Fair and Efficient Resource Allocation: Algorithms and Computational Complexity

A thesis submitted to the Indian Institute of Technology, Gandhinagar in partial fulfillment of the requirements for the degree of Doctor of Philosophy

by

Aditi Sethia

Computer Science and Engineering

Indian Institute of Technology, Gandhinagar

382355, Gujarat (India)

February 2025

Declaration

I hereby certify that the work presented in this thesis is the result of original research, is free of plagiarised materials, and has not been submitted earlier as a whole or in part for a degree/diploma at this or any other Institution or University.

Signature of the Author:

. . . .

Aditi Sethia
Computer Scince and Engineering
Indian Institute of Technology, Gandhinagar

Signature of the Thesis Supervisor:

. . . .

Neeldhara Misra
Associate Professor
Computer Scince and Engineering
Indian Institute of Technology, Gandhinagar

Acknowledgements

This piece of work is a collective effort of all my co-authors. Any of it would not have been possible without them, first and foremost of whom is my advisor, Prof. Neeldhara. I recall, ma'am, while starting, you said that we will have fun-filled academic journey. Today, while finishing this chapter, I take so much pleasure in re-iterating what you had said back then— I have had so much fun working under your guidance and learning from you, not only the nuances in research but also larger perspectives on life. I will cherish the journey these academic endeavours have enabled me to take on and I am so grateful to you for nudging me on this path. The kind of tireless efforts you put into your work and the evident joy it brings to you is so inspiring. Thank you for always leading by example. Thank you for all sorts of encouragements. Thank you for your patience. And special thanks for being so flexible with logistics. The world is a better and happier place because of mentors like you.

My amazing co-authors have contributed towards various interesting aspects of this thesis. Prof. Umang and Rohit — all the brainstorming sessions have been really helpful and thanks to them, I feel a little less afraid of the devil in the details. Special thanks to both of them for being great hosts for my respective short visits. Special mention to Prof. Hadi for hosting me at Pennsylvania State University. Thank you for making my visit a rich experience, both academically and otherwise. Prof. Balagopal, Jayakrishnan, Harshil, Saraswati, Anant, Suchismita, Sanjukta — thank you for all your time, efforts and discussions regarding the respective projects. Thanks to Dr. Prasant Misra for consistent guidance during my time at TCS Research.

To all the anonymous reviewers whose constructive comments have improvised the initial versions of the drafts, thank you so much!

To my Doctoral Studies Committee members, Prof. Bireswar Das and Prof. Akshaa Vatwani, I extend my sincere gratitude for their reviews and comments on the thesis narrative.

To the entire Computer Science and Mathematics faculty at IITGN, thank you for your amazing lectures and guidance that have shaped much of my fundamental understanding.

Many thanks to IITGN, PSU, USask, and CFD workshop for the travel grants to conferences and research visits that have inspired me and enhanced my perspectives. To the entire IITGN admin staff — thank you for your logistical support that facilitated my PhD journey to a great extent. Special mention to Mr. Vijay Meena and Mr. Raviraj Sukhadiya for guiding me with all the paperwork on multiple occasions.

Tanisha and Chinmay, thank you for your initial pep-talks. A little inspiration goes a long way! To my lab mates — Anant, Harshil, Shruti, Saraswati, Dhara, Akash, Jinia, Dipan and others, thanks for collectively creating a healthy and fun environment at the lab. The chai-breaks have made all the difference. Shraddha, Sanjukta, and all lab mates at PSU, thank you so much for making the place feel homely.

Harshil, it has been a pleasure to have known you in person. Thank you for being so generous, accommodating, and thoughtful, to not a few, but to everyone around you, both academically and otherwise. Regardless of how occupied you may have been, thank you for your willingness to proofread for me, very many times. I have gained a wealth of insights from you. I am both amazed by and envious of your precision and minute attention to detail. Your hunger for learning, coupled with your unparalleled humility is so inspiring. You would hear my clap the loudest on all your academic achievements.

To all my dear friends, in and outside IITGN, thank you for your company that I have always enjoyed and cherished so much. Akshita, Prachi, Megha, Yashika, Pragya, Harshil, Hardeep, Deepika, Stanzin, and others — you all have consistently added fun to life, time and again, in exceptional ways. This has implicitly helped me navigate my PhD journey as well. Thanks y'all!

To my entire family: all your endless love and unconditional support have been crucial for my life, in general, and this piece of work, in particular. To my parents, with all the freedom and love of the world, you have made me so happy and rich on all days and always. To mom, not a single day passes without your presence being felt in the little things. Cheers to the memories that still bring a smile in the void. To my brother, whose profound influence has shaped pretty much all my life decisions — big or small — the credit and fault will always be yours. To Bhabhi and my precious little niece, you have brought immeasurable joy to us all. To my *much* better half, how can I thank you enough for all that you bear? I am so glad to have you. Major kudos to my in-laws for being amazingly supportive in all my endeavors. I don't think I will ever have enough words to even remotely describe the amount of happiness and peace you all bring to the table!

Finally, to the readers: all the exciting stuff that you encounter in this piece of work is credited to the collective effort of my co-authors, and the mistakes, if any, are credited to my limitations. Happy reading!

In remembrance
Mrs. Sudha Jain

Abstract

Division of available resources among interested agents by taking into account and consolidating their individual choices has been a fundamental aspect of human survival mechanisms from the very beginning. In the current times also, world harmony hinges on the fact that the land is divided fairly across countries and each of the participating countries perceive the division to be fair to themselves. If this perception of fairness is absent, it often results in devastating wars and widespread chaos. From routine events like splitting the dinner bill to more crucial setting like matching organ donors to patients, algorithms for computing desirable allocations are highly solicited. This thesis discusses two dimensions of desirability among allocations—fairness and efficiency. The contribution of this thesis can be broadly divided into the following three parts:

1. First, we study the computational complexity of finding fair (approximately envy-free/equitable) allocations in various settings and for each of them, identify the domain restrictions where computational tractability holds. We also study the existence and complexity of these fairness notions in conjunction with efficiency notions.
2. Second, we quantify the loss in the various welfare measures due to the fairness constraints (approximately envy-free or equitable) and present tight bounds on the price of minimizing envy and the price of equitability. Rather than focussing on a single welfare measure, we give tight bounds for generalized p-mean welfares. We also identify structured instances where no price has to be paid in terms of welfare.
3. Third, we propose novel concepts of Secure and Abundant allocations, which in addition to being a relaxation of Consensus Allocations (Simmons and Su, 2003), also capture elements of human psychology molded and influenced by the perspectives of others.

Publications

Publications based on this thesis:

C1. (**ADT 2024**) "Envy-Free and Efficient Allocations on Graphical Valuations", *Neeldhara Misra, Aditi Sethia*

C2. (**SOFSEM 2021**) "Fair Division is Hard even for Amicable Agents", *Neeldhara Misra, Aditi Sethia*

C3. (**AAMAS 2023, Extended Abstract**) "The Complexity of Minimizing Envy in House Allocations", *Jayakrishnan Madathil, Neeldhara Misra, and Aditi Sethia*

C4. (**SAGT 2023**) "The Price of Equity with Binary Valuations and Few Agent Types", *Umang Bhaskar, Neeldhara Misra, Aditi Sethia, Rohit Vaish*

W5. (**IJCAI CFD Workshop 2024**) "Equitable Allocation of Mixtures of Goods and Bads", *Hadi Hosseini, Aditi Sethia*

R6. (**Under Review**) "Generalized Consensus Allocations: Valuing the Perspective of Others", *Neeldhara Misra, Aditi Sethia*

Publications outside the scope of this thesis:

C1. (**STACS 2023**) "Finding and counting patterns in sparse graphs", *Balagopal Komarath, Anant Kumar, Suchismita Mishra, Aditi Sethia.*

C2. (**CALDAM 2023**) "Diverse Fair Allocations: Algorithms and Complexity", *Harshil Mittal, Saraswati Girish Nanoti, Aditi Sethia*
(***To appear in the Journal of Discrete Applied Mathematics***)

C3. (**CCCG 2020**) "Red-blue point separation for points on a circle", *Neeldhara Misra, Harshil Mittal, Aditi Sethia*

Contents

Declaration **iii**

Acknowledgements **vii**

Abstract **ix**

Publications **xi**

Contents **xiii**

List of Figures **xvii**

List of Tables **xix**

1 Introduction **1**
- 1.1 Resource Allocation . . . 1
- 1.2 Preliminaries . . . 3
  - 1.2.1 Fair Allocations . . . 4
  - 1.2.2 Efficient Allocations . . . 6
  - 1.2.3 Trade-offs Between Fairness and Efficiency . . . 8
- 1.3 Recent Progress . . . 9
- 1.4 Our Contributions . . . 15

2 Envy-Free and Efficient Allocations for Graphical Instances **19**
- 2.1 Introduction . . . 19
- 2.2 Envy Free Allocations . . . 22
- 2.3 EFX and Welfare-Maximization . . . 30

2.4 Concluding Remarks. 36

3 Envy-Free and Efficient Allocations for Degenerate Instances 37
3.1 Introduction 37
3.2 Preliminaries 40
3.3 EF+fPO Allocations and Degeneracy 42
3.4 Consumption Graphs 56
3.5 The Case of Bounded Valuations 61
3.6 Concluding Remarks 62

4 Minimizing Envy in House Allocation 63
4.1 Introduction 63
4.2 Preliminaries 69
4.3 Pre-processing using Expansion Lemma 78
4.4 Optimal House Allocation 84
4.4.1 Polynomial Time Algorithms for OHA 84
4.4.2 Hardness Results for OHA 89
4.4.3 Parameterized Results for OHA 95
4.5 Egalitarian House Allocation 102
4.5.1 Polynomial Time Algorithms for EHA 102
4.5.2 Hardness Results for EHA 104
4.5.3 Parameterized Results for EHA 110
4.6 Utilitarian House Allocation 113
4.6.1 Polynomial Time Algorithms for UHA 113
4.6.2 Parameterized Results 117
4.7 Experiments 118
4.8 House Allocation on Single-Peaked/Dipped Rankings 119
4.8.1 Single-Peaked Preferences 120
4.8.2 Single-Dipped Preferences 128
4.9 Price of Fairness 130
4.10 Conclusion 138

5 Price of Equitability 139
5.1 Introduction 139
5.2 Preliminaries 146
5.3 Optimal Allocations for Binary Submodular Valuations 149
5.4 Lower Bounds on the PoE for Binary Additive Valuations 152
5.5 Upper Bounds on the PoE for Binary Additive Valuations 154
5.5.1 Upper bounds on the PoE for p=1 154
5.5.2 Upper bounds on the PoE for p < 1 155
5.6 PoE Bounds for Doubly Normalized Instances 164
5.7 PoE Bounds for Binary Submodular Valuations 171
5.8 Visualizing the PoE Bounds 173

5.9 Some Concluding Remarks on Chores . . . 174

6 Equitable and Efficient Allocations for Mixtures of Goods and Bads 177
6.1 Introduction . . . 177
6.2 Preliminaries . . . 182
6.3 EQ1 Allocations . . . 183
6.4 EQ1+PO Allocations . . . 195
6.5 EQX and Social Welfare . . . 202
6.6 EF+EQ+PO Allocations . . . 205
6.7 Concluding Remarks . . . 205

7 Generalized Consesus Allocations: Valuing the Perspectives of Others 207
7.1 Introduction . . . 207
7.2 Preliminaries . . . 213
7.3 Exact Equitable Allocations . . . 216
7.4 Secure Allocations . . . 219
7.4.1 Hardness Results for Secure Allocations . . . 219
7.4.2 Algorithms for Secure Allocations . . . 225
7.5 Abundant Allocations . . . 228
7.5.1 Hardness Results for Abundant Allocations . . . 228
7.5.2 Algorithms for Abundant Allocations . . . 230
7.6 Experiments . . . 231
7.7 Concluding Remarks . . . 231

Bibliography 233

List of Figures

1.1 A subset of real-world settings (house allocations [114], splitting dinner bills [89], inheritance division [116], organ transplantation [145], among others) where fair division algorithms are solitcited. 2
1.2 Desirable Allocations (Image Credits: [74]) 3
1.3 Containment of valuation functions. 4
1.4 If there is an allocation in the intersection, the price of fairness is said to be one. If there no allocation in the intersection, the price is strictly greater than one. . . 8

2.1 A schematic of reduced instance in the proof of Theorem 2.6. 26
2.2 A schematic of reduced instance in the proof of Theorem 2.11 32

3.1 The overall schematic of the construction in the proof of Theorem 1. The entries depicted by a $\star$ indicate small values. In this example, the literal corresponding to the agent b_i, i.e., u_i, belongs to the auxiliary clause C_j corresponding to the backup goods f_j^1 and f_j^2. 44

4.1 Single-peaked preferences with respect to the ordering $\rhd := h_7 \rhd h_5 \rhd h_3 \rhd h_1 \rhd h_2 \rhd h_4 \rhd h_6$. The house h_1 is a shared peak and h_7 is an individual peak. Notice that $peak(a_1) = h_7$ and $peak(a_2) = peak(a_3) = peak(a_4) = h_1$. Also, $base(h_1) = \{a_2, a_3, a_4\}$. And, $\text{span}(h_1) = 2$, which contains the houses h_1 and h_2 as these are the top 2 houses identically ranked by all the agents in $base(h_1)$. 120
4.2 A schematic of Case 2 in the proof of Lemma 4.59. 123
4.3 A schematic of Case 2(a) in Theorem 4.62 127

5.1 PoE as a function of r for $p = 1$ 174
5.2 PoE as a function of r for $p = 0$ 174

5.3 PoE as a function of r for $p = -1$ 174
5.4 PoE as a function of r for $p = -10$ 174

List of Tables

2.1 An EF allocation that allocates an item wastefully. 24

4.1 A partial summary of our results. Here, d denotes the maximum number of houses approved by any agent. The results marked with a $\star$ refer to reductions that imply hardness even when the standard parameter is a constant, while the result marked with a † is a FPT reduction and also implies W[1]-hardness in the standard parameter. 70
4.2 An example of an extremal instance $\mathcal{I} = (A, H, \mathcal{P}; k)$, where $D \subseteq H$ denotes the dummy houses. Once Reduction Rules 1, 2 and 3 are no longer applicable, then there are no dummy agents, but dummy houses must necessarily exist, i.e., $D \neq \emptyset$; and we must have $|H_L| < |A_L|$ and $|H_R| < |A_R|$. 82
4.3 The constraints of the ILP $\mathtt{P1}(\mathcal{I})$. 98
4.4 The constraints of the ILP $\mathtt{P2}(\mathcal{I})$. 112
4.5 A summary of the results, averaged over 100 instances of each type. The OHA column corresponds to the solution from OHA ILP and the max-envy and total envy in that column shows those values when the number of envious agents is minimized. Similarly for the EHA column. 118
4.6 Price of minimizing the number of envious agents, the maximum envy, and the total envy for binary valuations. 131

5.1 Summary of results for the price of equity (PoE). Each cell indicates either the lower or the upper bound (columns) on PoE for a specific welfare measure (rows) as a function of the number of *agent types* r. Our contributions are highlighted by shaded boxes. The lower bounds are from Theorem 5.2, while the upper bounds are shown in Theorem 5.3 and Theorem 5.4. Section 5.8 presents the upper and lower bounds graphically as a function of r, for $p = 1$, $p = 0$, $p = -1$, and $p = -10$. 143

6.1 A partial summary of our results. Each cell contains existence/computation results with ✓ implying existence, and × implying non-existence. The table entry, say '$\checkmark, \{-w, 0, w\}, \mathbb{P}'$ conveys that the corresponding fair and efficient allocation always exists for $\{-w, 0, w\}$ valuations and admits a polynomial time algorithm. The results marked '$\star'$ and '†' follow from Freeman et al. (2019) and Sun et al. (2023b) respectively. 181
6.2 Reduced instance as in the proof of Theorem 6.2. 184
6.3 Rich-poor and poor-rich transfers. In one step, only one of these transfers is executed, not both. 189
6.4 A schematic of Case 2(a) of the proof of Claim 6.7. r_1 and r_2 are two rich agents with utility $3c$ each and p_1 and p_2 are two poor agents with utility $2c$ each. The values in gray are forced, otherwise, we are done by Case 1. The values in red depict the contradiction that $E_c \in O^-$. 191
6.5 Leximin++ is not EQ1 even with {1, -1} normalized valuation. The highlighted allocation is leximin++ with the utility vector as $\{3, 1, 1, 1, 1\}$. This is not EQ1 since even if the last 5 agents choose to hypothetically ignore a good from $\Phi_1's$ bundle, they still fall short of the equitable utility. 195
6.6 Reduced instance as in the proof of Theorem 6.16 197

7.1 A partial summary of our results. The egalitarian version of the problem asks if every agent's utility is at least a given target, while the utilitarian version asks if the total utility derived by all agents meets a given target. For (left) extremal instances, we give efficient algorithms for all the capacitated problems. 209
7.2 Reduced Allocation Instance from Equitable 3-Coloring (Theorem 7.4), where G in the original instance is a cycle on 6 vertices with the edge set $\{(12), (23), (34), (45), (56), (61)\}$ 220
7.3 Reduced Allocation Instance from Vertex Cover (Theorem 7.7), where G in the original instance is a cycle on 4 vertices with the edge set $\{(12), (23), (34), (41)\}$. 223
7.4 Reduced Allocation Instance from Vertex Cover (Theorem 7.16), where G in the original instance is a cycle on 4 vertices with the edge set $\{(12), (23), (34), (41)\}$. 229

Chapter 1

Introduction

"Can the values which individual members of society attach to different alternatives be aggregated into values for society as a whole, in a way that is both fair and theoretically sound?"

- Amartya Sen, *Collective Choice and Social Welfare*

"Mirror, mirror on the wall, who is the fairest of them all?"

- The Queen, *Snow White and the Seven Dwarfs*

1.1 Resource Allocation

Division of available resources among interested parties by taking into account and consolidating their individual choices has been a fundamental aspect of human survival mechanisms from the very beginning. In the current times also, world harmony hinges on the fact that the land is divided fairly across countries and each of the participating countries perceive the division to be fair to themselves. If this perception of fairness is absent, it often results in devastating wars and widespread chaos.

From routine events and logistics like splitting dinner bills among a group of friends, deciding rent among flatmates (Gal et al., 2016; Velez, 2018; Edward Su, 1999), matching medical graduates to hospitals (Roth and Peranson, 1999), assigning students to public schools, splitting assets in bankruptcy, divorces or inheritance, allocating computational resources such as CPU, memory, storage among users and applications – to more critical life or death situations like – matching organ donors to patients, conducting kidney-exchanges (Freedman et al., 2020), dividing vaccines among states (Bertsimas et al., 2020), and dispute resolutions or land division among countries (recall the dramatic division of Indian subcontinent into India and Pakistan) – in all of the above problems, the aim is to arrive at a collective social decision by aggregating individual preferences and hence fair and efficient ways to do the same are highly solicited.

What is to be allocated? The set of resources could be of two types—one that can be assigned in fractions and the other that can not be assigned in fractions. The former are called *divisible resources* (like cake, rent, time-slots, land) and the latter *indivisible* resources (like organs, houses, university seats etc).

Figure 1.1: A subset of real-world settings (house allocations [114], splitting dinner bills [89], inheritance division [116], organ transplantation [145], among others) where fair division algorithms are solitcited.

What is a good way to consolidate different and probably orthogonal opinions into a collective decision that is acceptable by all the parties and agents involved is a burning question that social choice theory asks. There are two kinds of issues to be addressed here: one qualitative, which entails analyzing *what* constitutes a 'desirable' allocation, and another algorithmic, which entails *how* to find such an allocation. This study comes under what is called the interdisciplinary field of computational social choice, which aggregates individual preferences and opinions into a collective outcome such as an allocation. Due to the involvement of several rational agents, it is sometimes referred to as Multi-Agent Resource Allocation (MARA). This thesis takes a step forward in providing answers to some of the intriguing questions revolving around resource allocations, mainly revolving around indivisible resources.

What constitutes a desirable allocation? The two pillars of the desirability of an allocation are (a) Fairness (b) Efficiency (Social Welfare). The former entails that an individual feels that justice is being done to her while the latter ensures that the collective welfare of the society as a whole is maximized. We elaborate on these aspects in Section 1.2.1 and Section 1.2.2 and before that, we formalize the problem setting.

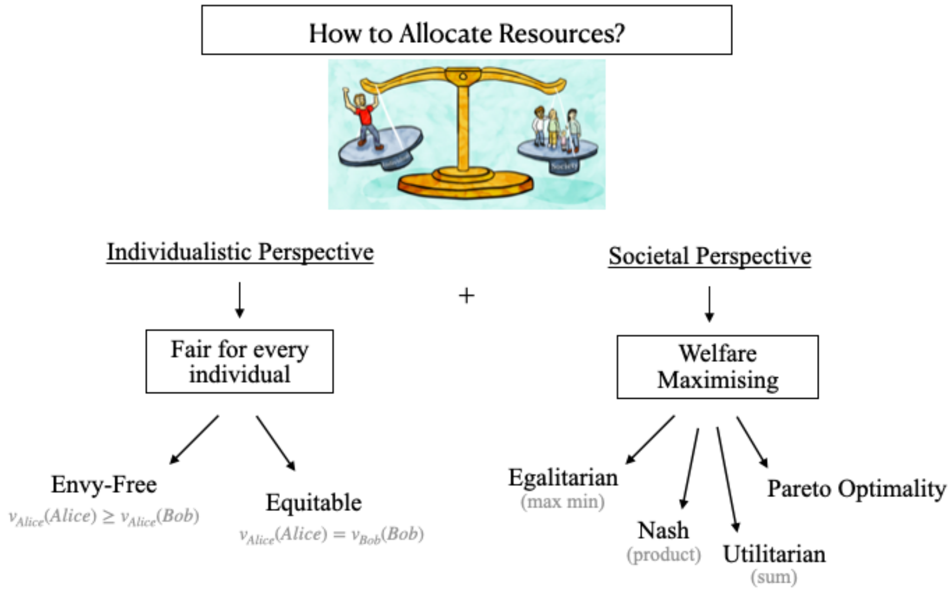

Figure 1.2: Desirable Allocations (Image Credits: [74])

1.2 Preliminaries

An instance of the fair division problem is specified by a tuple $\langle N, M, \mathcal{V} \rangle$, where $N = \{1, 2, \ldots, n\}$ is a set of $n \in \mathbb{N}$ *agents*, $M = \{g_1, \ldots, g_m\}$ is a set of m indivisible *items*, and $\mathcal{V} := \{v_1, v_2, \ldots, v_n\}$ is the *valuation profile* consisting of each agent's valuation function. For any agent $i \in N$, her valuation function $v_i : 2^M \to \mathbb{N} \cup \{0\}$ specifies her numerical value (or *utility*) for every subset of items in M. Further, an instance is *normalized* if for some constant W, $v_i(M) = W$ for all agents i. Normalization is a standard assumption in fair division literature which ensures that utility functions or resources are on a common scale, making it easier to compare the fairness of different allocations.

Valuation functions. A valuation function v is said to be:

- *monotone* if agents prefer more number of items, that is, for any two subsets of items S and T such that $S \subseteq T$, we have $v(S) \leqslant v(T)$,
- *monotone submodular* (or simply submodular) if it is monotone and for any two subsets of items S and T such that $S \subseteq T$ and any item $g \in M \setminus T$, we have $v(S \cup \{g\}) - v(S) \geqslant v(T \cup \{g\}) - v(T)$,
- *additive* if for any subset of items $S \subseteq M$, we have $v(S) = \sum_{g \in S} v(g)$,
- *binary submodular* (or *matroid rank*) if it is submodular and for any subset $S \subseteq M$ and any item $g \in M \setminus S$, we have $v(S \cup \{g\}) - v(S) \in \{0, 1\}$, and
- *binary additive* if it is additive and for any item $g \in M$, $v(\{g\}) \in \{0, 1\}$.

The containment relation between these classes is as follows:

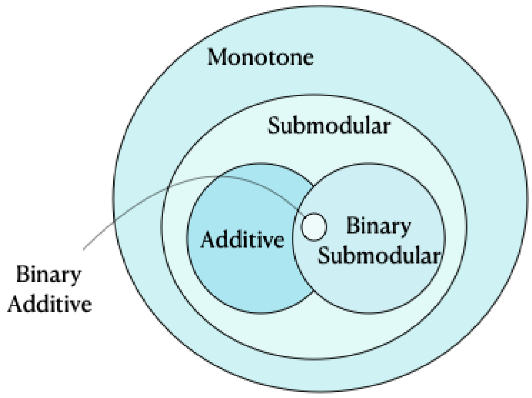

Figure 1.3: Containment of valuation functions.

Allocation. A *bundle* refers to any (possibly empty) subset of goods. An allocation $\Phi := (\Phi_1, \ldots, \Phi_n)$ is a partition of the set of items M into n bundles; here, Φ_i denotes the bundle assigned to agent i.

1.2.1 Fair Allocations

Envy-Freenes. Given an allocation, one basic thing to desire is that every agent is happy with her own share and does not wish to swap what she got with someone else. This is captured by the notion of *envy-freeness* (Gamow and Stern, 1958; Foley, 1967), wherein, no agent envies anyone else. In other words, everyone deems their own bundle of greater value than anyone else's bundle. The formal definition follows.

Definition 1.1. ***(Envy-Freeness.)*** *An allocation* $\Phi = (\Phi_1, \ldots, \Phi_n)$ *is said to be* envy-free *(*EF*) if for any pair of agents* $i, k \in N$*, we have* $v_i(\Phi_i) \geqslant v_i(\Phi_k)$*.*

Envy-freeness, as a fairness concept, is too much to ask for, in the sense that it is non-existential. There are instances where no allocation is EF. Given that the resources can be scarce and a substantial number of agents may compete for the same resource, hence, it is not always possible to achieve perfect envy-freeness. This has led to several workarounds in terms of relaxations (See Section 1.3). One such popular relaxation is *envy-freeness up to one item (*EF1*)* (Lipton et al., 2004; Budish, 2011) where an envious agent is no longer envious upon a hypothetical removal of *some* item from the envied agent's bundle. That is, the envy is there but it is bounded up to the removal of at most one item from each envied bundle. While EF1 surpasses the non-existence of EF and always exists, it can be argued that it is a fairly weak fairness notion in the sense it considers it fair to hypothetically remove a highly-valued item (like a diamond or a car) from the envied agent's bundle. This may not really serve the purpose since that item is the main cause of envy and it is unfair to ignore it. A stricter relaxation of envy-freeness up to *any* good (EFX) is thus considered. (Gourvès et al., 2014; Caragiannis et al., 2019b). An allocation is said to be EFX if the envy of an agent towards another agent can be eliminated by the hypothetical removal of any good in j's bundle. Formally,

Definition 1.2. ***(Relaxations of Envy-Freeness.)*** *An allocation is said to be* envy-free up to one good *(*EF1*) if for any pair of agents* $i, k \in N$ *such that* $\Phi_k \neq \emptyset$*, there is an item* $g \in \Phi_k$ *such that* $v_i(\Phi_i) \geqslant v_i(\Phi_k \setminus \{g\})$*. Further, an allocation is* envy-free up to any good *(*EFX*) if for any pair of agents* $i, k \in N$ *such that* $\Phi_k \neq \emptyset$*,* $v_i(\Phi_i) \geqslant v_i(\Phi_k \setminus \{g\})$ *for any item* $g \in \Phi_k$*.*

Example 1.3. *Below is an instance of a fair division problem where every agent values an item at either* 0 *or* 1*. The highlighted allocation* Φ *corresponds to allocating* $\{g_1, g_2, g_3\}$ *to Alice,* $\{g_4, g_5\}$ *to Bob and the remaining item* g_3 *to Carol. Note that* Φ *satisfies* EF1 *but it is not* EF*. Indeed, Carol does not envy Alice, but she envies Bob. However, if she chooses to ignore the item* g_5 *from Bob's bundle hypothetically, then she is no longer envious. Therefore, the envy in the system is subject to removing at most one item from the envied bundle. Moving to* EFX*, as long as the valuations are binary, and* EF1 *allocation also satisfies* EFX *property. So* Φ *is* EFX*. But this is not true in general. Suppose Carol valued* g_5 *at* 2*, then the same allocation no longer remains* EFX*. Indeed, the removal* *of g_4 from Bob's bundle does not get rid of the envy experienced by Carol.*

	g_1	g_2	g_3	g_4	g_5	g_6
Alice	(1)	(1)	(1)	0	0	0
Bob	0	0	0	(1)	(1)	1
Carol	0	0	0	1	1	(1)

Equitability. Another interpretation of fairness is that of *equitability* (Dubins and Spanier, 1961). This entails that no agent is discriminated and consequently, everyone derives equal value from what they have received. Unlike envy-freeness, here, agents are not concerned if they value the other bundle more, as long as the envied agent also derives the same amount of utility as they do. Formally,

Definition 1.4. ***(Equitability.)*** *An allocation $\Phi = (\Phi_1, \ldots, \Phi_n)$ is said to be* equitable (EQ) *if for any pair of agents $i, k \in N$, we have $v_i(\Phi_i) = v_k(\Phi_k)$.*

Similar to EF allocations, an EQ allocation may not always exist. Likewise, the following relaxations are proposed in the literature.

Definition 1.5. ***(Relaxations of Equitability.)*** *An allocation is said to be* equitable up to one good (EQ1) *if for any pair of agents $i, k \in N$ such that $\Phi_k \neq \emptyset$, there is an item $g \in \Phi_k$ such that $v_i(\Phi_i) \geqslant v_k(\Phi_k \setminus \{g\})$ (Gourvès et al., 2014; Freeman et al., 2019). Further, an allocation is* equitable up to any good (EQX) *if for any pair of agents $i, k \in N$ such that $\Phi_k \neq \emptyset$, $v_i(\Phi_i) \geqslant v_K(\Phi_k \setminus \{g\})$ for any item $g \in \Phi_k$.*

Note that the highlighted allocation in Example 1.3 neither satisfies EQ nor does it satisfy EQ1 (hence, EQX). Indeed, Carol's utility falls behind Alice's, even if she chooses to ignore an item from the latter's bundle. We remark here that EQ allocation may or may not exist but an EQ1 allocation always exists (under monotone valuations) and can be found efficiently by allocating the least happy agent her most favorite remaining item.

1.2.2 Efficient Allocations

A canonical example of a perfectly envy-free and equitable allocation is to not allocate anything at all. All the items are wasted. Indeed, no one is envious of anyone else as nobody owns anything. In addition, everyone derives zero value from this trivial allocation, which although makes it equitable but keeps it far from being an interesting allocation. This forces some sort of efficiency notions to be introduced in order to avoid the wastage of resources.

Completeness and Non-wastefulness. One basic notion is that of *completeness*, which forces that every item must be allocated to somebody. Along similar lines, we would want to avoid a situation where an item is allocated to an agent who does not value it. An allocation that satisfies this property is called a *non-wasteful* allocation. The allocation in Example 1.3 satisfies both completeness and non-wastefulness.

Utilitarian Social Welfare. Assuming that individual agents model their preferences using valuation functions that map subsets of resources to numbers, *Utilitarian Social Welfare* is defined as the sum of individual valuations. This is one way of measuring the quality of an allocation from the viewpoint of the system as a whole. Although it seems attractive, it might happen that there is one agent who values everything at a large number and hence, gets away with all the items thereby maximizing utilitarian welfare. This forces everyone else to end up with an empty bundle. Consider the allocation Φ in Example 1.3. For binary valuations, any non-wasteful allocation also maximizes the utilitarian welfare. Therefore, Φ is optimal in that sense. Consider a scenario where Alice valued all items at 1 each. Then, an equally optimal allocation would be to give all the items to Alice. This implies that a utilitarian maximal allocation alone does not distinguish how items are partitioned among the agents.

Egalitarian Social Welfare. An egalitarian society may want that to maximize what the poorest person can get. This is captured by *Egalitarian Social Welfare*, defined as the utility derived by the least happy agent. The allocation Φ in Example 1.3 is indeed an egalitarian optimal allocation since at least one of the agents, Bob or Carol, must end up with a utility of at most 1 under any allocation, hence, the egalitarian welfare is 1. Going beyond binary utilities, it may happen that insisting on egalitarian maximality leads to a decrease in the overall average utility. Suppose Alice had a value of 2 each for $g_4, g_5,$ and g_6, then allocating any item to Bob or Carol will decrease the utilitarian welfare while not allocating any item to either of them will decrease the egalitarian welfare.

Nash Social Welfare. Covering the sweet spot between average utility and minimum utility, a widely acknowledged welfare notion is that of *Nash Welfare* (Nash Jr, 1950; Eisenberg and Gale, 1959). It is defined as the geometric mean of the individual utilities. Again, note that Φ in Example 1.3 maximizes Nash welfare as well. Observe that the Nash can only provide a meaningful metric of social welfare if all individual utilities are non-negative.

Generalized p-mean welfares. One may view all of the above welfare notions as special cases of the p-mean welfare (where $p \leq 1$), which is defined as the generalized p-mean of utilities of agents under an allocation. Note that p-mean of a set of numbers $v_1, v_2, \ldots v_n$ is defined as $\left(\frac{1}{n}\sum_{i\in N} v_i^p\right)^{1/p}$. Note that for $p > 1$, the p-mean optimal allocation tends to concentrate the distribution among fewer agents. Consider the simple case of two identical agents with additive valuations who value each of two goods at 1. Then for $p = 2$, an allocation that gives both items to a single agent is optimal, which is contrary to the spirit of desirable allocations. Hence, we focus on $p \leqslant 1$. Moreover, p-mean welfare functions with $p \leq 1$ correspond to functions characterized by a set of natural axioms like the Pigou-Dalton transfer principle (Moulin, 2004). Hence, they are an axiomatically relevant set of welfare functions.

