nothing comes between you ...	Mac MacBook Pro - MacBook Air - iMac - Mac mini - Compare all Macs
iPhone Apple - Introducing iPhone 4S. With the dual-core A5 chip, all ...	Support iPad - iPhone Support - Complimentary Support - ...

Figure 1.2: An example of Internet advertisement shown above the standard searching results of the query “apple”.

over \$150 billion, largely on the strength of the Internet advertising market [104]. Roughly 85% of Google’s \$4.1 billion in 2005 revenue and roughly 45% of Yahoo!’s \$3.7 billion in 2005 revenue was likely attributed to Internet advertising auctions [104]. In 2011, full year internet advertising revenues totaled \$31.7 billion, up 22% from the \$26.04 billion reported in 2010 [76]. In every year since 2005, the annual growth rates of Internet advertising have exceeded those of other advertising media [76].

Revenue maximization mechanisms have an extensive history. In a remarkable paper [102], Myerson constructed a revenue maximization auction for a single item with risk-neutral consumers and independent private information. Riley and Samuelson had similar results [115]. Milgrom and Weber found that the market maker can receive a higher revenue by adopting some policies for certain auctions [99]. In a frequently cited paper, Myerson and Satterthwaite also designed a mechanism for a bargain game in bilateral trading in order to maximize sellers and broker’s expected revenue (if it exists) [103]. Hartline and Koltun constructed efficient approximation algorithms for several problems on computing the prices that maximize the revenue of the seller for a given set of items [71]. Guruswami et al. designed algorithms to approximate revenue-maximizing envy-free solution in combinatorial auctions under unit-demand consumers or single-minded consumers [69]. Abrams considered the problem of maximizing revenue for auctions with multiple units of goods where consumers have hard budget constraints [4]. Balcan et al. designed several mechanisms that target the seller’s revenue maximization in several circumstances [11, 12, 18, 19]. Briest and Krysta dealt with the problem of finding profit-maximizing prices for a finite number of distinct goods with single-minded consumers [29].

1.4 Problem Studied in This Thesis

In this thesis, we construct some solution concepts from the criteria above in order to satisfy the interests of different parties in different market models.

In Chapter 2, we consider a centralized combinatorial market where a market maker sells a given set of items to potential consumers, who are interested in purchasing their preferred items [74]. Formally, we have a set A with n consumers and a single market maker with a set B with m different items. For each consumer $i \in A$, there is a function $v_i(\cdot)$ that denotes the valuation of each subset of items by consumer i , i.e., the amount of

money that i is willing to pay for different combinations of items. An output $(\mathbf{p}, X)$ of the market is a combination of a price vector $\mathbf{p} = (p_j)_{j \in B}$ where p_j is the price of item j and an allocation vector $X = (X_i)_{i \in A}$ where X_i is the subset of items that i obtains. For any subset of items $S \subseteq B$, its price is defined as the sum of the prices of its elements. The utility of each consumer is her valuation of the set that she receives less the price of the set.

We study a solution concept called *Envy-free Pareto-efficient pricing*, which strikes a balance between the perspectives of consumers, the centralized authority, and the market maker. It captures fairness and balances the tradeoff between efficiency and revenue. An Envy-free Pareto-efficient pricing (i.e. *Envy-free Pareto-optimal pricing*) is a outcome $(\mathbf{p}, X)$ that is:

- (*Envy Free*) For any consumer i and any subset of items $S \subseteq B$, i prefers her allocation to S , i.e., $v_i(X_i) - p(X_i) \geq v_i(S) - p(S)$.
- (*Pareto Efficient*) An outcome of the market $(\mathbf{p}, X)$ is said to be a Pareto improvement of (or *Pareto-dominate*) another outcome $(\mathbf{p}', X')$ if
 1. each consumer i in $(\mathbf{p}, X)$ gets strictly higher or the same utility as in $(\mathbf{p}', X')$;
 2. the market maker's revenue in $(\mathbf{p}, X)$ is strictly higher or the same as in $(\mathbf{p}', X')$;
 3. at least one of the above two terms is strict.

We say that an outcome is *Pareto efficient* if it is not Pareto-dominated by another outcome.

Specifically, we will consider Envy-free Pareto-efficient pricing separately with two kinds of consumers.

- *Unit Demand Consumers*. Each consumer is interested in getting at most one item.
- *Single-minded Consumers*. Each consumer i is interested in getting only a set of items containing $S_i \in B$, S_i is called the *demand set* of consumer i .

Under unit demand consumers, we will prove the following theorem, which characterizes the economic property of Envy-free Pareto-efficient pricing [74].

Theorem 2.10. *In a market with unit demand consumers and ensuring envy-freeness, an output is Pareto efficient if and only if it maximizes social welfare.*

This theorem shows the advantage of our Envy-free Pareto-efficient pricing as compared to the *envy-free pricing* considered in [69]. Unfortunately, if our aim is to find *revenue-maximizing Envy-free Pareto-efficient pricing*, we have the following inapproximability result [74] showing that:

Theorem 2.12. *For any constant $c > 1$, there is no polynomial time algorithm for computing an Envy-free Pareto-efficient solution with revenue greater than $1/c$ times the revenue of a revenue maximizing Envy-free Pareto-efficient solution, unless $P=NP$.*

For single-minded consumers mainly, we have several hardness results.

In Chapter 3, we will consider a combination of max-min fairness and social welfare maximization. Since max-min fairness and optimal social welfare are both concrete quantities that can be optimized, it is very natural to ask the following question:

For any given r_f and r_w , does there exist an allocation that simultaneously achieves an approximation ratio r_f to max-min fairness and approximation ratio r_w to social welfare?

We consider this question in two classic resource allocation models: the indivisible-item allocation with subadditive functions and the multiple fractional knapsack allocation with capacities and demands. There is no market maker in these two models.

For the indivisible-item allocation problem, there is a set of indivisible items and a function $v(\cdot)$, which gives the valuation of different subsets of the items. We assume the valuation function $v(\cdot)$ is subadditive, i.e., $v(S) + v(T) \geq v(S \cup T)$ for any S and T (each consumer has the same valuation function). We need to partition these items into n disjoint sets and allocate them to n consumers. We mainly have the following two theorems.

Theorem 3.1. *For any subadditive valuation function $v(\cdot)$ and two positive constants r_f, r_w satisfying one of the following conditions*

- $0 < r_f \leq \frac{1}{2}$ and $r_w \leq \frac{1}{r_f+1}$
- $\frac{1}{2} < r_f \leq 1$ and $r_w \leq \frac{1}{3r_f}$

there always exists an allocation that simultaneously approximates fairness and welfare with the approximation ratios r_f and r_w , respectively.

Theorem 3.5. *For any given constants $0 \leq r_f, r_w \leq 1$ that satisfy $r_w > \frac{1}{r_f+1}$, there exists a subadditive valuation function $v(\cdot)$ such that no allocation can approximate fairness and welfare simultaneously with approximation ratios r_f and r_w , respectively.*

Theorem 3.1 shows the existence of solutions and theorem 3.5 shows the upper bound of these bifactor approximations. Notice that we do not consider computational issue in these two theorems, and in this thesis, we also determine what can be achieved in polynomial time:

Corollary 3.8. *For any given constant $r_f < \frac{1}{6}$, there is a polynomial time algorithm that has bifactor approximation ratios r_f and r_w with $r_w < \frac{1}{2} - 3r_f$ for subadditive valuation functions.*

For the multiple fractional knapsack allocation with capacities and demands, we have the following model: there are n consumers, each with a demand d_i , and m knapsacks, each with a capacity c_j . The resources are the knapsacks; an allocation consists of a mapping from consumers to knapsacks where every consumer is mapped to exactly one knapsack (i.e., the resource is unsplittable) and an assignment vector $X = (x_1, x_2, \dots, x_n)$, where x_i is the amount that consumer i receives from her corresponding knapsack.

Similarly to item allocation, we find the following results, which show the existence of the bifactor approximations.

Theorem 3.10. *For any given multiple fractional knapsack problem and two positive constants $0 < r_f, r_w \leq 1$ satisfying $r_w \leq \frac{1}{r_f+1}$, there always exists an allocation that simultaneously approximates fairness and welfare with approximation ratios r_f and r_w , respectively.*

In fact, in the multiple fractional knapsack problem, we also show that the bound given in Theorem 3.10 is tight, i.e., it is the best bifactor approximation that can be achieved. We will also consider a special case of the multiple fractional knapsack allocation problem and show an improved approximation ratio. For both the multiple knapsack problem and the special case, we provide polynomial time algorithms in order to compute allocations with constant bifactor approximations.

In Chapter 4, we will try to address the following question:

What is the computational complexity of computing an efficiency optimal allocation while ensuring fairness?

We study this question in an important resource allocation model – cake cutting [17]. Formally, a cake, which is represented by interval $[0, 1]$, is going to be divided among n consumers. Each consumer i has a density function $f_i(\cdot) : [0, 1] \rightarrow \mathbb{R}^+ \cup \{0\}$. The valuation of agent i for the interval $[a, b] \subseteq [0, 1]$ is denoted by $v_i(a, b) = \int_a^b f_i(x) dx$. There is no market maker in this model. Specifically, we require that every consumer receives a continuous piece of cake, while at the same time we try to output a welfare maximized division that ensures proportionality.

We consider two kinds of consumers separately:

- *Linear Valuations.* Each consumer i has a density function $f_i(x) = a_i \cdot x + b_i$, where a_i, b_i are constants, and $x \in [0, 1]$.
- *Piecewise Constant Valuations.* Each consumer i has a density function $f_i(x)$ such that we can partition the interval $[0, 1]$ into subintervals; $f_i(x)$ is constant on each interval.

For consumers with linear functions, we can actually achieve the best we can hope for in polynomial time: a Polynomial Time Approximation Scheme (PTAS) [17].

Theorem 4.4. *When all consumers have linear density functions, for any $\epsilon > 0$, there is an algorithm running in polynomial time in the input size and $\frac{1}{\epsilon}$, and outputs an $(1 - \epsilon)$ -fair allocation with efficiency within a factor of $(1 - \epsilon)$ to the optimum.*

$(1 - \epsilon)$ -fair solution means each consumer receives no less than $\frac{1-\epsilon}{n}$ share of her value for the entire cake.

For consumers with piecewise constant functions, we provide an inapproximability result [17] showing that:

Theorem 4.6. *When all consumers have piecewise constant density functions, it is NP-hard to approximate the optimum efficiency within a ratio of $\Omega(\frac{1}{\sqrt{n}})$.*

Chapter 2

Indivisible Resource Allocation with Payment

2.1 Introduction

In this chapter we will consider an indivisible resource allocation problem with payment, which is raised in the centralized combinatorial market. In a centralized combinatorial market, a market maker sells a set of m items to n potential consumers, where each consumer i has a valuation function $v_i(\cdot)$ measuring the maximum amount that i is willing to pay for different combinations of items. As an outcome of the market, the market maker specifies a price vector $\mathbf{p} = (p_j)$ for all items and an allocation vector $X = (X_i)$ which indicates the subset of items that every consumer i obtains. That is, for the given outcome $(\mathbf{p}, X)$, consumer i obtains subset X_i with payment $p(X_i) = \sum_{j \in X_i} p_j$ to the market maker; therefore his utility is defined to be $v_i(X_i) - p(X_i)$.

The centralized combinatorial market is one of the most fundamental market models that has received a lot of attention in the literature [46, 96]. It characterizes a number of applications, especially with the development of the Internet, e.g., Amazon's electronic market, Google and Yahoo's advertising markets, to name a few. A key question in such marketplaces is that for the given input information $v_i(\cdot)$, what kind of outcomes should the market output? In other words, what solution concepts should be selected in different applications with different focuses? In this chapter, for consumer's fairness perspective, we use envy-freeness as a criterion. From global point of view, we hope to output a welfare

	Fairness	Social welfare	Pareto efficiency	Revenue	Computation	Existence
WE	✓	✓	✓	small	easy	yes
REF	✓	×	×	large	hard	yes
REP	✓	✓	✓	medium	hard	yes

Table 2.1: Our results (last row) for unit-demand consumers.

maximization solution, or at bottom line, a Pareto efficient solution. On the other hand, the market maker wants to generate as much revenue as possible in the mean time.

Ideally, we would like to have a solution concept that captures all these three criteria simultaneously. However, due to the contrary interests amongst different parties in the market, such a solution does not exist in general. Therefore, any feasible solution concept can only have a partial focus on each of these properties. We will consider the following solution concepts in this chapter.

Walrasian equilibrium (WE) An outcome is called a *Walrasian equilibrium* if it satisfies, in addition to envy-freeness, a market clearance condition which says that every leftover item must be priced at zero. Walrasian equilibrium is one of the most important solution concepts in economics and has been studied extensively in the literature. It ensures individual fairness and overall market efficiency. In particular, the seminal first fundamental welfare theorem says that every Walrasian equilibrium maximizes social welfare; this implies that it is Pareto efficient as well.

However, it is well known that a Walrasian equilibrium may not exist in general; and even if one is guaranteed to exist (for example, when the valuation functions satisfy the gross substitutes condition [81]), the revenue it generates could be very small even for a revenue maximizing Walrasian equilibrium (RWE). These two issues largely limit the applicability of Walrasian equilibrium. For example, the market maker would like to seek to optimize his own objective, i.e., revenue maximization, which is one of the key business concerns in many applications. From this perspective, Example 2.1 shows that (revenue maximizing) Walrasian equilibrium is not a good solution concept in some applications.

Revenue maximizing envy-freeness (REF) The limitation of Walrasian equilibrium relies on the requirement of the market clearance condition. In applications like Google and Yahoo’s advertising markets, the market makers may hold some of their inventories unsold

	Fairness	Social welfare	Pareto efficiency	Revenue	Computation	Existence
WE	✓	✓	✓	small	hard	no
REF	✓	×	×	large	hard	yes
REP	✓	×	✓	medium	hard	no

Table 2.2: Our results (last row) for single-minded consumers.

(i.e., priced at infinity) to achieve their own objectives. Removing the market clearance condition leads to the solution concept of envy-freeness (EF), which is guaranteed to always exist. Finding an envy-free pricing solution that maximizes revenue (REF) for the seller has attracted a lot of attention recently in computer science [12, 37, 69]: The optimization problem is NP-hard to solve even when all consumers only desire one item or a fixed subset of items; approximation algorithms have been considered for special cases.

A remarkable property of a REF pricing solution is that there are instances in which the REF revenue can be arbitrarily larger than that of the RWE.

Example 2.1. (REVENUE COMPARISON: RWE vs REF). *There are two consumers i_1, i_2 and two items j_1, j_2 . Both consumers want to get only one item, and their valuations are $v_{i_1 j_1} = k$, $v_{i_1 j_2} = k + 2$, and $v_{i_2 j_1} = 0$ and $v_{i_2 j_2} = 1$. In an optimal RWE solution, i_1 wins j_2 and i_2 wins j_1 at a price vector $(0, 2)$. In an optimal REF solution, however, we may charge a price vector $(k, k + 2)$ with the same allocation. Hence, the total revenue is increased from 2 to $k + 2$. (Note that in the REF solution, we are free to set any large price for item j_1 , whereas in the RWE, its price has to be 0 by the market clearance condition.) Hence, the difference between an optimal RWE solution and REF solution can be arbitrarily large.*

While REF does generate the largest possible revenue provided the fairness condition, it may not maximize social welfare; and further, it may not even be Pareto efficient, which is arguably the weakest solution concept to ensure overall performance of a system.

Example 2.2 (REF is not Pareto efficient). *There are two consumers i_1, i_2 and two items j_1, j_2 . Again both consumers only want to get one item, and their valuations are $v_{i_1 j_1} = 100$, $v_{i_1 j_2} = 10$, and $v_{i_2 j_1} = 180$ and $v_{i_2 j_2} = 100$, as the following figure (A) shows: In this example, an optimal REF solution is given by (B) with total revenue 180 (note that we can set j_2 with any large price to ensure that i_2 is happy with his allocation j_1 at price 180). However, this is Pareto-dominated by the solution in (C) where both i_1 and i_2 obtain*

the same utility 0 and the market maker improves his revenue from 180 to 200. Note that the solution given by (C) is not envy-free. For this example, the optimal REP (defined below) is shown in (D).

Revenue maximizing envy-free Pareto efficiency (REP) Social welfare maximization, and the bottom line, Pareto efficiency, are critical solution conditions for a market to have healthy and long-term development. To capture these conditions, in particular, Pareto efficiency, we introduce a solution concept, envy-free Pareto efficiency (EP), which requires both envy-freeness and Pareto efficiency. A revenue maximizing envy-free Pareto-efficient (REP) outcome is one that maximizes revenue among all envy-free Pareto efficient solutions. The figure (D) in Example 2.2 gives an optimal REP solution for the given instance, which is different from the optimal REF solution shown in figure (B); this implies that an algorithm for finding a REF solution will not suffice for finding a REP solution.

By the definitions of the three solution concepts and the first fundamental theorem, we know that EP lies between WE and EF, that is, the solution space has the following structure:

$$\text{WE} \subseteq \text{EP} \subseteq \text{EF}$$

This further implies that the total revenue generated by REP also sits between RWE and REF, i.e.,

$$\text{RWE} \leq \text{REP} \leq \text{REF}$$

where both inequalities can be strict.

Including Pareto efficiency strikes a balance between system efficiency and revenue: REF solutions yield good revenue, but, as mentioned earlier, the overall performance of the system is not addressed well. On the other hand, WE , while fair and efficient, is not practical because the revenue of the market maker can be arbitrarily small. Our solution concept EP , in addition to satisfying fairness (in the form of envy-freeness), addresses both overall performance (in the form of Pareto efficiency) and revenue (which is sandwiched between the revenues guaranteed by REF and WE solutions).

Our Results

We will study these solution concepts in two domains that are, arguably, the most applicable: *unit-demand* consumers (where each consumer desires exactly one item, as shown in Example 2.1 and 2.2) and *single-minded* consumers (where each consumer desires a fixed subset of items).

For a market with unit-demand consumers, the existence of an EP solution follows from the existence of a WE immediately. Further, in addition to satisfying fairness and efficiency, we show that EP (and therefore REP) has the property that it maximizes social welfare (Theorem 2.9). This property of EP solutions is quite remarkable because it does not hold for REF solutions. Hence, our solution concept possesses all the nice economic properties that WE has. On the other hand, while approximating the revenue of an REP solution is NP-hard (Theorem 2.12), we show that for a special case which includes most instances of advertising markets, an optimal REP can be computed in polynomial time (Theorem 2.13).

For a market with single-minded consumers, our solution concept, REP , aligns with WE and REF . Similar to the complexity of determining the existence of a WE [36], we show that determining the existence of an EP solution is NP-hard. Even if an EP solution is guaranteed to exist, computing one that maximizes revenue is NP-hard as well (Theorem 2.16); this has the same computational complexity as REF [69]. These hardness results do not rule out the applicability of the solution concepts, but rather, illustrate their computational features. This is similar to Nash equilibrium, which is extensively applied in game theory, but shown to be hard to compute [38, 44, 51].

We summarize our results, and compare EP with WE and EF in Tables 2.1 and 2.2. Our results imply that EP captures fairness, and balances the tradeoff between efficiency and

revenue. Therefore, it provides a good alternative for solution concepts in combinatorial markets.

Related Work

Envy-freeness and Pareto efficiency are among the most well-studied solution concepts in economics. For instance, Pazner and Schmeidler [111] gave examples in which no Pareto efficient allocation is envy-free. Varian [128] considered general utility functions for the buyers and showed that if preferences of consumers are monotonic, then a fair allocation exists. For Walrasian equilibrium, Shapley and Shubik [122] showed that a solution always exists in unit-demand settings. The result was later improved to a more general class of functions [48, 52, 114]. Kelso and Crawford [81], as well as Gul and Stacchetti [68], proposed conditions (e.g., gross substitutes and single improvements) in which a Walrasian equilibrium always exists. More discussions are referred to the textbooks, e.g., [46, 96]. We note the major difference between our work and the literature is that we consider the centralized seller as a participant of the market as well; thus, in our model the market is composed of $n + 1$ entities.

From a computational point of view, Bouveret and Lang [21] studied the complexity of deciding whether there exists a Pareto efficient and envy-free allocation in several contexts when preferences are represented compactly. Guruswami et al. [69] showed that a revenue maximizing envy-free solution is NP-hard to approximate, even for unit-demand and single-minded consumers. The revenue maximizing envy-free pricing problem was later shown by Briest [28] that it is even NP-hard to approximate within a ratio of $O(\log^\epsilon n)$ for some $\epsilon > 0$. Our results, on the other hand, suggest that REP, as an alternative solution concept, could be more applicable from computational perspective.

2.2 Model and Solution Concepts

In a marketplace, we have a set A of n consumers and a single seller (i.e., the market maker) with a set B of m items, where each item has unit supply. For every consumer i and subset of items $S \subseteq B$, there is a value $v_i(S) \geq 0$ denoting the valuation that i obtains from S . We assume that $v_i(\emptyset) = 0$. An output of the market is given by a *price* vector $\mathbf{p} = (p_j)_{j \in B}$ where p_j is the price of item j , and an *allocation* vector $X = (X_i)_{i \in A}$ where X_i is the subset

of items that i obtains. For any $S \subseteq B$, its price is defined to be the sum of the prices of its elements, i.e., $p(S) = \sum_{j \in S} p_j$. Given an output $(\mathbf{p}, X)$, the *utility* of every consumer i is defined to be $u_i(\mathbf{p}, X) = v_i(X_i) - p(X_i)$. We assume all consumers are utility-maximizers, i.e., they would prefer outcomes with larger utilities.

In general, there could be many different outputs $(\mathbf{p}, X)$. A crucial study in economics and social choice is the selection of solution concepts that satisfy different criteria. All outputs considered in this chapter are assumed to be *feasible* (i.e., $X_i \cap X_{i'} = \emptyset$ for any $i \neq i'$) and *individually rational* (i.e., $u_i(\mathbf{p}, X) = v_i(X_i) - p(X_i) \geq 0$). One of the most notable solution concepts is that of Walrasian Equilibrium, which is defined formally as below.

Definition 2.3. (ENVY-FREENESS AND WALRASIAN EQUILIBRIUM). *An outcome $(\mathbf{p}, X)$ is called a Walrasian Equilibrium (WE) if it satisfies the following two conditions:*

- (Market clearance) *Every unsold item is priced at zero.*
- (Envy-freeness) *For any consumer i and any subset of items $S \subseteq B$, $v_i(X_i) - p(X_i) \geq v_i(S) - p(S)$.*

The first condition above is a market clearance condition, which requires that all unallocated items are priced at zero (or at any given reserve price). The second is a fairness condition, which says that at the given price vector, the utility of every consumer is maximized by the corresponding allocation, i.e., everyone is happy with his allocation.

While envy-freeness captures the interests of consumers, the seller, on the other hand, would like to maximize the total amount of payment received (i.e., *revenue*), defined by $r(\mathbf{p}, X) = \sum_{i \in A} p(X_i)$ for a given outcome $(\mathbf{p}, X)$. This motivates the following solution concepts:

- Revenue maximizing Walrasian equilibrium (RWE): As the name implies, it is a Walrasian equilibrium that offers the most revenue to the seller.
- Revenue maximizing envy-free solution (REF): These are envy-free solutions with maximum total revenue. Note that we do not require market clearance condition here, i.e., it is possible that an unallocated item has a positive price.

Although a (revenue maximizing) envy-free solution captures the interests of both the consumers and the seller to some extent, as Example 2.2 shows, it may not be Pareto efficient, which is defined formally as follows.