Below, we present the formal definitions of the above discussed welfare notion

Definition 1.6. *(**Welfare measures.**) Given an allocation* Φ,

- Utilitarian social welfare *is the sum of utilities of agents under* Φ, *i.e.,* $\mathcal{W}^{util}(\Phi) := \sum_{i \in N} v_i(\Phi_i)$.
- Egalitarian social welfare *is the utility of the least-happy agent under* Φ, *i.e.,* $\mathcal{W}^{egal}(\Phi) := \min_{i \in N} v_i(\Phi_i)$.
- Nash social welfare *is the geometric mean of utilities of agents under* Φ, *i.e.,* $\mathcal{W}^{Nash}(\Phi) := \left(\Pi_{i \in N} v_i(\Phi_i)\right)^{1/n}$, *and*
- *for any* $p \in \mathbb{R}$, *the* p-mean welfare *is the generalized* p*-mean of utilities of agents under* Φ, *i.e.,* $\mathcal{W}_p(\Phi) := \left(\frac{1}{n} \sum_{i \in N} v_i^p(\Phi_i)\right)^{1/p}$.

1.2.3 Trade-offs Between Fairness and Efficiency

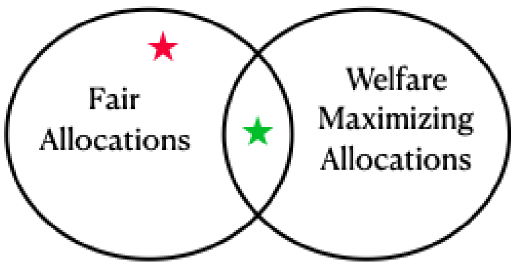

Figure 1.4: If there is an allocation in the intersection, the price of fairness is said to be one. If there no allocation in the intersection, the price is strictly greater than one.

Price of Fairness. Trade-offs are inevitable when we pursue multiple optimization objectives that are typically not simultaneously achievable. Quantifying such trade-offs is a fundamental problem in computation. How much Utilitarian welfare is lost when we insist on having envy-free allocations? How much Nash welfare is lost when we insist on having equitable allocations? For any welfare notion $\mathcal{W}$ and fairness notion $\mathcal{F}$, it is an imperative consideration to address the loss in W because of the constraint $\mathcal{F}$. This is captured by the *price of fairness*, as the "worst-case ratio" of the maximum welfare (measured by $\mathcal{W}$) that can be obtained by *any* allocation, to the maximum welfare that can be obtained among allocations that are fair according to $\mathcal{F}$. Formally,

Definition 1.7. *Let $\mathcal{I}_{n,m}$ denote the set of all fair division instances with n agents and m items. Let $\mathcal{A}(I)$ denote the set of all allocations in the instance I, and further let $\mathcal{A}_{\mathcal{F}}(I)$ denote the set of all allocations in the instance I that satisfy the fairness notion $\mathcal{F}$. Then, the price of fairness (*PoF*) of the fairness notion $\mathcal{F}$ with respect to the welfare measure $\mathcal{W}$ is defined as:*

$$\text{PoF}(\mathcal{F}, \mathcal{W}) := \sup_{I \in \mathcal{I}_{n,m}} \frac{\max_{\Phi^* \in \mathcal{A}(I)} \mathcal{W}(\Phi^*)}{\max_{\Phi \in \mathcal{A}_{\mathcal{F}}(I)} \mathcal{W}_p(\Phi)}.$$

For instance, consider Example 1.3. The highlighted allocation maximizes the utilitarian welfare. But this is clearly not EQ1 as Carol violates EQ1 with respect to Alice. In order to arrive at an EQ1 allocation, at least one item from Alice's bundle must be allocated wastefully to either Bob or Carol. Therefore, the maximum possible welfare under any EQ1 allocation in this instance is 5 while there is an allocation that achieves maximum welfare of 6. Therefore, for this instance, the price of EQ1 is $6/5$.

It is known from the work of Caragiannis et al. (2019b) that under additive valuations, any allocation that maximizes the Nash social welfare satisfies EF1. Thus, the price of fairness of EF1 with respect to Nash social welfare is 1. Further, Barman et al. (2020b) shows that the price of EF1 with respect to utilitarian welfare is $O(\sqrt{n})$ for normalized sub-additive valuations. We elaborate on price of fairness in Section 1.3.

1.3 Recent Progress

The fair division of indivisible items has garnered substantial attention in the past two decades. Unlike the divisible setting, where an EF allocation always exists, the indivisible items may not admit EF allocations under very simple settings like more agents and fewer items. Not only that, deciding whether an instance admits an EF allocation is computationally intractable even for very special and structured instances. In particular, it is NP-Complete even for binary valuations (Aziz et al., 2015) and weakly NP-Complete for two agents and identical valuations (Lipton et al., 2004).

Envy-Freeness and its Relaxations. The story does not end here with non-existence and intractability, but it paves way for a host of approximation algorithms and work-arounds.

- **Hypothetical Removal of an item.** Recall that if an envied agent can get rid of her envy by a hypothetical removal of *any* item from the envied bundle, then such allocations are EFX. This is a strong way to approximate EF, as also remarked by Plaut and Roughgarden (2020) that "Arguably, EFX is the best fairness analog of envy-freeness of indivisible items." There has been a significant amount of effort to explore the EFX landscape, gradually understanding the simpler setting of the small number of agents. Plaut and Roughgarden (2020) first showed that EFX exists for two agents and general valuations. Chaudhury et al. (2020a) extended this to three agents but additive valuations. For an arbitrary number of agents, Babaioff et al. (2021) showed that EFX exists for binary submodular valuations while Bu et al. (2023) extended it to general binary valuations.

  Although the intriguing question of whether an EFX allocation exists for additive valuations stands open, there have been natural relaxations of EFX, discussed as follows.

  - **Partial EFX Allocations (EFX with Charity)**. Caragiannis et al. (2019a) initiated the study of donating items to charity. Here, the property of completeness of allocations is relaxed in order to achieve EFX and the set of unallocated items are said to be donated to charity. Chaudhury et al. (2020b) showed that there always exists an EFX allocation with at most $n-1$ unallocated items. Berger et al. (2022) strengthened it to at most $n-2$ unallocated items for arbitrary agents and at most 1 unallocated item for 4 agents. Follow-up works contributed to decreasing the number of unallocated items and simplifying the previous algorithms (Akrami et al., 2022; Chaudhury et al., 2021; Berendsohn et al., 2022).
  - **Approximate EFX Allocations.** Plaut and Roughgarden (2020) showed that $1/2$-EFX allocations always exist for sub-additive valuations but it may require exponential time (in the number of items) to compute one. Chan et al. (2019) strengthened the above result by giving a polynomial time algorithm for the same. Amanatidis et al. (2020) improved the approximation guarantee from 0.5 to 0.618. Barman et al. (2024b) further imporved upon the approximation ratio in terms of an instance-dependent parameter that upper bounds, for each indivisible good in the given instance, the multiplicative range of nonzero values for the good across the agents.

– **Epistemic EFX Allocations.** Caragiannis et al. (2023) defined and proved the existence and polynomial-time computability of epistemic EFX for additive valuations where for every agent, there exists a way to shuffle the goods of the other agents such that she does not envy any other agent up to any good. Akrami and Rathi (2024) extended the existence to general monotone valuations.

EF1, the other relaxation of envy-freeness, weaker than EFX, comes with both existential and computational guarantees. An EF1 allocation can be found in polynomial-time by simple algorithms like *Envy-Cycle Elimination* and *Round-robin*. The former was proposed by Lipton et al. (2004) long before EF1 was formally defined by Budish (2011), and works for monotone valuations. Here, at every step, an unallocated item is given to an unenvied agent. Intially, every agent is unenvied, so the choice is arbitrary. At any point, if all the agents are envious, then there must be a directed cycle in associated envy graph, which is a directed graph in which each agent points to everyone it envies. The algorithm then breaks the cycles by a cyclic exchange of bundles and once all cycles are removed, the envy graph must have a source node with any outgoing edges. The agent corresponding to this node represents an unenvied agent, who can now receive the next unallocated item. Notice that the allocation is EF1 by construction, as every envious agent can choose to ignore the last added item in the envied bundle. Round-robin is a simpler algorithm that finds an EF1 allocation for additive valuations. It first fixes an ordering of the agents and, according to this ordering, it lets one agent at a time choose their favorite available item until all items have been allocated. To see why it is EF1, consider an agent i who comes before another agent j in the ordering. Then, i never envies j since it gets a chance to pick first in every round and stays ahead of j. On the other hand, j might be envious towards i but note that once i picks her first item say g, from that point onwards, j becomes the agent who gets to pick first. Therefore, j can choose to remove g and be envy-free. Therefore, the allocation is indeed EF1.

- **Minimizing the envy.** Apart from the relaxation in terms of the hypothetical removal of some items from the envied bundle, minimizing the degree of envy when envy-free allocations do not exist is also of interest. The degree of envy can be captured by various measures, for instance, the number of envious agents, the aggregate amount of envy experienced by all the agents, and so on. Chevaleyre et al. (2007) proposed a framework for defining the degree of envy of an allocation based on the degree of envy among individual agents, where the envy of agent i against j is

$\max(0, v_i(\Phi_j) - v_i(\Phi_i))$, while Nguyen and Rothe (2013) defined the amount of envy as $\max\left(1, \frac{v_i(\Phi_j)}{v_i(\Phi_i)}\right)$. Shams et al. (2021) focussed on choosing the allocation that is fair in the sense of the distribution of the envy among agents and used ordered weighted average of the envy vector as the measure of envy.

- **Subsidy.** The use of money for the fair division of indivisible items is well-studied (Svensson, 1983; Tadenuma and Thomson, 1993; Maskin and Feiwel, 1987; Edward Su, 1999; Klijn, 2000). Recently, Halpern and Shah (2019) showed that envy-freeness can be achieved with a small amount of money (divisible good) and gave a strongly polynomial time algorithm to compute the minimum subsidy. Brustle et al. (2020) strengthened the above result by showing that a subsidy of at most one dollar per agent is sufficient to guarantee the existence of an envy-free allocation (where the marginal value of each item is at most one dollar). Choo et al. (2024) looked at house allocations with minimum subsidy.
- **Sharing items.** Finding envy-free allocations with a bounded number of shared items was first studied by Brams and Taylor (1996) with two agents, additive utilities and at most 1 sharing. Sandomirskiy and Segal-Halevi (2019) showed that such an envy-free and fractionally Pareto optimal allocation with the smallest possible number of shared items can be found in polynomial time. Recently, Goldberg et al. (2022) considered consensus splitting (a partition of items into k subsets such that every subset is valued at same value by all the agents) and showed that computing a partition with at most $(k-1)n$ sharings can be done in polynomial time. Bismuth et al. (2024) provided bounds on number of sharings for various special cases. They showed that given the number of agents n and shared items s, for binary valuations, the existence of a fair allocation for n agents with s shared objects can be decided in polynomial time by a mixed integer linear program with a fixed number of variables. They also showed the hardness of finding envy-free allocations for non-degenerate instances.
- **Hiding items.** Hosseini et al. (2020) defined the notion of information withholding (envy-freeness upto k hidden items (HEF-k)) where an agent can choose to reveal only partial information about her allocated bundle. An agent can hide (or withhold) some of the goods in her bundle and reveal the remaining goods to the other agents. They show that deciding the existence of HEF-k is NP-Complete even for identical valuations, but in practice, envy-freeness can be achieved by hiding only a small number of items. Hosseini et al. (2023a) empirically demonstrated through crowd-sourcing experiments

that allocations achieved by withholding information are perceived to be fairer than EF1. Recently, Bliznets et al. (2024) considered the problem of finding the largest number of items to allocate to the agents in the given social network so that each agent hides at most one item and overall at most k items are hidden such that no one envies her neighbors. They showed that the problem admits an XP algorithm and is W[1]-hard parameterized by k.

Equitability and its Relaxations. Unlike EFX, it is known that EQX always exists for additive valuations and can be computed in polynomial time as well. It boils down to giving the least happy agent her most favorite remaining item (Gourvès et al., 2014). (They used the term near-jealousy freeness). As far as practicality, empirical relevance, and perceived fairness are considered, experiments suggest that equitability (or inequality aversion) is a preferable criterion over envy-freeness (Herreiner and Puppe, 2009, 2010; Gal et al., 2016). Freeman et al. (2019) explored EQX in conjunction with efficiency notions. They showed that an EQX and PO allocation always exists for strictly positive valuations but may fail to exist in the presence of 0-valued items. They gave a pseudo-polynomial time algorithm that always returns an EQ1+PO allocation for strictly positive valuations. They showed the (strong) hardness of finding an EQ+PO/EQX+PO/EQ1+PO allocation.

Freeman et al. (2020) did a similar analysis in the setting of chores (where agents derive negative value from the items). They demonstrate a set of differences between the goods and chores settings in the context of equitability. They showed that in the chores setting, Leximin does not even guarantee equitability up to one chore while in goods setting, leximin simultaneously satisfies EQX and PO. While they provide a pseudo-polynomial time algorithm for EQ1+PO, they showed that EQX+PO may no longer exist for chores. Given this non-existence, they defined equitability up to one/any duplicated chore (DEQ1/DEQX) properties. These entail that pairwise equitability can be restored by duplicating a chore from a poor agents's bundle and adding it to the rich agent's bundle, rather than only removing a chore from the less happy agent's bundle. They showed that the "duplicate" relaxations in conjunction with PO are satisfied by the Leximin allocation. Note that this is an existential result and the computational complexity of DEQX+PO allocations remains open.

Sun et al. (2023b) studied equitability for both goods and chores in conjunction with utilitarian and egalitarian welfare. They studied the computational complexity of deciding the existence of an EQX/EQ1 and welfare-maximizing allocation and computing a welfare maximizer among all EQX/EQ1 allocations. Recently, Barman et al. (2024a) considered EQX for the mixed instances with both goods and chores and showed that computing an EQX allocation (even without any efficiency requirement) in the mixed setting is weakly NP-Hard even for two agents and strongly NP-hard for more agents. They also extended the existence of EQX allocations to monotone valuations (not necessarily additive) and gave a pseudo-polynomial time algorithm for the same.

Fairness, Efficiency and Price of Fairness. Given the existence of EF1 allocations, a natural question is to find EF1 allocations in conjunction with efficiency notions. Caragiannis et al. (2019b) showed that any allocation that has the maximum Nash welfare is guaranteed to be EF1 and PO. However, maximizing the Nash social welfare over integral allocations is an NP-hard problem (Nguyen and Rothe, 2014). Additionally, the problem is known to be APX-hard Lee (2017). Therefore, although Nash guarantees EF1 and PO, it does not provide a tractable algorithm to find such an allocation. To that end, (Barman et al., 2018a) gave a pseudo-polynomial algorithm to find an allocation that is both EF1 and Pareto-optimal. It constructs integral Fisher markets wherein specific equilibria are not only efficient, but also fair. They also present a polynomial-time 1.45-approximation algorithm for the Nash social welfare maximization problem. Aziz et al. (2023b) looked at EF1 in conjunction with maximizing utilitarian welfare. They showed that among the utilitarian-maximal allocations, deciding whether there exists one that is also EF1 and among the EF1 allocations, computing one that maximizes the utilitarian welfare are both strongly NP-hard problems when the number of agents is variable. They design pseudo-polynomial time algorithms when the number of agents is fixed.

The study of efficiency loss due to fairness constraints in the fair division setting was initiated by Bertsimas et al. (2011) and Caragiannis et al. (2012), who studied the notion of envy-freeness, proportionality, and equitability in divisible and indivisible goods and chores. Later, Bei et al. (2021) considered fairness notions whose existences are guaranteed and presented lower-bounds of $O(\sqrt{n})$ and upper-bound of $O(n)$ on the price of EF1, which was then closed by Barman et al. (2020b) who showed that the price of envy-free up to one item and of (1/2)-approximate maximin share are both $O(\sqrt{n})$. Sun et al. (2023a) studied the efficiency loss of EF and its relaxations for chores. Sun et al. (2023b) studied the price of equitability with respect to both utilitarian and egalitarian welfare for indivisible goods and chores. Nicosia et al. (2017) studied efficiency loss due to maximin, Kalai–Smorodinski, and proportional fairness. Kurz (2016) highlighted the impact of the number of items on the price of fairness.

1.4 Our Contributions

The contribution of this thesis can be broadly divided into three parts:

1. First, we study the computational complexity of finding fair (approximately envy-free or equitable) allocations in various settings and for each of them, identify the domain restrictions where computational tractability holds. We analyse the tractability of EF allocations on graphical valuations and also present parameterized algorithms for the same. We analyze minimizing three different measures of envy in the context of house allocations and give algorithms for extremal instances, single-peaked and single-dipped instances. We study approximately equitable allocations for the setting of mixtures of goods and chores. We also study envy-freeness and equitability coupled with efficiency notions for each of the above areas.
2. Second, we quantify the loss in the various welfare measures due to the fairness constraints (approximately envy-free or equitable) and present bounds, tight up to a constant factor, on the price of minimizing envy and the price of equitability. We analyze the price of fairness from the perspective of agents-types—a smaller parameter than the number of agents. Rather than focussing on a single welfare measure, we give tight bounds for generalized p-mean welfares. We also identify structured instances where no price has to be paid in terms of welfare.
3. Third, we propose novel concepts of Secure and Abundant allocations, which in addition to being a relaxation of consensus allocations (Simmons and Su, 2003), also capture elements of human psychology molded and influenced by the perspectives of others.

We now highlight the main results below.

1. **Computational Complexity of finding (approximately) Envy-free or Equitable allocations**.
    - **Envy-Free Allocations:** Chapter 2 deals with allocation instances with graphical valuations. We show that, for graphical valuations, when agents have binary utilities over the items, the existence of EF allocations can be determined in polynomial time. In contrast, we show that allowing for even slightly more general utilities leads to intractability even for graphical valuations. In particular, we show that it is NP-complete to determine if an instance admits an EF allocation even when all agents value every item at either $0, 1$, or d. This motivates other approaches to tractability,

and to that end, we show that the problem is FPT when parameterized by the vertex cover number of the graph associated with the utilities when the number of distinct utilities is bounded.

Chapter 3 focusses on instances with low degeneracy (which is a property of the valuation matrix that captures the degree of similarity in the valuations). We show that for an arbitrary number of agents, finding an EF and fractionally Pareto optimal allocation remains hard even for instances with low degeneracy, establishing the limitation of degeneracy as a parameter.

Chapter 4 focuses on minimizing envy in the context of house allocations. We give a comprehensive picture of finding an allocation that minimizes the number of envious agents, the maximum envy experienced by any agent, and the total envy experienced by all the agents together. We give efficient algorithms for all the above three measures for extremal instances (Elkind and Lackner, 2015). Beyond extremal instances, we show the hardness of the first two measures even when every agent values at most two houses. Towards parameterized complexity, we formulate ILPs that lead to fixed-parameter tractability with respect to the parameter number of agent-types/house-types. We also show that all three measures admit linear kernel when parameterized by the number of agents, using the expansion lemma. We also give efficient algorithms for minimizing the number of envious agents for single-peaked and single-dipped preferences.

- **Equitable Allocations:** Chapter 6 focusses on equitability as the fairness measure in the context of mixed items. That is, agents could derive either positive or negative utility from a subset of items. Equitability for scenarios involving only goods (Freeman et al., 2019) and only chores (Freeman et al., 2020) had been previously explored, but its implications in a mixed setting remained unknown. To that end, for mixed items, we exhibit the hardness of finding an EQ1 allocation with non-normalized valuations. We complement the hardness by identifying some tractable restricted domains (objective valuations, $\{-w, 0, w\}$-valuations and such).

2. **Bounds on Price of EFX, Price of Minimizing Envy and Price of Equitability**.

   - **Price of EFX:** Section 2.3 discusses the price of EFX (for graphical instances) with respect to Utilitarian, Egalitarian, and Nash welfare. For binary graphical valuations, the "price of EF" relative to utilitarian social welfare is 1. Since EFX allocations are possibly wasteful, we also address the question of determining the price of fairness of EFX allocations in the context of graphical valuations. We show that the price of EFX with respect to utilitarian welfare in the context of graphical valuations is one for binary utilities, but can be arbitrarily large for instances where all items are valued at 0, 1, or d. We also show the hardness of deciding the existence of an EFX allocation which is also welfare-maximizing and of finding a welfare-maximizing allocation within the set of EFX allocations.

- **Price of Minimizing Envy:** Section 4.9 discusses the price of minimizing envy (in three different ways) with respect to Utilitarian welfare in the context of house allocations. We show that the price of minimizing envy is one for $m = n$ and binary valuations. Also, for such instances, there is an allocation (efficiently computable) that simultaneously minimizes the number of envious agents, the maximum and total envy, while maximizing Utilitarian welfare. Moreover, we show that for $m > n$, the price of fairness is $\Theta(n)$ for all the three envy optimization objectives.
- **Price of Equitability:** Chapter 5 discusses the price one has to pay in terms of welfare in order to arrive at approximately equitable allocations. We give tight bounds on the 'price of equitability' with respect to generalized p-mean welfare. We also identify instances for which no price has to be paid for any p-mean welfare measure (like doubly normalized binary valuations and identical binary submodular valuations).

3. **New notions of Secure and Abundant allocations**
   - **Generalized Consensus Allocations:** In the context of divisible items, the existence and complexity of "exact divisions" (where all agents agree on the value of the division) has been well-studied, usually referred to as consensus halving (Simmons and Su, 2003). In practice, a perfect consensus may, in general, be too much to ask for. A natural relaxation to ask for an approximate consensus: where all agents agree that all bundles have a value in some specified range, say $[p, q]$. In Chapter 7, we show that even in the setting of additive binary valuations, the problem of dividing m items into a collection of k bundles so that all k bundles are valued at either 0 or 1 is NP-complete. Moreover, even "almost" exact equitable divisions, where all but some c bundles are valued in the range $[p, q]$ is hard. Therefore, we consider other ways of relaxing the demands we make from a perfectly equitable consensus allocation. Here, we treat a common valuation as a

target lower or upper bound, instead of an exact goal. Now we consider allocations where all external valuations of any bundle are: (a) at least the base value; and (b) at most the base value. The former is called a *Secure* allocation while the latter is called an *Abundant* allocation. These concepts relate also to the quality of the valuations from a "user experience" perspective. Secure allocation captures the fact that everyone wants to feel validated/secure by the opinions of others. That is, the sense of security for an individual is linked to how much other agents appreciate her possessions. If an agent owns a bundle that she values highly but that everyone else deems worthless, then it is natural for the agent to be unhappy about the situation, or at any rate, be suspect about their own judgment. Likewise, we also want to avoid situations where agents underestimate their own bundles, which are captured by abundant allocations.

We show that both secure and abundant allocations always exist, but are not very interesting on their own. For instance, allocating all the items to one agent who values the entire bundle at the least values is trivially a secure but unfair allocation. So, we couple these notions with capacity and welfare constraints. We give efficient algorithms for extremal instances but show hardness for the general binary valuations.

Acknowledgements.

Chapter 2 and Chapter 3 are based on joint works with Neeldhara Misra. The former will appear at ADT 2024 and the latter appeared at SOFSEM 2021 (Misra and Sethia, 2021). Chapter 4, a joint work with Jayakrishnan Madathil and Neeldhara Misra appeared as an extended abstract at AAMAS 2023 (Madathil et al., 2023). Section 4.8 is based on a joint work (manuscript) with Hadi Hosseini and Sanjukta Roy. Chapter 6 is a joint work with Hadi Hosseini that will appear at CFD workshop at IJCAI 2024. Chapter 5 is based on a joint work with Umang Bhaskar, Rohit Vaish, and Neeldhara Misra which appeared at SAGT 2023 (Bhaskar et al., 2023). Chapter 7, joint work with Neeldhara Misra, is under review.

Chapter 2

Envy-Free and Efficient Allocations for Graphical Valuations

"If we can't be certain of mathematical truths, can we be certain of anything?"

- Chris Bernhardt, *The Birth of Computer Science*

2.1 Introduction

As discussed in the Chapter 1, finding allocations that are envy-free is one of the gold standards in allocation problems. This entails that no agent should feel envious of any other agent under the allocation. That is, every agent should value her own allocated bundle at least as much as it values anyone else's bundle. The problem with such envy-free allocations is two-fold: existential and computational. That is, they might not exist for many instances (say, when there are more agents than items), and deciding whether they exist is computationally intractable even for very special and structured instances. In particular, it is NP-complete even for binary valuations (where agents value items at either 0 or 1) (Aziz et al., 2015) and weakly NP-Complete for two agents and identical valuations (Lipton et al., 2004).

Motivated by these issues, in this work, we focus on envy-freeness in the context of graphical valuations. These are a recently introduced (Christodoulou et al., 2023) class of structured valuations which are interesting because they admit EFX allocations even for any number of agents. These are valuations where every item is valued by exactly two agents, lending a (simple) graph structure to the utilities, where the agents are associated with vertices and items with edges. Two agent-vertices are adjacent if and only if they value a (unique) common edge-item, represented by the edge between them. These kind of valuations capture scenarios where the number of agents with interest in any specific item is limited. Such valuations may arise in situations where agents only value the items that are geographically closer to them. For instance, in real estate allocation, potential buyers might only be interested in properties within a certain distance from their workplace or amenities; employees might value office spaces closer to their teams and likewise (Christodoulou et al., 2023).

Related Work.

Several special cases and approximations have been extensively studied in the fair division literature to understand the extent of tractability of EFX allocations: binary valuations (Bu et al., 2023); bounded number of agents (Plaut and Roughgarden, 2020; Chaudhury et al., 2020a; Akrami et al., 2023); and bounded number of unallocated items (Caragiannis et al., 2019a; Berger et al., 2022). Graphs have also been associated with fair division in various contexts and models. Allocations, where items allocated to each agent form a connected subgraph in a provided item graph, have been studied (Bouveret et al., 2017; Deligkas et al., 2021a; Bilò et al., 2022). In a different model, Payan et al. (2023) looked at graph-EFX which requires that an agent, represented by a vertex, satisfy EFX only against her adjacent vertices. Our work is closely aligned with that of Christodoulou et al. (Christodoulou et al., 2023) who introduced graphical valuations and showed the hardness of deciding the existence of an EFX orientation. Following this, Zeng and Mehta (Zeng and Mehta, 2024) characterized that graphs with chromatic number at most 2 admit EFX orientations for any given valuations, while graphs with chromatic number strictly greater than 3 may not admit such orientations for all valuations. They also characterized EFX orientability for binary valuations.

The quantification of welfare loss that is inevitable due to the fairness constraint has also been of interest in the literature. To capture this, the notion of *price of fairness* was proposed in the works of Bertsimas et al. (2011) and Caragiannis et al. (2012). Since then, various works have given bounds for the price of proportionality, envy-freeness, EF1, EFX, equitability, EQ1, maximum Nash welfare, and more (Aumann and Dombb, 2015; Bei et al., 2021; Sun et al., 2023a,b; Bhaskar et al., 2023).

Our Contributions.

We highlight our main contributions below and put them in context with the already-known results.

- We show that an EF allocation if it exists, can be found efficiently for graphical valuations where agents have binary ($\{0, 1\}$) valuations over the items (Theorem 2.5). This is in contrast to the intractability of EF allocation for binary utilities in general.
- We show that if we allow for even slightly more general valuations than binary, for instance, $\{0, 1, d\}$-valuations for some constant d, the problem again becomes intractable (Theorem 2.6).
- The above hardness motivates a parameterized approach towards tractability and towards that, we present a *fixed-parameter tractable*[1] algorithm for finding EF allocations for graphical instances with bounded number of distinct utilities, where the parameterization is in terms of the minimum vertex cover of the associated graph G (Theorem 2.8).
- We show that if there is an EF allocation for any graphical instance, then there is also an EF allocation that does not 'waste' any item, that is, it does not assign an item to an agent who derives 0 value from it. This shows that if there is an EF allocation, then there is an EF orientation of the graph G (Theorem 2.1). This result stands in contrast to the fact that an EFX allocation always exists but an EFX orientation may not exist (Christodoulou et al., 2023). In terms of the price of EF, this implies that for $\{0, 1\}$-graphical valuations, there is no loss in the welfare while achieving EF allocations, whenever they exist.
- Christodoulou et al. (2023) showed that EFX allocations not only always exist but can be found efficiently for graphical valuations. But this comes with a sacrifice in terms of welfare. In particular, there are cases where any EFX allocation must assign items to agents for which they are irrelevant (0-valued). In this work, we quantify the loss of welfare while achieving EFX allocations and show that for $\{0, 1\}$-graphical instances, the price of EFX for Utilitarian (sum of agents' utilities) welfare is 1 (Theorem 2.9). That is, restricted to binary graphical valuations, there is no loss in any of the welfare notions and an EFX allocation that maximizes the respective welfare can be found efficiently. On the other hand, we show that for slightly general valuations than binary, that is, for $\{0, 1, d\}$-valuations, there are instances with a huge loss in the utilitarian welfare and

[1] An algorithm that runs in time $f(k)poly(n, m)$ where f is some computable function of the parameter k.

consequently, price of EFX shoots up to ∞ (Theorem 2.10).

- On the computational side, we show that for general graphical valuations, finding EFX allocations that also maximize utilitarian welfare is NP-Hard (Theorem 2.11). It follows that finding a welfare-maximizing allocation within the set of EFX allocations is also hard.

Graphical Allocation Instance.

For standard fair allocation terminologies, we refer to Section 1.2. A graphical fair allocation instance $\mathcal{I} = \{G = (V, E), \mathcal{V}\}$ takes as input an undirected, simple graph G and a valuation function $\mathcal{V}$. The set of vertices V in G corresponds to n agents and the set of edges E in G corresponds to m items to be allocated. We will often use the terms "items" and "edges" interchangeably because of this correspondence. Every agent only values a subset of the incident edge-items. In addition, since an edge is incident on exactly two vertices, an edge-item is valued by exactly two agents. Note that every pair of agents value at most one edge together, the one which is incident on both of them. A $\{0,1\}$-graphical instance is one such that $v_i \in \{0,1\}\ \forall\, i \in N$. Given a graph G, an *orientation* O_G is an allocation with the additional property that every edge is assigned to one of the two endpoints. A directed graph that directs the edges of G towards the vertex that receives the edges is called an *orientation graph* of G. Note that every orientation corresponds to a complete allocation. An allocation is an orientation if it assigns the edges to the incident vertices. We say that an orientation satisfies a property if the corresponding allocation satisfies that property.

2.2 Envy Free Allocations

Although it is known that EFX allocations always exist on graphical valuations Christodoulou et al. (2023), an EF allocation may not exist on graphical instances as well, as illustrated by a simple example of a graph consisting of only one edge. Whichever incident vertex receives the edge, the other one is bound to be envious. We show that it is possible to determine if an EF allocation exists in polynomial time for $\{0,1\}$-graphical valuations, and in the event that the instance admits an EF allocation, such an allocation can be found in polynomial time. Before that, we present a series of structural results. The following result is in contrast to the EFX fairness, where the existence of an EFX allocation does not guarantee an EFX orientation but any EF allocation does guarantee an EF orientation.

Theorem 2.1. *Given a graphical allocation instance, there is an EF allocation if and only if there* *is an EF orientation.*

Proof. An orientation is EF if the corresponding allocation is EF, so the reverse direction holds. We argue the forward direction.

Suppose there is an EF allocation Φ for the given instance, which does not correspond to any EF orientation. We assume that everyone values at least one item, otherwise the agent can be removed from the instance. Since Φ is not an orientation, there must be some edges allocated to vertices that are not incident on them. All such edges are allocated wastefully as an agent does not value an edge that is not incident on itself. Consider the re-allocation Φ' such that all such wastefully allocated edges are re-allocated to one of their incident vertices, chosen arbitrarily. Say, edge $e = (ij)$ which was previously wastefully allocated to vertex k is now re-allocated to i, WLOG.

Under Φ', an agent who loses an item can not envy anyone, except possibly its neighbors, as its utility does not decrease. Any agent can potentially be envious of only those agents that are incident on it. Indeed, if i is not incident to k, then $v_i(\Phi'_k) = 0$ as k only receives the edges incident on itself, none of which are valued by i.

Moreover, suppose j is envious of i under Φ' as i receives the edge $e = (ij)$ that is also valued by j. Notice that e is the only item that is valued by j in the bundle Φ_i since it is the unique item valued by both i and j. Therefore, if j is envious of i, we have $v_j(\Phi'_i) = v_j(e) > v_j(\Phi'_j) \geq v_j(\Phi_j)$. The last inequality holds as no agent's utility decreases under the re-allocation Φ'. This implies that j valued e more than the bundle it got under the EF allocation Φ. But then, $v_j(\Phi_j) < v_j(e) \leq v_j(\Phi_k)$, where k is the recipient of e under Φ. This implies that j was envious of k in the allocation Φ, which is a contradiction to the fact that Φ was EF. Therefore, all the agents are EF under Φ', and Φ' assigns edges to only incident vertices. Therefore, Φ' corresponds to an EF orientation. □

Lemma 2.2. *Given any graphical allocation instance with general additive valuations, suppose v_i^{max} is the maximum value any agent i has for any item. Then, a non-wasteful allocation is EF if and only if i gets a utility of at least v_i^{max} $\forall\, i \in [n]$.*

Proof. Let Φ be any EF allocation. Let v_i^{max} be the maximum value an agent i has for an edge e. Suppose $v_i(\Phi_i) < v_i^{max}$, then clearly, $e \notin \Phi_i$. Let $e \in \Phi_j$ for some agent j. Then $v_i(\Phi_i) < v_i(e) = v_i(\Phi_j)$. Therefore, i is envious of j which is a contradiction to the fact that Φ is EF. Therefore, every agent i must get a utility of at least v_i^{max} under an EF allocation Φ. Conversely, suppose every agent gets a utility of at least v_i^{max} under a non-wasteful allocation Φ. Since Φ

	g_1	g_2	g_3	g_4
a_1	(1)	0	1	1
a_2	0	(1)	0	(0)
a_3	1	0	(1)	1

Table 2.1: An EF allocation that allocates an item wastefully.

is a non-wasteful allocation, it corresponds to an orientation in G. So every agent receives a subset of edges that are incident on it. Consider an agent i. We have $v_i(\Phi_i) \geq v_i^{max}$. Consider any other agent j incident on i. If the edge $e = (ij) \in \Phi_j$, then $v_i(\Phi_j) = v_i(e) \leq v_i^{max}$, else $v_i(\Phi_j) = 0$, as i does not value any edge incident on j except e. Also, for any agent j not incident on i, $v_i(\Phi_j) = 0$ as i does not value any edge which is not incident on itself. Therefore, we have that $v_i(\Phi_i) \geq v_i(\Phi_j)$ for all $1 \leq i \neq j \leq n$ and hence the orientation is EF. □

This gives us the following corollary.

Corollary 2.3. *For graphical instances, if an agent i gets a utility of at least v_i^{max} under a partial orientation O_P, then i remains EF under any extension of O_P.*

In particular, for binary valuations, there is an EF allocation where every item is allocated to an agent who values it at 1, therefore, we have the following result.

Corollary 2.4. *For $\{0,1\}$-graphical instances, the price of EF with respect to utilitarian social welfare is* 1.

Note that the above result is not true for binary valuations in general. Consider the instance in Table 2.1. It is not a graphical instance as a_1 and a_3 value 3 items positively. An EF allocation must allocate at least one item from $\{g_1, g_3, g_4\}$ wastefully. Indeed, if all of them are allocated non-wastefully, then the agent who ends up receiving two of them is envied by the other one. The highlighted allocation is one of the EF allocations.

We are now ready to present the algorithm for binary graphical instances.

Theorem 2.5. *For $\{0,1\}$-graphical instances, an EF allocation can be found efficiently, if it exists.*

Proof. Consider an instance $\mathcal{I}$ of $\{0,1\}$-graphical valuations. Since an EF allocation exists if and only if there is an EF orientation (Theorem 2.1), we will construct an EF orientation if it exists. For all asymmetric edges $e = (ij)$, we orient them towards the incident agent who values e at 1, say i. This does not create any envy in the graph as the only agent who values e is i. We call such vertices i as special vertices since they remain envy-free under any completion of the allocation and are not envied by anyone else. Once we orient all the asymmetric edges, we remove them from the graph. The edges which are valued at 0 by both end-points are oriented arbitrarily and removed from the graph. This gives us a collection of connected subgraphs $H = \{H_1, H_2, \ldots H_k\}$ such that all edges in H are symmetric and valued at 1 by both the end-points. For each $H_i \in H$, we consider the following cases:

1. H_i is a tree. Then, $V(H_i) = E(H_i) + 1$. By pigeonholing, at least one agent, say i, does not receive any edge item from $E(H_i)$. Such a vertex i is always envious under any allocation unless i is already a special vertex. In the former case, there is no complete EF allocation. Otherwise, if there is a special agent i, then we root H_i either on i and construct an orientation such that every vertex gets an edge item from her parent. This way, everyone except i receives a utility of at least 1 from the edges in H_i and hence is EF in any complete orientation. Also, i is EF since it is a special vertex.

2. H_i contains a cycle, say $C = \{u_1, u_2, \ldots u_c, u_1\}$. We orient the edges (u_i, u_{i+1}) towards u_i and (u_c, u_1) towards u_c. Then, every vertex in the cycle is EF as $v_i(\Phi_i) \geq 1$ and remains EF in any completion of this orientation (Corollary 2.3). Therefore, the edges inside the cycle can be oriented arbitrarily. We now remove the cycle C from H_i, replace it with a vertex c, and construct a spanning tree of H_i rooted at c. We then construct an orientation that allocates every vertex in the spanning tree, except c, an edge from her parent. This implies that every agent in the spanning tree except the root c ends up with a utility of at least 1. All agents corresponding to the root c already had a utility of at least 1. Since all the agents in H_i now have utility at least 1, therefore everyone is EF in any completion of the partial orientation. Therefore, the remaining edges in H_i can be oriented arbitrarily, and hence we get an EF allocation for H_i.

The algorithm loops over every H_i in H and if there is an EF allocation for every H_i, it corresponds to a complete EF allocation (since vertices across components do not envy each other). Else, if there is at least one H_i for which there is no EF allocation, then the algorithm outputs that no complete EF allocation exists. This is true because an envious agent in H_i can not be made EF by any of the edges in the other components, as it does not value them. This settles our claim. □

We now show in the following result that if we slightly generalize from binary to $\{0, 1, d\}$ graphical valuations, it becomes hard to decide if the graphical instance admits an EF allocation.

Theorem 2.6. *Deciding whether an EF allocation exists is NP-Hard even for symmetric $\{0, 1, d\}$-graphical valuations.*

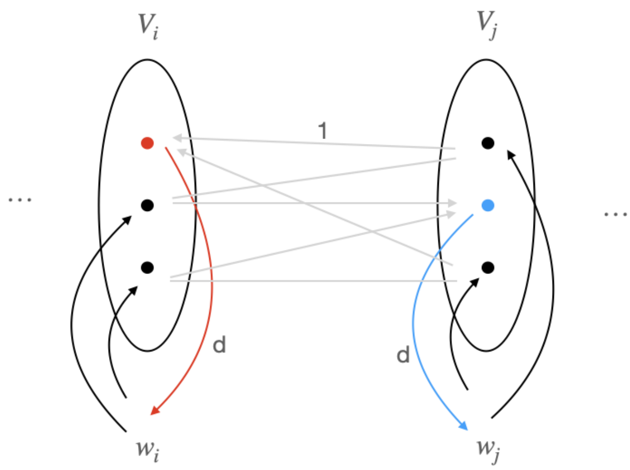

Figure 2.1: A schematic of reduced instance in the proof of Theorem 2.6.

Proof. We present a reduction from Multi-Colored Independent Set (MCIS) (Fellows et al., 2009), where given a regular graph $G = (V_1 \uplus \cdots \uplus V_k, E)$, the problem is to decide if there exists a subset $S \subseteq V(G)$ such that $G[S]$ is an independent set and $|V_i \cap S| = 1$ for all $i \in [k]$. We construct the graphical instance as follows. All vertices in $V(G)$ correspond to agents and all edges in $E(G)$ to items. Every agent $v \in V(G)$ values her incident edges at 1. That is, all edges in G are symmetric with a weight of 1. For every vertex partition V_i, we add a vertex-agent w_i adjacent to all the vertices in V_i. Every edge $\{(w_i, v) : v \in V_i\}$ is a symmetric edge such that w_i and v value it at d, where d is the degree of any vertex in the (regular) graph G. This completes the construction. A schematic of this construction is shown in Figure 2.1. We now argue the equivalence.

Forward Direction.

Suppose MCIS is a Yes-instance and there is an independent set $S = \{s_1, s_2, \ldots s_k\} \subseteq V(G)$ such that $G[S]$ is an independent set and $|V_i \cap S| = 1 \; \forall i \in [k]$. Then, we do the following orientation of $E(G)$ to get an EF allocation.

- $\{(s_i, w_i) : i \in [k]\}$ are oriented towards w_i.
- $\{(v, w_i) : v \in V_i \setminus \{s_i\}\}$ are oriented towards v.
- $\{(s_i, v) : v \in N(s_i) \setminus \{w_i\}\}$ are oriented towards s_i.
- All the remaining edges are oriented arbitrarily.

Let Φ be the allocation corresponding to the above orientation. Then, $u_{w_i}(\Phi_{w_i}) = d$. Note that since all the other edge-items valued at d by w_i are allocated to distinct agents in V_i, hence $u_{w_i}(\Phi_j) \leq d$ for any $j \in V(G)$, so all agents $\{w_i : i \in [k]\}$ are envy-free. Similarly all $v \in V_i \setminus \{s_i\}$ have a utility of d each for Φ_v and a utility of at most d for any other bundle. So, all such agents are envy-free. The remaining agents $\{s_i : i \in [k]\}$ get a utility of d from the d distinct edge-items ($|N(s_i) \setminus w_i| = d$) valued at 1 each in their respective bundles. Note that all s_i also value any other bundle at atmost d and hence $\{s_i : i \in [k]\}$ are envy-free. This implies that Φ is an EF allocation.

Reverse Direction.

Suppose there is an EF allocation Φ in the reduced instance. Under Φ, each of the $w_i's$ must get at least one incident edge-item to be envy-free. Otherwise, $u_{w_i}(\Phi_i) = 0$ but w_i values every bundle that ends up with any edge-item $\{(w_i, v) : v \in V(G)\}$ at d, and hence is envious. Also, since there are only $|V_i| - 1$ edge-items valued at d by all the $|V_i|$ agents in V_i, so by pigeon-holing, there is at least one agent in every partition V_i, say $s_i \in V_i$, which does not end up with a d-valued item. Since $u_{s_i}(\Phi_{w_i}) = d$, therefore, s_i must get a utility of at least d from the remaining items in order to be envy-free. This is feasible only if all the agents $\{s_i : i \in [k]\}$ get the respective d edge-items incident on them in the original graph G. This implies that $\{s_1, s_2, \ldots s_k\}$ must form an independent set in the original graph G. This settles the reverse direction. □

Given the hardness of finding EF allocation for $\{0, 1, d\}$-graphical valuations, we consider the parameterized tractability in this context. On a positive note, we show that the problem admits an FPT algorithm parameterized by the vertex cover number of the associated graph, which is the size of the smallest vertex cover (a set of vertices that includes at least one endpoint of every edge) of the graph. We will use the following classical result by Lenstra.

Theorem 2.7 (Lenstra (1983)). *An integer linear programming (ILP) instance of size L with p variables can be solved using $\mathcal{O}\left(p^{2.5p+o(p)} \cdot (L + \log M_x) \log (M_x M_c)\right)$ arithmetic operations and space polynomial in $L + \log M_x$, where M_x is an upper bound on the absolute value a variable can take in a solution, and M_c is the largest absolute value of a coefficient in the vector c.*

Theorem 2.8. *Given a graphical allocation instance with a 'bounded number of distinct utilities, the problem of finding an EF allocation admits an FPT algorithm parameterized by the Vertex Cover Number of the associated graph G.*

Proof. We formulate an ILP where the number of variables is bounded by a function of the size of the minimum vertex cover of G. We will show that the ILP is feasible if and only if there is an EF orientation in the allocation instance. Then, we invoke Theorem 2.7 to get a feasible solution of the ILP, if it exists, and hence, get the desired FPT algorithm parameterized by minimum vertex cover number.

Let B be the (bounded) set of distinct utilities. Let S be a minimum Vertex Cover of G and $|S| = k$. We have that $I = V(G) \setminus S$ is an independent set. We say that two vertices in I are of the 'same class' C_i if they are incident to the same subset of vertices in S. This partitions I into at most 2^k classes, corresponding to the subsets of S. That is, $I = \{C_1, C_2, \ldots C_{2^k}\}$. Further, for each class C_i, we say that two vertices have the 'same signature' σ_i if they value the subset in S in the same manner. That is, $\{u_1, u_2, \ldots u_s\} \in C_i$ have the same signature if their common neighborhood $\{s_1, s_2, \ldots s_t\} \in S$ is valued by all of them at $\{v_1, v_2, \ldots v_t\}$ such that $v_i \in B$. Since the degree of every vertex in I is at most k, this gives us at most $|B|^k$ many signatures for every class. All vertices of the same signature in a class are said to be of the same type. In aggregate, we have at most $2^k \cdot |B|^k$ many types of vertices in I.

For each vertex u in a type T, there are 2^d possible orientations of the edges incident on u, where d is the degree of u in G. Note that $d \leq k$, so there are at most 2^k such orientations. We say that an orientation is 'good' for the vertex u if it orients at least one of the highest-valued edges of u towards it. We denote the set of good orientations as O.

Towards formulating the ILP, for every type T and a good orientation o, we create the variables $x(T, o)$, which denote the number of vertices in the type T that are oriented according to the orientation o. Note that these are $f(k) = (2^k \cdot |B|^k) \cdot 2^k = 4^k \cdot |B|^k$ many variables.

We first describe the constraints to ensure the envy-freeness of the vertices in the independent set I. Let n_T be the number of vertices in the type T. Any vertex is EF if and only if it gets her highest valued edge oriented towards it (Corollary 2.3). Therefore, if the vertex u of type T ends up in a good orientation, it is EF. To ensure this, we add the constraints as described in Equation (2.1). Note that LHS of Equation (2.1) equals n_T only when every vertex in the type T is oriented according to some good orientation o. Indeed, if any vertex fails to end up in a good orientation, then it is not counted in the sum and hence RHS is strictly less than n_T, which then violates the constraint.

Now consider a vertex i in S. Let $v(i, T, o)$ denote the utility that an agent $i \in S$ gets when a vertex in type T is oriented according to the orientation o. Note that for a fixed orientation, $v(i, T, o)$ is a constant. If $x(T, o)$ many vertices in type T are oriented according to o, then the utility that agent i derives from the edge items across S and I from type T is precisely

$x(T,o) \cdot v(i,T,o)$. To capture the utility that i gets from edges in $E(S)$, we can do a brute-force search on which edges are allocated to i (since there are at most $\binom{k}{2}$ edges in $E(S)$). To that end, we create binary variables x_{ie} which take value 1 if the edge $e \in E(S)$ is allocated to i, otherwise 0. These are at most $g(k) = k \cdot k^2$ many variables. And, the utility that i derives from $E(S)$ is precisely $\sum_{e \in E(S)} v_i(e) x_{ie}$. Equation (2.2) ensures that every edge in S is allocated to at most 1 agent in S, while Equation (2.3) ensures that every edge in $E(S)$ is allocated. Finally, for i to be EF, it must get at least v_i^{max} utility under any allocation. This is captured by the constraints in Equation (2.4).

$$\sum_{o \in O} x(T,o) = n_T \quad \forall\, T \tag{2.1}$$

$$\sum_{i \in S} x_{ie} = 1 \quad \forall\, e \in E(S) \tag{2.2}$$

$$\sum_{i \in S} \sum_{e \in E(S)} x_{ie} = |E(S)| \tag{2.3}$$

$$\sum_{e \in E(S)} v_i(e) x_{ie} + \sum_{T,o} x(T,o) \cdot v(i,T,o) \geq v_i^{max} \quad \forall\, i \in S \tag{2.4}$$

$$x(T,o) \geq 0 \quad \forall\, T \;\&\; o \in O \tag{2.5}$$

$$x_{ie} \in \{0,1\} \quad \forall\, i \in S \;\&\; e \in E(S) \tag{2.6}$$

In aggregate, the number of variables created is $f(k) + g(k)$.

We now argue the correctness of the ILP. Let O be the orientation that corresponds to the values that $x(T,o)$ takes in some feasible solution of the ILP. For every vertex in the independent set I, Equation (2.1) ensures that it ends up in a good orientation, and therefore gets one of her highest valued edges oriented towards itself under O. This ensures the envy-freeness of vertices in I. The envy-freeness of vertices in S is ensured in Equation (2.4) via a brute-force search of an orientation that gives every vertex in S her highest valued edge. Therefore, a feasible solution to the ILP corresponds to an envy-free allocation of the original instance.

Conversely, suppose there is an EF allocation Φ in the original instance. Then, by Theorem 2.1, there is an EF orientation O. Let O_I and O_S be the restrictions of O for vertices in I and S respectively. Since O is EF, we have that both O_I and O_S are good orientations. This implies that there exist orientations under which every vertex ends up being in a good orientation. Hence, the constraints 2.1 and 2.4, which loop over all the good orientations, are satisfied when the variables $x(T, o)$ and x_{ie} correspond to the orientations O_I and O_S respectively. This implies that the ILP is feasible and this settles our claim. □

2.3 EFX and Welfare-Maximization

In this section, we discuss the price of EFX on graphical instances. Every graph may not admit an EFX orientation but does admit an EFX allocation, so it must be the case that some welfare is lost in the process of achieving EFX. We quantify this loss with respect to Utilitarian welfare.

Theorem 2.9. *For $\{0, 1\}$-graphical instances, a non-wasteful EFX allocation always exists and can be found in polynomial time. Therefore, the Price of EFX with respect to Utilitarian welfare is 1.*

Proof. Consider an instance $\mathcal{I}$ of $\{0, 1\}$-graphical allocations. For all the asymmetric edges $e = (ij)$, we orient them towards the incident agent who values e at 1, say i. We call i a special vertex, since i is now envy-free in any completion of this partial allocation and is also not envied by anyone else. Once we orient all the asymmetric edges, we remove them from the graph. The edges that are valued at 0 by both end-points can be allocated to the non-envied agents arbitrarily at the end of the algorithm, so for now, we remove them from the graph and consider a collection of connected subgraphs $H = \{H_1, H_2, \ldots H_k\}$ such that all edges in H are symmetric and valued at 1 by both the end-points. For each $H_i \in H$, we consider the following cases:

1. H_i is a Tree. Then, we claim that there is always a non-wasteful EFX allocation where there is at most one envious agent. Since H_i is a tree, we have $V(H_i) = E(H_i) + 1$. By pigeonholing, at least one agent, say i, does not receive any edge item from $E(H_i)$. Suppose there is a special agent i, then we root H_i on i and construct an orientation such that every vertex gets an edge item from her parent. This way, everyone except i receives a utility of at least 1 from the edges in H_i and hence is EF in any complete orientation. Also, i is EF since it is a special vertex. This gives us a non-wasteful EF (hence EFX) allocation such that every agent gets a utility of at least 1.

   Suppose there is no special vertex in the tree H_i. Then there is no complete EF allocation.

To find an EFX allocation, we root H_i on any vertex, say the least degree vertex i, and construct an orientation such that every vertex gets an edge item from her parent. Note that i leaves empty-handed and is envious of her neighbors in H_i. Since every envied agent (precisely, the children of i in the tree H_i) gets exactly one edge item (precisely the edge from i), the allocation is EFX.

2. H_i contains a cycle, say $C = \{u_1, u_2, \ldots u_c, u_1\}$. This case is the same as Case 2 in the proof of Theorem 2.5 and therefore a non-wasteful EF (hence EFX allocation exists) such that every agent receives a utility of at least 1.

Therefore, we get a non-wasteful EFX allocation Φ. For $\{0, 1\}$-valuations, a non-wasteful allocation is also utilitarian optimal, and hence Φ is also utilitarian optimal. Therefore, the price of EFX is 1 in this case. □

Theorem 2.10. *The price of EFX with respect to Utilitarian welfare is unbounded even for $\{0, 1, d\}$-graphical valuations.*

Proof. We construct an instance where the price of fairness is a function of the highest degree of a vertex in the graph. Consider a star graph G rooted at the vertex r which is incident to d many leaf vertices. The root vertex r values each of the d incident edges at d. All the leaf vertices value their incident edge at 1. A utilitarian welfare maximizing allocation gives all the edges to r, generating a welfare of d^2. Clearly, this allocation is not EFX since the envied agent r has multiple items and every leaf agent violates EFX. Under any EFX allocation, r can not receive more than 1 item, otherwise, the corresponding leaf vertex whose incident edge is allocated to r, violates EFX. Therefore, the maximum welfare under an EFX allocation is $d + (d - 1)$, where one d-valued edge is allocated to r and the rest all $(d - 1)$ edges are allocated to their corresponding leaf vertices, valued at 1 by each of them. Therefore, $PoF_{UM} = \frac{d^2}{d+(d-1)} > \frac{d^2}{2d} \approx d$. This implies that welfare loss can be as high as possible, and hence PoF is unbounded. □

Theorem 2.11. *Given an instance of graphical valuations, deciding the existence of a utilitarian welfare-maximizing and EFX allocation (UM+EFX) is NP-Hard.*

Proof. We present a reduction from Multi-Colored Independent Set (MCIS), where given a regular graph $G = (V_1 \uplus \cdots \uplus V_k, E)$ with degree d, the problem is to decide if there exists a subset $S \subseteq V(G)$ such that $G[S]$ is an independent set and $|V_i \cap S| = 1$ for all $i \in [k]$. We construct the graphical instance as follows. All vertices in $V(G)$ correspond to agents and all edges in $E(G)$ to items. Every agent $v \in V(G)$ values her incident edges at 1. That is, all edges

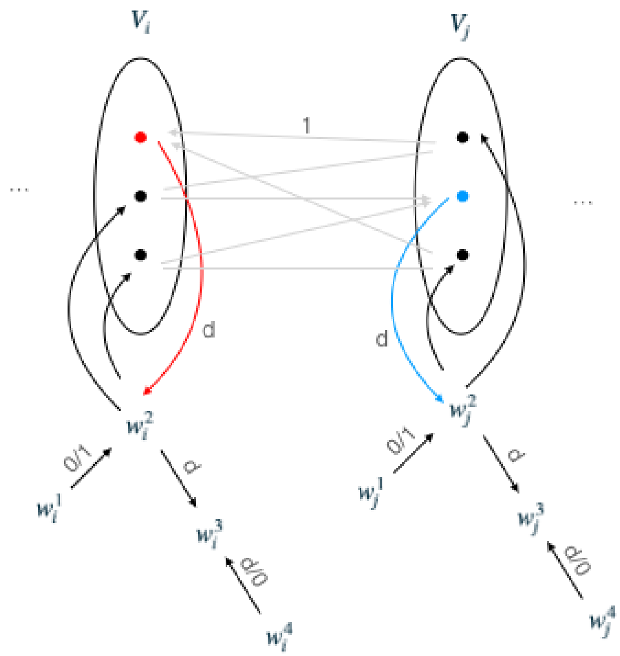

Figure 2.2: A schematic of reduced instance in the proof of Theorem 2.11

in G are symmetric with a weight of 1. For every vertex partition V_i, we add a path of three edges and four vertex-agents $\{w_i^1, w_i^2, w_i^3, w_i^4\}$ such that w_2^i is adjacent to all the vertices in V_i. All edges from w_i^2 to V_i are valued symmetrically at d by both end-points. The edge (w_i^1, w_i^2) is valued at 0 by w_i^1 and at 1 by w_i^2. The edge (w_i^2, w_i^3) is valued at d by both w_i^2 and w_i^3. Finally, the edge (w_i^3, w_i^4) is valued at d by w_i^3 and at 0 by w_i^4. This completes the construction. A schematic of this construction is shown in Figure 2.2. We now argue the equivalence.

Forward Direction.

Suppose MCIS is a Yes-instance and there is an independent set $S = \{s_1, s_2, \dots s_k\} \subseteq V(G)$ such that $|V_i \cap S| = 1$. Then, we do the following orientation of $E(G)$ to get an allocation that is welfare-maximizing and EFX.

- $\{(s_i, w_i^2)\}$ are oriented towards w_i $\forall\, i \in [k]$.
- $\{(v, w_i) : v \in V_i \setminus \{s_i\}\}$ are oriented towards v $\forall\, i \in [k]$.
- $\{(s_i, v) : v \in N(s_i) \setminus \{w_i\}\}$ are oriented towards s_i $\forall\, i \in [k]$.
- $\{w_i^1, w_i^2\}$ are oriented towards w_i^2 $\forall\, i \in [k]$.
- $\{w_i^2, w_i^3\}$ & $\{w_i^3, w_i^4\}$ are oriented towards w_i^3 $\forall\, i \in [k]$.

Let Φ be the allocation corresponding to the above orientation. Then, by construction, every edge is allocated to an agent who values it the most. Therefore, Φ is a (utilitarian) welfare-maximizing allocation. We now argue that Φ also satisfies EFX. The agents w_i^1 and w_i^4 do not value any item, so even though they are empty-handed under Φ, they do not envy any agent. All the agents in V_i except s_i get a utility of d each and they value every other bundle at most d, hence are envy-free. Likewise, w_i^2 is envy-free as it gets a utility of $d+1$ and values every other bundle no more than d. Also, w_i^3 gets all the edges it values, so there is no envy on her part. Lastly, each s_i gets d of her incident edges valued at 1 each, deriving a value of d, and hence they are envy-free. This implies that Φ is EF and hence, EFX.

Reverse Direction.

Suppose there is a welfare-maximizing allocation Φ which also satisfies EFX. Then because Φ maximizes welfare, it must satisfy the following partial allocation: w_i^3 must receive both her incident edges $\{w_i^2, w_i^3\}$ & $\{w_i^3, w_i^4\}$ as it values them highly at d, and $\{w_i^1, w_i^2\}$ must be allocated to w_i^2 as a utilitarian welfare maximizing allocation is also non-wasteful. This forces w_i^2 to be envious of w_i^3 even after one item is removed from the envied bundle. Therefore, w_i^2 must receive at least one item that it values at d incident to the partition V_i. This in turn forces at least one vertex from V_i to violate EFX with respect to w_i^2, hence it must receive at least d utility from the remaining items. This is feasible only when it is allocated all her d incident edges. Since this is true for at least one vertex in all V_i such that $i \in [k]$, it must be the case that all these k vertices form an independent set in G. This implies that MCIS is a yes-instance. This concludes the argument. □

We now present a polynomial-time reduction from deciding the existence of a welfare-maximizing and EFX allocation (UM+EFX) to finding a welfare-maximizing allocation within the set of EFX allocations (UM/EFX). Let $w^\star$ be the maximum utilitarian welfare ($w^\star$ can be computed in linear time by giving each item to an agent who values it the most). Now suppose the latter problem can be solved in polynomial time. Then, let w be the maximum welfare within EFX allocations. If $w = w^\star$, we have a "yes" instance of UM+EFX; else if $w \neq w^\star$, we have a "no" instance. Therefore, we get the following result.

Corollary 2.12. *Given an instance of graphical valuations, finding a utilitarian welfare maximizing allocation within the set of EFX allocations (UM/EFX) is NP-Hard.*

We now discuss the complexity and the loss in the egalitarian welfare due to the EFX constraint. The egalitarian welfare of an allocation Φ is defined as the minimum utility of any agent under

Φ. We say that an allocation is Egalitarian Maximal (EM) if it maximizes the minimum utility of any agent. In the following result, we exhibit the hardness of finding an egalitarian maximizing allocation within the set of EFX allocations (EM/EFX).

Theorem 2.13. *Given an instance of graphical valuations, finding an egalitarian welfare maximizing allocation within the set of EFX allocations (EM/EFX) is NP-Hard.*

Proof. We show that given a welfare threshold d, deciding the existence of an EFX allocation with egalitarian welfare at least d is NP-hard. To that end, we exhibit a reduction from Multi-Colored Independent Set (MCIS). The construction is similar as in the proof of Theorem 2.6.

Forward Direction.

Suppose MCIS is a Yes-instance and there is an independent set $S = \{s_1, s_2, \ldots s_k\} \subseteq V(G)$ such that $G[S]$ is an independent set and $|V_i \cap S| = 1$. Then, we do the following orientation of $E(G)$ to get an EF allocation.

- $\{(s_i, w_i) : i \in [k]\}$ are oriented towards w_i.
- $\{(v, w_i) : v \in V_i \setminus \{s_i\}\}$ are oriented towards v.
- $\{(s_i, v) : v \in N(s_i) \setminus \{w_i\}\}$ are oriented towards s_i.
- All the remaining edges are oriented arbitrarily.

Let Φ be the allocation corresponding to the above orientation. Then, for all agents a, we have $v_a(\Phi_a) = d$. By Lemma 2.2, Φ is an EF (hence, EFX) allocation. Therefore, there is an EFX allocation with egalitarian welfare of at least d.

Reverse Direction.

Suppose there is EFX allocation with egalitarian welfare at least d. Then, every w_i must receive at least one of her incident edges to derive a utility of at least d. This implies that there is at least one vertex v in every V_i does not get the edge (v, w_i). In order to get a utility of at least d, each of such vertices v must get all of her d incident edges in G. This is feasible if and only if they form an independent set in G. Therefore, if there is EFX allocation with egalitarian welfare at least d, then MCIS is a yes-instance and this settles the claim. □

Corollary 2.14. *For $\{0, 1\}$-graphical instances, the Price of EFX with respect to Egalitarian welfare is* 1.

Proof. If the optimal egalitarian welfare is 0, then there is nothing to prove. Otherwise, consider the case when the optimal egalitarian welfare is $k > 0$ and is achieved by the allocation Φ. We have $v_i(\Phi_i) \geq k \; \forall \; i \in [n]$. If Φ is a wasteful allocation, then consider the re-allocation of all wastefully allocated edges to one of her incident vertices, arbitrarily. Since this re-allocation does not bring down the utility of any agent, it remains egalitarian optimal but now corresponds to a non-wasteful allocation and hence, an orientation. By Lemma 2.2, we get that that this orientation is EF, since $v_i(\Phi_i) = k \geq 1 = v_i^{max} \; \forall \; i \in [n]$. Therefore, we get an egalitarian optimal EF (hence, EFX) allocation and this settles our claim. ☐

Nash welfare is defined as the geometric mean of the utilities of the agents. We show below that for binary graphical valuations, there is no loss in the Nash welfare as well.

Corollary 2.15. *For $\{0, 1\}$-graphical instances, the Price of EFX with respect to Nash welfare is 1.*

Proof. If the optimal Nash welfare is 0, then there is nothing to prove. Otherwise, consider the case when the optimal Nash welfare is $k > 0$ and is achieved by the allocation Φ. Then, we must have $v_i(\Phi_i) \geq 1 \; \forall \; i \in [n]$. Since every Nash optimal allocation is non-wasteful therefore, Φ is a non-wasteful allocation and hence, an orientation. By Lemma 2.2, we get that that this orientation is EF, since $v_i(\Phi_i) \geq 1 = v_i^{max} \; \forall \; i \in [n]$. Therefore, we get a Nash optimal EF (hence, EFX) allocation and this settles our claim. ☐

Since Nash optimal allocations can be found in polynomial time for general binary additive valuations (Darmann and Schauer, 2015), we can find an EFX and Nash welfare maximizing allocation in polynomial time.

We say that an allocation is a leximin allocation if, among all allocations, it lexicographically maximizes the utility profile, that is, maximizes the minimum utility, subject to that maximizes the second minimum, and so on. Clearly, leximin allocations are also egalitarian maximal allocations. Moreover, for general binary additive valuations, the set of leximin and Nash optimal allocations coincide Halpern et al. (2020), and hence a Nash optimal allocation is also maximizes egalitarian welfare. Since Nash optimal allocations are also non-wasteful, they also maximize the utilitarian welfare. This gives us the following result.

Corollary 2.16. *For $\{0, 1\}$-graphical instances, an EFX allocation that maximizes Utilitarian, Egalitarian, and Nash welfare always exists and can be found in polynomial time.*

2.4 Concluding Remarks.

We studied the complexity of finding envy-free allocations for graphical valuation and quantified the loss of welfare in the process of achieving approximate (i.e, EFX) envy-freeness, which was the original motivation for the study of the class of graphical valuations. We believe there are several directions of interest for future work that build on our preliminary line of inquiry here. For instance, for parameterized results, one could consider structural parameters that are smaller than vertex cover. Extending the PoF discussion beyond binary setting for other welfare notions would also be of interest. Finally, one might also generalize the class of graphical valuations in many ways. One generalization is to allow graphs with multiedges which then corresponds to instances where an item is liked by at most two agents but a pair of agents together can derive positive value from more than one item. The restricted setting of hypergraphs where every edge corresponds to multiple but same number of vertices is also an interesting direction to pursue.

Chapter 3

Fair and Efficient Allocations for Degenerate Instances

"Mathematics leaves no room for argument. If you made a mistake, that was all there was to it."

- Malcolm X, *The Autobiography of Malcolm X*

3.1 Introduction

The non-existence of envy-free allocations has led to several notions of "workarounds": approximate envy-freeness (e.g, requiring allocations to be envy-free up to the removal of one good (Budish, 2011; Lipton et al., 2004), or any good (Caragiannis et al., 2019b), or using hidden goods (Hosseini et al., 2020)), subsidy (introducing money to compensate for envy (Brustle et al., 2020)), donating items (this involves giving up on completeness, but to a limited extent (Chaudhury et al., 2020b)), and sharing (wherein we allow for some goods to be shared between agents (Bismuth et al., 2024; Sandomirskiy and Segal-Halevi, 2019)).

In this chapter, our focus is on the settings where introducing money is not feasible and hypothetical removal or actual donation of goods is not desirable, which typically turns out to be the case when several high-valued goods are involved. In such scenarios, sharing goods between agents appears to be the most reasonable of all workarounds, and the question of interest is to find allocations that meet our goals of fairness and efficiency with minimum sharing. We recall that finding complete envy-free allocations is hard already for the special case of no sharing even with just two agents with identical valuations: indeed, the problem is easily seen to be equivalent to Partition, a weakly NP-hard problem which asks if a set of numbers can be split into two parts of equal sum.

In a recent development, Sandomirskiy and Segal-Halevi (2019) show that there is a sense in which the case of identical valuations are in fact the "hardest" — they propose a notion of degeneracy which captures the degree of similarity across agent valuations and argue that the intractable cases are those that have a rather high degree of similarity. In retrospect, one might argue that similar valuations signal high conflict, and this possibly contributes to making this a hard scenario. We informally describe the notion of degeneracy here and refer the reader to Section 3.2 for the formal definitions. We say that a set of goods are valued similarly by two agents if the ratios of their values for all goods are the same (for example, two goods valued at 10 and 500 by agent A and at 20 and 1000 by agent B would be considered similar). The degree of similarity between two agents is one less than the largest number of goods that are valued similarly by them. The degeneracy of an instance with n agents is the highest degree of similarity across all pairs of agents. In particular, the degeneracy of an instance with identical valuations is $m - 1$, and this is one extreme example. On the other end of the spectrum, the degeneracy can be as small as zero, when all agents view all goods differently (more precisely, no pair of goods is valued similarly by any two agents).