Definition 2.4 (Pareto efficiency). *An outcome $(\mathbf{p}, X)$ is said to Pareto-dominate (or just dominate) another outcome $(\mathbf{p}', X')$ if (i) for all $i \in A$, $u_i(\mathbf{p}, X) \geq u_i(\mathbf{p}', X')$; (ii) $r(\mathbf{p}, X) \geq r(\mathbf{p}', X')$, and (iii) at least one of the above inequalities is strict. We say that an outcome $(\mathbf{p}, X)$ is Pareto efficient if it is not Pareto dominated by any other outcome.*

In this chapter, we will consider the following solution concept, combining revenue maximization (from the seller's perspective), envy-freeness (from the consumers' perspective), and Pareto efficiency (from the market's perspective).

Definition 2.5. (REVENUE MAXIMIZING ENVY-FREE PARETO EFFICIENCY). *An envy-free Pareto-efficient (EP) solution is one that is both envy-free and Pareto efficient. A revenue maximizing envy-free Pareto efficient (REP) solution is one with maximum revenue among all EP solutions.*

Another important property considered in the current chapter is social welfare, defined as below.

Definition 2.6 (Social welfare). *Given an outcome $(\mathbf{p}, X)$, the social welfare is defined to be the sum of utilities of all participants in the system, i.e., $\sum_{i \in A} u_i(\mathbf{p}, X) + r(\mathbf{p}, X)$.*

Observe that the utility of every consumer i is $u_i(\mathbf{p}, X) = v_i(X_i) - p(X_i)$ and the utility of the seller is $\sum_{i \in A} p(X_i)$. Hence, in our setting social welfare is equal to $\sum_{i \in A} v_i(X_i)$, which is independent of the selected price vector. A feasible allocation that maximizes social welfare is called *optimal*, i.e., it has the maximum total valuations $\sum_{i \in A} v_i(X_i)$. Therefore, an outcome maximizes social welfare if and only if the corresponding allocation is optimal.

2.3 Unit Demand Consumers

We say a consumer i has *unit demand* if $v_i(S) = 0$ for any $S \subseteq B$ with $|S| \geq 2$. That is, i is interested in obtaining at most one item. In this section, we assume all consumers have unit demand and denote $v_{ij} = v_i(\{j\})$. Note that when all consumers have unit demand, any feasible allocation $X = (X_i)_{i \in A}$ corresponds to a matching between consumers and items.

2.3.1 Pareto Efficiency and Social Welfare

In this section, we will examine the relation between Pareto efficiency and social welfare maximization (hence, optimal allocation) in the market with unit-demand consumers.

We begin with a useful lemma.

Lemma 2.7. *If $(\mathbf{p}, X)$ is not Pareto efficient, then there is another solution $(\mathbf{p}', X')$ that dominates $(\mathbf{p}, X)$, but the strict improvement is restricted to the seller's revenue. That is, $u_i(\mathbf{p}', X') = u_i(\mathbf{p}, X)$ for all consumers $i \in A$ and $r(\mathbf{p}', X') > r(\mathbf{p}, X)$.*

Proof. We provide a constructive proof. Since we are not concerned with envy-freeness, without loss of generality assume that all unsold items are priced at 0. Consider a solution $(\mathbf{q}, Y)$ that Pareto-dominates $(\mathbf{p}, X)$. If for all consumers $i \in A$, $u_i(\mathbf{q}, Y) = u_i(\mathbf{p}, X)$, then clearly $(\mathbf{q}, Y)$ dominates $(\mathbf{p}, X)$ at revenue, i.e., $r(\mathbf{q}, Y) > r(\mathbf{p}, X)$, and the claim follows. Hence, it remains to consider the case where there is $i \in A$ such that $u_i(\mathbf{q}, Y) > u_i(\mathbf{p}, X)$. For such case, we can simply increase the price of the item y_i allocated to i in $(\mathbf{q}, Y)$ till the inequality becomes tight. This procedure can be repeated for all such consumers with $u_i(\mathbf{q}, Y) > u_i(\mathbf{p}, X)$. Eventually all consumers' utilities in the resulting solution will match their utilities in the original solution $(\mathbf{p}, X)$, and the total revenue of the seller will be strictly larger. $\square$

Given the above characterization, we can show the following result.

Theorem 2.8. *Given an outcome $(\mathbf{p}, X)$, if X is an optimal allocation, then $(\mathbf{p}, X)$ is Pareto efficient.*

Proof. We prove by contradiction. Suppose that X is an optimal allocation but $(\mathbf{p}, X)$ is not Pareto efficient. Then there is a solution $(\mathbf{p}', X')$ that dominates $(\mathbf{p}, X)$, and by Lemma 2.7, the strict improvement is restricted to the seller's revenue. Notice that the total valuations given by an allocation is equal to the total revenue plus total utility. Hence,

$$\begin{aligned} \sum_{i \in A} v_{ix_i} &= \sum_{i \in A} (u_i(\mathbf{p}, X) + p_{x_i}) = r(\mathbf{p}, X) + \sum_{i \in A} u_i(\mathbf{p}, X) \\ &< r(\mathbf{p}', X') + \sum_{i \in A} u_i(\mathbf{p}', X') = \sum_{i \in A} (u_i(\mathbf{p}', X') + p'_{x'_i}) = \sum_{i \in A} v_{ix'_i}. \end{aligned}$$

Therefore, the total valuations given by X' is strictly greater than that given by X , which contradicts the fact that X is an optimal allocation. $\square$

We comment that the reverse direction of the above claim does not hold, i.e., Pareto efficiency may not imply optimal allocation. For example, there are two consumers i_1, i_2 , and one item j_1 , with $v_{i_1 j_1} = 10$ and $v_{i_2 j_1} = 8$. Assigning the item to i_2 at price 0 is not an optimal allocation, but it is Pareto efficient since it is not dominated by any other outcome given that the utility of i_2 should be at least 8. That is, optimal allocation is only a sufficient condition to ensure Pareto efficiency.

However, as the next theorem implies, given envy-freeness, optimal allocation is also a necessary condition for Pareto efficiency. We note that the following result is known for WE, but we consider the solution concept EP with $n + 1$ participating entities.

Theorem 2.9. *For any envy-free solution $(\mathbf{p}, X)$, if X is not an optimal allocation, then $(\mathbf{p}, X)$ cannot be Pareto efficient. Therefore any EP solution has the maximum social welfare and the corresponding allocation is optimal.*

Proof. Let Y be an arbitrary optimal allocation, i.e., a maximum weighted matching between consumers and items. For the given non-optimal allocation X , consider the allocation graph $G(X \cup Y)$. Then there is at least one connected component C in $G(X \cup Y)$ such that the total value of X is strictly less than that of Y restricted in C , i.e.,

$$\sum_{(i,j) \in C \cap X} v_{ij} < \sum_{(i,j) \in C \cap Y} v_{ij}.$$

Regarding the structure of the connected component C , we consider the following cases.

- The connected component C consists of the following alternating path of even length:

$$i_1 \xrightarrow{X} j_1 \xrightarrow{Y} i_2 \xrightarrow{X} j_2 \cdots \cdots i_{r-1} \xrightarrow{X} j_{r-1} \xrightarrow{Y} i_r$$

By envy-freeness of $(\mathbf{p}, X)$, we have

$$\begin{aligned} v_{i_1 j_1} - p_{j_1} &\geq 0, \\ v_{i_k j_k} - p_{j_k} &\geq v_{i_k j_{k-1}} - p_{j_{k-1}}, \quad 2 \leq k \leq r-1, \\ 0 &\geq v_{i_r j_{r-1}} - p_{j_{r-1}}. \end{aligned}$$

Adding these inequalities together yields

$$\begin{aligned} \sum_{(i,j) \in C \cap X} v_{ij} &= \sum_{1 \leq k \leq r-1} v_{i_k j_k} \geq \sum_{2 \leq k \leq r} v_{i_k j_{k-1}} \\ &= \sum_{(i,j) \in C \cap Y} v_{ij}, \end{aligned}$$

a contradiction.

- The connected component C consists of the following alternating path of even length:

$$j_1 \xrightarrow{Y} i_1 \xrightarrow{X} j_2 \xrightarrow{Y} i_2 \cdots \cdots j_{r-1} \xrightarrow{Y} i_{r-1} \xrightarrow{X} j_r$$

If for all $1 \leq k \leq r-1$, $p_{j_{k+1}} \geq \max\{v_{i_k j_{k+1}} - v_{i_k j_k}, 0\}$, then we define a new pricing $\mathbf{p}^*$ where $p_{j_k}^* = v_{i_k j_k} - (v_{i_k j_{k+1}} - p_{j_{k+1}})$ for all $1 \leq k \leq r-1$; and $p_{j_r}^*$ can be defined arbitrarily. Consider $\mathbf{p}^*$ with allocation Y : Note that each consumer has the same utility in $(\mathbf{p}^*, Y)$ as in $(\mathbf{p}, X)$, i.e., $u_i(\mathbf{p}^*, Y) = u_i(\mathbf{p}, X)$, and

$$\begin{aligned} \sum_{j \in C \cap Y} p_j^* &= \sum_{k=1}^{r-1} v_{i_k j_k} - (v_{i_k j_{k+1}} - p_{j_{k+1}}) \\ &= \sum_{(i,j) \in C \cap Y} v_{ij} - \sum_{(i,j) \in C \cap X} v_{ij} + \sum_{k=2}^r p_{j_k} \\ &> \sum_{j \in C \cap X} p_j. \end{aligned}$$

That is, the seller obtains more revenue in $(\mathbf{p}^*, Y)$ while ensuring that no consumer obtains less utility. This implies that $(\mathbf{p}, X)$ is not Pareto efficient, a contradiction.

Otherwise, let ℓ be the largest index such that

$$p_{j_\ell} < v_{i_\ell j_{\ell+1}} - v_{i_\ell j_\ell}.$$

By envy-freeness of $(\mathbf{p}, X)$, we have $p_{j_{k+1}} \leq v_{i_k j_{k+1}} - v_{i_k j_k} + p_{j_k}$ for all $1 \leq k \leq r-1$. Hence, we have

$$p_{j_{\ell+1}} < v_{i_\ell j_{\ell+1}} + v_{i_{\ell+1} j_{\ell+2}} - (v_{i_\ell j_\ell} + v_{i_{\ell+1} j_{\ell+1}}).$$

Continue with the same argument following $\ell, \ell + 1, \dots, r - 1, r$, eventually we can get

$$0 \leq p_{j_r} < \sum_{\ell \leq k \leq r-1} v_{i_k j_{k+1}} - \sum_{\ell \leq k \leq r-1} v_{i_k j_k}.$$

This implies that Y is not a maximum matching, a contradiction.

- The connected component C consists of the following alternating path of odd length:

$$i_1 \xrightarrow{X} j_1 \xrightarrow{Y} i_2 \xrightarrow{X} j_2 \cdots \cdots i_{r-1} \xrightarrow{X} j_{r-1} \xrightarrow{Y} i_r \xrightarrow{X} j_r$$

By envy-freeness, we have

$$\begin{aligned} v_{i_1 j_1} - p_{j_1} &\geq 0, \\ v_{i_k j_k} - p_{j_k} &\geq v_{i_k j_{k-1}} - p_{j_{k-1}}, \quad 2 \leq k \leq r. \end{aligned}$$

Adding these inequality together yields

$$\begin{aligned} \sum_{(i,j) \in C \cap X} v_{ij} &\geq \sum_{(i,j) \in C \cap X} v_{ij} - p_{j_r} \\ &= \sum_{1 \leq k \leq r-1} v_{i_k j_k} - p_{j_r} \geq \sum_{2 \leq k \leq r} v_{i_k j_{k-1}} \\ &= \sum_{(i,j) \in C \cap Y} v_{ij}, \end{aligned}$$

a contradiction.

- The connected component C consists of the following alternating path of odd length:

$$i_1 \xrightarrow{Y} j_1 \xrightarrow{X} i_2 \xrightarrow{Y} j_2 \cdots \cdots i_{r-1} \xrightarrow{Y} j_{r-1} \xrightarrow{X} i_r \xrightarrow{Y} j_r$$

If for all $2 \leq k \leq r$, $v_{i_k j_{k-1}} - p_{j_{k-1}} \leq v_{i_k j_k}$, then we can define a new pricing $\mathbf{p}^*$ where $p_{j_1}^* = v_{i_1 j_1}$ and $p_{j_k}^* = v_{i_k j_k} - (v_{i_k j_{k-1}} - p_{j_{k-1}})$ for $2 \leq k \leq r$. Consider $\mathbf{p}^*$ with allocation Y : Note that each consumer has the same utility in $(\mathbf{p}^*, Y)$ as in $(\mathbf{p}, X)$, i.e., $u_i(\mathbf{p}^*, Y) = u_i(\mathbf{p}, X)$, and

$$\begin{aligned}
\sum_{j \in C \cap Y} p_j^* &= v_{i_1 j_1} + \sum_{k=2}^r v_{i_k j_k} - (v_{i_k j_{k-1}} - p_{j_{k-1}}) \\
&= \sum_{(i,j) \in C \cap Y} v_{ij} - \sum_{(i,j) \in C \cap X} v_{ij} + \sum_{k=1}^{r-1} p_{j_k} \\
&> \sum_{j \in C \cap X} p_j.
\end{aligned}$$

That is, the seller obtains more revenue in $(\mathbf{p}^*, Y)$ while ensuring that no consumer obtains less utility. This implies that $(\mathbf{p}, X)$ is not Pareto efficient, a contradiction.

Otherwise, let ℓ be the smallest index such that $v_{i_\ell j_{\ell-1}} - p_{j_{\ell-1}} > v_{i_\ell j_\ell}$. Note that $\ell \geq 2$. By envy-freeness of $(\mathbf{p}, X)$, we have

$$\begin{aligned}
0 &\geq v_{i_1 j_1} - p_{j_1}, \\
v_{i_k j_{k-1}} - p_{j_{k-1}} &\geq v_{j_k i_k} - p_{j_k}, \quad 2 \leq k \leq \ell - 1.
\end{aligned}$$

Adding all these inequalities together yields

$$\sum_{2 \leq k \leq \ell} v_{i_k j_{k-1}} > \sum_{1 \leq k \leq \ell} v_{i_k j_k}.$$

This contradicts to the fact that Y is a maximum matching.

- The connected component C consists of the following alternating cycle:

$$i_1 \xrightarrow{Y} j_1 \xrightarrow{X} i_2 \xrightarrow{Y} j_2 \cdots \cdots j_{r-1} \xrightarrow{X} i_r \xrightarrow{Y} j_r \xrightarrow{X} i_1$$

We denote $j_{r+1} = j_1$. If for all $1 \leq k \leq r$, $p_{j_{k+1}} \geq \max\{v_{i_k j_{k+1}} - v_{i_k j_k}, 0\}$, then we define a new pricing $\mathbf{p}^*$ where $p_{j_k}^* = v_{i_k j_k} - (v_{i_k j_{k+1}} - p_{j_{k+1}})$ for all $1 \leq k \leq r$. Consider $\mathbf{p}^*$ with allocation Y : Note that each consumer has the same utility in

$(\mathbf{p}^*, Y)$ as in $(\mathbf{p}, X)$, i.e., $u_i(\mathbf{p}^*, Y) = u_i(\mathbf{p}, X)$, and

$$\begin{aligned} \sum_{j \in C \cap Y} p_j^* &= \sum_{k=1}^r (v_{i_k j_k} - (v_{i_k j_{k+1}} - p_{j_{k+1}})) \\ &= \sum_{(i,j) \in C \cap Y} v_{ij} - \sum_{(i,j) \in C \cap X} v_{ij} + \sum_{k=1}^r p_{j_k} \\ &> \sum_{j \in C \cap X} p_j. \end{aligned}$$

That is, the seller obtains more revenue in $(\mathbf{p}^*, Y)$ while ensuring that no consumer obtains less utility. This implies that $(\mathbf{p}, X)$ is not Pareto efficient, a contradiction.

Otherwise, let ℓ be an index such that $p_{j_\ell} < v_{i_\ell j_{\ell+1}} - v_{i_\ell j_\ell}$. Assume without loss of generality that $\ell \geq 2$. By envy-freeness of $(\mathbf{p}, X)$, we have $p_{j_{k+1}} \leq v_{i_k j_{k+1}} - v_{i_k j_k} + p_{j_k}$ for all $1 \leq k \leq r$. Hence, we have

$$p_{j_{\ell+1}} < v_{i_\ell j_{\ell+1}} + v_{i_{\ell+1} j_{\ell+2}} - (v_{i_\ell j_\ell} + v_{i_{\ell+1} j_{\ell+1}}).$$

Continue with the same argument following $\ell, \ell+1, \dots, r-1, r, 1, \dots, \ell-1$, eventually we can get

$$0 \leq p_{j_{\ell-1}} < \sum_{(i,j) \in C \cap X} v_{ij} - \sum_{(i,j) \in C \cap Y} v_{ij},$$

which is a contradiction.

Therefore, if any envy-free pricing is Pareto efficient, then the corresponding allocation must be a maximum matching, and therefore, optimal. Hence, for any REP solution, which is envy-free and Pareto efficient, it has the maximum social welfare. $\square$

The above claim implies that if we enforce both envy-freeness and Pareto efficiency, the solution is guaranteed to have maximum social welfare. We have the following (main) conclusion.

Theorem 2.10. *In a market with unit demand consumers, given envy-freeness, an output is Pareto efficient if and only if it maximizes social welfare.*

2.3.2 Determining Pareto Efficiency

In this section, we consider the following question: Given a market with unit demand consumers, can we determine whether an output $(\mathbf{p}, X)$ is Pareto efficient or not? While optimality of an allocation ensures Pareto efficiency (Theorem 2.8), the converse is not necessarily true. Therefore, our approach does not entail Theorem 2.8.

Next we will give an algorithm to determine Pareto efficiency of any given market output $(\mathbf{p}, X)$. For simplicity, denote $r = r(\mathbf{p}, X)$ and $u_i = u_i(\mathbf{p}, X)$. Let N be the number of consumers with positive utility u_i , i.e., $N = |\{i \in A \mid u_i > 0\}|$. We construct a weighted bipartite graph $G(\mathbf{p}, X) = (A, B; E)$ as follows: For each consumer i and each item j , if $v_{ij} - u_i \geq 0$, then we add an edge connecting i and j with weight $w_{ij} = v_{ij} - u_i$. Let W be a sufficiently large number, e.g., $W = n^2 \cdot \max_{ij} v_{ij}$. Further, we define another graph $G^*(\mathbf{p}, X) = (A, B; E)$ derived from $G(\mathbf{p}, X)$ where the only difference is on the weights of edges: In graph G^* , the weight of every edge $(i, j) \in E$ is defined to be $w_{ij}^* = w_{ij} + W$ if $u_i > 0$ and $w_{ij}^* = w_{ij}$ if $u_i = 0$. If the total weight of the maximum matching in the graph $G^*(\mathbf{p}, X)$ exceeds $NW + r$, the algorithm reports that the market output $(\mathbf{p}, X)$ is not Pareto efficient; otherwise, it is Pareto efficient.

Theorem 2.11. *The total weight of the maximum matching in the graph $G^*(\mathbf{p}, X)$ exceeds $NW + r$ if and only if the market output $(\mathbf{p}, X)$ is not Pareto efficient. Therefore, the above algorithm determines whether $(\mathbf{p}, X)$ is Pareto efficient or not in polynomial time.*

Proof. Let X^* be a maximum matching in the graph $G^*(\mathbf{p}, X)$ with total weight W^* ; assume that $W^* > NW + r$. By the construction of $G^*(\mathbf{p}, X)$, this implies that all N consumers with positive utilities u_i are matched in X^* . Consider the market output $(\mathbf{p}^*, X^*)$ to the market, where $p_j^* = w_{ij}^*$ if there is i such that $x_i^* = j$ and $p_j^* = \infty$ otherwise. Hence, $(\mathbf{p}^*, X^*)$ will have a total revenue of

$$\begin{aligned} r(\mathbf{p}^*, X^*) &= \sum_{(i,j) \in X^*} p_j^* = \sum_{(i,j) \in X^*} w_{ij}^* \\ &= \sum_{(i,j) \in X^*} w_{ij}^* - NW = W^* - NW > r. \end{aligned}$$

Further, for every assignment $(i, j) \in X^*$, we have

$$v_{ij} - p_j^* = v_{ij} - w_{ij}^* = v_{ij} - (v_{ij} - u_i) = u_i.$$

Hence, all consumers obtain the same utility in $(\mathbf{p}^*, X^*)$ as in $(\mathbf{p}, X)$ (for those who do not win any item in X^* , their utilities in both market outputs are 0). Therefore, $(\mathbf{p}^*, X^*)$ Pareto-dominates $(\mathbf{p}, X)$, which implies that $(\mathbf{p}, X)$ is not Pareto efficient.

On the other hand, suppose that there is another market output $(\mathbf{p}^*, X^*)$ with revenue $r^* = r(\mathbf{p}^*, X^*)$ that Pareto-dominates $(\mathbf{p}, X)$. By Lemma 2.7, without loss of generality, we can assume that (i) the domination occurs only at revenue, i.e., $r^* > r$; and (ii) all consumers obtain the same utility in $(\mathbf{p}, X)$ and $(\mathbf{p}^*, X^*)$, i.e., $u_i = v_{ij} - p_j^*$ for every $(i, j) \in X^*$. This implies, in particular, all consumers with positive u_i must be matched in X^* . Therefore, the matching X^* in the graph $G(\mathbf{p}, X)$ has total weight

$$\sum_{(i,j) \in X^*} w_{ij} = \sum_{(i,j) \in X^*} v_{ij} - u_i = \sum_{(i,j) \in X^*} p_j^* = r^* > r.$$

Hence, the total weight of X^* in the graph $G^*(\mathbf{p}, X)$ is at least

$$NW + \sum_{(i,j) \in X^*} w_{ij} = NW + r^* > NW + r.$$

This completes the proof. □

2.3.3 Complexity of Computing REP

In this section, we consider the problem of computing a revenue maximization solution that is both envy-free and Pareto efficient (i.e., an REP solution) in a unit-demand market. The following claim says that the problem in general does not admit a polynomial time algorithm. (Indeed, it does not admit a $(1 + \epsilon)$ approximation algorithm for an arbitrarily small $\epsilon > 0$.)

Theorem 2.12. *For any constant $c > 1$, there is no polynomial time algorithm for computing an EP solution with revenue greater than $1/c$ times the revenue of an REP solution, unless $P=NP$.*

Proof. Assume, for a contradiction, that there is a polynomial time c -approximation algorithm A for computing approximate REP solutions. We want to show that A can be used to give a constant approximation for maximum independent set problem. This means we will do a gap preserving reduction from independent set problem.

Given an instance of the independent set problem consisting of a graph $G = (V, E)$ with $|V| = n$ and $|E| = m$, we construct a unit-demand market. Our construction is similar to the one used in proving Theorem 3.1 in [69]. For each vertex v , construct a set of items M_v , $|M_v| = \delta v + 1$, where $\delta(v)$ is the degree of v in G . The set of all items in the markets is $M = \cup_{v \in V} M_v$; of course, $M = |2m + n|$. Define the set of buyers N as follows: for each vertex v , there is a buyer N_v who has a valuation of $c(m + n)$ for each item in set M_v and 0 for other items; for each edge (u, v) , there is a buyer N_{uv} who has a valuation 1 for each items in sets M_u and M_v and 0 for other items; therefore $N = (\cup_{v \in V} \{N_v\}) \cup (\cup_{(u,v) \in E} \{N_{uv}\})$ and $|N| = m + n$.

As we observed, if we have an independent set of size at least k , then the optimal solution of this market has a total revenue of at least $kc(m + n) + m + n - k > kc(m + n)$.

If we have an independent set of size less than k/c , then we can have at most k/c vertex buyers who buy items at price greater than 1 (otherwise there is a buyer N_{uv} who gets nothing in the output, the output is not Pareto efficient), then the optimal solution is at most $(k/c - 1) \cdot c(m + n) + m + n = k(m + n) - c(m + n) + m + n < k(m + n)$.