Informally speaking, we refer to the setting of low degeneracy, the ones where agent valuations over goods are generally dissimilar, as a scenario involving *amicable agents*. Unlike the case of identical valuations, we expect such valuations to invoke relatively "less conflict". One of the key results in Sandomirskiy and Segal-Halevi (2019) is that while finding EF allocations remains hard even with two amicable agents, finding allocations that are both fPO and EF is tractable for a constant number of amicable agents. In particular, the time to compute such allocations was shown to be $O(3^{\frac{n(n-1)}{2}d}m^{\frac{n(n-1)}{2}+2})$, where d is the degeneracy of the valuation matrix. We note that this running time is $3^{O(d)}m^{O(1)}$ for a constant number of agents, and $O(m^{O(1)})$[1] for a constant number of agents and valuations of degeneracy at most $O(\log m)$. In contrast, it was shown that the problem remains NP-hard for instances that have high degeneracy (informally, those that are closer to having the structure of

[1]We remark that this is a strongly polynomial running time.

instances with identical valuations) — specifically, if d is allowed to grow as a polynomial function of m, then it is hard to check if there exists a fPO and EF allocation with zero sharings even for instances with two agents.

Our Contributions.

The results of Sandomirskiy and Segal-Halevi (2019) nicely illustrate the influence of degeneracy on the complexity of finding fPO and EF allocations. Building on this line of work, we investigate the complexity from the perspective of the number of agents. For example, can this running time be improved to $(n+m)^{O(d)}$, which would increase the realm of tractability to scenarios with any number of agents and constant degeneracy, or more ambitiously, $O(2^{O(d)} \cdot (m+n)^{O(1)})$, which would make the problem tractable for instances with any number of agents and degeneracy logarithmic in $(n+m)$? Our main contribution here is to show that even the former goal is unlikely to be achievable: when the number of agents is unbounded, the problem of finding allocations that are fPO and EF remains *strongly* NP-complete for instances with degeneracy one, even for the specific question of allocations with no sharings.

Our result also has consequences for the problem of finding EF allocations. We recall that the problem of finding EF allocations is weakly NP-complete by a reduction from Partition (Sandomirskiy and Segal-Halevi, 2019). It turns out that the arguments in the reverse direction of our reduction do not require the allocation in question to be fPO. Since the valuation matrix of our reduced instance happens to only have values that are bounded by a polynomial function of n and m, we obtain a stronger hardness result for the problem of finding complete EF allocations for instances with constant degeneracy.

We also revisit the algorithm for finding fPO+EF allocations from Sandomirskiy and Segal-Halevi (2019). The algorithm relies on enumerating certain *consumption graphs* corresponding to fPO allocations that fix the sharing structure of a potential solution, after which the task of determining the exact proportions of sharing while respecting fairness constraints is outsourced to an ILP formulation. It is shown (Sandomirskiy and Segal-Halevi, 2019, Lemma 2.5) that there always exists a fPO allocation with at most $(n-1)$ sharings. We propose an alternate method for generating the relevant consumption graphs that takes advantage of the upper bound on the number of sharings upfront. This leads to a slightly different bound that leads to a better exponential term at the cost of a worse polynomial factor. Although the difference in the bound is not significant, we believe our approach lends additional understanding to the structure of class of graphs based on fPO allocations.

3.2 Preliminaries

Allocations and Sharing.

For standard fair allocation terminologies, we refer to Section 1.2. We elaborate on the ones we use in this chapter as follows.

A *bundle* of objects is a vector $\mathbf{b} = (b_j)_{j \in [m]} \in [0,1]^m$, where the component b_j represents the portion of g_j in the bundle. The total amount of each object is normalized to one. An *allocation* Φ, as previously, is a collection of bundles $(\Phi_i)_{i \in [n]}$, one for each agent, with the condition that all the objects are fully allocated. Note that an allocation can be identified with the matrix $\Phi := (\Phi_{i,j})_{i \in [n], j \in [m]}$ such that:

$$\text{all } \Phi_{i,j} \geq 0 \text{ and } \sum_{i \in [n]} \Phi_{i,j} = 1 \text{ for each } j \in [m].$$

Let $j \in [m]$ be arbitrary but fixed. If for some $i \in [n]$, $\Phi_{i,j} = 1$, then the object g_j is not shared – it is fully allocated to agent a_i. Otherwise, object g_j is *shared* between two or more agents. There are two natural measures quantifying the amount of sharing in an allocation Φ:

- The *number of shared objects* is given by the number of items that are shared:

$$\#s^{\dagger}(\Phi) = \left| \left\{ j \in [m] \,:\, \Phi_{i,j} \in (0,1) \text{ for some } i \in [n] \right\} \right|.$$

- The *total number of sharings* accounts for the number of times that an object is shared, i.e:

$$\#s^{\star}(\Phi) = \sum_{j \in [m]} \left(\left| \{ i \in [n] : \Phi_{i,j} > 0 \} \right| - 1 \right).$$

For allocations with no shared objects, both measures are zero, but they can differ by as much as $n - 2$ in general. Note that the number of sharings is at least the number of shared objects, since each shared object is shared at least once by definition. Unless mentioned otherwise, our measure for "extent of sharing" in the computational questions that we will shortly define will be the notion of the total number of sharings.

Value and Utility.

For every $i \in [n], j \in [m]$, $v_{i,j}$ denotes agent a_i's *value* for the entire object g_j. In the setting of *additive* utilities, the valuations naturally lead us to an utility function over bundles defined as follows:

$$u_i(\mathbf{b}) = \sum_{j \in [m]} v_{i,j} \cdot b_j.$$

The matrix $\mathbf{v} = (v_{i,j})_{i \in [n], j \in [m]}$ is called the *valuation matrix*; it encodes the information about the preferences of agents and is used as the input of fair division algorithms. We use $v^\star$ to denote the largest value in a valuation matrix v. We say that a class of inputs $\mathcal{C}$ has *bounded valuations* if there exists a polynomial $p(n, m)$ such that $v^\star \leq p(n, m)$ for all instances in $\mathcal{C}$.

We recall the notion of degeneracy that was proposed in (Sandomirskiy and Segal-Halevi, 2019). To this end, we say that two goods g_p, g_q are valued *similarly* by a pair of agents i, j if there exists a constant r such that:

$$v_{i,p} \cdot v_{j,q} = v_{i,q} \cdot v_{j,p} = r.$$

If all valuations in consideration are non-zero, then this is equivalent to the requirement that:

$$\frac{v_{i,p}}{v_{j,p}} = \frac{v_{i,q}}{v_{j,q}} = r.$$

Note that any collection of goods valued identically by a pair of agents would be pairwise similar with respect to the agents in question, but this definition generalizes the notion of "identical" to, roughly speaking, "identical up to a scaling factor".

Now, let us define the *similarity* between a pair of agents i and j as:

$$s_{\mathbf{v}}(i, j) = \max_{r > 0} \left| \left\{ k \in [m] \,:\, v_{i,k} = r \cdot v_{j,k} \right\} \right| - 1.$$

Note that the similarity of a pair of agents captures the notion of the largest number of goods that the agents value similarly when considered pairwise. This finally leads us the the notion of *degeneracy*, which is defined as:

$$d(\mathbf{v}) = \max_{i,j \in [n], i \neq j} s_{\mathbf{v}}(i,j).$$

Valuations for which $d(\mathbf{v}) = 0$ are called *non-degenerate*. Also, note that if any two agents have the same valuations for all goods, then $d(\mathbf{v}) = m - 1$.

Efficiency.

We now turn to notions of efficiency that will be relevant to our discussion. An allocation Φ is Pareto-dominated by an allocation $\mathbf{y}$ if $\mathbf{y}$ gives at least the same utility to all agents and strictly more to at least one of them. An allocation Φ is *fractionally Pareto-optimal (fPO)* if no feasible $\mathbf{y}$ dominates it. The following lemma provides a complete characterization of fPO allocations.

Lemma 3.1 (Sandomirskiy and Segal-Halevi (2019), Lemma 2.3)**.** *An allocation Φ is fractionally Pareto Optimal if and only if there exists a vector of weights $\lambda = (\lambda_i)_{i \in [n]}$ with $\lambda_i > 0$, such that for all agents $i \in [n]$ and goods $p \in [m]$, if $\Phi_{i,p} > 0$ then for any agent $j \in [n]$,*

$$\lambda_i \cdot v_{i,p} \geq \lambda_j \cdot v_{j,p}$$

Also, we mention here a related and weaker notion of efficiency: an allocation Φ is *discrete Pareto-optimal* if it is not dominated by any feasible $\mathbf{y}$ with $y_{i,j} \in \{0,1\}$.

Computational Questions.

We conclude this section with the computational questions that we address in this chapter. Our main focus is on the problem of determining a fair and efficient allocation that minimizes sharing.

Formally, for a fairness concept $\alpha \in \{\text{EF, EQ, Prop}\}$ and an efficiency concept $\beta \in \{\text{fPO, dPO}\}$, the (α, β)-Minimal Sharing problem is the following. Given $(\mathcal{A}, \mathcal{G}, \mathbf{v}, t \in \mathrm{N})$ as input, the question is if there exists an α, β allocation where the total number of sharings is at most t.

3.3 EF+fPO Allocations and Degeneracy

To prove Theorem 3.2, we will show a reduction from a structured version of Satisfiability problem called Linear Near-Exact Satisfiability (LNES) which is known to be NP-complete Dayal and Misra (2019). An instance of LNES consists of $5p$ clauses (where

$p \in \mathbb{N}$) denoted as follows:

$$\mathcal{C} = \{U_1, V_1, U_1', V_1', \cdots, U_p, V_p, U_p', V_p'\} \cup \{C_1, \cdots, C_p\}.$$

We will refer to the first $4p$ clauses as the *core* clauses, and the remaining clauses as the *auxiliary* clauses. The set of variables consists of p *main variables* $x_1, \ldots, x_p$ and $4p$ *shadow variables* $y_1, \ldots, y_{4p}$. Each core clause consists of two literals and has the following structure:

$$\forall\, i \in [p], U_i \cap V_i = \{x_i\} \text{ and } U_i' \cap V_i' = \{\bar{x}_i\}.$$

Each main variable x_i occurs exactly twice as a positive literal and exactly twice as a negative literal. The main variables only occur in the core clauses. Each shadow variable makes two appearances: as a positive literal in an auxiliary clause and as a negative literal in a core clause. Each auxiliary clause consists of four literals, each corresponding to a positive occurrence of a shadow variable. We will use u_i, v_i, u_i', and v_i' to refer to the shadow variables in the main clauses U_i, V_i, U_i', and V_i', respectively.

The LNES problem asks whether, given a set of clauses with the aforementioned structure, there exists an assignment τ of truth values to the variables such that *exactly one* literal in every core clause and *exactly two* literals in every auxiliary clause evaluate to TRUE under τ. The main result of this section is the following, and is established by a reduction from LNES.

Theorem 3.2. *(EF,fPO)-Minimal Sharing is NP-hard even when restricted to inputs with bounded valuations, degeneracy one, and no sharing.*

Proof. We reduce from LNES. Let

$$\mathcal{C} = \{U_1, V_1, U_1', V_1', \cdots, U_p, V_p, U_p', V_p'\} \cup \{C_1, \cdots, C_p\}.$$

be an instance of LNES as described above.

We begin with a description of the construction of the reduced instance. For each main variable x_i we introduce three agents: $\{a_i, \bar{a}_i, d_i\}$, and the goods $\{g_i, \bar{g}_i, h_i\}$. We refer to d_i as the *dummy* agent associated with x_i and a_i and $\bar{a}_i$ as the *key* agents associated with x_i. Also, we refer to h_i as the *trigger* good and g_i and $\bar{g}_i$ as *consolation* goods.

For the shadow variables u_i, v_i, u_i', v_i', we introduce four agents: b_i, c_i, b_i', c_i' which we simply refer to as *shadow agents* and four goods: r_i, s_i, r_i', s_i', which we refer to as the *essential* goods.

Finally, for each auxiliary clause C_j, we introduce the good f_j^1 and f_j^2. These goods are called *backup* goods.

Note that our instance consists of $n = 7p$ agents and $m = 9p$ goods. Thus the size of the valuation matrix is $N := 63 \cdot p^2$. We let $L = 4000 \cdot p^5$. We will use $\mathcal{A}$ and $\mathcal{G}$ to refer to the set of agents and goods that we have defined here.

Let $\mathbf{w} = (w_{i,j})_{i \in [n], j \in [m]}$ denote the $(7p \times 9p)$ matrix whose entries are given by $w_{i,j} = (i-1) \cdot m + j$. Intuitively, we can think of these values as being small enough to be negligible, and we will obtain our final valuation matrix by starting from $\mathbf{w}$ and "overwriting" some entries to reflect the fact that certain goods are valued highly by certain agents. This is done to ensure that the final valuation matrix has low degeneracy. We now describe the specific modifications that we have to make to $\mathbf{w}$.

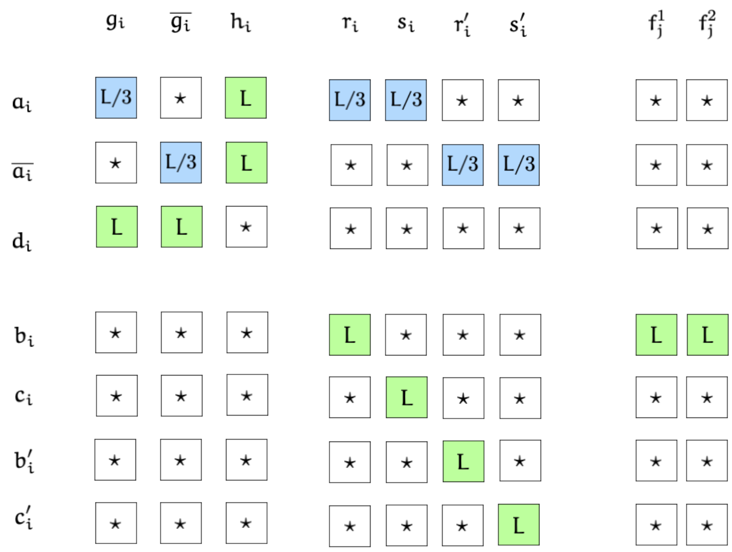

Figure 3.1: The overall schematic of the construction in the proof of Theorem 1. The entries depicted by a $\star$ indicate small values. In this example, the literal corresponding to the agent b_i, i.e., u_i, belongs to the auxiliary clause C_j corresponding to the backup goods f_j^1 and f_j^2.

To this end, let us define another set of values given by $\mathbf{w}^\star = (w^\star_{i,j})_{i \in [n], j \in [m]}$. Let $\pi : \mathcal{A} \to [n]$ and $\sigma : \mathcal{G} \to [m]$ be arbitrary but fixed orderings of the agents and goods, respectively.

- For $i \in [p]$, we have that the dummy agent corresponding to the main variable x_i has a high value for the consolation goods g_i and $\bar{g}_i$.

$$w^\star_{\pi(d_i),j} = \begin{cases} L & \text{if } \sigma^{-1}(j) \in \{g_i, \bar{g}_i\}, \\ 0 & \text{otherwise.} \end{cases}$$

- For $i \in [p]$, we have that the first key agent corresponding to the main variable x_i has a somewhat high value for the consolation good g_i and the essential goods r_i and s_i, and a high value for the trigger good h_i.

$$w^\star_{\pi(a_i),j} = \begin{cases} L/3 & \text{if } \sigma^{-1}(j) \in \{g_i, r_i, s_i\}, \\ L & \text{if } \sigma^{-1}(j) = h_i, \\ 0 & \text{otherwise.} \end{cases}$$

- For $i \in [p]$, we have that the second key agent corresponding to the main variable x_i has a somewhat high value for the consolation good $\bar{g}_i$ and the essential goods r'_i and s'_i, and also has a high value for the trigger good h_i.

$$w^\star_{\pi(\bar{a}_i),j} = \begin{cases} L/3 & \text{if } \sigma^{-1}(j) \in \{\bar{g}_i, r'_i, s'_i\}, \\ L & \text{if } \sigma^{-1}(j) = h_i, \\ 0 & \text{otherwise.} \end{cases}$$

- For $i \in [p]$ the shadow agents have a high value for their associated essential goods and the backup good which represents an auxiliary clause that contains the shadow variable associated with the shadow agent. Formally, we have:

$$w^\star_{\pi(b_i),j} = \begin{cases} L & \text{if } \sigma^{-1}(j) \in \{r_i, f^1_\ell, f^2_\ell\}, \\ 0 & \text{otherwise.} \end{cases}$$

  where ℓ is such that C_ℓ is the unique clause that contains the shadow variable u_i. The valuations for $w_{\pi(c_i),j}$, $w_{\pi(b'_i),j}$ and $w_{\pi(c'_i),j}$ are analogously defined, with r_i being replaced by s_i, r'_i, and s'_i, respectively, and ℓ would be such that C_ℓ is the unique clause that contains v_i, u'_i, and v'_i, respectively.

The final valuations that we will work with are obtained by taking a point-wise max of the two valuation matrices defined above with the following exceptions:

- Dummy agents continue to value the four essential goods associated with them at zero.

- The shadow agent b_i (respectively, c_i) values continues to value the consolation good g_i and the essential good s_i (respectively, r_i) at zero.
- The shadow agent b_i' (respectively, c_i') values continues to value the consolation good $\bar{g}_i$ and the essential good s_i' (respectively, r_i') at zero.

In particular, we propose the final valuation matrix $\mathbf{v} = (v_{i,j})_{i \in [n], j \in [m]}$ as follows:

$$v_{i,j} = \begin{cases} \min(w_{i,j}, w^{\star}_{i,j}) & \text{if } \pi^{-1}(i) = d_k \text{ and } \sigma^{-1}(j) \in \{r_k, s_k, r_k', s_k'\}, \\ & \text{or } \pi^{-1}(i) = b_k \text{ and } \sigma^{-1}(j) \in \{g_k, s_k\}, \\ & \text{or } \pi^{-1}(i) = c_k \text{ and } \sigma^{-1}(j) \in \{g_k, r_k\}, \\ & \text{or } \pi^{-1}(i) = b_k' \text{ and } \sigma^{-1}(j) \in \{\bar{g}_k, s_k'\}, \\ & \text{or } \pi^{-1}(i) = c_k' \text{ and } \sigma^{-1}(j) \in \{\bar{g}_k, r_k'\}, \\ & \text{for any } k \in [p] \\ \max(w_{i,j}, w^{\star}_{i,j}) & \text{otherwise.} \end{cases}$$

For convenience, we say an entry of $\mathbf{v}$ is *large* if it is at least $L/3$ and is *small* otherwise. For (i, j) which are such that $v_{i,j}$ is small, we introduce the notation $\varepsilon_{i,j}$ to denote $v_{i,j}$.

We ask if this instance admits an allocation with zero sharing. This completes the description of the construction. From here, we first argue the equivalence of the instances and then turn to demonstrating that the instance has degeneracy one. Note that the valuation matrix is clearly bounded.

Forward Direction.

Let τ be a boolean assignment for the variables of the LNES instance that we start with. Based on this, we will now propose an allocation $\Phi := (\Phi_{i,j})_{i \in [n], j \in [m]}$. For the formal expression of the allocation, we introduce a set of *compatible indices* I which is initialized to ϕ and will be developed further below. For all $i \in p$, we have the following:

- If $\tau(x_i) = 1$, then the first key agent a_i gets $\{g_i, r_i, s_i\}$, the second key agent $\bar{a}_i$ gets the trigger good h_i and the dummy agent d_i gets the consolation good $\{\bar{g}_i\}$. Based on this, we let:

$$I = I \cup \{(\pi(a_i), \sigma(g_i)), (\pi(a_i), \sigma(r_i)), (\pi(a_i), \sigma(s_i)),$$

$$(\pi(\bar{a}_i), \sigma(h_i)), (\pi(d_i), \sigma(\bar{g}_i))\}.$$

- If $\tau(x_i) = 0$, then the first key agent a_i gets the trigger good $\{h_i\}$, the second key agent $\bar{a}_i$ gets $\{\bar{g}_i, r_i', s_i'\}$ and the dummy agent gets the consolation good $\{g_i\}$. Based on this, we let:

$$I = I \cup \{(\pi(\bar{a}_i), \sigma(\bar{g}_i)), (\pi(\bar{a}_i), \sigma(r_i')), (\pi(\bar{a}_i), \sigma(s_i')),$$
$$(\pi(a_i), \sigma(h_i)), (\pi(d_i), \sigma(g_i))\}.$$

- If $\tau(x_i) = 1$, then the shadow agents b_i' and c_i' get the essential goods that they value highly, i.e, r_i' and s_i'. Based on this, we let:

$$I = I \cup \{(\pi(b_i'), \sigma(r_i')), (\pi(c_i'), \sigma(s_i'))\}.$$

- If $\tau(x_i) = 0$, then the shadow agents b_i and c_i get the essential goods that they value highly, i.e, r_i and s_i. Based on this, we let:

$$I = I \cup \{(\pi(b_i), \sigma(r_i)), (\pi(c_i), \sigma(s_i))\}.$$

Note that there are $2p$ shadow agents who have not been allocated any goods so far. It is easy to check that these shadow agents correspond exactly to shadow variables x for which $\tau(x) = 1$. Since τ is a satisfying assignment for the LNES instance, we know that each auxiliary clause C_ℓ contains exactly two shadow variables that which evaluate to true under τ. Let $\mu(C_\ell)$ denote the shadow agents corresponding to these shadow variables. Then, for each $j \in [p]$, the goods f_j^1 and f_j^2 are allocated arbitrarily, one each, to the two shadow agents in $\mu(C_j)$. In particular, if $\mu(C_j) = \{t_j^1, t_j^2\}$ we let:

$$I = I \cup \left\{(\pi(t_j^1), \sigma(f_j^1)) \mid j \in [p]\right\} \cup \left\{(\pi(t_j^2), \sigma(f_j^2)) \mid j \in [p]\right\}$$

Thus, our final allocation is given by:

$$\Phi_{i,j} = \begin{cases} 1 & \text{if } (i,j) \in I, \\ 0 & \text{otherwise.} \end{cases}$$

We now argue that the allocation has the desired efficiency and fairness properties. We first argue that the allocation is fPO.

Claim 3.3. *The allocation Φ defined above is an fPO allocation.*

Proof. By Lemma 1, the allocation will be fPO if there exists a vector of weights $\lambda = (\lambda_i)_{i \in [n]}$ with $\lambda_i > 0$, such that for all agents $i \in [n]$ and goods $p \in [m]$, if $\Phi_{i,p} > 0$ then for any agent $j \in [n]$,

$$\lambda_i \cdot v_{i,p} \geq \lambda_j \cdot v_{j,p}$$

We propose the following weight vector:

$$\lambda_i = \begin{cases} 3 & \text{if } \pi^{-1}(i) \text{ is a key agent,} \\ 1 & \text{otherwise.} \end{cases}$$

Consider a key agent a_i. Suppose a_i gets the trigger good, that is, $\Phi_{\pi(a_i),\sigma(h_i)} = 1$. In this case, recall that h_i is the only good allocated to the agent a_i, in other words, for any $j \neq \sigma(h_i)$, we have $\Phi_{\pi(a_i),j} = 0$. Then, since any other agent $j \in [n]$ values h_i at L or less, we have that for any agent j who is not a key agent:

$$\lambda_{\pi(a_i)} \cdot v_{\pi(a_i),\sigma(h_i)} = 3 \times L \geq \varepsilon_{j,\sigma(h_i)} = \lambda_j \cdot v_{j,\sigma(h_i)}.$$

Further, for the "partner" key agent, we have:

$$\lambda_{\pi(a_i)} \cdot v_{\pi(a_i),\sigma(h_i)} = 3 \times L \geq 3 \times L = \lambda_{\pi(\bar{a}_i)} \cdot v_{\pi(\bar{a}_i),\sigma(h_i)},$$

and for any other key agent j, we have:

$$\lambda_{\pi(a_i)} \cdot v_{\pi(a_i),\sigma(h_i)} = 3 \times L \geq 3 \times \varepsilon_{j,\sigma(h_i)} = \lambda_j \cdot v_{j,\sigma(h_i)}.$$

On the other hand, suppose a_i gets the bundle with a consolation good and two essential goods, that is, suppose:

$$\Phi_{\pi(a_i),\sigma(g_i)} = 1, \Phi_{\pi(a_i),\sigma(r_i)} = 1, \text{ and } \Phi_{\pi(a_i),\sigma(s_i)} = 1.$$

As before, recall that in this case, the only goods allocated to a_i are $\{g_i, r_i, s_i\}$. Note that the dummy agent d_i is the *only* other agent who has a large value for g_i and likewise, the shadow agents b_i and c_i are the *only* other agents who have a large value for the goods r_i and s_i, respectively. Since these values are all at most L, we have, for any $x \in \{\pi(d_i), \pi(b_i), \pi(c_i)\}$ and $y \in \{\sigma(g_i), \sigma(r_i), \sigma(s_i)\}$:

$$\lambda_{\pi(a_i)} \cdot v_{\pi(a_i),y} = 3 \times \frac{L}{3} = L \geq 1 \times L \geq \lambda_x \cdot v_{x,y},$$

and for any other agent j not accounted for above and any $y \in \{\sigma(g_i), \sigma(r_i), \sigma(s_i)\}$, we have:

$$\lambda_{\pi(a_i)} \cdot v_{\pi(a_i),y} = 3 \times \frac{L}{3} = L \geq \max(1,3) \times \varepsilon_{j,y} = \lambda_x \cdot v_{x,y}.$$

The case of the key agent $\bar{a}_i$ is symmetric to the discussion above. Now, consider the allocation of goods for all the other agents:

- Every dummy agent gets a consolation good, which is valued at L by the dummy agent, at $L/3$ by a key agent, and at a small number by any other agent.
- Every shadow agent either gets an essential good (which is valued at L by the shadow agent, at $L/3$ by a key agent, and at a small number by any other agent), or a backup good (which is valued at $2L$ by the shadow agent and three other shadow agents, and at a small number by any other agent).

Since the weights λ_i are one for any $i \in [n]$ who is not a key agent, it suffices to demonstrate that these agents have been allocated a good that they value at least as much as any other agent who is not a key agent, and at least three times as much the value ascribed by the key agent. However, note that this is evidently true from the description above.

Therefore, by Lemma 1, the allocation is indeed fPO. □

Next, we argue that Φ is an envy-free allocation.

Claim 3.4. *The allocation Φ defined above is an EF allocation.*

Proof. Recall that the instance has $2p$ key agents, p dummy agents, and $4p$ shadow agents. Note that everyone values their own bundle at L, which is a large value. We argue that the allocation is EF by case analysis.

Between key agents.

Recall that, depending on whether $\tau(x_i) = 1$ or 0, either $\Phi_{(\pi(\bar{a}_i),\sigma(h_i))} = 1$ or $\Phi_{(\pi(a_i),\sigma(h_i))} = 1$, that is, one of the two key agents a_i or $\bar{a}_i$ gets the trigger good h_i. The bundle of the agent who gets h_i, is valued at L by the other key agent. But, she also values her own bundle (which consists of two essential goods and one consolation good valued at $L/3$ each) also at L, so there is no envy between them. Also, note that any two key agents a_i and a_j, corresponding to two different main variables x_i and x_j do not envy each other, as the utility that they derive from their own bundle is large, while their valuation of the other bundle would be small.

Between dummy agents.

Any two dummy agents d_i and d_j get one of the consolation goods $\{g_i, \bar{g}_i\}$ and $\{g_j, \bar{g}_j\}$ respectively, which they highly value. Since:

$$u_{d_i}(b_{d_i}) = L \gg \varepsilon_{\pi(d_i),\star} = u_{d_i}(b_{d_j}),$$

where $\star \in \{g_j, \bar{g}_j\}$, and:

$$u_{d_j}(b_{d_j}) = L \gg \varepsilon_{\pi(d_j),\star} = u_{d_j}(b_{d_i}),$$

where $\star \in \{g_i, \bar{g}_i\}$, we conclude that they have little value for each others' bundle, and so there is no envy.

Between shadow agents.

The pair of shadow agents getting any essential good do not envy each other because of low utility for each others' bundle and high utility for their own. The pair of shadow agents getting backup goods value their partner's bundle either at L or at a small number. In any case, there is no envy. Finally between a shadow agent i getting an essential good, and j getting a backup good, note that the utility of i for j's bundle is at most L, and the utility of j for i's bundle is small, and therefore there is no envy.

Between key agents and dummy agents.

The dummy agent gets one of the consolation goods. Any key agent value these goods either at $\frac{L}{3}$ or at a small value. In any case, d_i is not envied. On the other hand, d_i values the bundle of the key agents either at L, which is equal to his value for his own bundle (when the key agent a_i gets a bundle of two essential goods and a consolation good — recall that the dummy agent has a zero value for these specific essential goods) or at a small value (when the key agent a_i gets the trigger good or if the agents under consideration are either a_j or $\bar{a}_j$ for some $j \neq i$). Therefore, the key agents are also not envied by d_i.

Between key agents and shadow agents.

We claim that any shadow agent values the key bundles allocated to key agents at at most L. Indeed, they have negligible value for trigger goods. Further, we claim that they have a total value of at most L for bundles comprising of three goods. Indeed, there are only two possible scenarios: either all goods have a negligible value for the shadow agent in question or the bundle contains an essential good that the shadow agent values at L, in which case the value that the agent would have for the other two goods is zero (by construction). Since all shadow agents derive a utility L from their own bundle, these agents are not envious of any key agents.

On the other hand, the key agents clearly do not envy the shadow agents who get the backup goods, as they have negligible value for all the backup goods. As for shadow agents whose bundle contains essential goods, recall that a key agent values any essential good at at most $L/3$, and her own bundle at L, so there is no envy between a key agent and shadow agents whose bundles comprise of essential goods.

Between shadow agents and dummy agents.

All the dummy agents get a subset of the consolation goods, which carry negligible value for the shadow agents. Likewise, all the shadow agents get a subset of the essential goods and the backup goods, both of little value to the dummy agents. So they do not envy each other. □

Reverse Direction.

For the discussion in the reverse direction, we say that an allocation is *valid* if it is EF and fPO and involves no sharing. Let $\Phi := (\Phi_{i,j})_{i \in [n], j \in [m]}$ be a valid allocation. First, we argue that Φ must have a certain structure.

We first claim that in the allocation Φ, any trigger good h_i must be allocated to one of the corresponding key agents $\{a_i, \bar{a}_i\}$.

Claim 3.5. *If Φ is a valid allocation, then:*

$$\max(\Phi_{\pi(a_i),\sigma(h_i)}, \Phi_{\pi(\bar{a}_i),\sigma(h_i)}) = 1,$$

Proof. Suppose that the trigger good h_i is not allocated to either a_i or $\bar{a}_i$, who are the only two agents who value it highly[2]. To compensate for the envy generated against the recipient of h_i, both of these agents must be given[3] their associated bundle of consolation and essential goods, i.e, a_i gets $\{g_i, r_i, s_i\}$ and $\bar{a}_i$ gets $\{\bar{g}_i, r'_i, s'_i\}$. This makes the dummy agent d_i envious of both a_i and $\bar{a}_i$. Now there is no way to satisfy d_i, since the goods valued by her, g_i and $\bar{g}_i$, are already taken. This contradicts the fact that the allocation Φ is EF. □

Next, we claim that every consolation good g_i is allocated to either to the key agent a_i or to the dummy agent d_i. Likewise, the good $\bar{g}_i$ is allocated to either to the key agent $\bar{a}_i$ or to the dummy agent d_i.

Claim 3.6. *If Φ is a valid allocation, then:*

$$\max(\Phi_{\pi(a_i),\sigma(g_i)}, \Phi_{\pi(d_i),\sigma(g_i)}) = 1 \textit{ and } \max(\Phi_{\pi(\bar{a}_i),\sigma(\bar{g}_i)}, \Phi_{\pi(d_i),\sigma(\bar{g}_i)}) = 1.$$

Proof. We know from the previous claim that h_i is allocated to one of the key agents, say a_i. To cater to now envious agent $\bar{a}_i$, she must be given the corresponding bundle $\{\bar{g}_i, r'_i, s'_i\}$. Now, if the dummy agent d_i's bundle does not contain the remaining consolation good g_i, d_i will envy $\bar{a}_i$, since all the other available goods carry a negligible value for d_i. Therefore, d_i's bundle must contain the good g_i. The argument is analogous when h_i is allocated to $\bar{a}_i$, and the claim follows. □

We now show that the consolation good g_i is allocated to a key agent a_i, then the shadow agents b_i and c_i must be allocated backup goods.

[2]Notice that if all other agents had a zero value for the trigger good then it would be straightforward to obtain a Pareto improvement here. However, recall that the other agents carry a small but non-zero value for the trigger good. Our argument, therefore, relies on the fact that $\mathbf{v}$ is EF.

[3]Note that the choice of L ensures that even the sum of utilities of all goods that carry a small utility would be significantly lower than $L/3$, and we use this fact implicitly in several arguments to say that certain allocations are forced.

Claim 3.7. *Let Φ be a valid allocation where $\Phi_{\pi(a_i),\sigma(g_i)} = 1$. Let C_x be the auxiliary clause containing u_i, and C_y be the auxiliary clause containing v_i. Then, there exists $\star \in \{1,2\}$ and $\dagger \in \{1,2\}$ such that:*

$$\Phi_{\pi(b_i),\sigma(f_x^\star)} = 1 \textit{ and } \Phi_{\pi(c_i),\sigma(f_y^\dagger)} = 1,$$

Proof. If the key agent a_i's bundle contains g_i, then from the proof of the previous claim we know that her entire bundle consists of the goods $\{g_i, r_i, s_i\}$. Clearly, the shadow agents b_i and c_i now envy a_i. If b_i is not allocated one of f_x^1 or f_x^2 then b_i derives a utility of less than L from her bundle, and therefore continues to envy a_i. The argument for the claim that c_i's bundle contains one of f_y^1 or f_y^2 is analogous, and the claim follows. □

By symmetry, we also have that if the consolation good $\bar{g}_i$ is allocated to a key agent $\bar{a}_i$, then the shadow agents b_i' and c_i' must be allocated backup goods. We state the following without proof since the argument is analogous.