This completes the reduction. $\square$

Although the above hardness result shows that in general we cannot expect a polynomial time algorithm, the following claim says that for certain special cases, where all items have positive prices in a revenue maximizing WE solution, we can compute an optimal REP solution efficiently. Note that in most applications like advertising markets and housing markets, the items on sale are quite competitive and almost surely all items will be priced positively in a RWE solution.

Theorem 2.13. *Given a market with unit demand consumers, let $(\mathbf{p}, X)$ be an optimal RWE solution. If we have $p_j > 0$ for all items, then $(\mathbf{p}, X)$ is an optimal REP solution as well. Hence, REP can be computed in polynomial time.*

Proof. By the definition of RWE, we know that $(\mathbf{p}, X)$ is one that maximizes the total revenue over all Walrasian equilibria. Consider any optimal REP solution $(\mathbf{q}, Y)$; assume to the contrary that the total revenue given by $(\mathbf{p}, X)$ is less than that given by $(\mathbf{q}, Y)$, i.e., $r(\mathbf{p}, X) < r(\mathbf{q}, Y)$. By Theorem 2.9, we know that Y is an optimal allocation. Hence, according to Lemma 6 of [68], which says that any Walrasian price vector and optimal allocation constitute a WE, $(\mathbf{p}, Y)$ is a WE as well.

Define

$$T = \{j \in B \mid p_j \leq q_j\} \text{ and } T' = \{j \in B \mid p_j > q_j\}.$$

By the assumption, we know that $T \neq \emptyset$ and there is at least one $j \in T$ such that $p_j < q_j$.

Define

$$S = \{i \in A \mid y_i \in T\} \text{ and } S' = \{i \in A \mid x_i \in T'\}$$

to be the subset of consumers that win items in T and T' in the allocations Y and X , respectively. Since $p_j > 0$ for all j and $(\mathbf{p}, X)$ and $(\mathbf{p}, Y)$ are Walrasian equilibria, all items are allocated out in both X and Y ; hence, $|S| = |T|$ and $|S'| = |T'|$. For any consumer $i \in S'$, since $x_i \in T'$, we have $p_{x_i} > q_{x_i}$; this implies that i must be a winner in $(\mathbf{q}, Y)$. Further, for any item $j \in T$, we have $v_{ix_i} - q_{x_i} > v_{ix_i} - p_{x_i} \geq v_{ij} - p_j \geq v_{ij} - q_j$; therefore, i strictly prefers x_i to j at price vector $\mathbf{q}$. These two facts imply that all consumers in S' still win all items in T' in the allocation Y . That is, we can rewrite S' as

$$S' = \{i \in A \mid x_i \in T'\} = \{i \in A \mid y_i \in T'\}.$$

Given $(\mathbf{p}, X)$ and $(\mathbf{q}, Y)$, we define another outcome $(\mathbf{r}, Z)$, where

- $r_j = q_j$ if $j \in T$, and $r_j = p_j$ if $j \in T'$;
- items in T are allocated according to Y , and items in T' are allocated according to X .

We claim that $(\mathbf{r}, Z)$ is a WE. By the above discussion, we know that Z is a feasible matching and all items are allocated out. Further, it can be seen that every winner obtains a non-negative utility at $(\mathbf{r}, Z)$. It remains to show that $(\mathbf{r}, Z)$ is envy-free. For any winner $i \in S$, i does not envy any other item in T by the envy-freeness of $(\mathbf{q}, Y)$; for any item $j \in T'$, we have $v_{iz_i} - r_{z_i} = v_{iy_i} - q_{y_i} \geq v_{ij} - q_j > v_{ij} - p_j = v_{ij} - r_j$, i.e., i does not envy j as well. For any winner $i \in S'$, i does not envy any other item in T' by the envy-freeness of $(\mathbf{p}, X)$; for any item $j \in T$, we have $v_{iz_i} - r_{z_i} = v_{ix_i} - p_{x_i} \geq v_{ij} - p_j \geq v_{ij} - q_j = v_{ij} - r_j$, i.e., i does not envy j as well. Finally, it can be seen that $(\mathbf{r}, Z)$ is envy-free for all losers as well.

Now both $(\mathbf{p}, X)$ and $(\mathbf{r}, Z)$ are Walrasian equilibria that allocate all items, where the former, by assumption, is a revenue maximizing WE. However, as discussed above, we know that $p_j \leq r_j$ for all j and there is $j \in T$ such that $p_j < q_j = r_j$. That is, the total revenue given by $(\mathbf{p}, X)$ is strictly less than that given by $(\mathbf{r}, Z)$, a contradiction. $\square$

2.4 Single-Minded Consumers

We say a consumer i is *single-minded* if there is a subset of items $S_i \subseteq B$, called the demand set of i , such that $v_i(S) = v_i(S_i) > 0$ if $S_i \subseteq S$ and $v_i(S) = 0$ otherwise. That is, i is only interested in obtaining a subset containing S_i . In this section, we assume that all consumers are single-minded and denote $v_i(S_i) = v_i$. Note that when all consumers are single-minded, we can assume without loss of generality that every consumer i either wins subset S_i or wins nothing. Hence, we can encode an allocation $X = (x_i)_{i \in A}$ to be a $(0, 1)$ -indicator vector where $x_i = 1$ implies that i wins subset S_i in X . Note that for any consumers $i \neq i'$ with $x_i = x_{i'} = 1$, we must have $S_i \cap S_{i'} = \emptyset$ in any feasible allocation since all items have unit supply. Given any solution $(\mathbf{p}, X)$, the revenue of the seller can be denoted by $r(\mathbf{p}, X) = \sum_{i \in A} x_i \cdot p(S_i)$.

Our goal is again to consider envy-free solutions which are Pareto efficient (EP) as well, and find one that maximizes the revenue among all such solutions (REP). Surprisingly, in a market with single-minded consumers, an EP solution may not exist at all; and even if one exists, an REP may not.

Example 2.14 (Non-existence of EP). *There are three items j_1, j_2 , and j_3 and six consumers $i_k, k = 1, \dots, 6$. The first three consumers i_1, i_2 , and i_3 are interested in $S_1 = \{j_1, j_2\}, S_2 = \{j_2, j_3\}$ and $S_3 = \{j_3, j_1\}$, each with valuation 2. The remaining consumers i_4, i_5 , and i_6 are interested in items j_1, j_2 , and j_3 , respectively, each with valuation 0.7. Suppose consumer i_1 is a winner in an envy free solution. Consumer i_6 is the only possible other winner. However, if i_6 is a winner, the solution will not be envy free as either i_2 or i_3 will have a valuation that exceeds the price of their demand set. Using similar argument for other cases, we can conclude that any envy free solution will have at most one winner.*

Consider any solution that sells at most two items. Clearly, we can increase the revenue by selling the third item to the appropriate consumer i_k , where $k \in \{4, 5, 6\}$. Therefore, the solution that sells to at most two items cannot be Pareto efficient. This immediately implies that no Pareto efficient solution can have less than two winners. Therefore, in this example, there is no EP solution.

Example 2.15 (Non-existence of REP when EP exists). *There are three items j_1, j_2, j_3 and four consumers $i_1, \dots, i_4$ with demand subsets $S_1 = \{j_1, j_2\}, S_2 = \{j_2, j_3\}, S_3 = \{j_3, j_1\}$, and $S_4 = \{j_1, j_2, j_3\}$, and valuations $v_{i_1} = v_{i_2} = v_{i_3} = 2$ and $v_{i_4} = 2.5$. It is easy to see*

that at most one consumer can be the winner. Further, it can be seen that i_4 cannot be the winner given envy-freeness since there is no solution to the equation system:

$$\begin{aligned} p_{j_1} + p_{j_2} + p_{j_3} &\leq v_{i_4} = 2.5, \\ p_{j_1} + p_{j_2} &\geq v_{i_1} = 2, \\ p_{j_2} + p_{j_3} &\geq v_{i_2} = 2, \\ p_{j_3} + p_{j_1} &\geq v_{i_3} = 2. \end{aligned}$$

Since i_1, i_2, i_3 are symmetric, assume without loss of generality that i_1 is the winner and $p_{j_3} = \infty$. Then i_2, i_3, i_4 all get utility zero. Due to Pareto efficiency, we cannot charge i_1 at price $p(S_1) = p_{j_1} + p_{j_2} = v_{i_1} = 2$, since otherwise it is dominated by another solution where i_4 wins all three items at a total price of 2.5 (the seller obtains more revenue in this solution). Therefore, to guarantee envy-freeness and Pareto efficiency, we have to charge i_1 at a price strictly less than $v_{i_1} = 2$. Hence, any price vector $(p_{j_1}, p_{j_2}, \infty)$ where $p_{j_1} + p_{j_2} < 2$ gives an EP solution. However, in this example, there is no exact REP solution since the revenue of the seller can be arbitrarily close to 2. This example further implies that even if an EP solution exists, it may not maximize social welfare since i_4 should be the winner in the optimal allocation.

Our computational results, however, are quite negative. In particular, we show in Theorem 2.16 that, in addition to computing an optimal REP solution, even testing Pareto efficiency of a given solution and determining the existence of one EP solution is hard.

Theorem 2.16. *In a market with single-minded consumers, the following results hold:*

- *Determining whether a given solution $(\mathbf{p}, X)$ is Pareto efficient is NP-complete.*
- *Determining the existence of an EP solution is NP-complete.*
- *Given the existence of an EP solution, computing (a revenue maximizing) one is NP-hard.*

Proof. We prove the claims one by one.

- We reduce from the Exact Cover problem, where we are given a set system with a universal set U and a collection C of subsets. For our purpose, we assume

without loss of generality that every element in C has cardinality at least 2. An exact cover of the set system is a collection $C' \subseteq C$ such that elements in C' are pairwise disjoint and the union of the sets in C' equals U . The Exact Cover problem, which is known to be NP-complete, asks whether the given set system has an exact cover.

We consider each item in U to be an item on sale in the market. For each set $S \in C$, we have a single-minded consumer who is interested in subset S with valuation $|S|$. Furthermore, we add an extra consumer i who is interested in $S_i = U \setminus \{u\}$ with valuation $|U| - 1$, where $u \in U$ is chosen arbitrarily. Consider the following solution $(\mathbf{p}, X)$, where i wins S_i and all others win nothing, and all items have price one. Hence, all consumers have utility zero and the market maker has revenue $|U| - 1$. Given the structure of the constructed market, it can be seen that $(\mathbf{p}, X)$ is not Pareto efficient if and only if the original instance of the exact cover problem has an exact cover, thus establishing the reduction.

- We use a reduction from Exact Cover by 3-sets. That is, given a family $S = (S_1, \dots, S_n)$ of subsets of $N = \{1, \dots, 3m\}$, where each S_i contains exactly three elements, we are asked whether there are m subsets of S covering all elements in N .

For the given Exact Cover problem, we construct a market with $6m + 2n + 2$ buyers and $3m$ items. Every element of N corresponds to an item for sale. The set of consumers are composed of the following groups:

- $A_0 = \{i_0\}$: i_0 desires all items at valuation $3m - 1$.
- A_1 : For each subset S_i , there is a corresponding consumer $i \in A_1$ with demand subset S_i and valuation 3. Note that $|A_1| = n$.
- A_2 : For each item j , there is a corresponding consumer $i \in A_2$ with demand subset $N \setminus \{j\}$ and valuation $3m - 1.5$. Note that $|A_2| = 3m$.
- A'_0, A'_1, A'_2 : For each consumer in A_0, A_1, A_2 , we make a duplicate consumer with the same demand set and valuation. Notice that the duplicate consumers ensure that for any $i \in A_0 \cup A_1 \cup A_2$, if i is a winner, the total price charged to him has to be his valuation, since we need to guarantee the envy-freeness of the duplicate consumer. Therefore, all consumers in the market will always have zero utility for any envy-free solution. Assume without loss of generality that

consumers in A'_0, A'_1, A'_2 are not winners.

We first characterize all EP solutions (if one exists) for the market constructed above. First, any consumer $i \in A_2$ cannot be a winner. This is because if any such i is a winner, we have to fully charge him at price $3m - 1.5$; this solution, however, is dominated by another one where i_0 is the winner with total payment $3m - 1$ since the seller can obtain more revenue. Second, consumer i_0 cannot be a winner. Otherwise, since the total price charged to i_0 is equal to his valuation $3m - 1$, there is at least one item priced at least $(3m - 1)/3m$; without loss of generality assume that item j is priced higher than or equal to $(3m - 1)/3m$. Then the subset $N \setminus \{j\}$ is priced at most $3m - 1 - (3m - 1)/3m = 3m - 2 + \frac{1}{3m}$. Since $m \geq 1$, we have $3m - 2 + \frac{1}{3m} < 3m - 1.5$; this means that the corresponding consumer in A_2 is not envy-free. Therefore, to guarantee envy-freeness and Pareto efficiency, consumers only from A_1 can be winners.

We claim that the Exact Cover problem has a solution if and only if there is an EP solution for the constructed market. If there are m subsets $\{S_{i_1}, \dots, S_{i_m}\} \subseteq S$ covering all elements in N , we may set the price of all items to 1, and allocate the subset S_{i_k} , $k = 1, \dots, m$, to the corresponding consumer in A_1 . In this case, all items are sold with a total revenue of $3m$. Hence, it is an EP solution, i.e., a solution exists. On the other hand, if there are no m subsets of S covering all elements in N , there are at least 3 items unsold and the total revenue is at most $3m - 3$. This solution, however, is dominated by selling all items to i_0 , making a total revenue of $3m - 1$. Hence, an EP solution does not exist.

In conclusion, the Exact Cover problem has a solution if and only if there is an EP solution in the market. Since the Exact Cover problem is NP-hard, determining whether there is an EP solution is NP-hard as well.

- To prove the final statement, we again reduce from the Exact Cover problem. We construct a market in which the items for sale are the elements of the universe U . Each set S in the family C corresponds to a consumer with demand set S . The market so constructed is guaranteed to have an EP solution — simply price all items at 1 and consider the optimal allocation for this pricing. To the end of computing an optimal REP solution, we may first ask if there is a Pareto efficient envy-free pricing

whose total revenue is $|U|$. Clearly, the answer to this question is yes if and only if the Exact Cover instance that we started with has an exact cover.

□

Despite the hardness results in Theorem 2.16, in a manner similar to Theorem 2.13, we present natural cases where optimal REP solutions can be easily computed.

Theorem 2.17. *Suppose a market with single minded consumers that admits an optimal RWE solution $(\mathbf{p}, X)$ in which $p_j > 0$ for all items j . Then, $(\mathbf{p}, X)$ is an optimal REF solution as well.*

Proof. Suppose there is another envy free solution $(\mathbf{q}, Y)$ which does not allocate all items, but $r(\mathbf{q}, Y) > r(\mathbf{p}, X)$. It is easy to see that there is a consumer i who is unallocated in X , but allocated in Y , and furthermore, $\mathbf{p}(S_i) < \mathbf{q}(S_i)$. However, such a consumer i will be envious in $(\mathbf{p}, X)$ because, had she been allocated in $(\mathbf{p}, X)$, her utility will exceed the utility she gained in $(\mathbf{q}, Y)$. This contradicts our assumption that $(\mathbf{p}, X)$ is envy free. □

Theorem 2.18. *Given a market with single minded consumers, let $(\mathbf{p}, X)$ be an optimal REF solution. If we have $p_j > 0$ for all items, then $(\mathbf{p}, X)$ is an optimal REP solution as well.*

Proof. It suffices to show that REF is Pareto efficient. There are three groups of participants. Group A_1 are the consumers getting their bundle in $(\mathbf{p}, X)$; Group A_2 are the consumers who do not get their bundle in $(\mathbf{p}, X)$; Finally, we have the seller. We analyze case by case to see that their utility can not be improved in any other solution $(\mathbf{q}, Y)$ without hurting others.

Case 1: If the utility of the seller is improved, $q(S_i) > p(S_i)$ for some allocated i in Y . If $i \in A_1$, then i 's utility is reduced; if i is in A_2 , then i is envious $(\mathbf{p}, X)$, contradiction.

Case 2: Suppose that the seller's revenue remains unchanged. If the utility of some consumer j is improved, then $q(S_j) < p(S_j)$. This means that some other consumer j' has to pay $q(S_{j'}) > p(S_{j'})$ in order to ensure that seller's revenue is not diminished. This will mean again that j' either has a reduced utility or was envious in $(\mathbf{p}, X)$, which is a contradiction.

This completes the proof. □

Theorems 2.17 and 2.18 immediately lead to the following corollary.

Corollary 2.19. *Given a market with single minded consumers, let $(\mathbf{p}, X)$ be an optimal RWE solution. If we have $p_j > 0$ for all items, then $(\mathbf{p}, X)$ is an optimal REP solution as well.*

2.5 Concluding Remarks

We study Envy-free Pareto-efficient pricing in the centralized combinatorial market and consider its application in two domains: unit demand consumers and single-minded consumers.

In unit demand case, although we show that it is impossible to approximate optimal Envy-free Pareto-efficient solution within any constant ratio unless $P=NP$, items are in short supply in most auction cases in the real world. Intuitively, that means all items will be allocated at a non-zero price, hence Theorem 2.13 can be applied and optimal Envy-free Pareto-efficient solution can be output in polynomial time.

In single-minded case, we note that the instances constructed in our reductions in Theorem 2.16 are somewhat artificial. Therefore, while the hardness results hold, we expect more natural instances that occur in real world scenarios to allow EP solutions. We believe that for most “natural” instances, an EP solution always exists and the revenue given by REP is close to that given by REF. We leave this as future work.

Chapter 3

Indivisible Resource Allocation without Payment

3.1 Introduction

In this chapter, we still face the situation: Individual participants may want to receive as much utility as they can, while centralized authority may have certain global objectives. This means that fairness and economic efficiency are the factors we need to take into consideration.

We will consider max-min fairness in this chapter, that is, the benchmark is the least allocation among all participants. In the mean time, we take social welfare (efficiency) into account, since it evaluates the performance of the whole system. Maximization of social welfare and max-min fairness yield two natural optimization problems, i.e., maximize the least individual allocation and maximize the total allocation. They are therefore concrete quantities that can be optimized or approximated.

Fairness and welfare are arguably the most critical solution criteria in many applications. An ideal solution would capture both of them simultaneously, and this does occur in some classic examples. For instance, in the bandwidth allocation problem, where in a network with edge capacities there are n source-sink pairs requesting routing flows, Megiddo [98] showed that if all sources are the same, there exist splittable fractional flows that optimize both fairness and efficiency simultaneously¹.

¹Here, ‘splittable’ means one can send flows between one source-sink pair through different paths and

In many other settings, however, one cannot achieve both fairness and welfare simultaneously: a fair allocation may result in a loss in social welfare efficiency, and an efficient allocation may incur unfairness. In this chapter, we will characterize the tradeoff and interconnection between fairness and welfare by quantifying their loss from a centralized viewpoint; our analysis is by the standard approach of approximation. Specifically, we will address the following question.

What is the best possible bifactor approximation ratios of fairness and welfare one can achieve? In other words, for any given approximation ratios r_f and r_w , $0 < r_f, r_w \leq 1$, does there exist an allocation with fairness approximation r_f and welfare approximation r_w , respectively?

We will study this question for two simple, but important, indivisible resource allocation problems: indivisible-item allocation and multiple fractional knapsack problem. Our results are on the positive side—we give constant bifactor approximation schemes for the two problems, and show that our ratios are (almost) optimal.

Indivisible-Item Allocation In this model, there are a set of indivisible items and a function $v(\cdot)$, which gives the valuation of different subsets of the items. We assume the valuation function $v(\cdot)$ is subadditive, i.e., $v(S) + v(T) \geq v(S \cup T)$ for any S and T [61]; subadditive functions are a pretty general class and are a superset of a few well studied functions, e.g., additive, gross substitutes, and submodular. An allocation corresponds to a partition of the items into n groups, denoted by $(A_1, \dots, A_n)$. Thus, its fairness is $\min_i v(A_i)$ and welfare is $\sum_i v(A_i)$.

We show that when $0 < r_f \leq \frac{1}{2}$ and $r_w \leq \frac{1}{r_f+1}$, as well as $\frac{1}{2} < r_f \leq 1$ and $r_w \leq \frac{1}{3r_f}$, there is an allocation that achieves fairness approximation r_f and welfare approximation r_w simultaneously. On the other hand, we show an impossibility result: for any $0 < r_f, r_w \leq 1$ that satisfy $r_w > \frac{1}{r_f+1}$, there exists a subadditive function such that no allocation can approximate fairness and welfare simultaneously with approximation ratios r_f and r_w , respectively. Our results give a tight characterization when $0 < r_f \leq \frac{1}{2}$, and leave a small region open when $\frac{1}{2} < r_f \leq 1$ (the grey area in Figure 3.1).

‘fractional’ means a source-sink demand can be fulfilled fractionally. In addition, precisely, the notion of fairness used in [98] is stronger than max-min fairness; more discussions on this definition are referred to Section 3.4.

Figure 3.1: Fairness and welfare bifactor approximations of indivisible items.

1. The region including and below the curves given by $r_w = \frac{1}{r_f+1}$ ($r_f \leq \frac{1}{2}$) and $r_w = \frac{1}{3r_f}$ ($r_f > \frac{1}{2}$) corresponds to allocations that can be achieved.
2. The region above the curve $r_w = \frac{1}{r_f+1}$ ($0 < r_f \leq 1$) corresponds to bifactor ratios that cannot be approximated theoretically.
3. The grid region given by $r_w < \frac{1}{2} - 3r_f$ corresponds to allocations that can be computed in polynomial time.

Our proof of the existence results, while constructive, may require exponential runtime (indeed, for subadditive functions, finding a fairness or welfare optimal allocation is NP-hard). A further question is therefore what bifactor approximations r_f and r_w are tractable in polynomial time². We show that if welfare can be approximated within a factor of α , there is a polynomial time algorithm that computes an allocation with bifactor approximations r_f and $r_w = \alpha - 3r_f$ for any $r_f < \frac{\alpha}{3}$. Given the known welfare $\frac{1}{2}$ -approximation algorithm [61], our result implies a $\frac{1}{6+\epsilon}$ -approximation algorithm for fairness, and constant bifactor approximation schemes for both fairness and welfare, illustrated in the grid region of Figure 3.1.

²Note that a subadditive function may require a representation that is exponential in the number of items. There are in general two different ways to deal with this issue: *value oracle* (which returns the value $v(S)$ for any subset S) and *demand oracle* (which, for any given price vector, returns a subset that maximizes the difference between the valuation and total price). In a polynomial time algorithm, we are allowed to ask the given oracle at most polynomial times.

Multiple Fractional Knapsack Problem In this model, there are a set N of n consumers, each with a demand d_i , and a set M of m knapsacks, each with a capacity c_j . The resources are the knapsacks; an allocation consists of a mapping $f : N \rightarrow M$ from consumers to knapsacks where every consumer is mapped to exactly one knapsack (i.e., unsplittable) and an assignment vector $X = (x_1, x_2, \dots, x_n)$, where x_i is the amount that consumer i receives from the corresponding knapsack. Note that within a knapsack, the assigned amount to each mapped consumer can be fractional, and must satisfy the demand and capacity constraints. The welfare of an allocation (f, X) is defined as $\sum_i x_i$ and fairness is defined as $\min_i x_i$.

The multiple knapsack problem is a natural generalization of the classic knapsack and bin packing problems, see, e.g., [35, 95] and the references within. Our specific model is motivated from applications like, for instance, each consumer subscribes to and obtains services (e.g., data usage) from one service provider (e.g., mobile phone or cable network carriers). Our model can be considered as an unsplittable fractional flow (i.e., bandwidth allocation) problem on a special network with two nodes s and t and m parallel edges between them, which is a basic resource allocation problem that involves fairness and welfare questions.