Claim 3.8. *Let Φ be a valid allocation where $\Phi_{\pi(\bar{a}_i),\sigma(\bar{g}_i)} = 1$. Let C_x be the auxiliary clause containing u_i', and C_y be the auxiliary clause containing v_i'. Then, there exists $\star \in \{1,2\}$ and $\dagger \in \{1,2\}$ such that:*

$$\Phi_{\pi(b_i'),\sigma(f_x^\star)} = 1 \textit{ and } \Phi_{\pi(c_i'),\sigma(f_y^\dagger)} = 1.$$

Now observe that the two claims above account for the allocation of $2p$ backup goods among $2p$ distinct shadow agents. Let us call these shadow agents *happy* and the remaining shadow agents *unhappy*. We claim that the bundle of every unhappy shadow agent must contain an essential good — this is because these are the only highly valued goods left in the pool and are the only way to eliminate the envy that the unhappy agents feel for the happy ones. Note that every unhappy agent values the bundle of exactly two happy shadow agents.

Claim 3.9. *Let Φ be a valid allocation. If $\Phi_{\pi(\bar{a}_i),\sigma(\bar{g}_i)}$, then:*

$$\Phi_{\pi(b_i),\sigma(r_i)} = 1 \textit{ and } \Phi_{\pi(c_i),\sigma(s_i)} = 1.$$

On the other hand, if $\Phi_{\pi(a_i),\sigma(g_i)} = 1$:

$$\Phi_{\pi(b_i'),\sigma(r_i')} = 1 \textit{ and } \Phi_{\pi(c_i'),\sigma(s_i')} = 1.$$

Proof. This follows from the fact that in the first case, b_i and c_i are unhappy agents, and in the second case, b'_i and c'_i are unhappy agents. For example, let C_z be the auxiliary clause that contains u_i. If the bundle of b_i does not contain the essential good r_i, then b_i envies the happy agents who have received the goods f_z^1 and f_z^2. Observe that there is no fix for this envy since all other goods carry negligible value for b_i. The argument is analogous for the remaining three cases. □

Based on this, we propose the following assignment of truth values to the variables of the LNES instance:

$$\tau(x_i) = \begin{cases} 1 & \Phi_{\pi(a_i),\sigma(g_i)} = 1, \\ 0 & \text{otherwise.} \end{cases}$$

We extend this assignment to shadow variables in the natural way: if $\tau(x_i) = 1$, then $\tau(u_i) = \tau(v_i) = 1$ and $\tau(u'_i) = \tau(v'_i) = 0$, while if $\tau(x_i) = 0$, then $\tau(u_i) = \tau(v_i) = 0$ and $\tau(u'_i) = \tau(v'_i) = 1$. We now argue that τ is a satisfying assignment for the original LNES instance.

Suppose g_i is allocated to a_i, that is $\Phi_{\pi(a_i),\sigma(g_i)} = 1$. We set $\tau(x_i) = 1$. This satisfies all the clauses containing the literal x_i, namely, U_i and V_i. Further, note that these clauses are satisfied exactly once, since we also set $\tau(u_i) = \tau(v_i) = 1$ (recall that u_i and v_i appear in these clauses with negative polarity). The other main clauses U'_i and V'_i are satisfied since we set $\tau(u'_i) = \tau(v'_i) = 0$, and these clauses are satisfied exactly once as well, since x_i appears in them with a negative polarity and we are in the case when $\tau(x_i) = 1$. The case when $\tau(x_i) = 0$ is analogous, and we see that all core clauses are satisfied exactly once by τ, as desired.

We now turn to the auxiliary clauses. Observe that $\tau(x_i) = 1$ if and only if $\Phi_{\pi(a_i),\sigma(g_i)} = 1$, that is, the key agent a_i gets the consolation good g_i. This implies that b_i and c_i are happy agents. On the other hand, recall that we also set $\tau(u_i)$ and $\tau(v_i)$ to one. Similarly, it can be argued that if $\tau(x_i) = 0$, then b'_i and c'_i are happy agents, and in this case, we had also set $\tau(u'_i)$ and $\tau(v'_i)$ to one. So we conclude that all happy agents correspond to variables that evaluate to one under τ. Along similar lines, it is easy to check that all unhappy agents who receive essential goods as explained in the last claim correspond to variables that are set to zero under τ.

Now consider an auxiliary clause C_ℓ. Notice that f_ℓ^1 and f_ℓ^2 have been allocated to happy agents that value these goods highly, so we know that C_ℓ contains at least two variables that evaluates to true. Now suppose there is some auxiliary clause that contains more than two variables that evaluate to true. This would imply the existence of more than $2p$ happy agents, which is a contradiction.

Constant Degeneracy.

We now argue that the valuation matrix $\mathbf{v}$ of the reduced instance has degeneracy one. We proceed by sketching the possible configurations of numbers that arise in 2×2 submatrices of the valuation matrix in a manner that draws attention to the case which witness a similarity value of two. A more tedious case analysis that accounts for every pair of agents and every pair of goods, but we omit that version of the argument for the sake of brevity. We first make some observations:

- Note that for shadow variables that appear together in a single auxiliary clause, we have two backup goods that are valued identically by the corresponding shadow agents. It is easy to verify that there is no other pair of goods that these agents value similarly.
- Between a shadow agent b_i and a dummy agent d_i, there is one good, namely s_i, that is zero-valued for both these agents. It is also easy to check that there is no pair of goods not involving s_i that these agents value similarly. This claim is analogous for agents c_i, b_i', and c_i' considered along with d_i.
- Let $\varepsilon_{i,j}$, $\varepsilon_{i',j}$, $\varepsilon_{i,j'}$, $\varepsilon_{i',j'}$ denote four small values from the matrix $\mathbf{w}$. It is straightforward to verify, by the structure of $\mathbf{w}$, that $\varepsilon_{i,j} \times \varepsilon_{i',j'} \neq \varepsilon_{i',j} \times \varepsilon_{i,j'}$.

It is easy to check that backup goods (f_j^1, f_j^2) constitutes a *maximal* pair of similar goods for the pairs of shadow agents corresponding to shadow variables that appear in C_j. Also, between d_i and b_i (respectively, c_i), we have that s_i (respectively, r_i) along with any other good constitutes a maximal pair of similar goods for these agents. Similarly, between d_i and b_i' (respectively, c_i'), we have that s_i' (respectively, r_i') along with any other good constitutes a maximal pair of similar goods for these agents.

Now, we claim that other than the scenarios considered above, any pair of agents do not value more than one good similarly. It is straightforward to verify that once the cases above are excluded, for any agents $i, j \in [n]$ and goods $p, q \in [m]$, it holds that at least one of the entries among $\mathbf{v}_{i,p}, \mathbf{v}_{i,q}, \mathbf{v}_{j,p}$, and $\mathbf{v}_{j,q}$ is small and non-zero. We first deal with the case when all four values are positive. Without loss of generality, let's say that $\mathbf{v}_{i,p} = \varepsilon_{i,p}$ is small. If this is the only small value then the term $v_{i,q} \cdot v_{j,p}$ dominates. If there are two small values, then they are either distributed on either side of the cross product and it is easy to see[4] that since the small numbers are distinct then the cross products are not equal, and in the other case again the term $v_{i,q} \cdot v_{j,p}$ dominates. If there are three small values then one term will clearly dominate the other. Finally, if there are four small values then the claim follows from the properties of small numbers as stated above.

It remains to discuss the case when some of the values can be zero, when considered to the exclusion of the cases that have already been accounted for. In particular, note that we may assume that no good is zero-valued for two agents, since the only situation in which this happens is if we have a dummy agent and a shadow agent, a case that has already been considered above. If there is exactly one zero value then it is easy to see that the cross product in question is not equal. On the other hand, the only way to have two zero-valued goods is if one agent values a pair of goods at zero each. Then, the other agent must have non-zero values for both goods, and while the cross products are equal to zero, such goods are not considered similar according to our definition. Finally, it is not possible for us to have three zero-valued goods since this will force a situation where one good is zero-valued by two agents, which has been ruled out already. This concludes our argument.

□

3.4 Consumption Graphs

In this section, we discuss bounds on the number of consumption graphs that can be associated with a valuation $\mathbf{v}$. For an allocation Φ, the allocation graph captures the structure of how the goods are shared among agents. More formally, for any allocation Φ, the *consumption graph* (CG_z) associated with it is a bipartite graph with agents in one part and goods in another and there is an edge between an agent i and a good j if and only if i gets a share of j under Φ, that is, $\Phi_{i,j} > 0$. These graphs are useful in the context of algorithms for the (α, fPO)-Minimal Sharing problem for the following reason. If we can afford to "guess" the consumption graph of the valuation that we are looking for, then it is possible to express the fairness constraints usually as a ILP — indeed, the consumption graph already tells us how the goods are shared between agents, and what remains to be addressed is the actual proportions, which can be captured by variables which can then be subjected to whatever fairness constraints we are

[4]We remark that there is one edge case here for the situation when the large values in question are $L/3$ and L and the small values are such that one is three times the other, but it is easy to check that even if these cases arise, they contribute a maximal pair of similar goods and do not affect the degeneracy argument. It is also possible to avoid the situation altogether by a careful ordering of the agents, but we omit the details here.

interested in. It turns out that the number of variables that are required is bounded by $2(n-1)$, so this gives us an efficient algorithm when the number of agents is fixed.

Therefore, for a valuation $\mathbf{v}$, it is useful to bound the number of fPO allocations with distinct consumption graphs as a function of its degeneracy and n, m. This will control the expense of the guess that we make before delegating the task of finding a specific allocation to an ILP. We mention that this is the overall approach followed by Sandomirskiy and Segal-Halevi (2019), and in this section, we only focus on the bound on the number of consumption graphs that need to be enumerated: any improvement here automatically improves the running time of the algorithm for computing EF+fPO allocations. While we obtain a slightly different bound (which is tighter than the one in Sandomirskiy and Segal-Halevi (2019) under some circumstances), in our discussion here our goal is to emphasize a slightly different way of looking at the set of consumption graphs: our approach is a direct brute-force enumeration of all relevant structures, as opposed to being recursive.

To this end, we first introduce some terminology. An object g_k is called a *good* for agent a_i if $v_{i,k} > 0$, and is called a *bad* if $v_{i,k} < 0$. Further, following the definitions of Sandomirskiy and Segal-Halevi (2019), we say that an object is a *good* if it is a good for at least one agent, and is a *bad* if it is a bad for every agent. It is also called a *pure good* if it is a good for every agent. Further, we say that an object g_k is *contentious* for agents a_i and a_j if $v_{i,k} \cdot v_{j,k} > 0$, i.e., both agents agree whether g_k is a good or a bad. On the other hand, we say that an object is *non-controversial* for agents a_i and a_j if $v_{i,k} \cdot v_{j,k} < 0$ and *useless* if $v_{i,k} = v_{j,k} = 0$. Finally, for an arbitrary but fixed "threshold" $t_{i,j} > 0$, we say that:

- *a_i values g_k strongly relative to the threshold $t_{i,j}$* if $\frac{|v_{i,k}|}{|v_{j,k}|} > t_{i,j}$,
- *a_i is ambivalent about g_k relative to the threshold $t_{i,j}$* if $\frac{|v_{i,k}|}{|v_{j,k}|} = t_{i,j}$,
- *a_i values g_k weakly relative to the threshold $t_{i,j}$* if $\frac{|v_{i,k}|}{|v_{j,k}|} < t_{i,j}$.

Building on Lemma 3.1, we recall the following useful thresholding property characterizing fPO allocations developed in Sandomirskiy and Segal-Halevi (2019). The following lemma basically says that if there is a fPO allocation, then there is a choice of thresholds $t_{i,j}$ for every pair of agents a_i and a_j which is such that:

- among the contentious objects g_k which are goods, if a_i values g_k strongly relative to $t_{i,j}$ then no part of g_k is allocated to a_j, while if a_i values g_k weakly relative to $t_{i,j}$ then no part of g_k is allocated to a_i, and
- among the contentious objects g_k which are bads, if a_i values g_k strongly relative to $t_{i,j}$

then no part of g_k is allocated to a_i, while if a_i values g_k weakly relative to $t_{i,j}$ then no part of g_k is allocated to a_j.

This is formalized below.

Lemma 3.10 (Sandomirskiy and Segal-Halevi (2019), Corollary 2.4). *For a fractionally Pareto-optimal allocation Φ and any pair of agents $a_i \neq a_j$, there is a threshold $t_{i,j} > 0$ ($t_{i,j} = \frac{\lambda_j}{\lambda_i}$ from Lemma 3.1) such that for any object g_k:*

- *if g_k is contentious for a_i and a_j (i.e, $v_{i,k} \cdot v_{j,k} > 0$), then:*
  - *if a_i values g_k strongly relative to $t_{i,j}$, i.e, $\frac{|v_{i,k}|}{|v_{j,k}|} > t_{i,j}$, we have $\Phi_{j,k} = 0$ in case of a good and $\Phi_{i,k} = 0$ in case of a bad,*
  - *if a_i values g_k weakly relative to $t_{i,j}$, i.e,* for *for $\frac{|v_{i,k}|}{|v_{j,k}|} < t_{i,j}$, we have $\Phi_{i,k} = 0$ in case of a good and $\Phi_{j,k} = 0$ in case of a bad,*
- *if g_k is non-controversial for a_i and a_j (i.e, $v_{i,k} \cdot v_{j,k} < 0$), then an agent with negative value cannot consume k.*

In particular, a_i and a_j can share only objects g_k that are useless or that both agents are ambivalent about, that is, $\frac{v_{i,k}}{v_{j,k}} = t_{i,j}$.

Based on this, we will now obtain the following alternate version of Proposition 3.8 in Sandomirskiy and Segal-Halevi (2019) for the number of consumption graphs associated with fPO allocations that have at most $(n-1)$ sharings.

Proposition 3.11. *For every fixed number of agents $n \geq 2$, the number of all fPO graphs*

$$CG(v) := \{CG(\Phi) \mid \Phi \text{ is fPO for } \mathbf{v} \text{ with at most } n-1 \text{ sharings}\}$$

satisfies the upper bound:

$$\left|CG(v)\right| \leq 2^{(1+d(\mathbf{v}))\frac{n(n-1)}{2}} \cdot m^{\frac{n(n-1)}{2}+cn}, \tag{3.1}$$

where c is a constant.

For comparison, we recall that the bound obtained in Sandomirskiy and Segal-Halevi (2019) was:

$$3^{(1+d(\mathbf{v}))\frac{n(n-1)}{2}} \cdot m^{\frac{n(n-1)}{2}},$$

and our bound can be viewed as a minor improvement in some scenarios (in particular, for any fixed n, we obtain an improvement in the exponential term at a cost of introducing a larger polynomial factor). The approach used in Sandomirskiy and Segal-Halevi (2019) started with building all possible consumption graphs for instances with two agents and then followed an iterative process to build the collection of consumption graphs for a larger number of agents. We describe a process that generates all the graphs for instances with n agents directly instead, which allows us to take advantage of the fact that the number of sharings is small.

We now turn to a description of our process.

Proof of Proposition 3.11. Our task here is to enumerate all consumption graphs associated with fPO allocations Φ. Recalling the property of fPO allocations from Lemma 3.10, we first guess the thresholds $t_{i,j}$ for all pairs[5] of agents (a_i, a_j) with $i < j$. Note that it suffices to consider threshold values that emerge from ratios in the valuation matrix, thus for any fixed pair of agents we may reasonably restrict our attention to at most m possible values of possible thresholds. Therefore, this guess requires us to examine at most $m^{\binom{n}{2}}$ possibilities.

Since we are only interested in allocations with at most $(n-1)$ sharings, we guess the set of edges in the consumption graph that are incident to shared goods. If there are s sharings, then it is easy to see that there are at most $2s$ such edges, and we have $\binom{mn}{2s} \leq m^{O(n)}$ choices.

Now, let $\tau = (t_{i,j})_{1 \leq i < j \leq n}$ be an arbitrary but fixed guess of thresholds. We now describe a procedure for generating all possible graphs G for which there exists an allocation Φ compatible with the thresholds τ as specified by Lemma 3.10 and whose consumption graph is given by G. The vertex set of our graphs will always be given by $\{w_i \mid i \in [n]\} \cup \{u_k \mid k \in [m]\}$. For each u_k, we will maintain a set F_k, which denotes the set of agents which do not share the object g_k, and a set E_k, which denotes the set of agents which do share the object g_k. To begin with, $F_k = \phi$ for all $k \in [m]$, and E_k is populated based on the guessed sharings in the first step above. Now, for every pair of agents a_i and a_j, we have the following:

- among the contentious objects g_k which are goods, if a_i values g_k strongly relative to $t_{i,j}$ then a_j is added to F_k, while if a_i values g_k weakly relative to $t_{i,j}$ then a_i is added to F_k, and
- among the contentious objects g_k which are bads, if a_i values g_k strongly relative to $t_{i,j}$ then a_i is added to F_k, while if a_i values g_k weakly relative to $t_{i,j}$ then a_j is added to F_k, and

[5] It is easy to check that the conclusions derived from Lemma 3.10 are identical when applied to (a_i, a_j) and (a_j, a_i), so it suffices to consider only ordered pairs.

- for any non-controversial object g_k, we add to F_k the agent who values g_k negatively.

Now, observe that the number of useless goods and the goods that the agents a_i and a_j are ambivalent about combined is at most $(1 + d(\mathbf{v}))$, by the definition of degeneracy. For these goods, we would like to try all possible combinations of allocations, namely: whether both agents get a share of the good under consideration, or whether the good is not shared by a_i or the good is not shared by a_j. Notice that we have $2^{\binom{n}{2}\cdot(1+d(\mathbf{v}))}$ possibilities to consider, because for a fixed pair of agents and a good, the upfront guess of the "shared edges" indicates if the good is to be shared between the said agents or not — if it is, then the choice is already predetermined, and if not, then we have only two possibilities to consider: denoting the good under consideration by g_k, in one scenario we add a_i to F_k and in the other we add a_j to F_k.

At the end of this process, goods for which $F_k \cup E_k = \mathcal{A}$, then the adjacencies of g_k in the consumption graph are fully determined. Now let us consider goods for which this is not the case. In particular, let i be the smallest index for which $a_i \notin F_k \cup E_k$. We claim that $a_j \in F_k$ for all $j \neq i$. Indeed, fix any $j \neq i$. Suppose, for the sake of contradiction, that $a_j \notin F_k$. If a_i and a_j are both ambivalent about g_k, or g_k is useless to both of them, then $a_j \notin F_k$ implies that either a_j and a_i are both in E_k or $a_i \in F_k$, both of which contradict the assumption that $a_i \notin F_k \cup E_k$. If g_k is non-controversial or contentious for a_i and a_j, then at least one of a_i or a_j belongs to F_k, which is again a contradiction. Therefore, $a_j \in F_k$ for all $j \neq i$. Since an allocation has to be complete, we know that it must be the case that g_k is allocated entirely to a_i now that all other agents are forbidden from consuming g_k according to our choices so far. So we add a_j to E_k in this situation.

It is straightforward to verify that the choices made so far completely determine the structure of a consumption graph, and that our guesses have accounted for all fPO allocations with at most $(n-1)$ sharings. That the number of possibilities is as desired follows from the fact that we have $m^{\binom{k}{2}}$ choices of thresholds, $m^{O(n)}$ choices for the shared edges, and $2^{\binom{k}{2}(1+d(\mathbf{v}))}$ choices for how goods that are ambivalent or useless for agent pairs are distributed between them. □

We note here that for additive valuations, EF1+fPO allocations always exist and an EF1+PO allocation can be computed in pseudo-polynomial time (Barman et al., 2018a). The proof relies on constructing Fisher markets along with an underlying integral equilibrium. In the process, bipartite graphs called Maximum bang per buck (MBB) graphs of a Fisher market instance are constructed and goods are exchanged only along the edges in the MBB graphs, which ensures that the allocation remains an equilibrium allocation. The first welfare theorem ensures that for a Fisher market with additive valuations, any equilibrium outcome is fPO (Mas-Colell et al., 1995), consequently, the modified allocation at every step remains fPO. Sandomirskiy and Segal-Halevi (2019) also characterizes fPO allocations and show that an allocation is fPO if and only if the associated consumption graphs satisfy certain properties. The interesting aspect is that the undirected consumption graphs of fPO allocations correspond to the maximal bang per buck (MBB) graphs of a Fisher market. Consequently, the algorithm for enumerating all consumption graphs can be used to find all MBB graphs.

3.5 The Case of Bounded Valuations

In this section, we obtain a bound on the size of the fair division instance in the case when valuation matrix has a certain structure. We argue that if the values that agents assign to the objects comprise of a small number of distinct non-zero values, then it can not be the case that we have arbitrarily large number of agents and objects when the valuation matrix also has bounded degeneracy. Formally, we have the following.

Proposition 3.12. *If the valuation matrix is such that all the entries are non-zero, number of distinct entries is at most t, and degeneracy $d(\mathbf{v}) < m-1$, then number of goods m and the number of agents n is bounded as a function of t and $d(\mathbf{v})$, and in particular, $m \leq d(\mathbf{v}) \cdot t^2$ and $n \leq t^{d(\mathbf{v}) \cdot t^2}$.*

Proof. Let $i \neq j \in [n]$ be arbitrary but fixed. Observe that:

$$\left| \left\{ \frac{v_{i,k}}{v_{j,k}} \;\middle|\; k \in [m] \right\} \right| \leq t^2$$

If $m > t^2 \cdot d(\mathbf{v})$, then by the pigeon-hole principle, there exist more than $d(\mathbf{v})$ choices of k for which the ratios $\frac{v_{i,k}}{v_{j,k}}$ are equal, which would contradict our assumption that the instance has degeneracy $d(\mathbf{v})$.

We say that any pair of agents i and j have same *type* if they value all the goods in the same way, that is, $v_{i,k} = v_{j,k}$ for all $k \in [m]$. Note that since $d(\mathbf{v}) < m-1$, no two agents i and j can have same type. Since the number of distinct entries is t, every agent has t many choices for assigning the value to any good $p \in [m]$. Every agent can assign t possible distinct values to each one of the m goods. This implies the number of types of agents is at most t^m, therefore, $n \leq t^{d(\mathbf{v}) \cdot t^2}$. □

We briefly justify the assumption that requires all entries to be non-zero. Indeed, even for binary values, i.e, when all entries in the valuation matrix are either 0 or 1, there may exist instances with small degeneracy and an arbitrary number of goods. For example, consider the example where one agent values all the goods at 1 and the other values all the goods at 0 except for one good which is valued at one. The degeneracy here is zero, and the number of goods can clearly be arbitrarily large. It would be intersting to explore if the bounds demonstrated by Proposition 3.12 are tight.

3.6 Concluding Remarks

We demonstrated the hardness of finding fPO+EF and EF allocations even for instances with constant degeneracy for instances with an unbounded number of agents. We note that running times of the form $d^{O(n)} \cdot poly(m, n)$ are "weakly ruled out" because of the hardness result in Sandomirskiy and Segal-Halevi (2019) which is based on a reduction from Partition. However, all the hardness results combined so far do not rule out the possibility of an algorithm with a running time of $c^{O(d+n)} \cdot m^{O(1)}$, which would imply strongly polynomial running times for instances where $(d + n)$ is bounded by $O(\log m)$. One framework to rule out such a possibility would be parameterized complexity, where one might attempt demonstrating W-hardness in the combined parameter (n, d).

Chapter 4

Minimizing Envy in House Allocation

I think perfect objectivity is an unrealistic goal; fairness, however, is not.

- Michael Pollan, *The Omnivore's Dilemma*

4.1 Introduction

The *house allocation problem* consists of n agents and m houses, where the agents have preferences over the houses, and we have to allocate the houses to the agents so that each agent receives exactly one house and each house is allocated to at most one agent.[1] The problem captures scenarios such as assigning clients to servers, employees to offices, families to government housing, and so on. Several variants of house allocation have been studied in the matching-under-preferences literature (Manlove, 2013; Hylland and Zeckhauser, 1979; Zhou, 1990), where the typical objectives have been economic efficiency requirements such as Pareto-optimality (Hylland and Zeckhauser, 1979; Abraham et al., 2004), rank maximality (Irving et al., 2006) or strategyproofness (Krysta et al., 2019). Another useful and desirable objective in any resource allocation setting is *fairness*, and we may equally well think of the house allocation problem as a special case of the fair division of indivisible goods setting, with

[1] Notice that these requirements immediately imply that we must have $m \geq n$, and under any allocation, exactly n houses are allocated and the remaining $m - n$ houses remain unallocated.

the additional requirement that each agent be allocated exactly one good. Given an allocation, say $\phi : A \rightarrow H$, an agent a envies a' if she values $\phi(a')$ more than $\phi(a)$. Finding "envy-free" allocations, i.e., ones where no agent envies another, is one of the main goals in fair division. Fairness in house allocation, in particular, has been a topic of interest in recent years (Gan et al., 2019; Kamiyama et al., 2021; Hosseini et al., 2023c, 2024b).

In the context of the house allocation problem, note that if $n = m$, an envy-free allocation exists if and only if there is a perfect matching in the following bipartite graph: introduce a vertex for every agent and every house, and let the vertex corresponding to an agent be adjacent to all the houses that she values no less than any other house. Since every house must be assigned when $n = m$, when an agent is assigned anything short of her best option, she will be envious. Therefore, the existence of an envy-free allocation can be determined efficiently in this situation using standard algorithms for checking if a perfect matching exists.

The question is less obvious when $n < m$, i.e., when there are more houses than agents. Indeed, one could work with the same bipartite graph, but it is possible for the house allocation instance to admit an envy-free allocation even though the bipartite graph does not have a perfect matching. Consider a situation with three houses and two agents, where both agents value one house above all else, and the other two equally. While the graph only captures the contention on the highly valued house, it does not lead us directly to the envy-free allocation that can be obtained by giving both agents the houses that they value relatively less (but equally). It turns out that the question of whether an envy-free allocation exists can be determined in polynomial time even when $n < m$, by an algorithm of Gan et al. (2019) that involves iteratively removing subsets of contentious houses.

When an envy-free allocation does not exist at all, the natural objective is to resort to a relaxation of the fairness objective. However, house allocation differs from the typical fair division setting, with an additional constraint that every agent receive exactly one item. This constraint renders the well-studied relaxed notions of fairness like envy-freeness up to 'some' good (where an agent chooses to hypothetically ignore one good from the envied bundle) futile. Hence we resort to a different kind of relaxation of envy-freeness: We *quantify* the envy involved in an allocation using different aggregate measures such as the number of envious agents, the maximum number of agents any agent envies, and look for allocations that minimize these "measures of envy." We note that Nguyen and Rothe (2013) and Shams et al. (2021) previously studied minimizing aggregate measures of envy in resource allocation problems. In the context of house allocation, Kamiyama et al. (2021) studied the problem of minimizing the number of envious agents. They showed that it is NP-complete to find allocations that minimize the number of envious agents, even for binary utilities, and this quantity is hard to approximate for general utilities. In this paper, we explore envy minimization in house allocation from a broader perspective and prove algorithmic results not only for minimizing the number of envious agents but for two other measures of envy as well—minimizing the amount of maximum envy experienced by any agent and minimizing the amount of total envy experienced by all the agents put together. We say that the *amount of envy* experienced by an agent a is the number of agents she is envious of.

We remark here that minimizing the number of envious agents may lead to a sub-optimal allocation in terms of maximum envy and vice-versa. For instance, consider an instance with 4 agents a, b, c, d and 4 houses h_1, h_2, h_3, h_4 with the following rankings:

$$\begin{aligned} a &: \mathbf{h_1} \succ h_2 \succ h_3 \succ h_4 \\ b &: h_1 \succ h_2 \succ h_3 \succ \mathbf{h_4} \\ c &: \mathbf{h_2} \succ h_3 \succ h_4 \succ h_1 \\ d &: \mathbf{h_3} \succ h_4 \succ h_1 \succ h_2 \end{aligned}$$

Consider the highlighted allocation, denoted as ϕ. Only agent b experiences envy under ϕ. This allocation effectively minimizes the number of envious agents. However, the envy experienced by the sole envious agent is substantial, as she envies all the other agents. Alternatively, the envy of agent b could have been reduced to just 1 under the allocation ϕ', where agent b receives house h_2 and agents c and d are allocated houses h_3 and h_4 respectively. While both ϕ and ϕ' are optimal in terms of minimizing overall total envy, ϕ falls short when it comes to addressing maximum envy, whereas ϕ' is not ideal for minimizing the number of agents who experience envy. This suggests that the three notions of envy minimization are not directly comparable in general and, therefore demand individual scrutiny and analysis.

When our focus is on minimizing the envy, there can be a trade-off with regards to *social welfare*, which is essentially the collective measure of individual agent utilities within any allocation. The method for aggregating these utilities can vary, including options such as the geometric mean (known as Nash), the summation of utilities, or the minimum utility of any agent, among others. We restrict our attention to the sum of the individual agent utilities as our measure of social welfare for this work. For this measure, the trade-off between welfare and envy-minimization is illustrated as follows. Consider an instance with $2n-1$ houses such that each of the n agents like only the first $n-1$ houses. Then the only envy-free allocation is to allocate the last n houses to everyone, resulting in zero social welfare. On the contrast, if we allow for envious agents, the above instance can achieve a social welfare of at least $n-1$. This potential loss in welfare due to fairness guarantees is captured by *price of* *fairness*, which is the worst-case ratio of the maximum social welfare in any allocation to that in a fair allocation. This notion was first proposed by Bertsimas et al. (2011), following which there has been substantial progress towards finding the bounds for the price for various combinations of fairness and welfare notions, specifically in resource allocation setting (Caragiannis et al., 2012; Bei et al., 2021; Barman et al., 2020b; Sun et al., 2023b; Bhaskar et al., 2023). To the best of our knowledge, this fairness-welfare trade-off has not been looked at in the house allocation setting previously. In this work, we give tight bounds for the price of fairness. Our investigation into PoF inspired us to look at simultaneously minimizing all three envy-minimization objectives while maximizing welfare. We show that we can indeed do this for $m = n$ and binary valuations. In this setting, we establish that there is an allocation that simultaneously maximizes welfare and minimizes the number of envious agents, maximum envy, and total envy.

Our Contributions.

We propose to study the issue of "minimizing envy" from a broader perspective, and to this end we consider three natural measures of the "amount of envy" created by an allocation: a) the total number of agents who experience envy (discussed above), b) the envy experienced by the *most* envious agent, where the amount of envy experienced by an agent is simply the number of agents that she is envious of, and c) the total amount of envy experienced by all agents. We refer to the questions of finding allocations that minimize these three measures of envy the Optimal House Allocation (OHA), Egalitarian House Allocation (EHA), and Utilitarian House Allocation (UHA) problems, respectively. A summary of our main results is in Table 4.1 and Table 4.6.

Hardness Results. We show the (parameterized) hardness of OHA and EHA even under highly restricted input settings. We show that OHA is NP-complete even on instances where every agent values at most two houses. Further, it is W[1]-hard when parameterized by k, the number of agents who are allowed to be envious (which implies that it is unlikely to admit a $f(k)(n+m)^{O(1)}$ time algorithm). As for EHA, we show that it is NP-complete, even on instances where every agent values at most two houses *and* every house is approved by a constant number of agents. In fact, we achieve this hardness even when the maximum allowed envy is just *one*, establishing that the problem is para-NP-hard when parameterized by k, the maximum envy (which implies that it is NP-hard even for a constant value of the parameter). The (parameterized) complexity of UHA, however, remains open.

Algorithmic and Experimental Results. Despite the hardness results even under the restricted input settings mentioned above, we explore tractable scenarios and prove a number of positive results. Observe that in a given instance, the number of houses(m) could be much larger than the number of agents(n). But we show that all three problems admit *polynomial-time pre-processing algorithms* that reduce the number of houses; in particular, after this pre-processing, we will have the guarantee that $m \leq 2(n-1)$. This result, in parameterized complexity parlance, means that all three problems admit *polynomial kernels* when parameterized by the number of agents. To prove this, we use a popular tool called the *expansion lemma*. While the kernels are interesting in their own right, we also leverage them to design polynomial-time algorithms for all three problems on binary *extremal* instances. An instance of OHA/EHA/UHA is extremal if the houses can be ordered in such a way that every agent values either the first few houses or the last few houses in the ordering; that is, there is an ordering $(h_1, h_2, \ldots, h_m)$ of the houses such that for every agent a, there is an index $i(a)$ with $0 \leq i(a) \leq m$ such that a either values the houses $\{h_1, h_2, \ldots, h_{i(a)}\}$ or $\{h_m, h_{m-1}, \ldots, h_{m-i(a)}\}$. We note here that extremal instances, although restrictive, form a non-trivial subclass for demonstrating tractability and have been studied in the literature (Elkind and Lackner, 2015). The hardness of the optimization problem even in the binary setting motivates to look at the structured binary preferences in the quest of tractability. In the context of house allocations, extremal instances appear where agents approvals are, for example, influenced by the distance of a house to either a hospital or a school, and the preferences decrease as the distance increases. In fact, despite the relatively simple structure of the preference profile, the obvious greedy approaches do not work and the three problems require three different lines of arguments.