Regarding fairness and welfare bifactor approximations, we show that when $0 < r_f \leq 1$ and $r_w \leq \frac{1}{r_f+1}$, there is an allocation that achieves r_f and r_w simultaneously. On the other hand, we show that all points on this curve are the best we can achieve. Note that the curve $r_w = \frac{1}{r_f+1}$ coincides with the one in the indivisible-item allocation model, but there is a slight difference: In both models our results are tight when $0 < r_f \leq \frac{1}{2}$, but on the domain $\frac{1}{2} < r_f \leq 1$ we only have a tight result for the knapsack model.

It can be noticed that there is a huge fairness loss if one wants to achieve welfare optimum. The reason of it is that one can (artificially) create many knapsacks with tiny capacities. For practical applications, in the literature it is commonly assumed that the knapsacks are *balanced*, i.e., the smallest knapsack is at least the largest³ demand (see, e.g., the survey [86]). For such balanced settings, we show an improved approximation scheme: when $0 < r_f \leq 1$ and $r_w \leq \frac{2}{2r_f+1}$, there is an allocation that achieves r_f and r_w simultaneously. Our results are illustrated in Figure 3.2.

We again give explicit algorithms to derive the above tight ratios. In the algorithm, we

³To derive our results, it is sufficient to assume that the smallest knapsack is at least the smallest demand, as the latter gives an upper bound on the fairness optimum.

Figure 3.2: Fairness and welfare bifactor approximations for multiple fractional knapsack.

1. The region including and below the curve given by $r_w = \frac{1}{r_f + 1}$ ($0 < r_f \leq 1$) corresponds to allocations that can be achieved. The region above this curve corresponds to bifactor ratios that cannot be approximated theoretically.
2. The region including and below the curves given by $r_w = 1$ ($0 < r_f \leq \frac{1}{2}$) and $r_w = \frac{2}{2r_f + 1}$ ($\frac{1}{2} < r_f \leq 1$) corresponds to allocations that can be achieved, provided that all knapsacks are balanced. The region above these curves corresponds to bifactor ratios that cannot be approximated theoretically.

need to solve an optimal solution of fairness and welfare. In our knapsack model, computing a fairness optimal allocation can be solved by a simple polynomial time algorithm; but the computation of a welfare optimal solution is NP-hard. However, we show that if welfare can be approximated within a ratio of α , there is a polynomial time algorithm to compute bifactor approximations $r_w \leq \frac{\alpha}{r_f + 1}$ and $r_w \leq \frac{2\alpha}{2r_f + 2}$ for the unbalanced and balanced models, respectively. That is, the two curves in Figure 3.2 are scaled down by a factor of α .

3.1.1 Related Work

In the indivisible resource allocation setting, the algorithmic study of maximizing fairness or welfare has a rich history in the literature. For the fairness maximization problem, Woeginger [132] studied additive functions and gave a polynomial time approximation scheme. As any allocation maximizes social welfare for additive functions, this result immediately

implies that there exist bifactor ratios $r_f = 1 - \epsilon$ and $r_w = 1$ for additive functions. To the best of our knowledge, our $\frac{1}{6+\epsilon}$ -approximation algorithm is the first constant approximation algorithm to maximize fairness with a subadditive function. A more generalized model is that different individuals may have different valuation functions. For this setting, the fairness maximization problem becomes notoriously challenging, and all known approximation ratios, e.g., additive functions [9, 15, 32, 50], submodular functions [66], and subadditive functions [82], are either polynomial or polylogarithmic.

Computing a welfare maximizing allocation is a classic combinatorial optimization problem. For submodular functions, Lehmann et al. [91] showed a $\frac{1}{2}$ -approximation greedy algorithm. Later, Dobzinski and Schapira [56] gave a $(1 - \frac{1}{e})$ -approximation algorithm given a demand oracle. Finally, Vondrak [130] showed a $(1 - \frac{1}{e})$ -approximation algorithm in the value oracle model. When the valuation functions are subadditive, Mirrokni et al. [100] gave a $\frac{1}{m^{1/2}-\epsilon}$ -approximation algorithm in the value query model, and Feige [61] showed a $\frac{1}{2}$ -approximation algorithm in the demand query model.

In bandwidth allocation, Megiddo [98] gave an allocation with optimal fairness and welfare for single source networks with splittable fractional flows. If all sources and sinks are not necessarily the same, one can easily show that there exist allocations that achieve bifactor ratios r_f and $r_w = 1 - r_f$, respectively, for any $0 \leq r_f \leq 1$ and these ratios are tight. When flows are unsplittable, there is a rich body of literature on its algorithmic studies, see, e.g., [83, 86] and the references within.

The tradeoff between fairness and welfare have been considered in many other problems, e.g., cake cutting [33, 43], congestion games [34], load balancing in job scheduling [8, 65, 85, 125] and bandwidth allocation [40, 64, 85, 90]. Some of these works considered an even larger scope of simultaneous multi-objective optimizations; for instance, in [64, 90] the goal is to compare with all possible feasible allocations, not only fairness and welfare optimum. We comment that an important feature of our algorithms, of which most previous work does not possess, is that the approximations consider all possible ratios of r_f and r_w , given by the curves in Figure 3.1. Such approximation schemes explicitly reveal the tradeoff between fairness and welfare, and provide a clear illustration for the centralized authority to choose a desirable allocation.

3.2 Indivisible-Item Allocation

We are given a set $M = \{a_1, \dots, a_m\}$ of indivisible items. There is a function $v(\cdot)$ which gives the valuation for any subset $S \subseteq M$. For simplicity, we denote $v(a) = v(\{a\})$ for any single item $a \in M$. We assume that $v(\emptyset) = 0$. In this section, we will consider monotone *subadditive* functions, i.e., $v(S) \leq v(T)$ for any $S \subset T \subseteq M$, and $v(A \cup B) \leq v(A) + v(B)$ for all $A, B \subseteq M$. By subadditivity, we can assume that $v(a) > 0$ for any $a \in M$, since otherwise, such items have no contribution to any subset and can be simply discarded.

An *allocation* of M is a partition into n disjoint subsets $A = (A_1, A_2, \dots, A_n)$, i.e., $A_i \subseteq M$, $A_1 \cup \dots \cup A_n = M$ and $A_i \cap A_j = \emptyset$ for any $i \neq j$. For any allocation $(A_1, A_2, \dots, A_n)$, its *social welfare* is defined as $\sum_i v(A_i)$, and its *fairness* is defined as $\min_i v(A_i)$. For a given instance, let

$$W_{opt} = \max \left\{ \sum_i v(A_i) \mid (A_1, A_2, \dots, A_n) \text{ is a valid allocation} \right\}$$

$$F_{opt} = \max \left\{ \min_i v(A_i) \mid (A_1, A_2, \dots, A_n) \text{ is a valid allocation} \right\}$$

denote the optimal welfare and optimal fairness value among all valid allocations, respectively. By the definition, a simple observation is that $W_{opt} \geq n \cdot F_{opt}$. It is well known that both welfare and fairness are NP-hard to optimize. We say an allocation $A = (A_1, A_2, \dots, A_n)$ approximates welfare with an approximation ratio of α if $\sum_i v(A_i) \geq \alpha \cdot W_{opt}$. Similarly, we say it approximates fairness with an approximation ratio of α if $\min_i v(A_i) \geq \alpha \cdot F_{opt}$. Note that $0 \leq \alpha \leq 1$, and our objective is to find allocations with as large approximation ratio as possible.

We will study the interconnection (or tradeoff) between fairness and welfare in terms of bifactor approximations r_f and r_w . That is, for any given two ratios r_f and r_w , $0 \leq r_f, r_w \leq 1$, does there always exist an allocation such that fairness and welfare can be simultaneously approximated with ratios r_f and r_w , respectively? And what is the largest possible values to which the question admits an affirmative answer? In the following subsections, we will study the questions from both positive and negative sides. Note that we will only consider those instances with $F_{opt} > 0$ (precisely, it should be $F_{opt} \geq \min_{a \in M} v(a)$ since every subset will have at least one item), since if $F_{opt} = 0$ (which is easy to verify), the problem

degenerates to a pure welfare maximization problem.

3.2.1 Existence

Our main theorem of this section is the following.

Theorem 3.1. *For any subadditive valuation function $v(\cdot)$ and two positive constants r_f, r_w satisfying one of the following conditions*

- $0 < r_f \leq \frac{1}{2}$ and $r_w \leq \frac{1}{r_f+1}$
- $\frac{1}{2} < r_f \leq 1$ and $r_w \leq \frac{1}{3r_f}$

there always exists an allocation that approximates fairness and welfare simultaneously with approximation ratios r_f and r_w , respectively.

Our proof is constructive: we will give an explicit algorithm that, starting from a welfare optimal solution, computes an allocation with the desired approximations by a number of reassignments. Note that the proof is about existence; thus, we will ignore computational issues.

We first give two simple algorithms which will be used in the main algorithm. In both of these two algorithms, we start with a given threshold t and an allocation with good welfare. Then we give a reassignment procedure that converts the given allocation to another one in which every agent's value is no less than t , and the reassignment will only lose a small fraction of welfare.

Given a threshold t , we can divide all the items in two categories: we call item a a *big item* if $v(a) \geq t$, otherwise we call it a *small item*. These two algorithms will deal with large and small items, respectively.

Lemma 3.2. *Given a threshold t and an initial allocation $(A_1, \dots, A_n)$, let $k = |\{a \in M \mid v(a) \geq t\}|$ be the number of big items with respect to t , and let $B_i = \{a \in A_i \mid v(a) \geq t\}$ and $S_i = A_i \setminus B_i$. Then ALG-BIG-ITEM can successfully return an allocation if and only if $|\{S_i \mid v(S_i) < t\}| \leq k$, and the returned allocation satisfies the following two properties:*

- *Every agent's valuation is no less than t .*
- *The loss of welfare is no more than nt (compared to the initial allocation).*

Algorithm 1 ALG-BIG-ITEM(n, t, M, A)

Input:

number of agents n ,
threshold t ,
item set $M = \{a_1, \dots, a_m\}$,
initial allocation $A = (A_1, A_2, \dots, A_n)$.

- 1: **while** there exists A_i such that $v(A_i) < t$ **do**
- 2: **if** there exist A_j and a big item $a \in A_j$ such that $v(A_j \setminus a) \geq t$ **then**
- 3: Let $A_j \leftarrow A_j \setminus \{a\}$.
- 4: Let $A_i \leftarrow A_i \cup \{a\}$.
- 5: **else**
- 6: Return *null*.
- 7: **end if**
- 8: **end while**
- 9: Return $(A_1, A_2, \dots, A_n)$.

Proof. First, notice that in the algorithm, items in each S_i will never be moved and only big items will be rearranged. If $|\{S_i \mid v(S_i) < t\}| > k$, since there are only k big items, at any step of the algorithm, at least one A_i will have $v(S_i) < t$ and there are no big items in it. Thus, $v(A_i) < t$. And according to the algorithm, it will never return an allocation in this case. On the other hand, assume that $|\{S_i \mid v(S_i) < t\}| \leq k$. Hence, at any step of the algorithm, if there is A_i with $v(A_i) < t$, then there must exist another agent j such that either A_j contains at least two big items or A_j contains one big item and $v(S_j) \geq t$. By the algorithm, it will move one big item from A_j to A_i . Therefore, at the end of the algorithm, we will have $v(A_i) \geq t$ for all A_i , and the algorithm will not return *null*.

Next we analyze the loss of welfare during the reassignment. Notice that in the algorithm, the only step that will change the allocation is the one where an item a is moved from A_j to A_i . We will call A_j a *big set* and A_i a *small set* for this operation. Note that after every such an operation, both the big and small set will have value no less than t . This means that every set A_i can be picked at most once as a small set; thus, there will be at most n operations performed in the course of the algorithm. And for each operation, our loss in welfare is

$$\begin{aligned} & (v(A_i) + v(A_j)) - (v(A_i \cup \{a\}) + v(A_j \setminus \{a\})) \\ \leq & (v(A_i) + v(A_j)) - (v(a) + v(A_j \setminus \{a\})) \leq v(A_i) + v(A_j) - v(A_j) = v(A_i) \leq t \end{aligned}$$

Hence the total welfare loss during the whole algorithm will be no more than nt . $\square$

As one can see, ALG-BIG-ITEM tries to produce our desired allocation from an arbitrary allocation by only rearranging big items. This is the first step we will try in our main algorithm. And if it does not successfully return an allocation, we will put all these big items aside to a set of agents, each one with one big item. Then we will use another algorithm ALG-SMALL-ITEM to allocate all remaining small items to the remaining agents. Thus, in ALG-SMALL-ITEM, we require that every item has its single value less than t .

Algorithm 2 ALG-SMALL-ITEM(n, t, M, A)

Input:

number of agents n ,
threshold t ,
item set $M = \{a_1, \dots, a_m\}$ where $v(a_i) < t$ for all $a_i \in M$,
initial allocation $A = (A_1, A_2, \dots, A_n)$.

- 1: **while** there exists A_i such that $v(A_i) < t$ **do**
- 2: **if** there exists A_j such that $v(A_j) \geq 3t$ **then**
- 3: Let $S = \emptyset$.
- 4: Add items from A_j to S one by one, until the point when $t \leq v(S) < 2t$.
- 5: Let $A_j \leftarrow A_j \setminus S$.
- 6: Let $A_i \leftarrow A_i \cup S$.
- 7: **else**
- 8: Return an arbitrary fairness optimal allocation.
- 9: **end if**
- 10: **end while**
- 11: Return $(A_1, A_2, \dots, A_n)$.

This algorithm has a similar flavor to ALG-BIG-ITEM for the part where items are moved from a big set to a small set. Our next lemma, similar to Lemma 3.2, says that when the threshold is small enough, algorithm ALG-SMALL-ITEM will always return a fair allocation with good welfare.

Lemma 3.3. *For any threshold $t \leq F_{opt}$, ALG-SMALL-ITEM will always return an allocation such that*

- *every agent's valuation is no less than t ,*
- *the loss of welfare is no more than $\max\{nt, n(3t - F_{opt})\}$ (compared to the initial allocation).*

In addition, if the initial allocation $(A_1, A_2, \dots, A_n)$ is an α -approximation of welfare and $t \leq \frac{\alpha F_{opt}}{3}$, then ALG-SMALL-ITEM will always return an allocation at Step 11.

Proof. First it is obvious to see that in the returned allocation (either by Step 8 or 11), every agent's valuation must be no less than t for any threshold $t \leq F_{opt}$.

We next examine at which step the reassignment procedure lose welfare. We focus on the step where set S is moved from A_j to A_i . Similar to the proof of Lemma 3.2, again we call A_j a *big set* and A_i a *small set* for this operation. After such an operation, we have

$$t \leq v(S) \leq v(A_i \cup S) \leq v(A_i) + v(S) < t + 2t = 3t$$

and

$$v(A_j \setminus S) \geq v(A_j) - v(S) > 3t - 2t = t.$$

Thus, both of these two agents will have value no less than t . Further, since $v(A_i \cup S) < 3t$, a small set performed in an operation will never be picked again as a small set or a big set in the future operations. Let $\delta_i = (v(A_i) + v(A_j)) - (v(A_i \cup S) + v(A_j \setminus S))$ denote the loss in welfare during this operation. Notice that we can use such notation δ_i without ambiguity because every A_i can be picked at most once as a small set. We have

$$\begin{aligned} \delta_i &= (v(A_i) + v(A_j)) - (v(A_i \cup S) + v(A_j \setminus S)) \\ &\leq (v(A_i) + v(A_j)) - (v(S) + v(A_j \setminus S)) \\ &\leq v(A_i) + v(A_j) - v(A_j) \\ &= v(A_i) \leq t \end{aligned}$$

and

$$\delta_i + v(A_i \cup S) = v(A_i) + v(A_j) - v(A_j \setminus S) \leq v(A_i) + v(S) < 3t.$$

On the other hand, if some set A_i has value at least t in the original allocation, we define $\delta_i = 0$ as it will never be chosen as a small set in any operation. Depending on the returned allocation, there are the following two cases.

- Algorithm returns at Step 11. Notice that $\sum_i \delta_i$ is the total welfare loss between the original allocation and the returned allocation. Since $\delta_i \leq t$ for all $1 \leq i \leq n$, we

have $\sum_i \delta_i \leq nt$.

- Algorithm returns at Step 8. Let us consider the step right before the execution of Step 8. Suppose the allocation at that moment is $A' = (A'_1, A'_2, \dots, A'_n)$. By the algorithm, we know that $v(A'_i) < 3t$ for all $1 \leq i \leq n$; hence, we have $\delta_i + v(A'_i) < 3t$ for all $1 \leq i \leq n$. This implies that $\sum_i v(A_i) = \sum_i v(A'_i) + \sum_i \delta_i < 3t \cdot n$. On the other hand, any fairness optimal allocation will have welfare no less than nF_{opt} . Thus, the social welfare loss from the original allocation will be no more than $3tn - nF_{opt} = n(3t - F_{opt})$.

Therefore, the total loss of welfare is upper bounded by $\max\{nt, n(3t - F_{opt})\}$.

Finally, from the above argument we know that if the algorithm returns at Step 8, we have $\sum_i v(A_i) < 3tn$. Thus, if the original allocation A is an α -approximation of welfare, we know $t > \frac{\sum_i v(A_i)}{3n} \geq \frac{\alpha W_{opt}}{3n} \geq \frac{\alpha F_{opt}}{3}$. This means that if $t \leq \frac{\alpha F_{opt}}{3}$, the algorithm will never return at Step 8. This completes the proof. $\square$

Having the above two algorithms handling different values of the items, now we can construct our main algorithm, denoted by ALG-MAIN-ITEM. Again given a threshold t , the main algorithm will transform a welfare optimal solution to an allocation such that fairness is no less than t without losing much value in welfare. The algorithm is described formally as follows.

Lemma 3.4. *For any given threshold t , if $t \leq F_{opt}$, ALG-MAIN-ITEM will always return an allocation such that*

- *every agent's valuation is no less than t ,*
- *the loss of welfare is no more than $\max\{nt, n(3t - F_{opt})\}$ compared with W_{opt} .*

Proof. First, if $|\{S_i \mid v(S_i) < t\}| \leq k$, the algorithm will use the output given by ALG-BIG-ITEM. By Lemma 3.2, we know that in this case ALG-BIG-ITEM will always successfully return an allocation which satisfies all the requirements. Thus, in the following it suffices to consider the case where $|\{S_i \mid v(S_i) < t\}| > k$.

In ALG-MAIN-ITEM we will first reassign all small items to $n - k$ agents such that welfare is maximized (Step 7); let W'_{opt} denote this maximum welfare value. Notice that if we discard the smallest k sets in $\{S_1, S_2, \dots, S_n\}$, the remaining $n - k$ sets will form a

Algorithm 3 ALG-MAIN-ITEM(n, t, M)

Input:

- number of agents n ,
 - threshold t ,
 - item set $M = \{a_1, \dots, a_m\}$.
 - 1: Let $A = (A_1, A_2, \dots, A_n)$ be a welfare optimal allocation.
 - 2: Let $B_i = \{a \in A_i \mid v(a) \geq t\}$ be the set of big items in A_i , and let $S_i = A_i \setminus B_i$ be the set of small items in A_i .
 - 3: Suppose without loss of generality that $\bigcup_i B_i = \{a_1, a_2, \dots, a_k\}$, where $k = |\{a \in M \mid v(a) \geq t\}|$ is the number of big items in M .
 - 4: **if** $|\{S_i \mid v(S_i) < t\}| \leq k$ **then**
 - 5: Return ALG-BIG-ITEM(n, t, M, A)
 - 6: **else**
 - 7: Let $A' = \{A'_1, \dots, A'_{n-k}\}$ be a welfare optimal allocation with item set $\bigcup_i S_i$ and $n - k$ agents.
 - 8: Let $(A''_1, \dots, A''_{n-k}) = \text{ALG-SMALL-ITEM}(n - k, t, \bigcup_i S_i, A')$.
 - 9: Return allocation $(\{a_1\}, \dots, \{a_k\}, A''_1, \dots, A''_{n-k})$.
 - 10: **end if**
-

feasible allocation with items from a subset of $\bigcup_i S_i$ into $n - k$ agents. Since the values of the k smallest sets in $\{S_1, S_2, \dots, S_n\}$ are all less than t , the social welfare of this specific allocation will be larger than $\sum_i v(S_i) - kt$. Hence, if we allocate all items in $\bigcup_i S_i$ to $n - k$ agents, the optimal welfare must have $W'_{opt} > \sum_i v(S_i) - kt$.

Since there are in total k big items with respect to threshold t , in a fairness optimal allocation, there are at least $n - k$ agents that only contain small items. Since the value of each agent is no less than F_{opt} , when assigning all small items $\bigcup_i S_i$ to $n - k$ agents, the optimal fairness value will be at least F_{opt} . Thus, when we run ALG-SMALL-ITEM($n - k, t, \bigcup_i S_i, A'$), since $t \leq F_{opt}$ and the initial allocation A' is a welfare optimal allocation, by Lemma 3.3, the algorithm will always return an allocation such that: (i) each agent's valuation is no less than t ; and (ii) the welfare of the allocation is no less than $W'_{opt} - \max\{(n - k)t, (n - k)(3t - F_{opt})\}$. Therefore, for the final returned allocation $(\{a_1\}, \dots, \{a_k\}, A''_1, \dots, A''_{n-k})$, every agent's valuation is no less than t ; and for welfare

we have

$$\begin{aligned}
\sum_{i=1}^k v(a_i) + \sum_{i=1}^{n-k} v(A_i'') &\geq \sum_i v(B_i) + W'_{opt} - \max\{(n-k)t, (n-k)(3t - F_{opt})\} \\
&\geq \sum_i v(B_i) + \sum_i v(S_i) - kt - (n-k) \cdot \max\{t, 3t - F_{opt}\} \\
&\geq \sum_i v(A_i) - \max\{nt, n(3t - F_{opt})\} \\
&= W_{opt} - \max\{nt, n(3t - F_{opt})\}.
\end{aligned}$$

This completes the proof. $\square$

Now we are ready to prove our main theorem.

Proof of Theorem 3.1. First, if $nF_{opt} \geq r_w W_{opt}$, then a fairness optimal allocation (whose fairness approximation ratio is 1 by definition) will have welfare approximation ratio at least r_w , which is just an allocation that we need. It suffices to consider the case where $nF_{opt} < r_w W_{opt}$. We will run algorithm $\text{ALG-MAIN-ITEM}(n, r_f F_{opt}, M)$ where the threshold $t = r_f F_{opt}$. By Lemma 3.4, we know that it will always return an allocation with fairness approximation ratio r_f . For welfare, we consider the following two cases.

- $0 < r_f \leq \frac{1}{2}$. In this case, we know that $nr_f F_{opt} \geq n(3r_f - 1)F_{opt} = n(3t - F_{opt})$. Since $r_w \leq \frac{1}{r_f + 1}$, we have $1 - r_f r_w \geq r_w$. Thus, by Lemma 3.4, the welfare is no less than

$$W_{opt} - nr_f F_{opt} > W_{opt} - r_f r_w W_{opt} = (1 - r_f r_w)W_{opt} \geq r_w W_{opt}.$$

- $\frac{1}{2} < r_f \leq 1$. In this case we have $nr_f F_{opt} < n(3r_f - 1)F_{opt} = n(3t - F_{opt})$. Because $r_f \leq \frac{1}{3r_w}$, we have $1 - 3r_f r_w \geq 0$. Thus, again by Lemma 3.4, the welfare is no less than

$$\begin{aligned}
W_{opt} - n(3r_f - 1)F_{opt} &> W_{opt} - (3r_f - 1)r_w W_{opt} \\
&= (1 - 3r_f r_w + r_w)W_{opt} \geq r_w W_{opt}.
\end{aligned}$$

Therefore, in both cases, the allocation always has a welfare approximation ratio of at least r_w , which completes the proof of the theorem. $\square$

3.2.2 Impossibility

In this subsection we will focus on upper bounds. We will show that actually $r_w \leq \frac{1}{r_f+1}$ is the best that we can do in this setting. More specifically, we show the following theorem.