Finally, we show that both OHA and EHA are fixed-parameter tractable (FPT)[2] when parameterized by the total number of house types or agent types; two agents are of the same type if they both like the same set of houses, and two houses are of the same type if they are both liked by the same set of agents. Notice that the number of (house or agent) types could potentially be much smaller than m or n. The FPT algorithms are obtained using ILP formulations with a bounded number of variables and constraints (for UHA, we obtain an integer quadratic program). The ILPs may also be of independent practical value. We implemented our ILPs for OHA and EHA over synthetic datasets of house allocation, generated uniformly at random. For a fixed number of houses and agents, the results show

[2]FPT w.r.t a parameter ℓ means the instance can be decided in time $f(\ell) \cdot (n+m)^{\mathcal{O}(1)}$, where f is an arbitrary computable function.

that the number of envious agents and the maximum envy decreases as the number of agent types (the maximum number of agents with distinct valuations) increases. Instances with identical valuations seem to admit more envy, attributed to the contention on the specific subset of houses. Also, when we increase the number of houses, for a constant number of agents and agent-types, the envy decreases, which is as expected, because of the increase in the number of choices and the fact that some houses (the more contentious ones) remain unallocated.

Price of Fairness. Along with different measures of envy, we also focus on the social welfare of an allocation, as captured by the sum of the individual agent utilities. Minimizing the measures of envy can lead to economically inefficient allocations with poor social welfare. For example, an allocation that only allocates the "dummy houses" (i.e., houses that no agent values) is trivially envy-free, but has zero social welfare. Quantifying this welfare loss, incurred as the cost of minimizing envy is, therefore, an imperative consideration. We quantify this trade-off between welfare and envy minimization using a metric called the *price of fairness (PoF)*; each measure of envy leads to its corresponding PoF, and we defer a formal definition of PoF to Section 4.9. We prove several tight bounds for PoF for binary valuations. In particular, we show that when $m > n$, PoF can be as large as the number of agents and that the bound is tight, for all the three envy measures of envy. We also identify the instances where no welfare has to be sacrificed in order to minimize envy, and hence no price has to be paid. In particular, we show that the price of fairness is 1 for $m = n$ and binary valuations and also for $m > n$ and binary doubly normalized valuations. We show in particular that when $m = n$, there is an allocation that simultaneously minimizes the number of envious agents, the maximum envy, and total envy while maximizing social welfare. Moreover, we can compute such an allocation in polynomial time.

Related Work.

Shapley and Scarf (1974) studied the house allocation model with existing tenants, which holds crucial applications in domains like kidney exchanges. Scenarios encompassing entirely new applicants, as well as mixed scenarios with a few existing tenants have also been studied in the literature (Hylland and Zeckhauser, 1979; Abdulkadiroğlu and Sönmez, 1999), and their practical implementations span diverse areas such as public housing and college dormitory assignments, among others.

The notion of fairness in house allocation setting was initiated by Beynier et al. (2019a) who studied a local variant for an equal number of agents and houses, where an agent can envy only those who are connected to her in a given social network. Gan et al. (2019) studied envy-freeness when the number of houses can be more than that of agents and gave an efficient algorithm that returns an envy-free solution if it exists. When such solutions do not exist, Kamiyama et al. (2021) initiated the study of finding allocations that minimize the number of agents who experience envy and showed that it is NP-complete to find allocations that minimize the number of envious agents, even for binary utilities, and this quantity is hard to approximate for general utilities. Further, Aigner-Horev and Segal-Halevi (2021) studied the relaxed variant of assigning at most one house to every agent and give an $O(m\sqrt{n})$ algorithm for finding an envy-free matching of maximum cardinality in the setting of binary utilities. Shende and Purohit (2020) studied envy-freeness in conjunction with strategy-proofness. In more recent work, Hosseini et al. (2023c, 2024b) have considered minimizing the sum of all pairwise envy values over all edges in a social network. They proved structural and computational results for various classes of underlying graphs on agents. Hosseini et al. (2024a) looked at the degree of fairness while maximizing the social welfare and the size of an envy-free allocation. Choo et al. (2024) discussed house allocations in the context of subsidies and showed that finding envy-free allocations with minimum subsidy is hard in general but tractable if agents have identical utilities or m differs from n by an additive constant.

The price of fairness was first proposed by Bertsimas et al. (2011), following which there has been substantial progress towards finding the bounds for the price for various combinations of fairness and welfare notions, specifically in resource allocation setting (Caragiannis et al., 2012; Bei et al., 2021; Barman et al., 2020b; Sun et al., 2023b; Bhaskar et al., 2023). Hosseini et al. (2024a) recently looked at the degree of fairness while maximizing the social welfare and the size of an envy-free allocation.

Organization of the paper.

We discuss the results for OHA, EHA, and UHA in Section 4.4, Section 4.5, and Section 4.6 respectively. We discuss the experiments in Section 4.7. Finally, we discuss the price of fairness in Section 4.9, which is largely independent of all the other sections.

4.2 Preliminaries

Let $[k]$ denote the set $\{1, 2, \ldots, k\}$ for any positive integer k.

	Cardinal					
	General	Binary				Rankings
		General	Extremal Intervals	d = 1	d = 2	
OHA	NP-Complete (by implication)	NP-Complete (†) from Clique (Theorem 4.20)	P (Theorem 4.15)	P (Theorem 4.17)	NP-Complete from Clique (Theorem 4.21)	NP-Complete (Theorem 4.20)
EHA	NP-Complete (by implication)		P (Theorem 4.37)	P (Theorem 4.38)	NP-Complete (⋆) from Independent Set (Theorem 4.39)	NP-Complete (⋆) from Multi-Colored Independent Set (Theorem 4.40)
UHA	?		P (Theorem 4.50)	P (Corollary 4.54)	?	

Table 4.1: A partial summary of our results. Here, d denotes the maximum number of houses approved by any agent. The results marked with a $\star$ refer to reductions that imply hardness even when the standard parameter is a constant, while the result marked with a † is a FPT reduction and also implies W[1]-hardness in the standard parameter.

An instance $\mathcal{I}$ of the House Allocation problem (HA) comprises a set $A = \{a_1, a_2, \ldots, a_n\}$ of agents, a set $H = \{h_1, h_2, \ldots, h_m\}$ of houses and a preference profile (rankings or cardinal utilities) that capture the preference of all agents over the houses. An assignment or house allocation is an injection $\Phi : A \to H$. Throughout this section, let Φ be an arbitrary but fixed allocation. While we make all our notation explicit with respect to Φ, during future discussions, the subscript Φ may be dropped if the allocation is clear from the context. There are a few different ways in which agents may express their preferences over houses, and we focus here on both linear orders as well as cardinal utilities.

Rankings

In this setting, each agent $a \in A$ has a linear order $\succ_a$ over the set of houses H. We will typically use $\succ_i$ to denote[3] the preferences of agent a_i. The *rank* of a house h in the order $\succ_a$ is one plus the number of houses h' such that $h' \succ_a h$. For example, if $m = 3$ and $\succ_a$ is given by $h_2 \succ h_3 \succ h_1$, then the houses h_2, h_3 and h_1 have ranks $1, 2,$ and 3 respectively. We denote the rank of a house h in an order $\succ$ by $rk(h, \succ)$. A ranking is said to have ties ($rk(h, \succeq)$) if there is an agent who ranks some pair of houses equally.

An agent $a \in A$ *envies* an agent $b \in A$ under the allocation Φ if $\Phi(a) \prec_a \Phi(b)$, which is to say that a perceives Φ to have allocated a house to b that she ranks more than the one allocated

[3] In some of the reductions, the indices of the agents in A are different from $[n]$, and we continue to adopt the convention that $\succ_\circ$ is used to describe the preferences of the agent $a_\circ$.

to her. We use $\mathcal{E}_\Phi(a)$ to denote the set of agents b such that a envies b in the allocation Φ.

An agent a is *envy-free* with respect to Φ if there is no agent b such that a envies b. In other words, $\mathcal{E}_\Phi(a) = \emptyset$. An allocation Φ is said to be *envy-free* if all agents a are envy-free with respect to Φ. The *amount of envy* experienced by an agent a is the number of agents she is envious of, that is, $|\mathcal{E}_\Phi(a)|$ and is denoted by $\kappa_\Phi(a)$. Note that if an agent is envy-free, then the amount of envy experienced by her is zero.

Binary Preferences

The utility that an agent a derives from a house h is denoted by $u_a(h)$. Preferences are said to be *binary* if $u_a(h) \in \{0, 1\}$ for all $a \in A$ and $h \in H$. We note here that binary utilities are a crucial subclass with simple elicitation and several works in fair division and voting literature have paid special attention to this case (Halpern et al., 2020; Barman et al., 2018b; Lackner and Skowron, 2023). A house h is called a *dummy* house if the utility of every agent for it is zero, that is, $u_a(h) = 0$ for all $a \in A$. The set of dummy houses is denoted by D. An agent a is called a *dummy* agent if it values every house at zero, that is, $u_a(h) = 0$ for all $h \in H$. The set of dummy agents is denoted by D'.

As previously, an agent $a \in A$ *envies* an agent $b \in A$ under the allocation Φ if $u_a(\Phi(a)) < u_a(\Phi(b))$, which is to say that a perceives Φ to have allocated a house to b that she values more than the one allocated to her. That is, $u_a(\Phi(a)) = 0$ but $u_a(\Phi(b)) = 1$. The definition of $\mathcal{E}_\Phi(a)$ and the notion of *envy-freeness* is the same as before. The amount of envy experienced by an agent is $|\mathcal{E}_\Phi(a)|$ and is denoted by $\kappa_\Phi(a)$. Just as with rankings, if an agent is envy-free, then the amount of envy experienced by her is zero.

Let $\mathcal{P}$ be a profile of binary utilities of agents A over houses H. For an agent $a \in A$, use $\mathcal{P}(a)$ to denote the set of houses h for which $u_a(h) = 1$. We say that these are houses that are valued by the agent a. For a subset $S \subseteq A$, we use $\mathcal{P}(S)$ to denote $\cup_{a \in S} \mathcal{P}(a)$. Similarly, for a house h, we use $\mathcal{T}(h)$ to refer to the set of agents who value h, and for a subset $S \subseteq H$, we use $\mathcal{T}(S)$ to denote $\cup_{h \in S} \mathcal{T}(h)$.

Two agents a_p and a_q are said to be of the same type if $\mathcal{P}(a_p) = \mathcal{P}(a_q)$ and two houses h_p and h_q are said to be of the same type if $\mathcal{T}(h_p) = \mathcal{T}(h_q)$. For an instance with n agents and m houses, we use $n^\star$ and $m^\star$ to denote the number of distinct types of agents and houses, respectively.

The *preference graph* G based on $\mathcal{P}$ is a bipartite graph defined as follows: the vertex set of G consists of one vertex v_a corresponding to every agent $a \in A$ and one vertex v_h corresponding to every house $h \in H$; and (v_a, v_h) is an edge in G if and only if a values h.

Extremal Interval Structure. We say that $\mathcal{P}$ has an *extremal interval structure with respect to houses* if there exists an ordering σ of the houses such that for every agent a, $\mathcal{P}(a)$ forms a prefix or suffix of σ (Elkind and Lackner, 2015). Further, we say that $\mathcal{P}$ has a ***left*** *(respectively,* ***right****) extremal interval structure with respect to houses* if there exists an ordering σ of the houses such that for every agent a, $\mathcal{P}(a)$ forms a prefix (respectively, suffix) of σ.
Analogously, $\mathcal{P}$ has an *extremal interval structure with respect to agents* if there exists an ordering π of the agents such that for every house h, the set of agents who value h forms a prefix or suffix of π. The notions of left and right extremal interval structures here are also defined as before. In our discussions, whenever we speak of an *extremal interval structure* without explicit qualification, it is with respect to houses unless mentioned otherwise.

Optimization Objectives

We focus on the following optimization objectives.

1. The *number of envious agents* in an allocation Φ is the number of agents $a \in A$ for which $\kappa_\Phi(a) \geq 1$ and will be denoted by $\kappa^\#(\Phi)$. Further, given an instance $\mathcal{I}$ of the house allocation problem, we use $\kappa^\#(\mathcal{I})$ to denote the number of envious agents in an optimal allocation, that is, $\kappa^\#(\mathcal{I}) := \min_\Phi(\kappa^\#(\Phi))$.
2. The *maximum envy* generated by Φ is $\max_{a \in A} \kappa_\Phi(a)$ and is denoted by $\kappa^\dagger(\Phi)$. As before, given an instance $\mathcal{I}$ of the house allocation problem, we use $\kappa^\dagger(\mathcal{I})$ to denote the maximum envy in an optimal allocation, that is, $\kappa^\dagger(\mathcal{I}) := \min_\Phi(\kappa^\dagger(\Phi))$.
3. The *total envy* generated by Φ is $\sum_{a \in A} \kappa_\Phi(a)$ and will be denoted by $\kappa^\star(\Phi)$. Again, given an instance $\mathcal{I}$ of the house allocation problem, we use $\kappa^\star(\mathcal{I})$ to denote the total envy in an optimal allocation, that is, $\kappa^\star(\mathcal{I}) := \min_\Phi(\kappa^\star(\Phi))$.

Computational Questions

We now formulate the computational problems that we would like to address.

Optimal House Allocation
Input: A set $A = \{a_1, a_2, \ldots, a_n\}$ of agents and a set $H = \{h_1, h_2, \ldots, h_m\}$ of houses, a preference profile describing the preferences of all agents over houses, and a non-negative integer $k \in \mathbb{Z}^+$.
Question: Determine if there is an allocation Φ with number of envious agents at most k, i.e, $\kappa^{\#}(\Phi) \leq k$.

Egalitarian House Allocation
Input: A set $A = \{a_1, a_2, \ldots, a_n\}$ of agents and a set $H = \{h_1, h_2, \ldots, h_m\}$ of houses, a preference profile describing the preferences of all agents over houses, and a non-negative integer $k \in \mathbb{Z}^+$.
Question: Determine if there is an allocation Φ with maximum envy at most k, i.e, $\kappa^{\dagger}(\Phi) \leq k$.

Utilitarian House Allocation
Input: A set $A = \{a_1, a_2, \ldots, a_n\}$ of agents and a set $H = \{h_1, h_2, \ldots, h_m\}$ of houses, a preference profile describing the preferences of all agents over houses, and a non-negative integer $k \in \mathbb{Z}^+$.
Question: Determine if there is an allocation Φ with total envy at most k, i.e, $\kappa^{\star}(\Phi) \leq k$.

We use $[\succ]$-OHA, $[\succeq]$-OHA, and $[0/1]$-OHA to denote the versions of the OHA problem when the preferences are given, respectively, by linear orders, rankings with ties, and binary utilities, respectively. We adopt this convention for EHA and UHA as well.

We note that for all three problems, the question of finding an envy-free allocation, i.e., one for which the optimization objective attains the value zero, is a natural special case. This amounts to finding an allocation where no agent has any envy for another and is therefore resolved (for both binary valuations and rankings) by the algorithm of Gan et al. (2019) which uses an approach based on iteratively eliminating subsets that violate Hall's condition in the preference graph.

We also observe here that all three problems are tractable for the special case when $m = n$. We assume without loss of generality, that every agent values at least one house. Observe that since $n = m$, all valid allocations have no unallocated houses.

Proposition 4.1 (folklore). *$[0/1]$-Optimal House Allocation and $[\succeq]$-Optimal House Allocation can be solved in polynomial time if $m = n$.*

Proof. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ be an instance of $[0/1]$-OHA and let $G = (A \cup H; E)$ be the associated preference graph. We obtain an optimal allocation in polynomial time, and we return an appropriate output based on how the value of the optimum compares with k.

Let M be a maximum matching in G. We claim that any optimal allocation for $\mathcal{I}$ has $|M|$ envy-free agents. To see that there exists an allocation that has at least $|M|$ envy-free agents, consider the allocation that gives the house $M(a)$ to every agent a saturated by M, and allocates the remaining houses arbitrarily among the agents not saturated by M. It is easy to see that this allocation has at least $|M|$ envy-free agents, namely the ones corresponding to those saturated by the matching M. On the other hand, suppose there is an allocation Φ with k envy-free agents, then the envy-free agents must have received houses that they value—indeed, consider any agent a, and let h be any house that a values. If a does not value the house $\Phi(a)$, then a envies the agent who received the house h. Therefore, the set:

$$M := \{(a, \Phi(a)) \mid a \text{ is envy-free with respect to } \Phi\}$$

corresponds to a matching with k edges in G, and this concludes the argument. Notice that this argument extends to weak orders by the natural extension of the notion of a preference graph: we have that an agent a is adjacent to all houses h that she prefers over all other houses. Therefore, we have the claim $[\succeq]$-OHA as well. □

Proposition 4.2. *$[0/1]$-Egalitarian House Allocation and $[\succ]$-Egalitarian House Allocation can be solved in polynomial time if $m = n$.*

Proof. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ be an instance of $[0/1]$-EHA. We show that this problem reduces to finding a perfect matching among high-degree agents in the preference graph, based on the following observations.

1. Suppose a is an agent who values at most k houses. Then, the amount of envy experienced by a is at most k in *any* allocation.
2. Suppose a is an agent who values at least $k+1$ houses. Then, if Φ is a valid solution, then a must value the house $\Phi(a)$.

It follows that $\mathcal{I}$ is a Yes-instance if and only if the projection of the preference graph G on $(A^\star \cup H)$ admits a perfect matching, where $A^\star$ is the subset of agents whose degree in G is at least $k+1$.

Now, let $\mathcal{I} := (A, H, \succ; k)$ be an instance of $[\succ]$-EHA. Consider the bipartite graph $G = (A \cup H; E^\star)$, where (a, h) is an edge if and only if the rank of h is at most $k+1$ in $\succ_a$. We claim that $\mathcal{I}$ is a Yes-instance if and only if G has a perfect matching.

Indeed, observe that the amount of envy experienced by any agent a with respect to an allocation Φ is exactly one less than the rank of $\Phi(a)$ in $\succ_a$. Therefore, if Φ is an allocation whose maximum envy is k, then the rank of $\Phi(a)$ in $\succ_a$ must be at most $k+1$ for all agents a. It is easy to check that such allocations are in one-to-one correspondence with perfect matchings in the graph G. □

Proposition 4.3. *[0/1]Utilitarian House Allocation and [≻]-Utilitarian House Allocation can be solved in polynomial time if $m = n$.*

Proof. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ be an instance of [0/1]-UHA. To begin with, let $d(a)$ denote the degree of a in the preference graph G of $\mathcal{I}$. Now consider a complete bipartite graph $G^\star$ with bi-partition $(A \uplus H)$ and a cost function c on the edges defined as follows:

$$c((a,h)) = \begin{cases} d(a) & \text{if } a \text{ does not value } h, \\ 0 & \text{otherwise.} \end{cases}$$

Let M be a minimum cost perfect matching in G with total cost t. We claim that $\mathcal{I}$ is a Yes-instance if and only if $t \leq k$. In the forward direction, if Φ is an allocation with total envy at most $\kappa^\star$, then consider the following perfect matching in G:

$$M := \{(a, \Phi(a)) \mid a \in A\}$$

Notice that the cost of M corresponds exactly to $\kappa^\star(\Phi)$, the total amount of envy in Φ. This shows that there is a perfect matching in $G^\star$ with cost at most k.

On the other hand, let M be a perfect matching in $G^\star$ with the cost at most k, and let $M(a)$ denote the house h such that $(a,h) \in M$. Then consider the allocation Φ given by $\Phi(a) = M(a)$ for all agents a. Notice that every zero-cost edge in M corresponds to an envy-free agent with respect to Φ, and every other edge $e = (a,h)$ corresponds to an agent in Φ who was allocated a house she did not value. Observe that the amount of envy experienced by a in Φ is the number of houses she values, in other words, $d(a)$; however, this is also exactly the cost of the edge e. Therefore, it follows that the amount of envy in Φ is exactly equal to the cost of the matching M, and this concludes the proof of our claim.

Now, let $\mathcal{I} := (A, H, \succ; k)$ be an instance of [≻]-UHA. As before, consider a complete bipartite graph $G^\star$ with bipartition $(A \uplus H)$. This time, we have the cost function c on the edges defined as $c((a,h)) = rk(h, \succ_a) - 1$. This cost reflects the envy experienced by the agent a if she were to be allocated the house h. Using arguments similar to the setting of binary valuations, it is easily checked that $\mathcal{I}$ is a Yes-instance if and only if there is a perfect matching in $G^\star$ whose cost is at most k. □

Parameterized Complexity

Parameterized algorithms or multi-variate analysis is a popular perspective in the context of "coping with computational hardness". The key idea here is to segregate the running time of our algorithms into two parts: one that is polynomially bounded in the size of the entire input so that it is efficient on a quantity that is expected to be large in practice and the other, a computable function of a carefully chosen *parameter*—and this component of the running time remains feasible in practice because the parameter is expected to be small. The parameterized perspective also allows us to formalize ideas about efficient preprocessing, and this is now an active subfield in its own right. We refer the readers to the books Cygan et al. (2015) and Downey and Fellows (2013) for additional background on this algorithmic paradigm, while recalling here only the key definitions relevant to our discussions.

Formally, a parameterized problem L is a subset of $\Sigma^* \times \mathbb{N}$ for some finite alphabet Σ. An instance of a parameterized problem consists of (x, k), where k is called the parameter. A central notion in parameterized complexity is fixed-parameter tractability (FPT), which means for a given instance (x, k) solvability in time $f(k) \cdot p(|x|)$, where f is an arbitrary function of k and p is a polynomial in the input size. The notion of kernelization is defined as follows.

Definition 4.4. *A kernelization algorithm, or in short, a kernel for a parameterized problem $Q \subseteq \Sigma^* \times \mathbb{N}$ is an algorithm that, given $(x, k) \in \Sigma^* \times \mathbb{N}$, outputs in time polynomial in $|x| + k$ a pair $(x', k') \in \Sigma^* \times \mathbb{N}$ such that (a) $(x, k) \in Q$ if and only if $(x', k') \in Q$ and (b) $|x'| + k' \leq g(k)$, where g is an arbitrary computable function. The function g is referred to as the size of the kernel. If g is a polynomial function then we say that Q admits a polynomial kernel.*

On the other hand, we also have a well-developed theory of parameterized hardness. We call a problem para-NP-hard if it is NP-hard even for a constant value of the parameter. Further, we have the notion of *parameterized reductions*, defined as follows.

Definition 4.5 (Parameterized reduction). *Let $A, B \subseteq \Sigma^* \times \mathbb{N}$ be two parameterized problems, A parameterized reduction from A to B is an algorithm that, given an instance (x, k) of A, outputs an instance (x', k') of B such that:*

1. *(x, k) is a yes-instance of A if and only if (x', k') is a yes-instance of B.*

2. $k' \leq g(k)$ for some computable function g, and

3. the running time is $f(k) \cdot |x|^{O(1)}$ for some computable function f.

A parameterized reduction from a problem known to be W[1]-hard is considered to be strong evidence that the target problem is not FPT. The formal definition of W[1]-hardness is beyond the scope of this discussion, but we state in this section the known W[1]-hard problems from which we will perform parameterized reductions to obtain our results.

We also note that an instance of integer linear programming is FPT in the number of variables.

Theorem 4.6. *(Lenstra, 1983) An integer linear programming instance of size L with p variables can be solved using $\mathcal{O}\left(p^{2.5p+o(p)} \cdot (L + \log M_x) \log (M_x M_c)\right)$ arithmetic operations and space polynomial in $L + \log M_x$, where M_x is an upper bound on the absolute value a variable can take in a solution, and M_c is the largest absolute value of a coefficient in the vector c.*

Finally, we state the problems that we will use in the reductions. We note that all the problems below are W[1]-hard when parameterized by k (Garey and Johnson, 1990; Fellows et al., 2009).

Clique (respectively, Independent Set)
Input: A graph G and an integer k.
Question: Does there exist a subset $S \subseteq V(G)$ such that $G[S]$ is a clique (respectively, independent set) and $|S| \geq k$?

Maximum Balanced Biclique
Input: A graph $G = (L \cup R, E)$ and an integer k.
Question: Does there exist a subset $S \subseteq L$ and $T \subseteq R$ such that $G[S \cup T]$ is a biclique and $|S| = |T| = k$?

Multi-Colored Independent Set
Input: A graph $G = (V_1 \uplus \cdots \uplus V_k, E)$.
Question: Does there exist a subset $S \subseteq V(G)$ such that $G[S]$ is an independent set and $|V_i \cap S| = 1$ for all $i \in [k]$?

4.3 Pre-processing using Expansion Lemma

In this section, we introduce the expansion lemma, a powerful and popular tool for kernelization, which we will use later to design our algorithms. Let G be a bipartite graph with vertex bi-partitions (A, B). A set of edges $M \subseteq E(G)$ is called an *expansion* of A into B if:

- every vertex of A is incident to exactly one edge of M;
- M saturates exactly $|A|$ vertices in B.

Note that an expansion saturates all vertices of A.

Lemma 4.7 (Expansion lemma (Cygan et al., 2015)**).** *Let G be a bipartite graph with vertex bi-partitions (A, B) such that*

1. *$|B| \geq |A|$, and*
2. *there are no isolated vertices in B.*

Then there exist non-empty vertex sets $X \subseteq A$ and $Y \subseteq B$ such that

- *there is an expansion of X into Y, and*
- *no vertex in Y has a neighbor outside X, that is, $N(Y) \subseteq X$.*

Furthermore, the sets X and Y can be found in time polynomial in the size of G (Thomassé (2010)).

We now consider an instance $\mathcal{I} = (A, H, \mathcal{P})$ of HA with binary valuations parameterized by k, where k is one of $\kappa^{\#}$, $\kappa^{\dagger}$ or $\kappa^{\star}$. We introduce the following reduction rules here whose implementation is parameter-agnostic.

Let $G = (A \cup H; E)$ denote the preference graph of $\mathcal{I}$. Note that we may assume that we have at most $(n-1)$ dummy houses in the instance $\mathcal{I}$, since instances with at least n dummy houses admit trivial envy-free allocations, where every agent is given a dummy house. We make this explicit in the following reduction rule:

Reduction Rule 1. *If $\mathcal{I}$ has at least as many dummy houses as agents, then return a trivial Yes-instance. The parameter k is unchanged.*

Let $G^{\star}$ denote the preference graph induced by $(A \setminus D') \cup (H \setminus D)$, where D denotes the vertices corresponding to dummy houses in $\mathcal{I}$ and D' denotes the vertices corresponding to the dummy agents in $\mathcal{I}$. We now propose the following reduction rule based on the expansion lemma. The safety of the following reduction rule is argued separately for the three parameters in Theorems 4.23, 4.43 and 4.55, respectively.

Reduction Rule 2. *In $\mathcal{I}$, if $|H \setminus D| \geq |A \setminus D'|$, then let (X, Y) be as given by Lemma 4.7 applied to $G^\star$, and let $M \subseteq E(G^\star)$ be the associated expansion. Proceed by eliminating all agents and houses saturated by M. The parameter k is unchanged.*

Note that once Reduction Rules 1 and 2 are applied exhaustively, we have that:

$$|H| = |H \setminus D| + |D| \leq |A \setminus D'| + |D| \leq |A| + |A| - 1 \leq 2 \cdot (|A| - 1),$$

where we are slightly abusing notation and using H and A to denote the houses and agents in the reduced instance. Thus, once the safety of these reduction rules is established, we conclude that all the three problems under consideration—OHA, EHA and UHA—admit polynomial kernels with $\mathcal{O}(|A|)$ houses when parameterized by the number of agents.

After the above two reduction rules have been applied, we have $|H| \geq |A|$ and $|H \setminus D| \leq |A \setminus D'|$, we get $|H| - |H \setminus D| \geq |A| - |A \setminus D'|$ and hence $|D| \geq |D'|$. Therefore, the number of dummy houses are at least as much as the number of dummy agents. We then apply the following reduction rule.

Reduction Rule 3. *In a reduced instance $\mathcal{I}$ with respect to Reduction Rules 1 and 2, proceed by allocating a dummy house to each of the dummy agents and eliminate $|D'|$ dummy agents and $|D'|$ dummy houses. The parameter k is unchanged.*

Towards establishing the safety of the above reduction rule, we argue that in the reduced instance, there is an optimal allocation where all the dummy agents receive a dummy house each. Indeed, if not, say a dummy agent a_d receives a house $h \in H \setminus D$ in an optimal allocation. Since $|A \setminus D'| > |H \setminus D|$, there is an agent $a \in A \setminus D'$ who received a dummy house h_d. We re-allocate h_d to a_d and allocate all the houses in $|H \setminus D|$ to agents in $|A \setminus D|$ by finding a maximum matching in the associated preference graph G, restricted to these houses and agents. This re-allocation either does not create any new envy (a_d is indifferent and no one else if worse-off) or makes an additional agent in $A \setminus D$ envy-free (because of the allocation of h to an agent who values it). Hence, either we get the desired optimal allocation under which all the dummy agents receive a dummy house or we contradict the fact that we started with an optimal allocation. So we can allocate dummy houses to the dummy agents and assume going forward that there are no dummy agents. That is, $|D'| = \phi$.

We now make a claim here that will be useful in the arguments that we make later about the safety of Reduction Rule 2.

Definition 4.8. *Let $(A, H, \mathcal{P})$ and X, Y and M be as in the premise of Reduction Rule 2. An* *allocation Φ is said to be* good *if $\Phi(a) = M(a)$ for all $a \in X$, where $M(a)$ denotes the unique vertex $h \in Y$ such that $(a, h) \in M$.*

Claim 4.9. *Let $(A, H, \mathcal{P})$ and X, Y and M be as in the premise of Reduction Rule 2. There is a good allocation that minimizes the number of envious agents.*

Proof. Let Φ be an allocation that minimizes the number of envious agents. If Φ is already good then there is nothing to prove. Otherwise, suppose $\Phi(a) \neq M(a)$ for some $a \in X$. If Φ does not assign the house corresponding to $M(a)$ to any agent, then consider the modified allocation Φ' where we assign $M(a)$ to a while letting $\Phi(a)$ become an unassigned house, that is:

$$\Phi'(c) = \begin{cases} M(a) & \text{if } c = a, \\ \Phi(c) & \text{otherwise.} \end{cases}$$

On the other hand, suppose b is such that $\Phi(b) = M(a)$. Then consider the modified allocation Φ' where we swap the houses of a and b, that is:

$$\Phi'(c) = \begin{cases} M(a) & \text{if } c = a, \\ \Phi(a) & \text{if } c = b, \\ \Phi(c) & \text{otherwise.} \end{cases}$$

We keep modifying the original allocation Φ in the manner described above until we arrive at a good allocation. Let $\Phi^\star$ denote this final allocation.

Now consider an agent $c \notin X$. We claim that the amount of envy experienced by c does not increase at any step of the process of morphing Φ to $\Phi^\star$. Consider the following cases that arise at any step, where $a \in X$ and by a slight abuse of notation, we use Φ to denote the allocation that is being modified:

1. Suppose $M(a)$ is assigned to a and $\Phi(a)$ is unassigned. We know that c does not value $M(a)$ since $c \notin X$. If c valued $\Phi(a)$, then her envy with respect to the new allocation will be one less than her envy with respect to the previous allocation.
2. Suppose $M(a)$ and $\Phi(a)$ are swapped between agents a and b; and $c \neq b$. Then the amount of envy experienced by c does not change.
3. Suppose $M(a)$ and $\Phi(a)$ are swapped between agents a and b; and $c = b$. If c valued

$\Phi(a)$, then her envy with respect to the new allocation will be less than her envy with respect to the previous allocation. On the other hand, if c does not value $\Phi(a)$ then the amount of envy experienced by c does not change.

Also, in the final allocation $\Phi^\star$, all agents in X are envy-free, since they are assigned houses that they value via the expansion M. Therefore, the total number of envious agents in the final allocation $\Phi^\star$ is the same as the number of envious agents in the original allocation Φ—recall that Φ minimized the number of envious agents. □

The following claims can be shown by the same argument that was used for Claim 4.9, since for any agent a, the amount of envy experienced by a respect to in $\Phi^\star$ is at most the amount of envy experienced by a with respect to Φ.

Claim 4.10. *Let $(A, H, \mathcal{P})$ and X, Y and M be as in the premise of Reduction Rule 2. There is a good allocation that minimizes the maximum envy.*

Claim 4.11. *Let $(A, H, \mathcal{P})$ and X, Y and M be as in the premise of Reduction Rule 2. There is a good allocation that minimizes total envy.*

Applications to Extremal Instances

We observe several properties of instances $\mathcal{I} = (A, H, \mathcal{P}; k)$ of [0/1]-HA, where $\mathcal{P}$ has an extremal interval structure (with respect to the houses). The properties are parameter-agnostic, and therefore, hold for instances of [0/1]-OHA, [0/1]-EHA and [0/1]-UHA. As a shorthand, we say that $\mathcal{I} = (A, H, \mathcal{P}; k)$ is an extremal instance (or simply extremal) if $\mathcal{P}$ has the extremal interval structure.

Consider an extremal instance $\mathcal{I} = (A, H, \mathcal{P}; k)$ of [0/1]-HA. That is, there is an ordering σ of the houses, say, $\sigma = (h_1, \ldots, h_m)$ such that for every agent $a \in A$, either $\mathcal{P}(a) = \phi$ or there exists an index $i(a)$ such that $1 \leq i(a) \leq m$ and either $\mathcal{P}(a) = \left\{h_1, h_2, \ldots, h_{i(a)}\right\}$ or $\mathcal{P}(a) = \left\{h_m, h_{m-1}, \ldots, h_{i(a)}\right\}$. (See Table 4.2 for an example.) If $\mathcal{P}(a) = \phi$ or $\left\{h_1, h_2, \ldots, h_{i(a)}\right\}$ for every $a \in A$, then we say that the instance $\mathcal{I}$ is left-extremal. If $\mathcal{P}(a) = \phi$ or $\left\{h_m, h_{m-1}, \ldots, h_{i(a)}\right\}$ for every $a \in A$, then we say that $\mathcal{I}$ is right-extremal.