Theorem 3.5. *For any given constants $0 \leq r_f, r_w \leq 1$ that satisfy $r_w > \frac{1}{r_f+1}$, there exists a subadditive valuation function $v(\cdot)$ such that no allocation can approximate fairness and welfare simultaneously with approximation ratios r_f and r_w , respectively.*

Proof. Let $r_w(r_f + 1) = 1 + x$, where $x > 0$. Pick n to be a sufficiently large number such that $r_f > \frac{1}{n-1}$ and $n > \frac{1+x}{x}$. Assume there are n agents and $2n - 1$ items $M = \{a_1, a_2, \dots, a_{n-1}, b_1, b_2, \dots, b_n\}$. For any nonempty subset $S \subseteq M$, if $S \cap \{b_1, b_2, \dots, b_n\} = \emptyset$, let $v(S) = r_f - \frac{1}{n-1}$; otherwise, let $v(S) = |S \cap \{b_1, b_2, \dots, b_n\}|$. It is easy to verify that $v(\cdot)$ is a subadditive function with following properties:

- The unique welfare optimal allocation is given by $A_i = \{a_i\}$ for $1 \leq i \leq n - 1$ and $A_n = \{b_1, b_2, \dots, b_n\}$, with $W_{opt} = (n - 1)(r_f - \frac{1}{n-1}) + n = (n - 1)(r_f + 1)$.
- A fairness optimal allocation is given by $b_i \in A_i$ for all $1 \leq i \leq n$ ($\{a_1, a_2, \dots, a_{n-1}\}$ can be arbitrarily assigned), with $F_{opt} = 1$.

In order to have an allocation $(A_1, A_2, \dots, A_n)$ that approximates fairness by a factor of r_f , it needs to have $v(A_i) \geq r_f$ for all $1 \leq i \leq n$. In this case, we have $v(A_i) \geq 1$, which means that $|A_i \cap \{b_1, b_2, \dots, b_n\}| = 1$ and $v(A_i) = 1$ for all $1 \leq i \leq n$. Hence, the welfare of this allocation will be n , and the social welfare approximation ratio satisfies

$$\begin{aligned} \frac{n}{(n-1)(r_f+1)} &= \frac{n}{(n-1)(1+x)} \cdot r_w = \frac{1}{(1-\frac{1}{n})(1+x)} \cdot r_w \\ &< \frac{1}{(1-\frac{x}{1+x})(1+x)} \cdot r_w = r_w \end{aligned}$$

This implies that there is no allocation that can approximate fairness and social welfare simultaneously within factor r_f and r_w , respectively, with respect to the constructed valuation function. $\square$

3.2.3 Polynomial Time Algorithms

While the existence result characterized by Theorem 3.1 illustrates interconnections between approximation ratios of fairness and welfare, one important problem left is whether we can find such allocations in polynomial time for subadditive functions. Unfortunately, Theorem 3.1 does not imply that we can find those allocations in polynomial time, as in the ALG-MAIN-ITEM we need to compute welfare optimal solutions, which is NP-hard. Further, in the algorithm, we use $t = r_f F_{opt}$ as a threshold to distinguish big and small items, but F_{opt} is NP-hard to compute as well.

To fix these problems, we will use approximate welfare and fairness solutions to replace optimal allocations in the algorithm and to estimate the value of F_{opt} . Given an approximation algorithm for welfare as a black box, we can first use it to find a fairness approximation value, and then use them together to get a polynomial time bifactor approximation algorithm.

Theorem 3.6. *For any subadditive function $v(\cdot)$, assume that there exists a polynomial time α -approximation algorithm for welfare, where $0 < \alpha \leq 1$ is a constant. Then for any constants $r_f < \frac{\alpha}{3}$ and $r_w < \alpha - 3r_f$, there exists a polynomial time algorithm that computes an allocation with fairness and social welfare approximation ratios r_f and r_w , respectively.*

Proof. First, we make the following three changes to ALG-MAIN-ITEM and ALG-SMALL-ITEM:

- (1) At the first step of ALG-MAIN-ITEM, we start from an α -approximate welfare allocation A instead of an optimal one.
- (2) In Step 7 of ALG-MAIN-ITEM, instead of using a welfare optimal solution, we use the one with larger welfare between the following two allocations:
 - $(S_1 \cup S_2 \cup \dots \cup S_{k+1}, S_{k+2}, \dots, S_n)$ (assume that $v(S_1) \leq v(S_2) \leq \dots \leq v(S_n)$),
 - an α -approximate welfare allocation.
- (3) In Step 8 of ALG-SMALL-ITEM, return *null* instead of a fairness optimal allocation.

It's easy to check that after making these changes, ALG-MAIN-ITEM can now be done in polynomial time. Note that change (2) is to make sure that the initial allocation used

to call ALG-SMALL-ITEM is not only an α -approximate welfare allocation, but also with social welfare no less than $\sum_i v(S_i) - kt$, as both of these two conditions are required in the proof of Lemma 3.4. Thus, by Lemma 3.3 we know that for any given threshold $t \leq \frac{\alpha F_{opt}}{3}$, ALG-SMALL-ITEM will never return *null* at Step 8. This further implies that, by a similar argument to the proof of Lemma 3.4, in the returned allocation every agent's valuation is no less than t and the loss of welfare is no more than nt (compared to an optimal solution) when $t \leq \frac{\alpha F_{opt}}{3}$.

The next step is to approximate the optimal fairness value F_{opt} . We can do a polynomial time binary search to find the largest integer threshold t' such that the modified ALG-MAIN-ITEM always successfully returns an allocation⁴. It can be seen that through the binary search, we can have $t' \geq \frac{\alpha F_{opt}}{3+\epsilon}$, where $\epsilon > 0$ is any given small constant. Denote the returned allocation for the final value of t' by A' . Note that in the allocation A' , by Lemma 3.4, every agent's valuation is no less than t' ; thus, its welfare is no less than nt' .

Having this t' as an approximation of F_{opt} , we now let $t = \frac{(3+\epsilon)r_f t'}{\alpha}$, where we choose ϵ to be sufficiently small such that $t \leq t'$, and run the modified ALG-MAIN-ITEM again with the following further change:

(1') At the first step of ALG-MAIN-ITEM, we choose an initial allocation according to the following rule:

- if $\alpha W_{opt} \geq nt'$, use the α -approximate welfare allocation A ,
- if $\alpha W_{opt} < nt'$, use the allocation A' .

By the defined value t and the above discussions, our algorithm will always successfully return an allocation, and the fairness of the final returned allocation is no less than $t = \frac{(3+\epsilon)r_f t'}{\alpha} \geq r_f F_{opt}$. For welfare, if $\alpha W_{opt} \geq nt'$, since allocation A has welfare no less than αW_{opt} , we will use A and the final returned allocation is no less than

$$\alpha W_{opt} - nt = \alpha W_{opt} - \frac{(3+\epsilon)r_f nt'}{\alpha} \geq (\alpha - (3+\epsilon)r_f)W_{opt} \geq r_w W_{opt},$$

where the last inequality follows, again, by choosing ϵ to be sufficiently small. If $\alpha W_{opt} <$

⁴Note that the given subadditive function $v(\cdot)$ may have non-integer values. To implement the binary search on the range $[\min_{a \in M} v(a), v(M)]$, we can multiply a large number to the value of each subset to ensure that the value of every subset is larger than a fixed integer; this can be done easily in the queries to the oracle. It can be seen that the runtime of the binary search is bounded by $\log v(M)$, thus, in polynomial time.

nt' , we will use A' whose welfare is at least nt' and the final returned allocation is no less than

$$nt' - nt = nt' \left(1 - \frac{(3 + \epsilon)r_f}{\alpha} \right) > (\alpha - (3 + \epsilon)r_f)W_{opt} \geq r_w W_{opt}.$$

In either case, we will have fairness at least $r_f F_{opt}$ and welfare at least $r_w W_{opt}$; this yields the desired allocation. $\square$

As a corollary from the proof of the above Theorem 3.6, we directly get a polynomial time fairness approximation algorithm for any subadditive function.

Corollary 3.7. *For any subadditive function $v(\cdot)$, assume that there exists a polynomial time α -approximation algorithm for welfare. Then there exists a polynomial time $\frac{\alpha}{3+\epsilon}$ -approximation algorithm for fairness for any constant $\epsilon > 0$.*

For welfare, there are a $\frac{1}{2}$ -approximation algorithm for subadditive functions [61] and a $(1 - \frac{1}{e})$ -approximation algorithm for submodular functions [56, 130], given a demand oracle. Combining these algorithms with the above results, we have the following claim.

Corollary 3.8. *For any given constant $r_f < \frac{1}{6}$, there is a polynomial time algorithm that has bifactor approximation ratios r_f and any $r_w < \frac{1}{2} - 3r_f$ for subadditive functions in the demand query model. Further, the ratio is $r_f < \frac{e-1}{3e}$ and $r_w < \frac{e-1}{e} - 3r_f$ for submodular functions.*

3.2.4 Heterogeneous Agents

A very natural generalization of our problem is to consider the case where all agents are heterogeneous and have different valuation functions. Suppose that agent i has a valuation function $v_i(\cdot)$. Given an allocation $A = (A_1, A_2, \dots, A_n)$, welfare is defined as $\sum_i v_i(A_i)$, and fairness is $\min_i v_i(A_i)$ ⁵.

Next we present a negative result, showing that one cannot hope to have an allocation that approximate both social welfare and fairness well, even for additive functions.

Theorem 3.9. *For any given constants $r_f > 0$ and $r_w > 0$, there exist n additive functions $v_1, v_2, \dots, v_n$ such that there is no allocation that approximates fairness and welfare with approximation ratio r_f and r_w , respectively, even when all functions are normalized.*

⁵Given that now the agents may have different maximum values $v_i(M)$ over all items, one may define the measure of fairness as $\min_i \left\{ \frac{v_i(A_i)}{v_i(M)} \right\}$. The negative result of Theorem 3.9 applies to this definition as well.

Proof. Assume there are n items, i.e., $M = (a_1, a_2, \dots, a_n)$. Since these n valuation functions are all additive, we only need to define their values on each single item. For $1 \leq i \leq n-1$, let $v_i(a_i) = \epsilon$, $v_i(a_{i+1}) = 1 - \epsilon$, and $v_i(a_j) = 0$ for all other $j \neq i, i+1$. For $i = n$, let $v_n(a_n) = \epsilon$, $v_n(a_{n-1}) = 1 - \epsilon$ and $v_n(a_j) = 0$ for all other $j \neq n, n-1$.

Since there are equal number of items and agents, the only way to have a positive approximate fairness solution is to let each agent have exactly one item with positive value. Thus, we must allocate item a_i to agent i for all $1 \leq i \leq n-2$, and allocate items a_{n-1} and a_n to agents $n-1$ and n . Thus, the total welfare given by this solution will be no more than $(n-2)\epsilon + 2$. However, if we allocate item a_{i+1} to agent i for all $1 \leq i \leq n-1$, the welfare is at least $(n-1)(1-\epsilon)$. Hence, if we want to have an allocation with positive fairness approximation ratio, then welfare approximation ratio will be no more than $\frac{(n-2)\epsilon+2}{(n-1)(1-\epsilon)}$. Pick $n = \lceil \frac{3}{r_f} + 2 \rceil$ and $\epsilon = \frac{1}{n-1}$, we have

$$\frac{(n-2)\epsilon+2}{(n-1)(1-\epsilon)} = \frac{\frac{n-2}{n-1}+2}{n-2} < \frac{3}{n-2} \leq r_f.$$

This finishes the proof of the theorem. $\square$

3.3 Multiple Fractional Knapsack Problem

In this problem, we have a set N of n consumers, each with a demand $d_i \in \mathbb{N}^+$ for $i = 1, \dots, n$, and a set M of m knapsacks, each with a capacity $c_j \in \mathbb{N}^+$ for $j = 1, \dots, m$. Each consumer is assigned to exactly one knapsack; within a knapsack, the allocation that each assigned consumer obtains can be arbitrary and fractional, subject to the capacity constraint. Formally, a valid allocation is an assignment function $f : N \rightarrow M$ and a vector $X = (x_1, x_2, \dots, x_n)$ under the constraints that (i) $x_i \leq d_i$ for all $1 \leq i \leq n$, and (ii) $\sum_{i:f(i)=j} x_i \leq c_j$ for all $1 \leq j \leq m$. Here, $f(i)$ denotes the knapsack that consumer i is assigned to, and x_i is the amount of resources that i receives from knapsack $f(i)$. The welfare of an allocation (f, X) is defined as $\sum_i x_i$ and fairness is defined as $\min_i x_i$.

Similar to the previous section, define

$$W_{opt} = \max \left\{ \sum_i x_i \mid (f, X) \text{ is a valid allocation} \right\}$$

$$F_{opt} = \max \left\{ \min_i x_i \mid (f, X) \text{ is a valid allocation} \right\}$$

to be the optimal welfare and optimal fairness values, respectively. Again given bifactor ratios r_f and r_w , we will consider if there exists an allocation with fairness of at least $r_f F_{opt}$ and welfare of at least $r_w W_{opt}$.

3.3.1 Tight Ratios

We have the following result, which gives an upper bound on the existence of bifactor approximations.

Theorem 3.10. *For any given multiple fractional knapsack problem and two positive constants $0 < r_f, r_w \leq 1$ satisfying $r_w \leq \frac{1}{r_f+1}$, there always exists an allocation that approximates fairness and welfare simultaneously with approximation ratios r_f and r_w , respectively.*

Similar to the previous section, our proof is by constructing an algorithm that produces such an allocation. In the algorithm, similar to ALG-MAIN-ITEM, we again start with a threshold and an initial allocation, then try to perform a sequence of reassignments to convert this initial allocation to another one with desired fairness, but with only a small fraction of welfare loss.

Algorithm 4 ALG-KNAPSACK($t, n, m, (f, X)$)

Input:

- threshold t ,
 - n agents each with demand at least t ,
 - m knapsacks,
 - initial allocation (f, X) .
 - 1: **while** there exists $x_i < t$ **do**
 - 2: **if** there exist $x_j \geq 2t$ **then**
 - 3: Let $f(i) = f(j)$ and $x_i = t$.
 - 4: Let $x_j = x_j - t$.
 - 5: **else**
 - 6: Return an arbitrary fairness optimal allocation.
 - 7: **end if**
 - 8: **end while**
 - 9: Return (f, X) .
-

Lemma 3.11. *For any threshold $t \leq F_{opt}$ and initial allocation (f, X) , ALG-KNAPSACK always returns an allocation such that*

- *fairness is no less than t ,*
- *the loss of welfare is no more than nt (compare to the initial allocation).*

Proof. Given the initial allocation (f, X) , let $k = |\{x_i \mid x_i < t\}|$ be the number of consumers whose initial allocation is less than t . Notice that during each reassignment procedure (step 3 and step 4 of the algorithm), the number of such consumers will be reduced by 1, and the welfare loss during this reassignment procedure is exactly x_i , which is less than t . Hence, if algorithm returns at step 9, the total welfare loss will be less than kt , which is no more than nt .

If ALG-KNAPSACK returns at step 6. Consider the step right before the execution of step 6, let (f', X') be the allocation at that moment. According to the algorithm, no consumer gets an allocation larger than $2t$ in (f', X') . Let k' be the number of consumers whose assigned amount is larger than t in (f', X') . Note that if a consumer receives an allocation less than t in the initial allocation, she can never has an allocation strictly more than t at any step of the algorithm. Thus we have $k' \leq n - k$, and the welfare of allocation (f', X') satisfies $\sum_i x'_i < (n - k') \cdot t + k' \cdot 2t = nt + k't$. Further, since this allocation has a welfare loss no more than kt compared to the initial allocation (f, X) , we know that the welfare of the initial allocation is no more than $nt + k't + kt \leq 2nt$. On the other hand, any fairness optimal allocation has a welfare of at least $nF_{opt} \geq nt$. This means that if the algorithm returns at step 6, the welfare loss will be no more than nt compared to the initial allocation.

Finally, it is easy to check that in either case, the returned allocation has fairness at least t . This completes the proof. $\square$

Given the above lemma, we are ready to prove Theorem 3.10.

Proof of Theorem 3.10. First, if $nF_{opt} \geq r_w W_{opt}$, then any fairness optimal allocation yields the desired welfare approximation ratio. If $nF_{opt} < r_w W_{opt}$, let (f, X) be a welfare optimal allocation, we run $\text{ALG-KNAPSACK}(r_f F_{opt}, n, m, (f, X))$ where the threshold $t = r_f F_{opt}$. By Lemma 3.11, the returned allocation of the algorithm must have a fairness

approximation ratio no less than r_f , and its welfare approximation ratio is no less than

$$\frac{W_{opt} - nr_f F_{opt}}{W_{opt}} > \frac{W_{opt} - r_f r_w W_{opt}}{W_{opt}} = 1 - r_f r_w \geq r_w.$$

This completes the proof of the theorem. $\square$

On the other hand, we have the following lower bound result, which, together with Theorem 3.10, gives tight approximation ratios.

Theorem 3.12. *For any given constants $0 < r_f, r_w \leq 1$ that satisfy $r_w > \frac{1}{r_f+1}$, there exists a multiple fractional knapsack problem instance such that no allocation can approximate fairness and welfare with approximation ratios r_f and r_w , respectively.*

Proof. For any given constants $0 < r_f, r_w \leq 1$ that satisfy $r_w > \frac{1}{r_f+1}$, let ϵ be a sufficiently small constant such that $r_f - \epsilon$ is a rational number. Let $n = \lfloor \frac{1}{\epsilon} \rfloor$. Let p be an integer such that $(r_f - \epsilon)p$ is also integral. Assume that there are n consumers with demand $d > np$ each, and n knapsacks with capacities $c_1 = np, c_2 = c_3 = \dots = c_n = (r_f - \epsilon)p$. It is easy to check the following observations.

- Welfare optimal allocation maps each consumer to a different knapsack and assigns all capacities of the knapsack to that consumer. This gives $W_{opt} = np + (n-1)(r_f - \epsilon)p$.
- Fairness optimal allocation maps all consumers to the first knapsack, which gives $F_{opt} = p$.

In order for an allocation to approximate fairness by a factor of r_f , i.e., each consumer receives an allocation no less than $r_f \cdot p$, no consumer can be mapped to any of the knapsacks $2, 3, \dots, n$. Thus, the welfare of this allocation will be no more than $c_1 = np$, which means that the welfare approximation of this allocation is no more than

$$\frac{np}{np + (n-1)(r_f - \epsilon)p} = \frac{n}{n + (n-1)(r_f - \epsilon)} \leq \frac{n}{n + (n-1)(r_f - 1/n)}$$

where the last inequality is because $\epsilon \leq \frac{1}{n}$. Since $r_w > \frac{1}{r_f+1}$, the above formula is strictly less than r_w for a sufficiently large n . This implies that in the constructed knapsack instance, no allocation can approximate fairness and social welfare simultaneously within factors r_f and r_w , respectively. $\square$

3.3.2 Balanced Knapsacks

In this subsection we consider a special case when all knapsacks are *balanced*, i.e., $\max_i d_i \leq \min_j c_j$. (Technically, a weaker condition $\min_i d_i \leq \min_j c_j$ suffices for our results in this section.) This is a pretty common assumption used in the literature for practical applications [86]. For balanced knapsacks, we have a stronger upper bound on bifactor approximations.

Theorem 3.13. *For any given balanced multiple fractional knapsack problem and two constants $0 < r_f, r_w \leq 1$ satisfying $r_w \leq \frac{2}{2r_f+1}$ when $r_f \geq \frac{1}{2}$, there always exists an allocation that approximates fairness and welfare simultaneously with approximation ratios r_f and r_w , respectively.*

Note that an implication of the above result is that there is an allocation with approximation ratios $r_f = \frac{1}{2}$ and $r_w = 1$. Thus, the same existence result holds for all $0 < r_f \leq \frac{1}{2}$ and $0 < r_w \leq 1$.

The proof is again constructive. Given an allocation (f, X) , we first consider the following algorithm, which reassigns allocations within a given knapsack.

Algorithm 5 RELLOCATE-BIN($j, n, (f, X)$)

Input:

knapsack j ,
 n agents,
allocation (f, X) .

- 1: Let $S_j = \{i \in N \mid f(i) = j\}$ and $\hat{d}_j = \min_{i \in S_j} d_i$.
 - 2: **if** $\frac{c_j}{|S_j|} \leq \hat{d}_j$ **then**
 - 3: Let $x_i = \frac{c_j}{|S_j|}$ for $i \in S_j$.
 - 4: **else**
 - 5: Let $z = c_j - |S_j| \cdot \hat{d}_j$.
 - 6: **for each** $i \in S_j$ **do**
 - 7: Let $y = \min\{d_i - \hat{d}_j, z\}$.
 - 8: Let $x_i = \hat{d}_j + y$ and $z = z - y$.
 - 9: **end for**
 - 10: **end if**
 - 11: Return (f, X) .
-

In the reallocation, we first try to divide the whole capacity evenly among all consumers. If this average share is larger than any demand, based on such a shared value we allocate

the remaining capacity to consumers arbitrarily, subject to their demand constraints. Let X be the allocation returned by RELLOCATE-BIN, since the size of each knapsack is at least $\min_i d_i$ which is always larger than or equal to F_{opt} , then for any $i \in S_j$, either $x_i \geq F_{opt}$ or $x_i = \frac{c_j}{|S_j|}$. This observation is crucial to our analysis below.

Our main algorithm ALG-BALANCED-KNAPSACK is given below.

Algorithm 6 ALG-BALANCED-KNAPSACK($t, n, m, (f, X)$)

Input:

threshold t ,
 n agents each with demand at least t ,
 m knapsacks,
initial allocation (f, X) .

- 1: **for** $j = 1, \dots, m$ **do**
- 2: RELLOCATE-BIN($j, n, (f, X)$).
- 3: Let $S_j = \{i \in N \mid f(i) = j\}$.
- 4: **end for**
- 5: **while** there exists $x_i < t$ **do**
- 6: Pick a knapsack j that satisfies $c_j \geq (|S_j| + 1)t$.
- 7: Let $f(i) = j$, $S_{f(i)} = S_{f(i)} \setminus \{i\}$ and $S_j = S_j \cup \{i\}$.
- 8: RELLOCATE-BIN($j, n, (f, X)$).
- 9: RELLOCATE-BIN($f(i), n, (f, X)$).
- 10: **end while**
- 11: **Return** (f, X) .

Lemma 3.14. *For any threshold $t \in [\frac{F_{opt}}{2}, F_{opt}]$ and an initial allocation (f, X) , ALG-BALANCED-KNAPSACK always returns an allocation such that*

- *fairness is no less than t ,*
- *the loss of welfare is no more than $\frac{n(2t - F_{opt})}{2}$ compared to the initial allocation.*

Proof. For any given initial allocation (f, X) , let $k = |\{x_i \mid x_i < t\}|$ be the number of consumers whose initial allocation is less than t . First we show that whenever the algorithm gets to step 6, such a knapsack always exists.

Assume to the contrary that for any $j \neq f(i)$, $c_j < (|S_j| + 1)t$. Then we have $\lfloor \frac{c_j}{t} \rfloor < |S_j|$. Since $x_i < t \leq F_{opt}$, by the above observation, the allocation of all consumers in the

knapsack $f(i)$ equals x_i . Therefore, we have $x_i = \frac{c_{f(i)}}{|S_{f(i)}|} < t$, which implies that

$$\sum_{j=1}^m \left\lfloor \frac{c_j}{t} \right\rfloor < \sum_{j=1}^m |S_j| = n.$$

On the other hand,

$$\sum_{j=1}^m \left\lfloor \frac{c_j}{t} \right\rfloor \geq \sum_{j=1}^m \left\lfloor \frac{c_j}{F_{opt}} \right\rfloor \geq n,$$

where the second inequality is because a fairness optimal solution always exists. This gives a contradiction.

In addition, for each run of the “while” loop from step 5 to step 10, the number of consumers whose allocation is less than t is reduced by 1. Therefore, the algorithm terminates in at most n steps with fairness at least t for each consumer.

Next we examine welfare loss in the algorithm. First, note that in step 8, since we move consumer i to knapsack j , there is one more consumer to share the capacity c_j and no welfare loss. Thus, welfare loss can only occur in step 9 at knapsack $f(i)$. Consider any execution of step 9. Since $x_i < t \leq F_{opt}$, by the balanced condition and the above observation, before running the corresponding “while” loop, we have

$$\sum_{i' \in S_{f(i)}} x_{i'} = \frac{c_{f(i)}}{|S_{f(i)}|} \cdot |S_{f(i)}| = c_{f(i)} \geq \min_{i \in N} d_i \geq F_{opt}.$$

Therefore, at that moment $S_{f(i)}$ contains at least two consumers.

After running step 9, if there is any welfare loss at knapsack $f(i)$, then each consumer in $S_{f(i)} \setminus \{i\}$ will receive an allocation at least F_{opt} (otherwise their total allocation is still $c_{f(i)}$). Therefore, the amount of welfare loss at knapsack $f(i)$ is at most

$$c_{f(i)} - (|S_{f(i)}| - 1) \cdot F_{opt} < t \cdot |S_{f(i)}| - (|S_{f(i)}| - 1) \cdot F_{opt} \leq 2t - F_{opt}.$$

Finally, for each such knapsack $f(i)$, welfare loss occurs at most once (the moment when the total capacity $c_{f(i)}$ is not fully fulfilled). Further, each of these knapsacks contains at least two consumers in the initial allocation, and at least one will never be moved to other knapsacks in the algorithm. Thus, welfare loss at step 9 occurs at most $\frac{n}{2}$ times, which leads to a total welfare loss of at most $\frac{n(2t - F_{opt})}{2}$. $\square$

Now we are ready to prove the main result.

of Theorem 3.13. First, if $nF_{opt} \geq r_w W_{opt}$, then we can simply output a fairness optimal allocation, which has the desired welfare approximation ratio. If $nF_{opt} < r_w W_{opt}$, let (f, X) be a welfare optimal allocation, we run $\text{ALG-BALANCED-KNAPSACK}(r_f F_{opt}, n, m, (f, X))$ with threshold $t = r_f F_{opt}$ and $r_f > \frac{1}{2}$. By Lemma 3.14, the returned allocation of the algorithm must have a fairness approximation ratio no less than r_f , and its welfare approximation ratio is no less than

$$\frac{W_{opt} - n \frac{2r_f F_{opt} - F_{opt}}{2}}{W_{opt}} > \frac{W_{opt} - (r_f - 1/2)r_w W_{opt}}{W_{opt}} = 1 - (r_f - 1/2)r_w \geq r_w.$$

This completes the proof of the theorem. $\square$

We finally give a lower bound on the bifactor approximation ratios, showing that the ratios given in Theorem 3.13 are tight.

Theorem 3.15. *For any given constant $0 < r_w \leq 1$ that satisfies $r_w > \frac{2}{2r_f+1}$ when $r_f > \frac{1}{2}$, there exists a balanced knapsack problem instance such that no allocation can approximate fairness and welfare with approximation ratios r_f and r_w , respectively.*

Proof. Consider any constants r_f and r_w satisfying $r_w > \frac{2}{2r_f+1}$. Let ϵ be a sufficiently small constant such that $2r_f - \epsilon > 1$ and is a rational number. Let $2r_f - \epsilon = \frac{p}{q}$ where p and q are two integers; note that $2q > p > q$. Now consider the following balanced knapsack instance: there are $n = 3m + 2$ consumers with demand q each, and $m + 1$ knapsacks with capacities $c_1 = c_2 = \dots = c_m = p$ and $c_{m+1} = (m + 1)q$. It is easy to see that the constructed instance is balanced. We have the following observations for fairness and welfare optimum.

- If m consumers use the first m knapsacks and other consumers share the last one, each consumer receives value q , we get $F_{opt} = q$.
- If $2m$ consumers use the first m knapsacks and other consumers share the last one, we get $W_{opt} = mp + (m + 1)q$.

To approximate fairness by a factor of $r_f > \frac{1}{2}$, each consumer will receive a value no less than $r_f \cdot q$. Therefore at most 1 consumer can be allocated in each of the first

m knapsacks. In this case, maximum welfare is obtained by allocating m consumers to the first m knapsacks and others to the last knapsack, which is $(2m + 1)q$. Note that $\frac{(2m+1)q}{mp+(m+1)q} = \frac{2m+1}{m(2r_f-\epsilon)+(m+1)}$. Since $r_w > \frac{2}{2r_f+1}$, for sufficiently large m and sufficiently small ϵ , we have $r_w > \frac{2m+1}{m(2r_f-\epsilon)+(m+1)}$. This implies that for the above balanced knapsack instance, no allocation can simultaneously approximate fairness and welfare within ratios r_f and r_w , respectively. $\square$

3.3.3 Polynomial Time Computation

In our algorithms ALG-KNAPSACK and ALG-BALANCED-KNAPSACK, we use both fairness and welfare optimum. In this subsection, we address the computational issue of these two problems.

For fairness, its optimum can be computed in polynomial time using a standard binary search approach: Suppose an instance has n consumers with demands $d_1, \dots, d_n$ and m knapsacks with capacities $c_1, \dots, c_m$. First, for any value t , the problem admits an allocation with fairness t if and only if $t \leq \min_i d_i$ and $\sum_{j=1}^m \lfloor \frac{c_j}{t} \rfloor \geq n$. Further, notice that F_{opt} is always in the set $\{\frac{c_j}{i} \mid 1 \leq i \leq n, 1 \leq j \leq m, \frac{c_j}{i} \leq \min_i d_i\} \cup \{\min_i d_i\}$, which is of size $O(mn)$. Thus we can perform a binary search over this set to find the optimal fairness value F_{opt} , and a fairness optimal allocation can be produced accordingly. The total running time is $O((n + m) \cdot \log mn)$.

For welfare, its optimum is NP-hard to compute. This can be seen by a simple reduction from the partition problem: There are two knapsacks with capacity $\sum_i d_i/2$ each. Then the optimal value of social welfare is $\sum_i d_i$ if and only if the partition instance given by $d_1, \dots, d_n$ has a positive answer. However, one can show that a simple greedy algorithm gives 0.5-approximation. Thus, similar to previous section, given a welfare approximation algorithm as a blockbox, we are able to get a polynomial time bifactor approximations algorithm.

Theorem 3.16. *If there is a polynomial time α -approximation algorithm for welfare, then*

- *for any unbalanced knapsack instance, and any constants $0 < r_f, r_w \leq 1$ satisfying $r_w \leq \frac{\alpha}{r_f+1}$;*
- *for any balanced knapsack instance, and any constants $r_f \geq \frac{1}{2}$ and $0 < r_w \leq 1$ satisfying and $r_w \leq \frac{2\alpha}{2r_f+1}$,*

there exists a polynomial time algorithm that computes an allocation with fairness and welfare approximation ratios r_f and r_w , respectively.

Proof. We only prove the first claim, and the second one follows similarly. For any constants r_f and r_w satisfying $r_w \leq \frac{\alpha}{r_f+1}$, in order to have polynomial runtime, we use a welfare α -approximation allocation (f, X) and run $\text{ALG-KNAPSACK}(r_f F_{opt}, n, m, (f, X))$. In addition, after the algorithm returns an allocation, we add one more step to compare the returned allocation to a fairness optimal allocation, and use the one with larger welfare as the final output.

It is easy to check that the whole process can be done in polynomial time and the output allocation always has fairness approximation ratio no less than r_f . Regarding welfare, if $nF_{opt} \geq r_w W_{opt}$, it is obvious that the output allocation has a welfare approximation ratio no less than r_w . If $nF_{opt} < r_w W_{opt}$, by Lemma 3.11, the returned allocation has a welfare approximation ratio no less than

$$\frac{\alpha W_{opt} - nr_f F_{opt}}{W_{opt}} > \frac{\alpha W_{opt} - r_f r_w W_{opt}}{W_{opt}} = \alpha - r_f r_w \geq r_w.$$

This completes the proof. □

3.4 Concluding Remarks

We consider the resource allocation problem in two domains, indivisible-item allocation and multiple fractional knapsack, and give nearly optimal bifactor approximation schemes for fairness and welfare. As mentioned in Introduction, our approximation schemes consider all possible ratios of r_f and r_w , and explicitly reveal the tradeoff between fairness and welfare. It is an intriguing future direction to study the tradeoff between fairness and welfare using such approximation schemes in other resource allocation problems, e.g., load balancing, facility location and bandwidth allocation.

Our focus is on *existence*, that is, for given ratios r_f and r_w , if there is an allocation achieving these ratios for fairness and welfare, respectively; thus, randomization cannot help. In particular, notice that one can easily have a random allocation with expected fairness $\alpha \cdot F_{opt}$ and expected welfare $(1 - \frac{\alpha}{2}) \cdot W_{opt}$ for any $0 \leq \alpha \leq 1$, i.e., the straight line between $(0, 1)$ and $(1, 0.5)$ in Figure 3.2; but such a (fixed) allocation does not exist.

In addition, we also consider computational complexity of finding fairness and welfare bicriteria allocations. The two lines of Figure 3.1 indicate the gap between efficient computation and theoretical limit. Closing the gap is an interesting question; specifically, what is the extreme one can achieve to approximate fairness and welfare in polynomial time?

Finally, we remark that the fairness considered in our work is the least allocation, a concrete quantity. A more generalized fairness notion is that of max-min allocation vector. That is, considering an allocation as a vector, ordering all coordinates non-decreasingly, a fairness optimal allocation is one that maximizes this vector lexicographically. In other words, we cannot increase one's allocation without decreasing those who already have a smaller allocation. In our indivisible-item allocation model, it can be shown that any allocation vector with F_{opt} being the least allocation is a coordinate-wise 2-approximation to this stronger fairness. Hence, our algorithms still yield constant approximations. A further question for future study is to consider different fairness measures and their tradeoffs with welfare using the approximation scheme framework.

Chapter 4

Divisible Resource Allocation without Payment

4.1 Introduction

Fairness and efficiency together are among the most critical factors considered in many resource allocation problems. An ideal solution is to have an optimum for both, but this may not necessarily be achievable. In this chapter, we will consider economic efficiency optimization given the required fairness condition. Specifically, we address following question which captures economic efficiency and fairness simultaneously:

What is the computational complexity of computing an efficiency optimal allocation while ensuring fairness?

We will study this question in an divisible resource allocation model – cake cutting, where the resource is a continuous cake and can be divided into pieces and allocated to n individuals. Cake cutting has been a central problem studied in computer science and social science; see, e.g., [118]. An important feature of cake cutting is that the cake is divisible and individuals may have different values for different portions of the cake. The problem is therefore represented by a valuation measure function for each of the individuals.

Our interest is to find a division of the cake that maximizes efficiency given fairness. As we see in Chapter 1, there are mainly several notions that captures fairness in resource allocations: proportionality, envy-freeness, exact division, equitable division and max-min fairness. We will focus on proportionality in this chapter.

Cake cutting to maximize efficiency given fairness was first considered in [43] where the authors gave a linear program to compute an envy-free allocation with the maximum possible efficiency for piecewise constant measure functions. Their algorithm can be generalized easily to a polynomial time approximation scheme for any functions that can be represented concisely¹, and immediately implies the same result for proportional fairness.

The solution generated by the linear program, however, involves a large number of cuts and the allocation to each participant consists of many separate ‘tiny’ pieces, which is not desirable in some applications (e.g., division of a territory). We therefore focus on divisions with exactly $n - 1$ cuts, that is, everyone obtains a connected region. Note that with $n - 1$ cuts, a proportional allocation always exists, e.g., the well known moving knife procedure [58] generates a fair allocation with $n - 1$ cuts for n participants.

Our Results

We will study two types of measure functions: linear and piecewise constant. For linear functions, we give a polynomial time approximation scheme (PTAS) that computes an $(1 - \epsilon)$ -approximate efficiency optimal allocation. That is, for any constant $\epsilon > 0$, the solution given by the scheme is within a factor of $(1 - \epsilon)$ to the optimum. Note that this is the best we can hope for as for linear functions, an optimal solution may involve irrational numbers. Our algorithm is based on a critical ingredient of determining an optimal order of participants to cut the cake: While the intuition is that the optimal order would be according to the slopes of the linear functions, we show that this is actually not the case and give a precise characterization of the order.

For piecewise constant functions, finding an efficiency optimal allocation turns out to be much harder. While for any given fixed order of participants, an exactly optimal solution can still be found using a dynamic programming approach, the problem is NP-hard to solve without knowing an optimal order, even for normalized valuations. For general piecewise constant functions, we show that it is even NP-hard to find an allocation whose efficiency is within a factor of $\Omega(\frac{1}{\sqrt{n}})$ to the optimum. (We refer to [109] as a classic textbook for computational complexity and approximations.)

¹We can divide the cake into small pieces and approximate each small piece by a constant; this leads to piecewise constant functions.

Related Work

Maximizing efficiency in cake cutting with fairness has been studied extensively in the past. [116] studied the computation of a Pareto-efficient envy-free allocation. [106] provided an envy-freeness preserved procedure to improve social efficiency. [33] introduced the concept of price of fairness to characterize the loss of efficiency at a cost of fairness. [7] followed the notion of price of fairness and focused on divisions with only connected pieces. In contrast with our work which is to find exact allocations, the work [7, 33] addressed on the worst case ratio between the optimum efficiency without and with fairness conditions. Note that their benchmark is optimum efficiency without fairness, whereas in our efficient algorithm design the benchmark is optimum efficiency with fairness. The work closest to ours was by [43], where the authors considered efficiency optimization given the envy-free fairness condition, and studied several valuation functions. Note that in the work of [43], there is no restriction on the number of cuts, whereas in ours we only allow $n - 1$ cuts.

4.2 Preliminaries

There are n consumers to share a cake, which is represented by the interval $[0, 1]$; we will call ‘0’ the left side and ‘1’ the right side. Each consumer i has a density function $f_i(\cdot) : [0, 1] \rightarrow \mathbb{R}^+ \cup \{0\}$. The valuation of consumer i for the interval $[a, b] \subseteq [0, 1]$ is denoted by $v_i(a, b) = \int_a^b f_i(x) dx$. Assume without loss of generality that $v_i(0, 1) > 0$ (otherwise, we can simply remove such consumers). For any $S \subseteq [0, 1]$, let $v_i(S)$ denote the valuation of consumer i for S , i.e., the sum of the valuations over disjoint intervals in S .

An *allocation* (i.e., division) of the cake is a partition into n disjoint subsets, denoted by $A = (A_1, A_2, \dots, A_n)$. We require allocations to be fair, captured by *proportionality*: $v_i(A_i) \geq \frac{1}{n}v_i(0, 1)$ for any i , i.e., every consumer gets at least an average share with respect to its own measure. Our objective is to find a proportional allocation that maximizes *efficiency*, defined as $\sum_i v_i(A_i)$. We will restrict our attention on exactly $n - 1$ cuts. That is, we can divide the cake by only $n - 1$ cuts; thus, each consumer will obtain a connected interval.

Our interest is to study efficiency maximization from an algorithmic viewpoint. We say an algorithm is an α -approximation if the efficiency of its solution is always within a factor of α to the optimum efficiency, both conditioned on proportionality. In other words, the

compared benchmark of an algorithm is the maximum efficiency among all proportional allocations with $n - 1$ cuts.

In this chapter we will consider two special families of density functions that can be represented concisely: linear valuations and piecewise constant valuations.

4.3 Linear Functions

In this section we will consider linear density functions, i.e., each $f_i(\cdot)$ can be represented as $f_i(x) = a_i \cdot x + b_i$, where a_i, b_i are two given numbers. We assume that $f_i(x) \geq 0$ for any $x \in [0, 1]$; this implies in particular that $b_i \geq 0$. Further, we have $v_i(0, 1) = \frac{a_i}{2} + b_i > 0$. Linear function is one of the most fundamental and well studied valuation functions. It reflects gradually increasing or decreasing interests of consumers for the resource.

We will give a polynomial time approximation scheme (PTAS) to compute an efficiency optimal allocation. First we will characterize the order of allocations to the consumers in an optimal solution. Having this characterization, we will then use a dynamic programming algorithm to compute a nearly optimal solution.

4.3.1 Optimal Order

As there are in total $n - 1$ cuts, an allocation corresponds to a permutation of the n consumers. To design an algorithm to compute an efficiency optimal allocation, we first need to decide what order of the consumers should be. For the considered linear functions, one might expect that the optimal order, from left to right, would be according to the increasing order of the slope a_i . Counter-intuitively, this is not the case, as the following example shows.

Example 4.1. Consider two consumers $\{1, 2\}$ with density functions $f_1(x) = x$ and $f_2(x) = 2x + 10000$. Clearly, in an optimal allocation, consumer 2 will get as much as he can to maximize efficiency. If the order is 1, 2 from left to right according to slope, then consumer 1 needs to get $[0, \frac{\sqrt{2}}{2}]$ to ensure his proportional share and consumer 2 gets the rest. The efficiency in this order is

$$\frac{1}{4} + \int_{\frac{\sqrt{2}}{2}}^1 (2x + 10000) dx = 10000 + \frac{3}{4} - 5000\sqrt{2} \approx 2930.75.$$

If the order is reversed, then consumer 1 gets $[\frac{\sqrt{2}}{2}, 1]$ due to proportionality, and consumer 2 gets the rest. The efficiency in this order is

$$\frac{1}{4} + \int_0^{\frac{\sqrt{2}}{2}} (2x + 10000) dx = \frac{3}{4} + 5000\sqrt{2} \approx 7070.75,$$

which is larger. Thus, ordering consumers according to their slopes does not yield an optimal solution.

As we show in the following, the optimal order is actually determined by a quantity d_i , called *order determinant*: for consumer i with density function $f_i(x) = a_i x + b_i$, define

$$d_i = \frac{a_i}{a_i + 2b_i}.$$

We first have the following claim.

Proposition 4.2. *For any consumers i and j and any $0 \leq p < q < r \leq 1$, the order determinant $d_i < d_j$ if and only if $\frac{v_i(p,q)}{v_i(p,r)} > \frac{v_j(p,q)}{v_j(p,r)}$.*

Proof. We first prove the case where $p = 0$ and $r = 1$. For consumer i , we have $v_i(0, 1) = \frac{1}{2}a_i + b_i$ and $v_i(0, q) = \int_0^q a_i x + b_i dx = \frac{1}{2}a_i q^2 + b_i q$. Thus, we have

$$\frac{v_i(0, q)}{v_i(0, 1)} = \frac{\frac{1}{2}a_i q^2 + b_i q}{\frac{1}{2}a_i + b_i} = q + (q^2 - q) \frac{a_i}{a_i + 2b_i} = q + (q^2 - q)d_i.$$

Similarly, we have $\frac{v_j(0,q)}{v_j(0,1)} = q + (q^2 - q)d_j$. Because $q^2 - q < 0$, it is obvious that $d_i < d_j$ if and only if $\frac{v_i(0,q)}{v_i(0,1)} > \frac{v_j(0,q)}{v_j(0,1)}$.

For the case where $p \neq 0$ or $r \neq 1$, we apply a transformation to the density function $f_i(\cdot)$: let

$$\begin{cases} a'_i = a_i \cdot (r - p) \\ b'_i = a_i p + b_i \end{cases}$$

and define a new density function $f'_i(x) = a'_i x + b'_i$. Let $y = \frac{q-p}{r-p}$; one can verify that

$$\frac{v_i(p, q)}{v_i(p, r)} = \frac{\frac{1}{2}a_i(q^2 - p^2) + b_i(q - p)}{\frac{1}{2}a_i(r^2 - p^2) + b_i(r - p)} = \frac{\frac{1}{2}a'_i y^2 + b'_i y}{\frac{1}{2}a'_i + b'_i} = \frac{v'_i(0, y)}{v'_i(0, 1)}.$$

Using a similar transformation to $f_j(\cdot)$, we can reduce the problem to the case where $p = 0$

and $r = 1$. Hence, we have

$$\begin{aligned}
\frac{v_i(p, q)}{v_i(p, r)} > \frac{v_j(p, q)}{v_j(p, r)} &\Leftrightarrow \frac{v'_i(0, y)}{v'_i(0, 1)} > \frac{v'_j(0, y)}{v'_j(0, 1)} \\
&\Leftrightarrow \frac{a'_i}{a'_i + 2b'_i} < \frac{a'_j}{a'_j + 2b'_j} \\
&\Leftrightarrow \frac{a_i(r - p)}{a_i(r + p) + 2b_i} < \frac{a_j(r - p)}{a_j(r + p) + 2b_j} \\
&\Leftrightarrow \frac{a_i}{a_i + 2b_i} < \frac{a_j}{a_j + 2b_j} \\
&\Leftrightarrow d_i < d_j
\end{aligned}$$

This completes the proof. $\square$

The order determinant precisely characterizes the order of the consumers in an optimal allocation, shown by the following lemma.

Lemma 4.3. *There exists an efficiency optimal allocation in which the allocation (from left to right) is according to the non-decreasing order of d_i .*

Proof. Given an efficiency optimal allocation, assume that the allocation to the consumers are in the order $1, 2, \dots, n$. Then for any $1 \leq i < n$, it suffices to show the following:

- $d_i \leq d_{i+1}$;
- if $d_i = d_{i+1}$, then switching the order of i and $i + 1$ will lead to an efficiency optimal allocation as well.

We will prove the claims by considering each $i = 1, \dots, n$ in the sequence.

Suppose that in an efficiency optimal allocation, consumer i is assigned interval $[p, q]$, and consumer $i + 1$ is assigned to $[q, r]$. Let $\alpha_i = \frac{v_i(p, q)}{v_i(p, r)}$ and $\alpha_{i+1} = \frac{v_{i+1}(q, r)}{v_{i+1}(p, r)}$ denote the ratio of the valuation of consumer i and $i + 1$ between their own allocation and the union of their allocations, respectively. We next consider two cases.