We can check in polynomial time whether a given instance is (left/right)-extremal, and if so, then find the ordering σ on the houses. Also, note that removing a subset of houses and agents from an extremal instance does not destroy the extremal property. That is, if $\mathcal{I} = (A, H, \mathcal{P}, k)$

		Left-houses (H_L)			Dummy houses (D)				Right-houses (H_R)		
		h_1	h_2	h_3	h_4	h_5	h_6	h_7	h_8	h_9	h_{10}
Left-agents (A_L)	a_1	1	0	0	0	0	0	0	0	0	0
	a_2	1	1	0	0	0	0	0	0	0	0
	a_3	1	1	0	0	0	0	0	0	0	0
	a_4	1	1	1	0	0	0	0	0	0	0
	a_5	1	1	1	0	0	0	0	0	0	0
Right-agents (A_R)	a_6	0	0	0	0	0	0	0	1	1	1
	a_7	0	0	0	0	0	0	0	1	1	1
	a_8	0	0	0	0	0	0	0	0	1	1
	a_9	0	0	0	0	0	0	0	0	0	1

Table 4.2: An example of an extremal instance $\mathcal{I} = (A, H, \mathcal{P}; k)$, where $D \subseteq H$ denotes the dummy houses. Once Reduction Rules 1, 2 and 3 are no longer applicable, then there are no dummy agents, but dummy houses must necessarily exist, i.e., $D \neq \emptyset$; and we must have $|H_L| < |A_L|$ and $|H_R| < |A_R|$.

is extremal, then so is $\mathcal{I}' = (A', H', \mathcal{P}')$, where $A' = A \setminus X$ and $H' = H \setminus Y$ and $\mathcal{P}'$ is the restriction of $\mathcal{P}$ to $(A \cup H) \setminus (X \cup Y)$. So, in particular, we can safely apply Reduction Rules 2 and 3 to extremal instances.

Now, consider an extremal instance $\mathcal{I} = (A, H, \mathcal{P}; k)$, which is irreducible with respect to Reduction Rules 1, 2 and 3. Let $\sigma = (h_1, h_2, \ldots, h_m)$ be an extremal ordering on the houses. Let D' be the set of dummy agents. Due to Reduction Rule 3, we have $|D'| = \phi$ in the $\mathcal{I}$. Let D be the set of dummy houses in $\mathcal{I}$. Then, we have seen that $|H \setminus D| \leq n - 1$. Therefore, $D \neq \emptyset$. Let $h_d, h_{d'} \in D$ be such that h_d is the first dummy house and $h_{d'}$ is the last dummy house in the ordering σ. It may be the case that $d = d'$. Then, for every i, where $d \leq i \leq d'$, the house h_i is a dummy house. Let $H_L = \{h_1, h_2, \ldots, h_{d-1}\}$ and $H_R = \{h_{d'+1}, h_{d'+2}, \ldots, h_m\}$. We call the houses in H_L the left-houses and the houses in H_R the right-houses. Assume that the reduced instance contains both left and right houses, that is, $|H_L| > 0$ and $H_R > 0$. For every agent $a \in A$, either $\mathcal{P}(a) \subseteq H_L$, in which case we call the agent a a left-agent, or $\mathcal{P}(a) \subseteq H_R$, in which case we call the agent a a right-agent. Let A_L and A_R respectively denote the set of left and right agents. See Table 4.2. Thus, we have a partition of A into A_L and A_R and a partition of H into H_L, H_R and D. Notice now that if $|H_L| \geq |A_L|$ or $|H_R| \geq |A_R|$, then Reduction Rule 2 would apply. We thus have $|A_L| > |H_L|$ and $|A_R| > |H_R|$. For an allocation $\Phi : A \to H$ and an agent $a \in A$, we say that a is extremality-respecting under Φ if either (a) $a \in A_L$ and $\Phi(a) \in H_L \cup D$ or (b) $a \in A_R$ and $\Phi(a) \in H_R \cup D$. We say that Φ is extremality-respecting if every agent in $A_L \cup A_R$ is extremality-respecting under Φ. We now claim that there exists an extremality-respecting optimal allocation, irrespective of whether we are dealing with an instance of [0/1]-OHA, [0/1]-EHA or [0/1]-UHA.

Claim 4.12. *There exists an extremality-respecting optimal allocation for any instance irreducible with respect to the Reduction Rules 1, 2 and 3.*

Proof. Let $\Phi : A \to H$ be an optimal allocation that maximizes the number of extremality-respecting agents. If Φ is extremality-respecting, then the claim trivially holds. So, assume not. Then, assume without loss of generality that there exists $a \in A_L$ with $\Phi(a) \in H_R$. (The case when $a \in A_R$ with $\Phi(a) \in H_L$ is symmetric.)

Since $|H_R| < |A_R|$, there exists an agent $a' \in A_R$ such that $\Phi(a') \notin H_R$. Then, $\Phi(a') \in H_L \cup D$. Let Φ' be the allocation obtained by swapping the houses of a and a'. This house-swapping between a and a' does not cause any increase in the number of envious agents or the envy experienced by any agent. So we have $\kappa^{\#}(\Phi') \leq \kappa^{\#}(\Phi)$, $\kappa^{\dagger}(\Phi') \leq \kappa^{\dagger}(\Phi)$ and $\kappa^{\star}(\Phi') \leq \kappa^{\star}(\Phi)$. Therefore, Φ' is optimal. In addition, note that $\Phi'(a) = \Phi(a') \in H_L \cup D$, and $\Phi'(a') = \Phi(a) \in H_R$, and hence, both a and a' are extremality-respecting under Φ. That is, the number of extremality-respecting agents under Φ' is strictly greater than that in Φ, a contradiction. □

Remark 4.13. *Claim 4.12 shows that whenever dealing with an instance $\mathcal{I} = (A, H, \mathcal{P}, k)$ of [0/1]-HA, where $\mathcal{I}$ is extremal, we only need to look for an extremality-respecting optimal allocation, say Φ. Note that there are no dummy agents in the reduced instance (Reduction Rule 3). Now suppose n_L dummy houses get allocated to left-agents under Φ, and n_R dummy houses get allocated to right-agents under Φ. Hence, we can guess the pair (n_L, n_R) and split $\mathcal{I}$ into two instances, $\mathcal{I}_L$ and $\mathcal{I}_R$, where $\mathcal{I}_L$ consists of the left-agents, the left-houses and n_L dummy houses, and $\mathcal{I}_R$ consists of the right-agents, the right-houses and n_R dummy houses. Thus, $\mathcal{I}_L$ is left-extremal and $\mathcal{I}_R$ is right-extremal. We only need to solve the problem separately on $\mathcal{I}_L$ and $\mathcal{I}_R$. Notice that the number of guesses for the pair (n_L, n_R) is at most n^2. By reversing the ordering on the houses in the instance $\mathcal{I}_R$, we can turn $\mathcal{I}_R$ into a left-extremal instance as well. So, it suffices to solve the problem for left-extremal instances. Hence, whenever dealing with an extremal instance $\mathcal{I}$, we assume without loss of generality that $\mathcal{I}$ is left-extremal.*

Remark 4.14. *Consider a left-extremal instance $\mathcal{I} = (A, h, \mathcal{P}, k)$. Let $\sigma = (h_1, h_2, \ldots, h_m)$ be a left-extremal ordering on the houses. So for every agent $a \in A$, there exists $i(a) \in [n]$ such that $\mathcal{P}(a) = \left\{h_1, h_2, \ldots, h_{i(a)}\right\}$. Notice that σ imposes an ordering on the agents, say $\sigma_A = (a_1, a_2, \ldots, a_n)$ so that $\mathcal{P}(a_1) \subseteq \mathcal{P}(a_2) \subseteq \cdots \subseteq \mathcal{P}(a_n)$. Notice also that we can check* *in polynomial time if a given instance is left-extremal, and if so, find a left-extremal ordering σ on the houses and then the ordering σ_A on the agents. So, whenever dealing with a left-extremal instance $\mathcal{I}$, we assume without loss of generality that σ and σ_A are given. Whenever we talk about, for example, the "first/last house," we always mean the first/last house with respect to the ordering σ. Same with the ordering σ_A and the "first/last agent."*

4.4 Optimal House Allocation

In this section, we deal with the Optimal House Allocation problems, where the goal is to minimize the number of envious agents. We start by discussing the cases for which we have polynomial time algorithms for Optimal House Allocation.

4.4.1 Polynomial Time Algorithms for OHA

We first prove that $[0/1]$-Optimal House Allocation is polynomial-time solvable on instances with an extremal structure.

Theorem 4.15. *There is a polynomial-time algorithm for $[0/1]$-Optimal House Allocation when the agent valuations have an extremal interval structure.*

Proof. In light of Remark 4.13, it suffices to prove for the case when agent valuations are left extremal. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ denote an instance of HA with left extremal valuations, which is irreducible with respect to Reduction Rule 1, 2 and 3.

We first make the following claim:

Claim 4.16. *Given an instance of $[0/1]$- Optimal House Allocation when the agent valuations are left extremal, there exists an optimal allocation where the set of allocated houses form an interval.*

Proof. Suppose we start with an optimal allocation Φ under which the set of allocated houses does not form an interval. Let h_u be an unallocated house such that h_l and h_r are allocated, where $l < u < r$. We will show that we can allocate h_u instead of h_r without any increase in the envy, and hence iteratively convert Φ to another optimal allocation Φ' under which the set of allocated houses do form an interval.

Suppose no agent is envious of the allocation of h_r. Then, we can allocate h_u to the recipient of h_r without increasing the envy. Indeed, everyone who values h_r also values h_u because of the interval structure, so the envy of the recipient does not increase. Since no one envies the allocation of h_r, no one should become newly envious of the allocation of h_u as well, since it lies to the left of h_r. If any agent is envious in this re-allocation because of h_u, it must be already envious due to the allocation of h_l which lies to the left of h_u. So, whoever gets h_r can get h_u without any increase in the envy.

On the other hand, suppose we have an agent a who is envious due to the allocation of h_r. Since a values h_r, due to the interval structure, she must value h_u as $u < r$. If we swap $\Phi(a)$ with h_u, then a ceases to be envious. Also, note that allocation of h_u does not create any new envious agent—indeed if an agent a' was not envious before h_u was allocated, it means either she got a house she likes or her interval ended before h_l. In either case, the allocation of h_u can not be a cause of envy to her. This contradicts the fact that we started with an optimal allocation. Also, notice that all the dummy houses (if any), in the reduced instance, lie to the extreme right of all the houses that are valued. The allocation of a dummy house can always be done respecting the interval property, as agents do not distinguish between any two dummy houses. □

Based on the above claim, our algorithm Alg works as follows. It enumerates over all possibilities of the first allocated house in the optimal allocation where the set of allocated houses forms an interval. There are at most $m - n$ such choices, in particular, the first $m - n$ houses. For each such h_i, Alg chooses the next consecutive n houses to be allocated. This reduces the instance to the one where $m = n$ and by Proposition 4.1, this can be done in polynomial time.

To see the correctness, in one direction, if under any iteration i, the allocation Φ constructed by Alg has at most k envious agents, then it returns Yes, and the allocation is the witness that $\mathcal{I}$ is a Yes instance. In the other direction, if $\mathcal{I}$ is a Yes instance, then there exists an optimal allocation Opt such that $\kappa^{\#}(\text{Opt}) \leq k$ and it allocates the consecutive houses, say $[h_i, h_{i+n}]$. The algorithm captures this optimal allocation when it iterates over the house h_i, and hence returns Yes. □

We now present our algorithms for the cases when (a)every agent approves exactly one house and (b)every house is approved by almost two agents. To that end, we first present a reduction rule that will be helpful in the following two results.

Reduction Rule 4. *For every house h such that $d(h) = 1$, we allocate h to $N(h)$.*

The above reduction is safe.[4] Indeed, if $d(h) = 1$, then no one except exactly one agent a values

[4]This rule is subsumed by Reduction Rule 2 when $|H \setminus D| \geq |A|$.

the house h, and so allocating h to a does not generate any envy.

Theorem 4.17. *There is a polynomial-time algorithm for* $[0/1]$-*Optimal House Allocation when every agent approves exactly one house.*

Proof. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ denote an instance of HA such that every agent approves exactly one house. Consider the associated preference graph $G = (A \cup H, E)$. We first apply the reduction rules 1, 2 and 4. In the reduced instance, we have $|A| \leq |H| \leq 2(|A| - 1)$ and the degree of any house h in G is strictly greater than 1, that is, $d(h) > 1$.

We order the houses in the reduced instance in $H \setminus D$ as $\{h_1, h_2, \ldots h_t\}$ such that $d(h_1) \geq d(h_2) \geq \ldots d(h_t)$, where t denotes the number of houses in $H \setminus D$. (Note that $t \geq 1$, since $|D| \leq n - 1$.) The last $|A| - |D|$ of these houses are then allocated to their neighbors (chosen arbitrarily). The remaining $|D|$ agents get the $|D|$ dummy houses.

We argue the correctness of the above algorithm Alg. We show that the number of envious agents under the allocation Φ returned by Alg is equal to the number of envious agents under some optimal allocation. We first claim that in any optimal allocation Opt, all the dummy houses must be allocated. Suppose $h^\star \in D$ is unallocated. Consider a house h_i allocated to some agent a who approves it. This causes $d(h_i) - 1$ agents to have envy, and since every agent likes exactly one house, there is no way these $d(h_i) - 1$ agents can become envy-free once h_i is allocated. Since $h^\star$ is available, we can allocate it to a and add h_i to the set of unallocated houses. This decreases the number of envious agents by $d(h_i) - 1$, without adding to the envy of anyone else. This contradicts the fact that we started with an optimal allocation. Therefore all the dummy houses must be allocated.
Now, consider the set S of houses in $H \setminus D$ that are allocated by Opt. Let T be the set of such houses allocated under Φ. S and T both contain exactly $|A| - |D|$ many houses from $H \setminus D$. Consider any two house h_i and h_j in $H \setminus D$ such that $d(h_i) > d(h_j)$. Note that the allocation of h_i creates $d(h_i) - 1$ many envious agents, strictly greater than the number of envious agents created by the allocation of h_j, which is $d(h_j) - 1$. Therefore, S contains $|A| - |D|$ houses of the least degree from among the set of $H \setminus D$ houses. Since T also contain the $|A| - |D|$ houses of the least degree, therefore $\kappa^\#(\Phi) = \kappa^\#(\text{Opt})$. □

We now state the algorithm when the house degree is bounded, while the agent degree is not.

Theorem 4.18. *There is a polynomial-time algorithm for* $[0/1]$-*Optimal House Allocation when every house is approved by at most two agents.*

Proof. Consider the agent-house bipartite preference graph G. We first apply the reduction rules 1, 2, 3 and 4. Then, we apply the following reduction rule.

Reduction Rule 5. *For every cycle $C = (h_1, a_1, h_2, a_2, \ldots a_i, h_1)$ in G, allocate h_i to the agent a_i.*

The above reduction is safe. Indeed, since $d(h) \leq 2 \; \forall \; h \in G$, a house that participates in a cycle C is valued *only* by the agents participating in the same cycle. Since G is a bipartite graph, all cycles are of even length, so the number of agents in $C(= |A_C|)$, is equal to the number of houses in C $(= |H_C|)$. This implies that every agent in C can get a house she values from C. This does not make any agent outside C envious.

Consider the remaining graph G after the application of reduction rules. Note that G is a collection of trees, that is, $G = T_1, T_2, \ldots T_r$. Let D be the set of dummy houses in G. We now describe the algorithm Alg. First, sort the trees in increasing order of sizes, that is, $|T_1| \leq |T_2|, \ldots \leq |T_r|$. For each tree T_i, we root T_i at some leaf agent a_j such that $d(a_j) = 1$. (After Reduction Rule 4, such a leaf agent in T_i always exists, since every leaf house is allocated to the parent agent by the above reduction rule.) Let $n_1, n_2, \ldots n_r$ be the number of agents in the trees $T_1, T_2, \ldots T_r$ respectively. Let j be the first index such that $(n_1 + n_2 \ldots + n_j) + (r - j) > |D|$. Then, $(n_1 + n_2 \ldots + n_{j-1}) + (r - (j-1)) \leq |D|$. We allocate all the agents in $T_1, T_2, \ldots T_{j-1}$ a dummy house. Then, the number of dummy houses that remain is $|D| - (n_1 + n_2 + \ldots + n_{j-1}) \geq (r - (j-1))$. For the remaining trees, $T_j, \ldots T_r$, we match the non-root agents to their parent house and allocate the root agent a dummy house. There are $r - j + 1$ root agents and there are at least so many dummy houses.

Note that only the root agents in the tree $T_j, T_{j+1}, \ldots T_r$ are the envious ones — indeed, such a root agent gets a dummy house but a house valued by her is allocated. Notice that every non-root agent in the above trees got a house she valued, and all the agents in the trees $T_1, T_2, \ldots T_{j-1}$ got a dummy house, and none of the houses they valued got allocated.

Therefore, under Alg, the number of envious agents = $r - j + 1$ = total amount of envy.

Claim 4.19. *Alg returns an allocation that minimizes the number of envious agents.*

Proof. Let G be the reduced graph after the reduction rules. Note that $|H_{T_i}| = |A_{T_i}| - 1$ for any $T_i \in G$. Indeed, root T_i at a house say h_1. Let $N(h_1) = a_1$ and a_2. Since there is no leaf house, and $d(h) = 2 \; \forall \; h$, every $h \neq h_1$ is a parent to a unique agent $a \neq a_1, a_2$. This gives a bijection from $H_i \setminus h_1$ to $A_i \setminus \{a_1, a_2\}$.

Let Opt be the allocation that minimizes the number of envious agents in the reduced graph

G. Let l and l' be the number of envious agents under Opt and the allocation returned by Alg respectively.

If $l = l'$, we are done. We will now show that $l \not< l'$. Suppose, $l < l'$. Notice that every tree T_i has either no envious agents (in the case when every agent in T_i gets a dummy house) or exactly one envious agent (in the case when houses from T_i are allocated). Indeed, every tree is rooted at a house vertex which is of degree at most two. Also, if any of the leaf vertex in the tree T_i is a house h, then by the structure of the tree, there is only one agent a (namely, the parent of h in the tree T_i) that values the house h. Therefore, h can be safely assigned to a, without causing envy to anyone else. Therefore, we can assume that every leaf vertex in the tree is an agent. This implies that in any tree, there are more agents than houses. That is why, in case, any of the house from T_i is allocated, then at least one agent will be envious. The only case where no agent from T_i is envious is when none of the houses from T_i are allocated and every agent in T_i gets a dummy house.

Now consider that there are at least two envious agents in a tree T_i under OPT. Then, houses from T_i must have been allocated. But if so, then consider the following re-allocation where every agent vertex receives its parent house vertex and one agent from the two that are incident to the root house h, say a, receives a dummy house. Note that a is now the only envious agent in T_i. The agents, say in T_j, ($\neq T_i$) who previously received houses from T_i are re-allocated houses from T_j. If T_j has more houses than agents, then everyone is allocated a house they value from T_j. Otherwise, agents receive their respective parent houses, and either one of the two agents incident to the root vertex ends up with a dummy house.

The agents not in T_i who might have previously received houses from T_i under OPT are now allocated their respective parent houses, and either one of the two agents incident to the root vertex ends up with a dummy house. Note that such agents who are not in T_i do not value the houses in T_i and therefore do not differentiate between them and dummy houses. This implies that there is an allocation where at most one agent is envious in every tree. Therefore, if there are at least two envious agents in a tree T_i in the OPT, then one of them can be made envy-free without an increase in the envy of any other agent. This would contradict the fact that the allocation was OPT to begin with.

So, under Opt, there are exactly l trees that have exactly one envious agent. Since these l agents did not get what they value, they must have got a dummy house because of the reduction rules 4, 5, and the fact that every tree has one agent more than the number of houses. Also, since there are no envious agents in the remaining $r - l$ trees (say, $T_{i_1}, T_{i_2}, \ldots T_{i_{r-l}}$), all the agents in these trees must have got dummy houses each. Therefore the number of dummy houses allocated under Opt are:

$$(n_{i_1} + n_{i_2} + \dots n_{i_{r-l}}) + l \geq (n_1 + n_2 + \dots + n_{r-l}) + l \tag{4.1}$$

Note that $l' = r - j + 1$ where j is the first index such that $n_1 + n_2 \dots + n_j + r - j > |D|$. Now since $l < l'$, therefore,

$$(r - l) > (r - l') \Rightarrow (r - l) \geq (r - l' + 1) = j \Rightarrow (r - l) \geq j \tag{4.2}$$

This implies that there are at least j envy-free trees under Opt. Suppose WLOG there are exactly j envy-free trees under Opt. Then the number of dummy houses allocated under Opt: $(n_{i_1} + n_{i_2} + \dots n_{i_j}) + (r - j) \geq (n_1 + n_2 + \dots + n_j) + (r - j) > |D|$ which is a contradiction. □

This concludes the argument. □

4.4.2 Hardness Results for OHA

In this section, we design three different reductions that will establish the hardness of Optimal House Allocation (both 0/1-OHA and $[\succeq]$-OHA), including in restricted settings. The first is a parameterized reduction from Clique to 0/1-OHA.

Theorem 4.20. *[0/1]-Optimal House Allocation is* NP*-complete when every house is approved by at most three agents.*

Proof. We sketch a reduction from Clique. Let $\mathcal{I} := (G = (V, E); k)$ be an instance of Clique. Let $|V| = n'$ and $|E| = m'$. We first describe the construction of an instance of OHA based on G:

- We introduce a house h_e for every edge $e \in E$. We call these the *edge houses.*
- We also introduce $m' + n' - \binom{k}{2}$ dummy houses.
- We introduce an agent a_v for every vertex $v \in V$ and an agent a_e for every edge $e \in E$. We refer to these as the vertex and edge agents, respectively.
- Every edge agent a_e values the house h_e.
- For every vertex $v \in V$, the vertex agent a_v values the edge house h_e if and only if e is

incident to v in G.

Note that in this instance of OHA, there are $m' + n'$ agents, m' edge houses and $m' + n' - \binom{k}{2}$ dummy houses, and every house is approved by at most three agents. We let k be the target number of envious agents. This completes the construction of the reduced instance. We now turn to a proof of equivalence.

The forward direction.

Let $S \subseteq V$ be a clique of size k. Then consider the allocation Φ that assigns h_e to a_e for all $e \in E(G[S])$ and dummy houses to all other agents. Note that no edge agent is envious in this allocation, and the only vertex agents who are envious are those that correspond to vertices of S. Since $|S| = k$, this establishes the claim in the forward direction.

The reverse direction.

Let Φ be an allocation that has at most k envious agents. We say that Φ is nice if every edge house is either unallocated by Φ or allocated to an edge agent who values it. If Φ is not nice to begin with, notice that it can be converted to a nice allocation by a sequence of exchanges that does not increase the number of envious agents. In particular, suppose h_e is allocated to an agent $a \neq a_e$. Then we obtain a new allocation by swapping the houses h_e and $\Phi(a_e)$ between agents a and a_e. This causes at least one envious agent to become envy-free (i.e., a_e) and at most one envy-free agent to become envious (i.e., a), and therefore the number of envious agents does not increase. Based on this, we assume without loss of generality, that Φ is a nice allocation.

Now note that any nice allocation Φ is compelled to assign n dummy houses among the n' vertex agents, and this leaves us with $m' - \binom{k}{2}$ dummy houses that can be allocated among m' edge agents. Therefore, at least $\binom{k}{2}$ edge agents are assigned edge houses.

Let $F \subseteq E$ be the subset of edges corresponding to edge agents who were assigned edge houses by Φ. Let $S \subseteq V$ be the set of vertices in the span of F, that is:

$$S := \bigcup_{e=(u,v) \in F} \{u, v\}.$$

Note that for all $v \in S$, a_v is envious with respect to Φ, since—by the definitions we have so far—a_v valued an assigned house and was assigned a dummy house. Since Φ admits at most k envious agents, we have that $|S| \leq k$. However, S is also the span of at least $\binom{k}{2}$ distinct edges, so it is also true[5] that $|S| \geq k$. Therefore, we conclude that $|S| = k$, and since every edge in F belongs to $G[S]$ and F has $\binom{k}{2}$ edges, it follows that S is a clique of size k in G. This concludes the argument in the reverse direction. □

Theorem 4.21. *[0/1]-Optimal House Allocation is* NP*-complete even when every agent approves at most two houses.*

Proof. Let $\mathcal{I} := (G = (V, E); k)$ be an instance of Clique where G is a d-regular graph. Let $|V| = n'$ and $|E| = m'$. (The problem of finding a clique restricted to regular graphs is also NP-complete (Mathieson and Szeider, 2012).) We first describe the construction of an instance of OHA based on G:

- We introduce a house h_v for every vertex $v \in V$. We call these the *vertex houses.*
- We also introduce $m' + n' - k$ dummy houses.
- We introduce an agent a_v for every vertex $v \in V$ and an agent a_e for every edge $e \in E$. We refer to these as the vertex and edge agents, respectively.
- Every vertex agent a_v values the house h_v.
- For every edge $e = (u, v)$ in E, the edge agent a_e values the houses h_u and h_v.

Note that in this instance of OHA, there are n' vertex houses, $m' + n' - k$ dummy houses, and every agent approves at most two houses. We let $kd - \binom{k}{2}$ be the target number of envious agents. This completes the construction of the reduced instance. We now turn to a proof of equivalence.

The forward direction.

Let $S \subseteq V$ be a clique of size k. Then consider the allocation Φ that assigns h_v to a_v for all $v \in S$ and dummy houses to all other agents. Note that no vertex agent is envious in this allocation, and the only edge agents that are envious are those that correspond to edges in G that have at least one of their endpoints in S. The total number of distinct edges incident on S is at most kd, but since $G[S]$ induces a clique, the exact number of distinct edges incident on S is $kd - \binom{k}{2}$, and this establishes the claim in the forward direction.

[5]Intuitively, a smaller set of vertices would not be able to accommodate as many edges; and specifically a subset of at most $k-1$ vertices can account for at most $\binom{k-1}{2} < \binom{k}{2}$ edges.

The reverse direction.

Let Φ be an allocation that has at most $kd - \binom{k}{2}$ envious agents. We say that Φ is nice if every vertex house is either unallocated by Φ or allocated to a vertex agent who values it. If Φ is not nice to begin with, notice that it can be converted to a nice allocation by a sequence of exchanges that does not increase the number of envious agents. In particular, suppose h_v is allocated to an agent $a \neq a_v$. Then we obtain a new allocation by swapping the houses h_v and $\Phi(a_v)$ between agents a and a_v. This causes at least one envious agent to become envy-free (i.e., a_v) and at most one envy-free agent to become envious (i.e., a), and therefore the number of envious agents does not increase. Based on this, we assume without loss of generality, that Φ is a nice allocation.

Now note that any nice allocation Φ is compelled to assign m' dummy houses among the m' edge agents, and this leaves us with $n' - k$ dummy houses that can be allocated among n' vertex agents. Therefore, at least k vertex agents are assigned vertex houses. We may assume that exactly k vertex agents are assigned vertex houses—indeed if more than k vertex agents are assigned vertex houses, these houses can be swapped with dummy houses without increasing the number of envious agents, and we perform these swaps until we run out of dummy houses to swap with. Finally, observe that the set of k vertex agents (say, S) who are assigned vertex houses induce a clique of size k in G. Indeed, if not:

$$\#\text{ of edges incident on } S = kd - |E(G[S])| > kd - \binom{k}{2}.$$

The claim follows from the fact that every edge incident on S in G corresponds to a distinct edge agent who experiences envy in the reduced instance, and this would contradict our assumption that the number of agents envious with respect to Φ was at most $kd - \binom{k}{2}$. □

We now exhibit the hardness of approximation of finding the maximum number of envy-free agents. Computing a maximum balanced biclique is known to be hard to approximate within a factor of $n^{(1-\gamma)}$ for any constant $\gamma > 0$, where n is the number of vertices (Manurangsi, 2018), assuming the Small Set Expansion Hypothesis (Raghavendra and Steurer, 2010). The reduction below shows that any $f(n)$ approximation to the maximum number of envy-free agents (where n is the number of agents) implies a $2(1+\epsilon) \cdot f(|L|)$ approximation to the maximum balanced biclique where $|L|$ is the size of the left bi-partition. The argument is similar to that of Kamiyama et al. (2021), but we produce here for the sake of completeness.

Theorem 4.22. *If the Small Set Expansion Hypothesis holds, then finding the maximum number of envy-free agents for instances with weak rankings is hard to approximate within a factor of $n^{1-\gamma}$ for any constant $\gamma > 0$.*

Proof. We will first show that there is a polynomial-time reduction that takes an instance $G = (L, R, E)$ of maximum balanced biclique and produces an allocation instance such that if there is a bi-clique of size at least k in G, then there exists an allocation Φ in the reduced allocation instance such that Φ admits at least k envy-free agents. On the other hand, given any allocation Φ with at least k envy-free agents in the reduced instance, then there is a bi-clique of size $k/2$ in G, which can be found in polynomial time. Additionally, the number of agents n in the reduced instance is exactly $|L|$.

Consider an instance $G = (L, R, E)$ of maximum balanced biclique such that $L = \{b_1, \ldots, b_{n'}\}$ and $R = \{c_1, \ldots, c_{m'}\}$, we create an allocation instance as follows: an agent a_i for each $b_i \in L$ and a house h_j for each $c_j \in R$. We also have n' additional starred houses $\{h_1^\star, \ldots, h_{n'}^\star\}$. This amounts to a total of n' agents and $m' + n'$ houses. An agent a_i ranks the houses as follows:

$$\succeq_{a_i} = \begin{cases} h_j \succ h_l, \text{ if } j > l; (b_i, c_j) \notin E \text{ and } (b_i, c_l) \notin E \\ h_j \succ h_l, \text{ if } (b_i, c_j) \notin E \text{ and } (b_i, c_l) \in E \\ h_j = h_l, \text{ if } (b_i, c_j) \in E \text{ and } (b_i, c_l) \in E \\ h_j^\star \succ h_l^\star, \text{ if } j > l \\ h_i \succ h_j^\star, \; \forall\, i \in [m'], j \in [n'] \end{cases}$$

Essentially, every agent ranks his non-neighbor houses first in a fixed strict order, then ranks all his neighbors equally, and at last, ranks all the additional starred houses, again, in some fixed strict order.

We now establish the correctness of the above reduction. In the forward direction, suppose there is a k-sized balanced biclique in G, consisting of vertices $\{b_{i_1}, b_{i_2}, \ldots b_{i_k}\}$ from L and $\{c_{i'_1}, \ldots c_{i'_k}\}$ from R. Then consider the following assignment:

$$\Phi(a_i) = \begin{cases} h_{i'_l} \text{ if } i = i_l \text{ for some } l \in [k] \\ h_{i^\star} \text{ otherwise} \end{cases}$$

That is, the agents corresponding to the clique vertices get one of their adjacent houses, again corresponding to the clique vertices. Note that each of the $\{a_{i_1}, \ldots a_{i_k}\}$ rank the houses

$\{h_{i'_1}, \ldots h_{i'_k}\}$ equally, so they do not envy each other. Also, they do not envy the remaining agents who get $h^\star_j$, as they all have the ranking $h_{i_l} \succ h^\star_j$ for all i and j.

For the reverse direction, suppose there exists an assignment Φ of houses such that there are k envy-free agents. We claim that there is a balanced biclique of size $\frac{k}{2}$ in G. Let A_{EF} denote the set of envy-free agents with respect to Φ. Notice that none of the agents in A_{EF} owns a starred house under Φ. If not, suppose some $a \in A_{EF}$ gets $h^\star_j$. Consider another agent a' in A_{EF} such that $a' \neq a$. In case a' gets a house h_j for some $j \in [m']$, then a would be envious as she ranks all h_j better than the starred houses. Else, if a' gets a starred house, say $h^\star_l$, then depending on whether $j > l$ or not, one of these two agents experiences envy, as they both rank all the starred houses in the same manner. Hence, $\Phi(A_{EF}) \subseteq \{h_1, \ldots, h_{m'}\}$.

Now, let the k houses under $\Phi(A_{EF})$ be $\{h_{j_1}, \ldots h_{j_k}\}$ such that $j_1 < j_2 < \ldots < j_k$. Let $a_{i_l} = \Phi^{-1}(h_{j_l})$. Then consider the set S, consisting of the vertices in L corresponding to the first half agents in $\Phi^{-1}(A_{EF})$ and the set T, consisting of the vertices in R, corresponding to the remaining half house vertices in $\Phi(A_{EF})$. Precisely, $S = \{b_{i_1}, \ldots b_{i_{\frac{k}{2}}}\}$ and $T = \{c_{j_{\frac{k}{2}+1}}, \ldots, c_{j_k}\}$. We claim that S and T together induce a biclique of size $\frac{k}{2}$ in G. Suppose $(b_{i_l}, c_{j_{l'}}) \notin E$ for some $b_{i_l} \in S$ and some $c_{j_{l'}} \in T$. By the choice of S and T, notice that $l < l'$. This means that a_{i_l} ranks the house $h_{j_{l'}}$ strictly better that h_{j_l}, therefore envies $a_{i_{l'}}$, who is assigned the house $h_{j_{l'}}$. In order for a_{i_l} to not envy the owner of $h_{j_{l'}}$, it must be the case that $(b_{i_l}, c_{j_{l'}}) \in E$. Therefore, S and T form a biclique of size $\frac{k}{2}$. This completes the correctness of the reduction.

With this reduction in hand, we will now argue that a polynomial time $f(n)$-approximation to finding the maximum number of envy-free agents gives a $2(1+\epsilon)f(|L|)$-approximation algorithm for maximum balanced biclique. But assuming the Small Set Expansion Hypothesis, this contradicts the inapproximability result for maximum balanced biclique by Manurangsi (2018).