- (1) $\alpha_i \geq \alpha_{i+1}$. We can find a point $q' \in [p, q]$, such that $\frac{v_i(p, q')}{v_i(p, r)} = \alpha_{i+1}$. Now we consider a new allocation, in which consumer $i + 1$ gets $[p, q']$, consumer i gets $[q', r]$, and the allocation of all other consumers remains unchanged. The social welfare difference

between the original allocation and the new one is

$$\begin{aligned}\delta &\triangleq (v_i(p, q) + v_{i+1}(q, r)) - (v_{i+1}(p, q') + v_i(q', r)) \\ &= (v_{i+1}(q, r) - v_{i+1}(p, q')) + (v_i(p, q') - v_i(q, r))\end{aligned}$$

If $d_i > d_{i+1}$, by Proposition 4.2, we have $\frac{v_{i+1}(p, q')}{v_{i+1}(p, r)} > \frac{v_i(p, q')}{v_i(p, r)} = \alpha_{i+1}$. Hence

$$\begin{aligned}v_{i+1}(p, q') &> \alpha_{i+1} \cdot v_{i+1}(p, r) \\ &= \frac{v_{i+1}(q, r)}{v_{i+1}(p, r)} \cdot v_{i+1}(p, r) = v_{i+1}(q, r).\end{aligned}$$

Similarly, because $\frac{v_{i+1}(p, q)}{v_{i+1}(p, r)} > \frac{v_i(p, q)}{v_i(p, r)}$, we have

$$\begin{aligned}\frac{v_i(q, r)}{v_i(p, r)} &= 1 - \frac{v_i(p, q)}{v_i(p, r)} \\ &> 1 - \frac{v_{i+1}(p, q)}{v_{i+1}(p, r)} = \frac{v_{i+1}(q, r)}{v_{i+1}(p, r)} = \alpha_{i+1}.\end{aligned}$$

Hence, $v_i(q, r) > \alpha_{i+1} \cdot v_i(p, r) = v_i(p, q')$. This means that $\delta < 0$, i.e., the new allocation has larger efficiency than the original allocation, which contradicts the fact that the original one is an efficiency optimal allocation. Thus, we have $d_i \leq d_{i+1}$.

On the other hand, if $d_i = d_{i+1}$, then by the calculation same as above, we can get $v_i(q, r) = v_i(p, q')$ and $v_{i+1}(p, q') = v_{i+1}(q, r)$, which implies that $\delta = 0$. Hence, we get another optimal allocation with the order of i and $i + 1$ switched.

(2) $\alpha_i < \alpha_{i+1}$. In this case we can find a point $q' \in [q, r]$ such that $\frac{v_{i+1}(q', r)}{v_{i+1}(p, r)} = \alpha_i$. The remaining details of the proof are similar to the first case and are omitted.

Therefore, by considering each $i = 1, \dots, n$ sequentially, the lemma follows. $\square$

4.3.2 Polynomial Time Approximation Scheme

Having a characterization of the order of allocations in an optimal solution, we will next compute such an allocation. However, there is one issue: For linear functions, a fair allocation may require to cut the cake at irrational points. For instance, in Example 4.1, the unique fair allocation requires to cut the cake at point $\frac{\sqrt{2}}{2}$. This barrier prevents us from finding a tractable algorithm that always outputs a precisely fair allocation. As a result, we

relax our requirement by considering $(1 - \epsilon)$ -fair allocation, in which each consumer gets no less than $\frac{1-\epsilon}{n}$ share of its value for the entire cake.

We will use the approach of dynamic programming. For any given $\epsilon > 0$, we divide the interval $[0, 1]$ evenly into $K = \lceil \frac{2n}{\epsilon} \rceil$ subintervals; and we will keep each subinterval as a whole in the allocation. Let $H(i, k)$ denote the maximum efficiency one can get if we allocate the interval $[0, \frac{k}{K}]$ to the first i consumers in the given optimal order, under the condition that each consumer j , $1 \leq j \leq i$, gets a share of at least $\frac{1-\epsilon}{n}$. The dynamic programming algorithm is described as follows.

ALG-LINEAR

- 1: Sort all consumers in a non-decreasing order with respect to their order determinant d_i .
- 2: Let $H(0, 0) = 0$ and $H(0, k) = -\infty$ for all $0 < k \leq K$.
- 3: **for** $1 \leq i \leq n$, $1 \leq k \leq K$ **do**
- 4: Among all $0 \leq k' \leq k$ such that $v_i(\frac{k'}{K}, \frac{k}{K}) \geq (\frac{1-\epsilon}{n})v_i(0, 1)$, let

$$H(i, k) = \max_{k'} \left\{ H(i-1, k') + v_i\left(\frac{k'}{K}, \frac{k}{K}\right) \right\}.$$

- 5: **end for**
 - 6: Return $H(n, K)$ and its corresponding allocation.
-

Theorem 4.4. *When all consumers have linear density functions, for any $\epsilon > 0$, ALG-LINEAR runs in time polynomial in the input size and $\frac{1}{\epsilon}$, and outputs an $(1 - \epsilon)$ -fair allocation with efficiency within a factor of $(1 - \epsilon)$ to the optimum.*

Proof. It is easy to check that this algorithm can be done in $O(n^3(\frac{1}{\epsilon})^2)$ time, and it will always output an efficiency maximal ϵ -fair allocation in which the cake is cut only at the given separations, assuming there exists one. Let $A = (A_1, \dots, A_n)$ be an efficiency optimal fair allocation. If we round each cutting point of A to the closest smaller separation given by K , in the new allocation every consumer i will get a value no less than $v_i(A_i) - \frac{\epsilon}{2n}(a_i + b_i) \geq v_i(A_i) - \frac{\epsilon}{n}v_i(0, 1) \geq (1 - \epsilon)v_i(A_i)$. Thus, it is an ϵ -fair allocation, and the efficiency of this allocation will be at least $(1 - \epsilon) \sum_i v_i(A_i)$, which is always within a factor of $(1 - \epsilon)$ to the optimum. This finishes the proof of the theorem. $\square$

4.4 Piecewise Constant Functions

In this section, we consider piecewise constant functions. A density function $f(\cdot)$ is called *piecewise constant* if one can partition the whole interval $[0, 1]$ into subintervals such that f has a constant value on each subinterval. Piecewise constant functions are a representative class because it is easy to describe algorithmically and can approximate any density functions with arbitrarily good accuracy.

For a given piecewise constant function $f(\cdot)$, we say it contains k segments if there exist

$$0 \leq a_1 < b_1 \leq a_2 < b_2 \leq \dots \leq a_k < b_k \leq 1$$

such that (i) f has a constant positive value on each $[a_i, b_i]$, (ii) f has value 0 on all remaining regions, and (iii) for each $[a_i, b_i]$ and $[a_{i+1}, b_{i+1}]$, either $b_i < a_{i+1}$ or f has different values on these two segments. We say $[a_i, b_i]$ a *segment* and a_i, b_i its endpoints.

Next we will consider computational complexity of computing an efficiency optimal allocation for piecewise constant functions. As can be seen, its complexity is completely different from the linear function case.

4.4.1 Fixed Order

For a piecewise constant function, its input consists of its value and two endpoints (rational numbers) of all of its (finite) segments. It can be seen (below) that there always exists an optimal solution whose cutting points are rational. Indeed, if the winning order of the consumers is explicitly given, such an optimal solution can be found in polynomial time.

Suppose the order of the consumers, from left to right, is fixed as $1, 2, \dots, n$. Let $0 \leq p_1 \leq p_2 \leq \dots \leq p_K \leq 1$ be the lists of endpoints of all segments of all consumers. While in an optimal allocation, the points of cuts are not necessarily in $\{p_1, \dots, p_K\}$, those cutting points that are not in the set can be nicely characterized. This leads to the following dynamic programming, which computes an *exactly* optimal allocation.

Let $H(i, p_k)$ denote the maximum efficiency one can get if we allocate the interval $[0, p_k]$ to the first i consumers. The dynamic programming is described as follows.

Theorem 4.5. *Given a fixed winning order of consumers, ALG-PIECEWISE-CONSTANT computes an efficiency optimal proportional allocation in polynomial time.*

ALG-PIECEWISE-CONSTANT

1: Let $H(0, p_k) = 0$ for all $1 \leq k \leq K$.

2: **for** $1 \leq i \leq n, 1 < k \leq K$ **do**

3: Let

$$H(i, p_k) = \max_{\substack{1 \leq j \leq i \\ 1 \leq k' < k}} \left\{ H(j-1, p_{k'}) + S(j, i, p_{k'}, p_k) \right\},$$

where $S(j, i, p_{k'}, p_k)$ is the maximum efficiency of proportional allocation one can get by allocating interval $[p_{k'}, p_k]$ to consumers $j, j+1, \dots, i$ under the condition that at most one consumer can have valuation more than $1/n$ of its share (if no such allocation exists, define it to be 0).

4: **end for**

5: Return $H(n, p_K)$ and its corresponding allocation.

Proof. In order to prove correctness of the algorithm, we first show the following characterization, which immediately implies that the dynamic programming always outputs an optimal solution. That is, to allocate interval $[p_{k'}, p_k]$ to consumers $j, j+1, \dots, i$ under the fairness condition, there always is an efficiency optimal allocation such that one of the following holds:

- Some consumer's allocation is cut at one endpoint of its own segments.
- At most one consumer has valuation more than $1/n$ of its total share.

To prove this, given an optimal allocation, assume that no consumer's allocation is cut at any endpoint of its segments and there are two consumers x, y both getting more than $1/n$ of their total share. If we keep the allocation valuations of all other consumers fixed, except x, y , and add a disturbance Δ_x to consumer x 's valuation, it will sequentially result in a $-\Delta_y$ change in consumer y 's valuation, where $\Delta_x, \Delta_y > 0$. (That is, slightly move the cutting points between x and y .) Similarly, if we disturb x 's valuation by $-\Delta_x$, y 's valuation will be changed by Δ_y . This means that if $\Delta_x \neq \Delta_y$, we can always perform such a disturbance to get a proportional allocation with larger efficiency. If $\Delta_x = \Delta_y$, we can keep performing this disturbance, until either one of x and y has reached exactly $1/n$ of its share, or one consumer has reached an endpoint cut. In this process, the efficiency of the allocation remains unchanged. The claim then follows by applying this process repeatedly until the number of consumers whose share is more than $1/n$ is 1 in the final allocation.

Finally we prove that the algorithm runs in polynomial time. It suffices to show that

$S(j, i, p_{k'}, p_k)$ can be polynomially computed. In fact, one can just enumerate the consumer whose valuation is (potentially) more than $1/n$ of its total share. Once that consumer is fixed, all other consumers should have exactly $1/n$ of their total shares. Thus the allocation can be uniquely determined, and can be computed in polynomial time. This finishes the proof of the theorem. $\square$

4.4.2 NP-Hardness

While the above algorithm gives a hope to find an optimal solution in polynomial time, we do not necessarily know the optimal order of the consumers in an optimal allocation. In general, we have the following hardness result, which says that computing an allocation whose efficiency is within a factor of $\Omega(\frac{1}{\sqrt{n}})$ to the optimum, conditioned on proportionality, is NP-hard.

Theorem 4.6. *When all consumers have piecewise constant density functions, it is NP-hard to approximate the optimum efficiency within a ratio of $\Omega(\frac{1}{\sqrt{n}})$, even if all functions have at most 2 segments.*

Proof. Our reduction is from the partition problem. Consider the following partition problem instance: Given a set A' of $2k$ positive integers $A' = (a'_1, a'_2, \dots, a'_{2k})$; let $2m' = \sum_i a'_i$. It is NP-complete to decide whether A' can be divided into two subsets such that both of them have exactly k integers, and the sum of the numbers in each subset equals m' [62]. Without loss of generality, we assume that all a'_i are even numbers. Further, we add an integer $2km'$ to each a'_i and get a new set $A = (a_1, a_2, \dots, a_{2k})$ where $a_i = a'_i + 2km'$ for each i . It is easy to see that A has a partition of size

$$m \triangleq 2k^2m' + m'$$

if and only if A' has a partition of size m' . Let $M = \max_i a_i$.

For the instance A , we have the following observations, which can be verified easily.

1. If A can be partitioned into A_1 and A_2 with equal sums, then $|A_1| = |A_2| = k$.
2. If the absolute value of the difference between the total values of A_1 and A_2 is less than 1, then they have the same total value (as all numbers are even).

We next construct a cake-cutting instance from the partition instance A . Assume that the cake is modeled as an interval $[0, 3m(2k - 1) + 2k(n + 1)]$ (which can be easily normalized to $[0, 1]$). There are in total $n = 4k^2$ consumers, divided into 3 groups: A, B, C .

- Group A has $2k - 1$ consumers. We label them by

$$1, 1 + (2k + 1), 1 + 2(2k + 1), \dots, 1 + (2k - 2)(2k + 1),$$

and define their valuation density functions to be

$$f_{1+t(2k+1)}(x) = \begin{cases} 1 & x \in [m + 3mt, m + 3mt + 1] \\ 0 & \text{otherwise} \end{cases}$$

for $t = 0, 1, 2, \dots, 2k - 2$. Note that the total value of each consumer is 1.

- Group B has in total $2k(2k - 1)$ consumers, which are further divided into $2k - 1$ subgroups $B_0, B_1, \dots, B_{2k-2}$. For each subgroup B_t , there are $2k$ consumers, which are labeled by

$$1 + t \cdot 2k, 2 + t \cdot 2k, 3 + t \cdot 2k, \dots, 2k + t \cdot 2k.$$

Define two intervals on the cake as follows:

$$I_1 \triangleq [3mt, 3mt + 2m + 1]$$

and

$$I_2 \triangleq [3m(2k - 1) + t(n + 1) + n, 3m(2k - 1) + t(n + 1) + n + 1]$$

The valuation functions of the consumers in B_t are defined to be

$$f_{i+t \cdot 2k}(x) = \begin{cases} 1 & x \in I_1 \\ na_i - (2m + 1) & x \in I_2 \\ 0 & \text{otherwise} \end{cases}$$

where $i = 1, 2, \dots, 2k$. Note that it is guaranteed that $na_i - (2m + 1) > 0$.

It can be seen that the total value of each consumer $i + t(2k - 1)$ is na_i ; thus, in a proportional allocation, its share is at least a_i . In addition, the positive valuation

intervals are disjoint for consumers in different subgroups B_t . For each B_t , it contains exactly one consumer from A .

- Group C contains only one consumer, and his valuation density function is

$$f_n(x) = \begin{cases} M^c & x \in [3m(2k-1), \\ & 3m(2k-1) + 2k(n+1) - 1] \\ 0 & \text{otherwise} \end{cases}$$

where c is a large integer.

The figure below shows the construction of the cake-cutting instance. Note that all consumers have piecewise constant valuations and at most two segments.

Figure 4.1: Construction of the cake-cutting instance.

For each consumer in group B_t , he can get at least $1/n$ of his total valuation only if he gets an allocation in T_1 of length at least a_i or in I_2 of at least certain length. If there is a valid partition in the partition instance A , then, for each t , consumers in subgroup B_t can share the interval $[3mt, 3mt + m]$ and $[3mt + m + 1, 3mt + 2m + 1]$ completely to get $1/n$ of their total valuation. Thus, we can allocate consumers in group A and group C all their interested intervals. The total valuation in this case is therefore

$$M^c(2k(n+1) - 1) + (2k - 1)(2m + 1).$$

On the other hand, if the partition instance does not have a valid solution, by the two observations established above, we cannot partition A to two subsets such that the difference between their sums is less than 1. This means that each subgroup B_t has at least one consumer allocated inside the interval I_2 . As there are $2k - 1$ subgraphs and the allocation must be continuous, consumer C can get only one of the $2k$ leftover segments:

$$\begin{aligned}
& [3m(2k - 1), 3m(2k - 1) + n] \\
& [3m(2k - 1) + n + 1, 3m(2k - 1) + 2n + 1] \\
& \dots\dots \\
& [3m(2k - 1) + (2k - 2)(n + 1) + n + 1, 3m(2k - 1) + 2k(n + 1) - 1]
\end{aligned}$$

where each segment has length n . The total valuation in this case is then at most

$$M^c(n + 1) + (2k - 1)(2m + 1) + (2k - 1)nM,$$

where the first term is an upper bound on the valuation for the consumer in group C , and the last two terms are for the consumers in groups A and B (note that $na_i - (2m + 1) < nM$).

Therefore, we have

$$\begin{aligned}
& \frac{M^c(n + 1) + (2k - 1)(2m + 1) + (2k - 1)nM}{M^c(2k(n + 1) - 1) + (2k - 1)(2m + 1)} \\
& < \frac{M^c(n + 1) + (2k - 1)(2m + 1) + (2k - 1)nM}{M^c(2k(n + 1) - 1)} \\
& = \frac{n + 1}{2k(n + 1) - 1} + \frac{(2k - 1)(2m + 1) + (2k - 1)nM}{M^c(2k(n + 1) - 1)} \\
& \leq \Omega\left(\frac{1}{2k}\right) = \Omega\left(\frac{1}{\sqrt{n}}\right)
\end{aligned}$$

where the last inequality follows from the observations that $m \leq kM$ and we can pick c to be a sufficiently large number. Therefore, it is NP-hard to approximate efficiency within a ratio of $\Omega\left(\frac{1}{\sqrt{n}}\right)$. $\square$

4.4.3 Normalized Valuations

We say a density function f *normalized* if $\int_0^1 f(x) dx = 1$. For such functions, computing an efficiency optimal allocation is NP-hard as well².

Theorem 4.7. *When all consumers have normalized piecewise constant density functions, it is NP-hard to compute the optimum efficiency, even if all functions have exactly 1 segment.*

Similar to the proof of Theorem 4.6, the proof of the above result uses a reduction from partition as well. But here, because all functions are normalized, we have to balance the valuation and length of a segment in the constructed instance, and cannot obtain a stronger hardness of approximation result.

Proof. Consider the following partition instance: There are a set of $2k$ positive integers $A' = \{a'_1, a'_2, \dots, a'_{2k}\}$ with $\sum_{i=1}^{2k} a'_i = 2m'$. The problem is to partition these integers to two subsets such that both sets have total value m' . Without loss of generality, we assume that $a'_1 \leq a'_2 \leq \dots \leq a'_{2k}$ and each a'_i is a multiple of 3. We construct a new set $A = \{a_1, a_2, \dots, a_{2k}\}$, where each $a_i = a'_i + (4k + 8)m'$. It is easy to see that A has a partition of total value $m \triangleq m' + k(4k + 8)m'$ each if and only if A' has a valid partition. Denote $a_1 = a$ and $\Delta = a_{2k} - a$. Note that $\Delta = a'_{2k} - a'_1$ and $ka \leq m \leq k \cdot (a + \Delta)$.

We have the following observations.

1. If A can be partitioned into A_1 and A_2 with equal sums, then $|A_1| = |A_2| = k$.
2. If the absolute value of the difference between the total values of A_1 and A_2 is less than 2, then they have the same total value (as all numbers are multiples of 3).
3. $(2k + 3)a_1 > 2m + 2$. This is because

$$\begin{aligned}
 (2k + 3)a_1 &= (2k + 3)(a'_1 + (4k + 8)m') \\
 &> (2k + 3)(4k + 8)m' \\
 &\geq 2(m' + k(4k + 8)m') + 2 \\
 &= 2m + 2
 \end{aligned}$$

²Note that Theorem 4.6 does not imply the claim here, as in its proof consumers' valuations are not normalized. On the other hand, the NP-hardness result established by Theorem 4.7 applies to non-normalized setting as well; but the statement of Theorem 4.6, $\Omega(\frac{1}{\sqrt{n}})$ -hardness, is much stronger.

4. $\frac{\Delta}{a} < \frac{2k-1}{2k(2k+4)}$. This is because

$$\begin{aligned} \frac{\Delta}{a} &= \frac{\Delta}{a'_1 + (4k+8)m'} < \frac{\Delta}{(4k+8)m'} \\ &< \frac{1}{4k+8} < \frac{k}{2k(2k+4)} < \frac{2k-1}{2k(2k+4)} \end{aligned}$$

5. $\frac{\Delta}{a} < \frac{1}{2k+3}$. To see this, we have

$$\begin{aligned} \frac{\Delta}{a} &< \frac{2k-1}{2k(2k+4)} \\ &= \frac{1}{2k+4} - \frac{1}{2k(2k+4)} < \frac{1}{2k+4} < \frac{1}{2k+3} \end{aligned}$$

6. Let $b_i = \frac{1}{a_i(2k+3)}$, for $i = 1, \dots, 2k$, and

$$b = \frac{1}{a_{2k}(2k+3) - (2m+2)}.$$

It can be seen that $b > b_1 \geq b_2 \geq \dots \geq b_{2k}$. Further, the above inequalities and a careful calculation show that $2m(b_1 - b_{2k}) < a_1(b - b_1)$ (details omitted).

The Figure 4.2 below shows the construction of the reduction.

Figure 4.2: Construction of the cake-cutting instance (normalized valuations).

Now we construct a cake-cutting instance with $2k+3$ consumers. For the first $2k$ consumers, their density functions are

$$v_i(x) = \begin{cases} b_i & x \in [0, a_i(2k+3)] \\ 0 & \text{otherwise} \end{cases}$$

For the last three consumers, let

$$v_{2k+1}(x) = \begin{cases} 1 & x \in [m, m+1] \\ 0 & \text{otherwise} \end{cases}$$

$$v_{2k+2}(x) = \begin{cases} 1 & x \in [2m+1, 2m+2] \\ 0 & \text{otherwise} \end{cases}$$

$$v_{2k+3}(x) = \begin{cases} b & x \in [2m+2, a_{2k}(2k+3)] \\ 0 & \text{otherwise} \end{cases}$$

In this construction, note that each consumer has the same unit total valuation, and each of the first $2k$ consumers need a length of a_i to get its proportional share. Moreover, the two intervals $[m, m+1]$ and $[2m+1, 2m+2]$ are ‘reserved’ to consumers $2k+1$ and $2k+2$ due to their large density values and the second observation above.

It can be seen that (detailed omitted) if the partition instance has a valid solution, the total valuation V_{yes} is at least

$$V_{yes} = \sum_{i=1}^{2k} a_i b_i + 2 + 1 \geq 2mb_{2k} + 3.$$

And if the partition instance does not have a valid solution, the total valuation V_{no} is at most

$$V_{no} < 2mb_1 + 2 + a_1 b_1 + (a_{2k}(2k+3) - (2m+2) - a_1)b.$$

Therefore,

$$\begin{aligned} & V_{yes} - V_{no} \\ & > 1 - a_1 b_1 - [a_{2k}(2k+3) - (2m+2) - a_1]b - 2m(b_1 - b_{2k}) \\ & = b[a_{2k}(2k+3) - (2m+2)] - a_1 b_1 \\ & \quad - [a_{2k}(2k+3) - (2m+2) - a_1]b - 2m(b_1 - b_{2k}) \\ & = a_1(b - b_1) - 2m(b_1 - b_{2k}) > 0 \end{aligned}$$

Since the partition problem is NP-complete and $V_{yes} > V_{no}$, computing an optimal allocation for the constructed normalized cake-cutting problem is NP-hard. $\square$

4.5 Concluding Remarks

We study efficiency optimization in cake cutting. The solution conditions are that of proportional fairness and $n - 1$ cuts. While the fairness condition is relatively easy to handle (see, e.g., [43]), the requirement on the number of cuts makes the problem much more difficult to analyze. For instance, some powerful tools like convex programming do not apply. We give an optimal design for linear valuation functions, and show hardness results for piecewise constant functions. A natural direction for future research is to design (approximate) algorithms, especially for (normalized) piecewise constant functions, to maximize efficiency given fairness and $n - 1$ cuts.

Chapter 5

Conclusions

Resource allocation is an extraordinarily crucial and broad topic in computer science and economics. It covers a wide range of topics including computer scheduling, portfolio selection, load distribution, social welfare distribution, production control, product pricing, tax policy making and so on. In this thesis, we select several typical resource allocation problems in different social welfare distribution scenarios, including indivisible resource allocation with payment, indivisible and divisible resource allocation without payment. We try to make different parties happy in different scenarios. These results can be adopted to different resource allocations: the indivisible resource allocation problem with payments can be applied to electronic commerce such as online auctions; the indivisible and the divisible resource allocation problem without payments can be applied to social goods distribution. Problems in divisible resource allocation with payment is the famous market equilibrium problem, which has been sufficiently studied in recent years; we do not include this part in the thesis.

In the combinatorial market (indivisible resource allocation problem with payment), the market maker will suffer from a revenue loss if envy-freeness is required [69]. On top of that, if Pareto efficiency is also required, the market maker will lose even more. The social welfare is maximized though in this case, which makes the results desirable to the society. My hope of finding a good approximation scheme for optimal Envy-free Pareto-efficient pricing is smashed by the establishing of the inapproximability results; however, when items are rare and all items can be sold out, an optimal Envy-free Pareto-efficient pricing can be output in polynomial time. In the future, discovering the domains in which

optimal Envy-free Pareto-efficient pricing can be computed or approximated in polynomial time is a quite attractive and practical direction.

Open Question 5.1. *Is there any domain of the combinatorial market in which optimal Envy-free Pareto-efficient pricing can be computed or approximated in polynomial time?*

Moreover, in the indivisible-item allocation problem with subadditive valuations and the multiple fractional knapsack problem, bifactor approximations and upper bounds are constructed, which provides a clear illustration of the tradeoff between fairness and welfare. In the case of the indivisible-item allocation problem with subadditive valuations, there are gaps indicated among efficient computation, theoretical existence and the upper bound. Closing the gaps will be a natural direction for future research. In fact, the bifactor approach in Chapter 3 is also an intriguing future direction to study the tradeoff between fairness and welfare. Applying such approximation schemes in other resource allocation problems, (e.g., cake cutting, load balancing, facility location and bandwidth allocation) will be very interesting. We leave the following open questions:

Open Question 5.2. *In the indivisible-item allocation problem with subadditive valuations and the multiple fractional knapsack problem, what is the best bifactor approximation that we can achieve in polynomial time?*

Open Question 5.3. *For these problems, if we consider other fairness criteria (e.g. envy-free), what can we achieve for efficiency (in polynomial time)?*

Open Question 5.4. *How can this bifactor scheme be applied to other resource allocation problems?*

Additionally, in the cake-cutting problem, the complexity to obtain an efficient proportional solution is considered. We identify that when consumers have linear density functions, a nearly optimal solution can be output in polynomial time. Although an inapproximability result for general case is established, we still hope to find an approximation algorithm for (normalized) piecewise constant density functions case in the future.

Open Question 5.5. *Is there any approximation scheme to maximize efficiency, given fairness and $n - 1$ cuts?*

As a conclusion, we can observe the tradeoff among fairness, welfare and revenue. In most of the cases, we cannot simultaneously optimize all the three factors. In both the combinatorial market and the general case of cake cutting, we have strong inapproximability results. Fortunately, in both cases, there are some special cases, in which optimal solutions can be output in polynomial time. Hence, it greatly increases the applicability of our solution concepts.

From a higher-level point of view, our work purely focuses on efficiency and fairness—we do not attempt to consider strategic behaviors of the consumers in the model. However, in many applications, individual consumers are likely to maximize their own utility by adapting certain strategies (e.g. by reporting a fraud valuation function). Taking this into consideration, the research will go into a new dimension—designing an incentive compatible, efficient and fair protocol for resource allocation problems. We leave this for future research.

Publications Related to This Thesis

X. Bei, N. Chen, X. Hua, B. Tao, E. Yang, *Optimal Proportional Cake Cutting with Connected Pieces*, AAAI 2012, 1263-1269.

X. Hua, *On Envy-Free Pareto Efficient Pricing*, FAW-AAIM 2012, 70-81.

Bibliography

- [1] A. Abdulkadiroglu, T. Sonmez, *Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems*, *Econometrica*, V.66(3), 689-701, 1998.
- [2] A. Abdulkadiroglu, T. Sonmez, *House Allocation with Existing Tenants*, *Journal of Economic Theory*, V.88(2), 233-260, 1999.
- [3] A. Abdulkadiroglu, T. Sonmez, *School Choice: A Solution to the Student Assignment Problem*, Mimeo, Columbia University, 2000.
- [4] Z. Abrams, *Revenue Maximization When Bidders Have Budgets*, SODA 2006, 1074-1082 .
- [5] K. Arrow, *The Organization of Economic Activity: Issues Pertinent to the Choice of Market versus Nonmarket Allocation*, In *The Analysis and Evaluation of Public Expenditures: The PPB System*, 1. U.S. Joint Economic Committee, 91st Session. U.S. Government Printing Office, Washington, DC, 59-73, 1969.
- [6] K. Arrow, D. Debreu, *Existence of Equilibrium for a Competitive Economy*, *Econometrica*, V.22(3), 265-290, 1954.
- [7] Y. Aumann, Y. Dombb, *The Efficiency of Fair Division with Connected Pieces*, WINE 2010, 26-37.
- [8] J. Aslam, A. Rasala, C. Stein, N. Young, *Improved Bicriteria Existence Theorems for Scheduling*, SODA 1999, 846-847.
- [9] A. Asadpour, A. Saberi, *An Approximation Algorithm for Max-Min Fair Allocation of Indivisible Goods*, STOC 2007, 114-121.
- [10] M. Bagnoli, M. Mckee, *Voluntary Contribution Games: Efficient Private Provision of Public Goods*, *Economic Inquiry*, V.29(2), 351-366, 1991.

- [11] M. Balcan, A. Blum, *Approximation Algorithms and Online Mechanisms for Item Pricing*, EC 2006, 29-35.
- [12] M. Balcan, A. Blum, Y. Mansour, *Item Pricing for Revenue Maximization*, EC 2008, 50-59.
- [13] S. Baliga, E. Maskin, *Mechanism Design for the Environment*, Chapter 7, Handbook of environmental economics, 2003 .
- [14] J. Balkin, *Too Good to Be True: The Positive Economic Theory of Law*, Columbia Law Review V.87(7), 1447-1489, 1987.
- [15] N. Bansal, M. Sviridenko, *The Santa Claus Problem*, STOC 2006, 31-40.
- [16] J. Barbanel, S. Brams, W. Stromquist, *Cutting a Pie Is Not a Piece of Cake*, American Mathematical Monthly, V.116(6), 496-514, 2009.
- [17] X. Bei, N. Chen, X. Hua, B. Tao, E. Yang, *Optimal Proportional Cake Cutting with Connected Pieces*, AAAI 2012, 1263-1269.
- [18] A. Blum, J. Hartline, *Near-optimal online auctions*, SODA 2005, 1156-1163.
- [19] A. Blum, V. Kumar, A. Rudra, F. Wu, *Online learning in online auctions*, SODA 2003, 202-204.
- [20] R. Boadway, M. N. Bruce, *Welfare Economics*, Oxford: Basil Blackwell, 1984.
- [21] S. Bouveret, J. Lang, *Efficiency and Envy-Freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity*, IJCAI 2005, 935-940.
- [22] S.J. Brams, *Fair Division*, The Encyclopedia of Public Choice, V.2, 572-574, 2003.
- [23] S.J. Brams, P. Fishburn, *Fair-division of Indivisible Items between Two People with Identical Preferences: Envy-freeness, Pareto-optimality, and Equity*, Social Choice and Welfare, V.17(2), 247-267, 2000.
- [24] S.J. Brams, M.A. Jones and C. Klamler, *Perfect Cake-Cutting Procedures with Money*, Preprint, Department of Politics, New York University, 2004.
- [25] S.J. Brams, M.A. Jones and C. Klamler, *Proportional Pie-cutting*, International Journal of Game Theory, V.36(3-4), 353-367, 2008.

- [26] S.J. Brams, D. L. King, *Efficient Fair Division: Help the Worst Off or Avoid Envy?*, Rationality and Society, V.17(4), 387-421, 2005 .
- [27] S.J. Brams, A.D. Taylor, *An Envy-Free Cake Division Protocol*, American Mathematical Monthly V.102(1), 9-18, 1995.
- [28] P. Briest, *Uniform Budgets and the Envy-Free Pricing Problem*, ICALP 2008, 808-819.
- [29] P. Briest, P. Krysta, *Single-minded Unlimited Supply Pricing on Sparse Instances*, SODA 2006, 1093-1102.
- [30] P. Briest, P. Krysta, B. Vocking, *Approximation Techniques for Utilitarian Mechanism Design*, STOC 2005, 39-48.
- [31] J. Buchanan, *Positive Economics, Welfare Economics, and Political Economy*, Journal of Law and Economics, V.2, 124-138 1959.
- [32] D. Chakrabarty, J. Chuzhoy, S. Khanna, *On Allocating Goods to Maximize Fairness*, FOCS 2009, 107-116.
- [33] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, *The Efficiency of Fair Division*, WINE 2009, 475-482.
- [34] D. Chakrabarty, A. Mehta, V. Nagarajan, *Fairness and Optimality in Congestion Games*, EC 2005, 52-57.
- [35] C. Chekuri, S. Khanna, *A Polynomial Time Approximation Scheme for the Multiple Knapsack Problem*, SIAM Journal on Computing , V.35(3), 713-728, 2005.
- [36] N. Chen, X. Deng, X. Sun, *On Complexity of Single-Minded Auction*, Journal of Computer and System Sciences, V.69(4), 675-687, 2004.
- [37] N. Chen, A. Ghosh, S. Vassilvitskii, *Optimal Envy-Free Pricing with Metric Substitutability*, SIAM Journal on Computing, V.40(3), 623-645, 2011.
- [38] X. Chen, X. Deng, S. H. Teng, *Settling the Complexity of Computing Two-Player Nash Equilibria*, Journal of the ACM, V.56(3), 2009.
- [39] Y. Chen, J. Lai, D. Parkes, A. Procaccia, *Truth, justice and cake cutting*, AAAI 2010, 756-761.

- [40] S. Cho, A. Goel, *Pricing for Fairness: Distributed Resource Allocation for Multiple Objectives*, STOC 2006, 197-204.
- [41] E. H. Clarke, *Multipart Pricing of Public Goods*, Public Choice, V.11, 17-33, 1971.
- [42] B. Cohen, *Incentives build robustness in BitTorrent*, Workshop on Economics of Peer-to-Peer systems.
- [43] Y. Cohler, J. Lai, D. Parkes, A. Procaccia, *Optimal Envy-Free Cake Cutting*, AAAI 2011, 626-631.
- [44] V. Conitzer, T. Sandholm, *Complexity Results about Nash Equilibria*, IJCAI 2003, 765-771.
- [45] Cormen, Thomas H.; Leiserson, Charles E.; Rivest, Ronald L.; Stein, Clifford , *Introduction to Algorithms (Second ed.)*, MIT Press and McGraw-Hill, 651-664, 2001.
- [46] P. Cramton, Y. Shoham, R. Steinberg, *Combinatorial Auctions*, MIT Press, 2006.
- [47] V. Crawford, *A Procedure for Generating Pareto-Efficient Egalitarian-Equivalent Allocations*, Econometrica, Vol. 47(1), 49-60, 1979.
- [48] V. Crawford, E. Knoer, *Job Matching with Heterogeneous Firms and Workers*, Econometrica, V.49(2), 437-450, 1981.
- [49] M. Dall'Aglio, T. Hill, *Maximin share and minimax envy in fair-division problems*, Journal of mathematical analysis and applications, V.281(1), 346-361, 2003.
- [50] M. Dall'Aglio, R. Mosca, *How to Allocate Hard Candies Fairly*, Mathematical Social Sciences, V.54(3), 218-237, 2008.
- [51] C. Daskalakis, P. Goldberg, C. Papadimitriou, *The complexity of computing a Nash equilibrium*, SIAM Journal on Computing, V.39(1), 195-259, 2009.
- [52] G. Demange, D. Gale, *The Strategy of Two-Sided Matching Markets*, Econometrica, V.53, 873-888, 1985.
- [53] Y. Dinitz, N. Garg, M. Goemans, *On the Single-Source Unsplittable Flow Problem*, FOCS 1998, 290-299.
- [54] S. Dobzinski, N. Nisan, M. Schapira, *Approximation Algorithms for Combinatorial Auctions with Complement-Free Bidders*, STOC 2005, 610-618.

- [55] S. Dobzinski, N. Nisan, M. Schapira, *Truthful randomized mechanisms for combinatorial auctions*, STOC 2006, 644-652.
- [56] S. Dobzinski, M. Schapira, *An Improved Approximation Algorithm for Combinatorial Auctions with Submodular Bidders*, SODA 2006, 1064-1073.
- [57] A. Drazen, V. Grilli, *The Benefits of Crises for Economic Reforms*, American Economic Review, V.83(3), 598-607, 1993.
- [58] L. Dubins, E. Spanier, *How to Cut a Cake Fairly*, American Mathematical Monthly, V.68(1), 1-17, 1961.
- [59] P. Edelman, P. Fishburn, *Fair division of indivisible items among people with similar preferences*, Mathematical Social Sciences, V.41(3), 327-347, 2001.
- [60] S. Even, A. Paz, *A Note on Cake Cutting*, Discrete Applied Mathematics, V.7(3), 285-296, 1984.
- [61] U. Feige, *On Maximizing Welfare when Utility Functions are Subadditive*, STOC 2006, 41-50.
- [62] M. Garey, D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, 1979.
- [63] D. Gillies, *Solutions to general non-zero-sum games*, Contributions to the Theory of Games IV, Annals of Mathematics Studies, V.40, 47-85, 1959.
- [64] A. Goel, A. Meyerson, S. Plotkin, *Combining Fairness with Throughput: Online Routing with Multiple Objectives*, Journal of Computer and System Sciences, V.63(1), 62-79, 2001.
- [65] A. Goel, A. Meyerson, S. Plotkin, *Approximate Majorization and Fair Online Load Balancing*, SODA 2001, 384-390.
- [66] M. Goemans, N. Harvey, S. Iwata, V. Mirrokni, *Approximating Submodular Functions Everywhere*, SODA 2009, 535-544.
- [67] T. Groves, *Incentives in Teams*, Econometrica, V.41(4), 617-631, 1973.
- [68] F. Gul, and E. Stacchetti. *Walrasian Equilibrium with Gross Substitutes*. Journal of Economic Theory, V.87(1), 95-124, 1999.

- [69] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, F. McSherry, *On Profit-Maximizing Envy-Free Pricing*, SODA 2005, 1164-1173.
- [70] E. Hahne, *Round-robin Scheduling for Max-min Fairness in Data Networks*, Selected Areas in Communications, IEEE Journal on, V.9(7), 1024 - 1039, 1991.
- [71] J. Hartline, V. Koltun, *Near-Optimal Pricing in Near-Linear Time*, Lecture Notes in Computer Science, V.3608/2005, 422-431, 2005.
- [72] J. D. Herreiner, C. Puppe, *A Simple Procedure for Finding Equitable Allocations of Indivisible Goods*, Social Choice and Welfare, V.19(2), 415-430, 2002.
- [73] J. R. Hicks, *The Rehabilitation of Consumers' Surplus*, The Review of Economic Studies, V.8(2), 108-116, 1941.
- [74] X. Hua, *On Envy-Free Pareto Efficient Pricing*, FAW-AAIM 2012, 70-81.
- [75] T. Ichiishi, A. Idzik, *Equitable allocation of divisible goods*, Journal of Mathematical Economics, V.32(4), 389-400, 1999.
- [76] Interactive Advertising Bureau, *IAB Internet Advertising Revenue Report*, http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_FY_2011.pdf, 2012.
- [77] M. A. Jones, *Equitable, Envy-Free, and Efficient Cake Cutting for Two People and Its Application to Divisible Goods*, Mathematics Magazine, V.75(4), 275-283, 2002.
- [78] L. Kalampoukas, A. Varma, K. Ramakrishnan, *An Efficient Rate Allocation Algorithm for ATM Networks Providing Max-min Fairness*, In Proceedings of the 6th IFIP International Conference on High Performance Networking , 1995, 143-154.
- [79] Kaldor, Nicholas, *Welfare Propositions of Economics and Interpersonal Comparisons of Utility*. The Economic Journal V.49(195), 549-552, 1939.
- [80] M. Keen, D. Wildasin, *Pareto-Efficient International Taxation*, The American Economic Review, V.94(1), 259-275, 2004.
- [81] A. Kelso, V. Crawford, *Job Matching, Coalition Formation, and Gross Substitutes*, Econometrica, V.50(6), 1483-1504, 1982.

- [82] S. Khot, A. Ponnuswami, *Approximation Algorithms for the Max-Min Allocation Problem*, Approximation, Randomization, and Combinatorial Optimization. Algorithms AND Techniques, V.4627/2007 204-217, 2007.
- [83] J. Kleinberg, *Approximation Algorithms for Disjoint Paths Problems*, Ph.D Thesis, MIT, 1996.
- [84] J. Kleinberg, *Single-Source Unsplittable Flow*, FOCS 1996, 68-77.
- [85] J. Kleinberg, Y. Rabani, E. Tardos, *Fairness in Routing and Load Balancing*, Journal of Computer and System Sciences, V.63(1), 2-20, 2001.
- [86] S. Kolliopoulos, *Edge Disjoint Paths and Unsplittable Flow*, Handbook on Approximation Algorithms and Metaheuristics (editor: T. Gonzalez), Chapman and Hall, 2007.
- [87] S. Kolliopoulos, C. Stein, *Improved Approximation Algorithms for Unsplittable Flow Problems* FOCS 1997, 426-435.
- [88] A. Kothari, D. Parkes, S. Suri, *Approximately Strategy Proof and Tractable Multi-unit Auction*, Decision Support Systems, V.39(1), 105-121, 2005.
- [89] A. Kronman, *Wealth Maximization as a Normative Principle*, The Journal of Legal Studies, V.9, 227-239, 1980.
- [90] A. Kumar, J. Kleinberg, *Fairness Measures for Resource Allocation*, SIAM Journal on Computing, V.36(3), 657-680, 2006.
- [91] B. Lehmann, D. Lehmann, N. Nisan, *Combinatorial Auctions with Decreasing Marginal Utilities*, EC 2001, 18-28.
- [92] R. Lipton, E. Markakis, E. Mossel, A. Saberi, *On Approximately Fair Allocations of Indivisible Goods*, EC 2004, 125-131.
- [93] Q. Ma, P. Steenkiste, H. Zhang, *Routing High-bandwidth Traffic in Max-min Fair Share Networks*, SIGCOMM 1996, 206-217.
- [94] J. MacKie-Mason, H. Varian, *Generalized Vickrey Auctions*, Technical report, Department of Economics, University of Michigan, 1994.
- [95] S. Martello, P. Toth, *Knapsack Problems: Algorithms and Computer Implementations*, John Wiley and Sons, 1990.

- [96] A. Mas-Colell, M. Whinston, J. Green, *Microeconomic Theory*, Oxford University Press, 1995.
- [97] L. McKenzie, *On the Existence of general equilibrium for a Competitive Market*, *Econometrica*, V.27(1), 54-71, 1959.
- [98] N. Megiddo, *Optimal Flows in Networks with Multiple Sources and Sinks*, *Mathematical Programming*, V.7(1), 97-107, 1974.
- [99] P. Milgrom, R. Weber, *A Theory of Auctions and Competitive Bidding*, *Econometrica*, V.50(5), 1089-1122, 1982.
- [100] V. Mirrokni, M. Schapira, J. Vondrak, *Tight Information-Theoretic Lower Bounds for Welfare Maximization in Combinatorial Auctions*, EC 2008, 70-77.
- [101] A. Mu'alem, N. Nisan, *Truthful Approximation Mechanisms for Restricted Combinatorial Auctions*, *Games and Economic Behavior*, V.64(2), 612-631, 2008.
- [102] R. Myerson, *Optimal Auction Design*, *Mathematics of operations research*, V.6(1), 58-73, 1981.
- [103] R. Myerson, M. Satterthwaite, *Efficient Mechanisms for Bilateral Trading*, *Journal of Economic Theory*, V.29(2), 265-281, 1983.
- [104] N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, *Algorithmic Game Theory*, Cambridge University Press, 2007.
- [105] N. Nisan, I. Segal, *The Communication Requirements of Efficient Allocations and Supporting prices*, *Journal of Economic Theory*, V.129(1), 192-224, 2006.
- [106] S. Nuchia, S. Sen, *Improving Optimality of n Agent Envy-Free Divisions*, ATAL 2001, 277-289.
- [107] F. Parisi. *Positive, Normative and Functional Schools in Law and Economics*, *European Journal of Law and Economics*, V.18(3), 259-272, 2004.
- [108] D. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*, Ph.D. thesis, University of Pennsylvania, 2001.
- [109] Papadimitriou, *Computational Complexity*, Addison Wesley, 1993.

- [110] V. Pareto, *Manual of Political Economy*, Trans. A. Schwier, Scholars Book Shelf, 1971.
- [111] E. Pazner, D. Schmeidler, *A Difficulty in the Concept of Fairness*, The Review of Economic Studies, V.41(3), 441-443, 1974.
- [112] O. Pikhurko, *On Envy-Free Cake Division*, The American Mathematical Monthly, V.107(8), 736-738, 2000.
- [113] R. Posner. *Ethical and Political Basis of the Efficiency Norm in Common Law Adjudication*. Hofstra Law Review, V.8(3), 487-508, 1980.
- [114] M. Quinzii, *Core and Competitive Equilibria with Indivisibilities*, International Journal of Game Theory, V.13(1), 41-60, 1984.
- [115] J. Riley, W. Samuelson, *Optimal Auctions*, The American Economic Review, V.71(3) 381-392, 1981.
- [116] J. Reijnierse, J. Potters, *On Finding an Envy-Free Pareto-Optimal Division*, Mathematical Programming, V.83(1-3), 291-311, 1998.
- [117] J. Robertson, W. Webb, *Near Exact and Envy Free Cake Division*, Ars Combinatoria, V.45, 97-108, 1997.
- [118] J. Robertson, W. Webb, *Cake-Cutting Algorithms: Be Fair If You Can*, A K Peters/CRC Press, 1998.
- [119] S. Ross, *The Economic Theory of Agency: The Principal's Problem*, The American Economic Review, V.63(2), 134-139, 1973.
- [120] D. Salvatore, *Microeconomics, Theory and Applications*, Macmillan Publishing Company, a division of Macmillan, Inc, 1986.
- [121] F. Schroyen, *Pareto Efficient Income Taxation under Costly Monitoring*, Journal of Public Economics, V.65(3), 343-366, 1997.
- [122] L. Shapley, M. Shubik, *The Assignment Game I: The Core*, International Journal of Game Theory, V.1(1), 111-130, 1971.
- [123] A. Smith, *Wealth of Nations*, edited by C. Bullock, Vol. X, The Harvard Classics, New York: P.F. Collier & Son, 1909-14.

- [124] R. Stavins, A. Wagner, G. Wagner, *Interpreting Sustainability in Economic Terms: Dynamic Efficiency Plus Intergenerational Equity*, Economics Letters,, V.79(3), 339-343, 2003.
- [125] C. Stein, J. Wein, *On the Existence of Schedules That Are Near-Optimal for Both Makespan and Total Weighted Completion Time*, Operations Research Letters, V.21(3), 115-122, 1997.
- [126] J. Stiglitz, *Pareto Efficient and Optimal Taxation and the New Welfare Economics*, Handbook of Public Economics, V.2, 991-1042, 1987.
- [127] H. Tzeng, K. Sin, *On Max-min Fair Congestion Control for Multicast ABR Service in ATM*, Selected Areas in Communications, IEEE Journal on, V.15(3), 545 - 556, 1997.
- [128] H. Varian, *Equity, Envy, and Efficiency*, Journal of Economic Theory, V.9, 63-91, 1974.
- [129] W. Vickrey, *Counterspeculation, Auctions and Competitive Sealed Tenders*, Journal of Finance, V.16(1), 8-37, 1961.
- [130] J. Vondrak, *Optimal Approximation for the Submodular Welfare Problem in the Value Oracle Model*, STOC 2008, 67-74.
- [131] L. Walras, *Elements of Pure Economics*, 1877, Harvard University Press, 1954.
- [132] G. Woeginger, *A Polynomial Time Approximation Scheme for Maximizing the Minimum Machine Completion Time*, Operations Research Letters, V.20(4), 149-154, 1997.