Consider an instance G of maximum balanced biclique. We first construct the reduced allocation instance $\mathcal{I}$. Suppose there is a polynomial time $f(n)$-approximation to finding the maximum number of envy-free agents that outputs an allocation Φ for $\mathcal{I}$. We first find the biclique (S, T) corresponding to Φ in G. Let $\beta = 2(1 + \frac{1}{\epsilon})$. We enumerate all subsets of size 2β in G and consider the largest biclique (S', T') (of size at most 2β). We then output the largest of the two bicliques (S, T) and (S', T'). Let Opt be the size of optimal biclique in G. If $Opt \leq \beta \cdot f(|L|)$, then we have that the brute force biclique (S', T') has size at least $\frac{Opt}{f(|L|)}$ and we are done. Otherwise, suppose $Opt > \beta \cdot f(|L|)$. By the reduction and

$f(n)$-approximation, we have that the number of envy-free agents under Φ is at least $\frac{Opt}{f(n)} = \frac{Opt}{f(|L|)}$. Then we have the size of the biclique (S, T) as

$$|S| = |T| = \lfloor \frac{Opt}{2f(|L|)} \rfloor > \frac{Opt}{2f(|L|)} - 1 > \frac{Opt}{2f(|L|)} - \frac{Opt}{\beta f(|L|)} = \frac{(\beta - 1)Opt}{\beta(2f(|L|))} = \frac{Opt}{2f(|L|)(1+\epsilon)}$$

Therefore, we get an approximation ratio of $2f(|L|)(1+\epsilon)$. This settles the claim. □

Note that $[\succeq]$-OHA remains NP-Complete from as a corollary to Theorem 4.20, which establishes the hardness for binary valuations, a specific case of rankings with ties. We mention here that the complexity of $[\succ]$-OHA remains open.

4.4.3 Parameterized Results for OHA

We now turn to the parameterized complexity of OHA and first present a linear kernel parameterized by the number of agents.

Theorem 4.23. *$[0/1]$-Optimal House Allocation admits a linear kernel parameterized by the number of agents. In particular, given an instance of $[0/1]$-Optimal House Allocation, there is a polynomial time algorithm that returns an equivalent instance of $[0/1]$-Optimal House Allocation with at most twice as many houses as agents.*

Proof. It suffices to prove the safety of Reduction Rule 2. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ denote an instance of HA with parameter k. Further, let $\mathcal{I}' = (H' := H \setminus X, A' := A \setminus Y, \mathcal{P}'; k)$ denote the reduced instance corresponding to $\mathcal{I}$. Note that the parameter for the reduced instance is k as well.

If $\mathcal{I}$ is a Yes-instance of OHA, then there is an allocation $\Phi : A \to H$ with at most k envious agents. By Claim 4.9, we may assume that Φ is a good allocation. This implies that the projection of Φ on $H' \cup A'$ is well-defined, and it is easily checked that this gives an allocation with at most k envious agents in $\mathcal{I}'$.

On the other hand, if $\mathcal{I}'$ is a Yes-instance of OHA, then there is an allocation $\Phi' : A' \to H'$ with at most k envious agents. We may extend this allocation to $\Phi : A \to H$ by allocating the houses in Y to agents in X along the expansion M, that is:

$$\Phi(a) = \begin{cases} \Phi'(a) & \text{if } a \notin X, \\ M(a) & \text{if } a \in X. \end{cases}$$

Since all the newly allocated houses are not valued by any of the agents outside X and all agents in X are envy-free with respect to Φ, it is easily checked that Φ also has at most k envious agents. □

The following results follow from using the algorithm described in Proposition 4.1 after guessing the allocated houses, which adds a multiplicative overhead of $\binom{m}{n} \leq 2^m$ to the running time.

Corollary 4.24. *$[0/1]$-Optimal House Allocation is fixed-parameter tractable when parameterized either by the number of houses or the number of agents. In particular, $[0/1]$-Optimal House Allocation can be solved in time $O^\star(2^m)$.*

Corollary 4.25. *$[\succ]$-Optimal House Allocation is fixed-parameter tractable when parameterized by the number of houses and can be solved in time $O^\star(2^m)$.*

The next two results follow from Theorem 4.20.

Corollary 4.26. *$[0/1]$-Optimal House Allocation is W[1]-hard when parameterized by the solution size, i.e, the number of envious agents, even when every house is approved by at most three agents.*

Corollary 4.27. *$[\succeq]$-Optimal House Allocation is W[1]-hard when parameterized by the solution size, i.e., the number of envious agents.*

We now show that $[0/1]$-OHA is fixed-parameter tractable when parameterized by the number of types of houses or the number of types of agents. To that end, we formulate $[0/1]$-OHA as an integer linear program and then invoke Theorem 4.6.

Theorem 4.28. *$[0/1]$-Optimal House Allocation is fixed-parameter tractable when parameterized by either the number of houses types or the number of agents types.*

Consider an instance $\mathcal{I} = (A, h, \mathcal{P}, k)$ of $[0/1]$-OHA. Recall that we use n^* to denote the number of types of agents in $\mathcal{I}$ and the m^* to denote the number of types of houses in $\mathcal{I}$. With a slight abuse of notation, for $i \in [n^*]$, we use $\mathcal{P}(i)(\subseteq [m^*])$ to denote the set of types of houses that each agent of type i values. Also, for $i \in [n^*], j \in [m^*]$ and an allocation $\Phi : A \to H$, let $A(\Phi, i, j) \subseteq A$ be the set of agents of type i who receive a house of type j under Φ.

Observation 4.29. *Consider an instance $\mathcal{I} = (A, h, \mathcal{P}, k)$ of $[0/1]$-OHA. Then, (1) $n^* \leq 2^{m^*}$ and (2) $m^* \leq 2^{n^*}$. To see (1), for each $i \in [n^*]$, the agents of type i are uniquely identified by the types of houses they prefer, and the number of distinct choices for the types of houses is at most 2^{m^*}. Similarly, to see (2), observe that for each $j \in [m^*]$, the houses of type j are uniquely identified by* *the types of agents who prefer houses of type j; and the number of distinct choices for the types of agents is at most 2^{n^*}.*

Lemma 4.30. *Consider an instance $\mathcal{I} = (A, h, \mathcal{P}, k)$ of [0/1]-OHA and an allocation $\Phi : A \to H$. Consider any fixed pair (i, j), where $i \in [n^*]$ and $j \in [m^*]$. Then, for every $a \in A$ and for every $a', a'' \in A(\Phi, i, j)$, either both a' and a'' envy a, or neither a' nor a'' envies a.*

Proof. Consider $a \in A$ and $a', a'' \in A(\Phi, i, j)$. First, if $j \in \mathcal{P}(i)$, then neither a' nor a'' envies any agent. So, assume that $j \notin \mathcal{P}(i)$. Let $\Phi(a)$ be of type ℓ, for some $\ell \in [m^*]$. If $\ell \notin \mathcal{P}(i)$, then, neither a' nor a'' envies a. If $\ell \in \mathcal{P}(i)$, then both a' and a'' envy a. □

We now move to formulate the [0/1]-OHA problem as an integer linear program (ILP). The number of variables in our ILP will be $\mathcal{O}(n^* \cdot m^*)$. By Observation 4.29, the number of variables will be bounded separately by both $\mathcal{O}(n^* \cdot 2^{n^*})$ and $\mathcal{O}(m^* \cdot 2^{m^*})$. That is, the number of variables will be bounded separately by both the number of house types and the number of agent types. The result will then follow from Theorem 4.6.

Consider an instance $\mathcal{I} = (A, h, \mathcal{P}, k)$ of [0/1]-Optimal House Allocation. For each $i \in [n^*]$ and $j \in [m^*]$, let n_i be the number of agents of type i and m_j the houses of type j. To define our ILP, we introduce the following variables. For each $i \in [n^*]$, $j \in [m^*]$, we introduce four variables: x_{ij}, z_{ij}, d_{ij} and d'_{ij}. Here, x_{ij} and z_{ij} are integer variables, and d_{ij} and d'_{ij} are binary variables. The semantics of the variables are as follows. (1) We want x_{ij} to be the number of agents of type i who receive houses of type j. Equivalently, we want x_{ij} to be the number of houses of type j that are allocated to agents of type i. (2) And we want z_{ij} to be the number of *envious* agents of type i who receive houses of type j. By Lemma 4.30, either all type i agents who receive type j houses are envious, or none of them is envious. That is, we must have either $z_{ij} = x_{ij}$ or $z_{ij} = 0$. Notice that type i agents who receive type j houses are envious if and only if $j \notin \mathcal{P}(i)$, and for some $j' \in \mathcal{P}(i)$, at least one house of type j' has been allocated (to, say, an agent of type i' for some $i' \in [n^*]$). That is, we must have $z_{ij} > 0$ if and only if $j \notin \mathcal{P}(i)$, $x_{ij} > 0$ and $x_{i',j'} > 0$ for some $i' \in [n^*]$ and $j' \in \mathcal{P}(i)$. Hence, for $j \notin \mathcal{P}(i)$, either $x_{ij} = 0$ or $z_{ij} > 0$ if $\sum_{i' \in [n]} \sum_{j' \in \mathcal{P}(i)} x_{i'j'} > 0$. (3) The variables d_{ij}, d'_{ij} are only dummy variables that we use to enforce the "either or" constraints.

We now formally describe our ILP. Minimize $\sum_{i \in [n^*]} \sum_{j \in [m^*]} z_{ij}$ subject to the constraints in Table 4.3.

For convenience, we name this ILP $\texttt{P1}(\mathcal{I})$, and denote the set of variables of $\texttt{P1}(\mathcal{I})$ by $Var(\texttt{P1}(\mathcal{I}))$ and the optimum value of $\texttt{P1}(\mathcal{I})$ by $opt(\texttt{P1}(\mathcal{I}))$. Constraint C1.$i$ ensures that for

(C1.i).	$\sum_{j\in[m^*]} x_{ij} = n_i$	$\forall i \in [n^*]$
(C2.j).	$\sum_{i\in[n^*]} x_{ij} \leq m_j$	$\forall j \in [m^*]$
(C3.a.i.j).	$x_{ij} \leq nd'_{ij}$	
(C3.b.i.j).	$\sum_{i'\in[n^*]} \sum_{j'\in\mathcal{P}(i)} x_{i'j'} \leq nmz_{ij} + nm(1-d'_{ij})$	$\forall i \in [n^*], j \in [m^*] \setminus \mathcal{P}(i)$
(C3.c.i.j).	$z_{ij} \leq n_i \sum_{i'\in[n^*]} \sum_{j'\in\mathcal{P}(i)} x_{i'j'}$	
(C4.a.i.j).	$z_{ij} \leq n_i d_{ij}$	
(C4.b.i.j).	$x_{ij} - z_{ij} \leq n_i(1-d_{ij})$	$\forall i \in [n^*], j \in [m^*]$
(C4.c.i.j).	$z_{ij} \leq x_{ij}$	
(C5.i.j).	$z_{ij} = 0$	$\forall i \in [n^*], j \in \mathcal{P}(i)$
(C6.a.i.j).	$x_{ij} \geq 0$	
(C6.b.i.j).	$z_{ij} \geq 0$	
(C6.c.i.j).	$d_{ij} \in \{0,1\}$	$\forall i \in [n^*], j \in [m^*]$
(C6.d.i.j).	$d'_{ij} \in \{0,1\}$	

Table 4.3: The constraints of the ILP $\texttt{P1}(\mathcal{I})$.

each $i \in [n^*]$, the number of houses allocated to agents of type i is exactly n_i. In other words, all agents of type i receive houses. Constraint C2.j ensures that for each $j \in [m^*]$, the number of houses of type j that are allocated does not exceed m_j. For $i \in [n^*]$ and $j \in [m] \setminus \mathcal{P}(i)$, constraints C3.a.$i$.$j$ and C3.b.i.j together ensure that depending on whether $d'_{ij} = 0$ or $d'_{ij} = 1$, we have either $x_{ij} = 0$ or $z_{ij} > 0$ if $\sum_{i'\in[n^*]} \sum_{j'\in[\mathcal{P}(i)]} x_{i'j'} > 0$. Constraint C3.c.$i$.$j$ ensures that if $x_{i'j'} = 0$ for every $i' \in [n^*], j' \in [m] \setminus \mathcal{P}(i)$, then $z_{ij} = 0$. Constraints C4.a.i.j-C4.c.i.j together ensure that depending on $d_{ij} = 0$ or $d_{ij} = 1$, we have either $z_{ij} = 0$ or $z_{ij} = x_{ij}$.

To establish the correctness of $\texttt{P1}(\mathcal{I})$, we prove the following two claims.

Claim 4.31. *For any allocation $\Phi : A \rightarrow H$, there exists a feasible solution $f_\Phi : Var(P1(\mathcal{I})) \rightarrow \mathbb{Z}$ for $P1(\mathcal{I})$ such that $\kappa^\#(\Phi) = \sum_{i\in[n^*]} \sum_{j\in[m^*]} f_\Phi(z_{ij})$.*

Claim 4.32. *For every optimal solution $f : Var(P1(\mathcal{I})) \rightarrow \mathbb{Z}$ for $P1(\mathcal{I})$, there exists an allocation $\Phi_f : A \rightarrow H$ such that $\kappa^\#(\Phi_f) = \sum_{i\in[n^*]} \sum_{j\in[m^*]} f(z_{ij})$.*

Remark 4.33. *Notice that Claim 4.31, in fact, proves that ILP $P1(\mathcal{I})$ is always feasible as there is always an allocation. Moreover, for an allocation Φ, since $0 \leq \kappa^\#(\Phi) \leq n$, and since $\kappa^\#(\Phi) = \sum_{i\in[n^*]} \sum_{j\in[m^*]} f_\Phi(z_{ij})$, we can conclude that $P1(\mathcal{I})$ has a bounded solution. Therefore,* opt$(P1(\mathcal{I}))$ *is well-defined.*

Assuming Claims 4.31 and 4.32 hold, we now prove the following claim.

Claim 4.34. *We have* $\kappa^{\#}(\mathcal{I}) = \mathsf{opt}(P1(\mathcal{I}))$.

Proof. To prove the claim, we will prove that (1) $\kappa^{\#}(\mathcal{I}) \geq \mathsf{opt}(\mathtt{P1}(\mathcal{I}))$ and (2) $\mathsf{opt}(\mathtt{P1}(\mathcal{I})) \geq \kappa^{\#}(\mathcal{I})$.

To prove (1), consider an optimal allocation $\Phi : A \to H$. That is, $\kappa^{\#}(\mathcal{I}) = \kappa^{\#}(\Phi)$. By Claim 4.31, we have $\kappa^{\#}(\Phi) = \sum_{i \in [n^*]} \sum_{j \in [m^*]} f_\Phi(z_{ij}) \geq \mathsf{opt}(\mathtt{P1}(\mathcal{I}))$, where f_Φ is as defined in Claim 4.31. We thus have $\kappa^{\#}(\mathcal{I}) \geq \mathsf{opt}(\mathtt{P1}(\mathcal{I}))$.

Now, to prove (2), consider an optimal solution f for $\mathtt{P1}(\mathcal{I})$. That is,

$\mathsf{opt}(\mathtt{P1}(\mathcal{I})) = \sum_{i \in [n^*]} \sum_{j \in [m^*]} f(z_{ij})$. By Claim 4.32, we have

$\mathsf{opt}(\mathtt{P1}(\mathcal{I})) = \sum_{i \in [n^*]} \sum_{j \in [m^*]} f(z_{ij}) = \kappa^{\#}(\Phi_f) \geq \kappa^{\#}(\mathcal{I})$, where Φ_f is as defined in Claim 4.32. We thus have $\mathsf{opt}(\mathtt{P1}(\mathcal{I})) \geq \kappa^{\#}(\mathcal{I})$. □

We are now ready to prove Theorem 4.28.

Proof of Theorem 4.28. Given an instance $\mathcal{I}$ of $[0/1]$-OHA, observe that we can construct the ILP $\mathtt{P1}(\mathcal{I})$in polynomial time. The number of variables in $\mathtt{P1}(\mathcal{I})$ is bounded by $4n^*m^*$. The number of constraints in $\mathtt{P1}(\mathcal{I})$ is also bounded by $\mathcal{O}(n^*m^*)$. The maximum value of any coefficient or constant term in $\mathtt{P1}(\mathcal{I})$ is bounded by nm. So, $\mathtt{P1}(\mathcal{I})$ can be encoded using $\mathrm{poly}(n^*, m^*) \cdot \mathcal{O}(\log(nm))$ bits. The result then follows from Theorem 4.6 and Claim 4.34. □

We now only have to prove Claims 4.31 and 4.32.

Proof of Claim 4.31. Consider any allocation $\Phi : A \to H$. Recall that $A(\Phi, i, j)$ is the set of agents of type i who receive houses to type j under Φ.

We define a solution $f_\Phi : Var(\mathtt{P1}(\mathcal{I})) \to \mathbb{Z}$ for $\mathtt{P1}(\mathcal{I})$ as follows. For each $i \in [n^*], j \in [m^*]$, we set (1) $f_\Phi(x_{ij}) = |A(\Phi, i, j)|$; (2) $f_\Phi(z_{ij}) = f_\Phi(x_{ij})$ if there exists an envious agent $a \in A(\Phi, i, j)$ and $f_\Phi(z_{ij}) = 0$ otherwise; (3) $f_\Phi(d_{ij}) = 0$ if $f_\Phi(z_{ij}) = 0$ and $f_\Phi(d_{ij}) = 1$ otherwise; and (4) $f_\Phi(d'_{ij}) = 0$ if $f_\Phi(x_{ij}) = 0$ and $f_\Phi(d'_{ij}) = 1$ otherwise.

To see that f_Φ satisfies all the constraints of $\mathtt{P1}(\mathcal{I})$, observe first that $|A(\Phi, i, j)|$ is the number of agents of type i who receive houses of type j under Φ; equivalently, $|A(\Phi, i, j)|$ is the number of houses of type j that have been allocated to agents of type i under Φ. Therefore, $\sum_{j \in [m^*]} |A(\Phi, i, j)| = n_i$ and $\sum_{i \in [n^*]} |A(\Phi, i, j)| \leq m_j$. Hence, (1) $\sum_{j \in [m^*]} f_\Phi(x_{ij}) = \sum_{j \in [m^*]} |A(\Phi, i, j)| = n_i$ and $\sum_{i \in [n^*]} f_\Phi(x_{ij}) = \sum_{i \in [n^*]} |A(\Phi, i, j)| \leq m_j$.

Thus f_Φ satisfies constraints C1.i and C2.j for every $i \in [n^*]$ and $j \in [m^*]$. Also, note that $\sum_{i' \in [n^*]} \sum_{j' \in \mathcal{P}(i)} f_\Phi(x_{i'j'}) \leq \sum_{i' \in [n^*]} \sum_{j' \in [m^*]} |A(\Phi, i, j)| \leq n$.

Now, consider $i \in [n^*]$, $j \in [m^*]$. Suppose first that $f_\Phi(x_{ij}) = 0$. Then, by the definition of f_Φ, we have $|A(\Phi, i, j)| = 0$, and hence $f_\Phi(z_{ij}) = 0$, which implies $f_\Phi(d_{ij}) = 0$; and $f_\Phi(x_{ij}) = 0$ implies that $f_\Phi(d'_{ij}) = 0$. Note that in this case, f_Φ satisfies all the constraints. In particular, constraint C3.b.$i.j$ is satisfied because $f_\Phi(z_{ij}) = f_\Phi(d'_{ij}) = 0$ implies that the right side of constraint C3.b.$i.j$ is exactly equal to nm, and the left side of the constraint is at most n. Since $f_\Phi(x_{ij}) = f_\Phi(z_{ij}) = 0$, all the other constraints corresponding to the pair (i, j) are also satisfied.

Suppose now that $f_\Phi(x_{ij}) > 0$. Then by the definition of f_Φ, we have $f_\Phi(d'_{ij}) = 1$. There are two possibilities: (1) $f_\Phi(z_{ij}) = 0$ and (2) $f_\Phi(z_{ij}) > 0$.

Assume first that $f_\Phi(z_{ij}) = 0$. Again, by the definition of f_Φ, we have $f_\Phi(d_{ij}) = 0$. Notice that this choice of values satisfies all the constraints corresponding to the pair (i, j), except possibly C3.b.$i.j$. To see that C3.c.$i.j$ is also satisfied, assume that $j \in [m^*] \setminus \mathcal{P}(i)$. Since $f_\Phi(z_{ij}) = 0$, the definition of f_Φ implies that the agents of type i who receive houses of type j are not envious. Hence we can conclude that none of the houses of type j' have been allocated under Φ, for any $j' \in \mathcal{P}(i)$. That is, $A(\Phi, i', j') = \emptyset$ for every $i' \in [n^*]$ and $[j'] \in \mathcal{P}(i)$. Thus the left side of constraint C3.b.$i.j$ is 0; and the right side is 0 as well, as $f_\Phi(z_{ij}) = 1$ and $f_\Phi(d'_{ij}) = 1$.

Finally, assume that $f_\Phi(z_{ij}) > 0$. Then, by the definition of f_Φ, we have $f_\Phi(z_{ij}) = f_\Phi(x_{ij}) = |A(\Phi, i, j)| \leq n_i$ and $f_\Phi(d_{ij}) = 1$. Notice that this choice of values satisfies constraints C4.a.$i.j$-C4.c.$i.j$. From the definition of f_Φ, we can also conclude that the agents of type i who receive houses of type j under Φ are envious, which implies that $j \notin \mathcal{P}(i)$ and $|A(\Phi, i', j')| = f_\Phi(x_{i'j'}) > 0$ for some $i' \in [n^*]$, $j' \in \mathcal{P}(i)$. Thus the right side of constraint C3.c.$i.j$ is strictly positive; and since $f_\Phi(z_{ij}) \leq n_i$, constraint C3.c.$i.j$ is satisfied. Since $f_\Phi(d'_{ij}) = 1$ and $f_\Phi(x_{ij}) \leq n$, constraint C3.a.$i.j$ is satisfied. Finally, constraint C3.b.$i.j$ is satisfied because its left side is at most n, and the right side is at least nm as $f_\Phi(z_{ij}) > 0$. Notice also that in this case $\mathtt{P1}(\mathcal{I})$ does not contain constraint C5.$i.j$ as $j \notin \mathcal{P}(i)$.

We have thus shown that f_Φ satisfies all the constraints of $\mathtt{P1}(\mathcal{I})$.

Consider $i \in [n^*], j \in [m^*]$. Suppose that $A(\Phi, i, j) \neq \emptyset$. By Lemma 4.30, either all agents in $A(\Phi, i, j)$ are envious or none of them are. By the definition of f_Φ, we have $f_\Phi(z_{ij}) = f_\Phi(x_{ij}) = |A(\Phi, i, j)|$ if and only if the agents in $A(\Phi, i, j)$ are envious; and $f_\Phi(z_{ij}) = 0$ otherwise. Therefore, the number of envious agents under Φ, $\kappa^\#(\Phi) = \sum_{i \in [n^*]} \sum_{j \in [m^*]} f_\Phi(z_{ij})$. □

To prove Claim 4.32, we first prove two preparatory claims below. In both these claims, $f : Var(\texttt{P1}(\mathcal{I})) \to \mathbb{Z}$ is a feasible solution for $\texttt{P1}(\mathcal{I})$.

Claim 4.35. *For every $i \in [n^*], j \in [m^*]$, either $f(z_{ij}) = 0$ or $f(z_{ij}) = f(x_{ij})$.*

Proof. Fix $i \in [n^*], j \in [m^*]$. Since constraint C6.c.$i.j$ is satisfied, we have $f(d_{ij}) \in \{0, 1\}$. If $f(d_{ij}) = 0$, then constraints C4.a.$i.j$ implies that $f(z_{ij}) \leq 0$ and then constraint C6.b.$i.j$ implies that $z_{ij} = 0$. Instead, if $f(d_{ij}) = 1$, then, constraint C4.b.$i.j$ implies that $f(x_{ij}) - f(z_{ij}) \leq 0$, which implies that $f(x_{ij}) \leq f(z_{ij})$. But then constraint C4.c.$i.j$ implies that $f(z_{ij}) = f(x_{ij})$. □

Claim 4.36. *For every $i \in [n^*], j \in [m^*]$, $f(z_{ij}) > 0$ if and only if $j \notin \mathcal{P}(i)$, $f(x_{ij}) > 0$ and $f(x_{i'j'}) > 0$ for some $i' \in [n^*]$ and $j' \in \mathcal{P}(i)$.*

Proof. Fix $i \in [n^*], j \in [m^*]$. Assume first that $f(z_{ij}) > 0$. Then, constraint C5.$i.j$ implies that $j \notin \mathcal{P}(i)$. Constraint C4.c.$i.j$ implies that $f(x_{ij}) \geq f(z_{ij}) > 0$. Now, if $f(x_{i'j'}) = 0$ for every $i' \in [n], j' \in \mathcal{P}(i)$, then constraint C3.c.i.j would imply that $f(z_{ij}) \leq 0$, which is not possible. Hence, we have $f(x_{i'j'}) > 0$ for some $i' \in [n], j' \in \mathcal{P}(i)$.

Assume now that $j \notin \mathcal{P}(i)$, $f(x_{ij}) > 0$, and $f(x_{i'j'}) > 0$ for some $i' \in [n^*]$ and $j' \in \mathcal{P}(i)$. Since $f(x_{ij}) > 0$, constraint C3.a.$i.j$ implies that $f(d'_{ij}) > 0$. By constraint C6.d.$i.j$, we then have $f(d'_{ij}) = 1$. Then, constraint C3.b.$i.j$,along with the fact that $f(x_{i'j'}) > 0$ for some $i' \in [n^*]$ and $j' \in \mathcal{P}(i)$, implies that $0 < \sum_{i' \in [n^*]} \sum_{j' \in \mathcal{P}(i)} f(x_{i'j'}) \leq nmf(z_{ij})$, which, then implies that $f(z_{ij}) > 0$. □

Proof of Claim 4.32. Given $f : Var(\texttt{P1}(\mathcal{I})) \to \mathbb{Z}$, we define $\Phi_f : A \to H$ as follows. For each $i \in [n^*], j \in [m^*]$, we allocate $f(x_{ij})$ houses of type j to agents of type i (one house per agent). Thus, we have $|A(\Phi_f, i, j)| = f(x_{ij})$. Notice that as f satisfies constraints C1.i and C2.j, the allocation Φ_f is valid.

To complete the proof of the claim, we only need to prove that for every $i \in [n^*], j \in [m^*]$, the number of envious agents of type i who received houses of type j under Φ_f is exactly equal to $f(z_{ij})$. Fix $i \in [n^*], j \in [m^*]$. By Lemma 4.30, either all agents in $A(\Phi_f, i, j)$ are envious or none of them is envious. Notice that the agents in $A(\Phi_f, i, j)$ (if they exist) are envious if and only if $|A(\Phi_f, i, j)| > 0$, $j \notin \mathcal{P}(i)$ and $A(\Phi_f, i', j') \neq \emptyset$ for some $i' \in [n^*], j' \in [m^*] \setminus \mathcal{P}(i)$. That is, the agents in $A(\Phi_f, i, j)$ are envious if and only if $f(x_{ij}) = |A(\Phi_f, i, j)| > 0, j \notin \mathcal{P}(i)$ and $f(x_{i'j'}) = |A(\Phi_f, i', j')| > 0$ for some $i' \in [n^*], j' \in [m^*] \setminus \mathcal{P}(i)$. On the other hand, by Claim 4.36, $f(z_{ij}) > 0$ if and only if $f(x_{ij}) > 0$, $j \notin \mathcal{P}(i)$ and $f(x_{i'j'}) > 0$ for some

$i' \in [n^*], j' \in [m^*] \setminus \mathcal{P}(i)$. We can thus conclude that the agents in $A(\Phi_f, i, j)$ are envious if and only if $f(z_{ij}) > 0$. By Claim 4.35, we also have $f(z_{ij}) = 0$ or $f(z_{ij}) = f(x_{ij})$. This implies that the agents in $A(\Phi_f, i, j)$ are envious if and only if $f(z_{ij}) = f(x_{ij}) > 0$. By Lemma 4.30, if the agents in $A(\Phi_f, i, j)$ are envious, then the number of envious agents in $A(\Phi_f, i, j)$ is exactly $|A(\Phi_f, i, j)| = f(x_{ij}) = f(z_{ij})$. We thus have $\kappa^{\#}(\Phi_f) = \sum_{i \in [n^*]} \sum_{j \in [m^*]} f(z_{ij})$. □

4.5 Egalitarian House Allocation

In this section, we deal with the EHA problems, where the goal is to minimize the maximum envy experienced by any agent. We first discuss the polynomial time algorithms for EHA.

4.5.1 Polynomial Time Algorithms for EHA

Theorem 4.37. *There is a polynomial-time algorithm for* [0/1]-*Egalitarian House Allocation when the agent valuations have an extremal interval structure.*

Proof. In light of Remark 4.13, it suffices to prove for the case when agent valuations are left-extremal. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ be an instance of EHA with left-extremal valuations. Consider an envious agent a_l in the allocation, who approves the interval $[h_1, h_i]$. Note that if $\mathcal{I}$ is a Yes instance, no more than k houses can be allocated from $[h_1, h_i]$, else the envy experienced by a_l will be more than k.

Based on this observation, the algorithm works as follows. We order the agents in the increasing order of the length of their intervals, that is, a appears before a' if $\mathcal{P}(a) \subseteq \mathcal{P}(a')$. We guess the last envious agent a_l. There are at most n such guesses. Suppose the interval that a_l approves is $[h_1, h_j]$ for some $j \in [m]$. Since any valid allocation is bound to allocate at most k houses from $[h_1, h_j]$, we iterate over the number of houses i allocated from $[h_1, h_j]$. (Note that $i \leq k$.) For each such i, we construct the associated preference graph $G_i = (A_{\geq l+i+1} \cup H_{\geq j+1}; E)$ where $A_{\geq l+i+1} = \{a_{l+i+1}, a_{l+i+2}, \ldots a_n\}$ and $H_{\geq j+1} = \{h_{j+1}, \ldots h_m\}$. We then find a matching in G_i that saturates $A_{\geq l+i+1}$, and if it does not exist, then the iteration is discarded. The algorithm then constructs the allocation Φ_i as follows. The matched houses under M_i are allocated to the matched agents. The first i houses from $[h_1, h_j]$ are allocated to the first i agents whose interval ends after h_j, particularly, to $\{a_{l+1}, a_{l+2}, \ldots, a_{l+i}\}$. The remaining unmatched houses are assigned arbitrarily to the unmatched agents. If the number of remaining houses is less than the remaining agents, then the iteration is discarded and the algorithm moves to the next iteration. The algorithm returns Yes if for some iteration i, Φ_i is a complete allocation, that is, every agent gets a house under Φ_i. Else, it returns No.

To argue the correctness of the above algorithm, we show that if $\mathcal{I}$ is a Yes instance, if and only if for some iteration i, there exists a complete allocation Φ_i. In the forward direction, suppose $\mathcal{I}$ is a Yes instance. There must exist some optimal allocation Opt such that the maximum envy under Opt is at most k. Among all the envious agents under Opt, consider the agent a whose interval $[h_1, h]$ is longest, that is, for all envious agents $a', \mathcal{P}(a') \subseteq \mathcal{P}(a)$. Suppose Opt allocates k' houses from $[h_1, h]$. Then consider the iteration in A that iterates over the agent $a_l = a$ and $i = k'$. (This implies $[h_1, h_j] = [h_1, h]$.) Note that under Opt, all the agents whose interval ends after h must get a house that they like. This implies that for all the agents $\{a_{l+1}, a_{l+2}, \ldots a_n\}$, there exists a house that they like and that can be allocated to them. As Opt allocates exactly $k'(= i)$ houses from $[h_1, h_j]$, at most i of the agents among $\{a_{l+1}, a_{l+2}, \ldots a_n\}$ can get a house they like from $[h_1, h_j]$. For the remaining ones, there must exist at least one house $[h_{j+1}, h_m]$ that they like and can be allocated to them, which implies that there must exist a matching saturating the said agents under the said iteration i. Therefore all the agents $\{a_{l+1}, a_{l+2}, \ldots a_n\}$ get a house that they like under Φ_i. Also, since OPT is a complete allocation, so there are enough remaining houses for the agents $\{a_1, \ldots a_l\}$ that can be allocated to them, once $\{a_{l+1}, a_{l+2}, \ldots a_n\}$ get what they value. This implies that there are enough houses remaining under Φ_i as well to be allocated to the remaining agents. (Note that even if all the agents $\{a_1, \ldots a_l\}$ are envious, their envy is bounded by at most $i \leq k$.) This implies that the allocation Φ_i constructed in the iteration i is indeed complete, and the algorithm returns Yes.

In the reverse direction, suppose there exists a complete allocation Φ_i. Note that all the envious agents under Φ_i do not value any house outside $[h_1, h_j]$ and exactly $i(\leq k)$ houses are allocated from $\{h_1, h_j\}$. Therefore the maximum envy is at most k, and $\mathcal{I}$ is a Yes instance. This concludes the proof. □

We now state the algorithm for the restricted setting when every agent approves exactly one house.

Theorem 4.38. *There is a polynomial-time algorithm for* $[0/1]$*-Egalitarian House Allocation when every agent approves exactly one house.*

Proof. Notice that when every agent approves exactly one house, then in any allocation, the maximum envy $\kappa^{\dagger}(\Phi)$ is bounded by 1. Given an instance $\mathcal{I} := (A, H, \mathcal{P}; k)$ of HA, we invoke the Hall's Violators Algorithm by Gan et al. (2019). The above algorithm checks whether there is an envy-free allocation and returns one if it exists. Therefore, if it returns an allocation, then

$\kappa^\dagger(\Phi) = 0$, else it has to be at least 1 under any allocation. □

4.5.2 Hardness Results for EHA

In this section, we establish the hardness of EHA for both binary and linear orders, even under restricted settings.

Theorem 4.39. *There is a polynomial-time reduction from* Independent Set *to* $[0/1]$*-*Egalitarian House Allocation*, where every agent approves at most two houses and every house is approved by at most ten agents. This reduction shows that* $[0/1]$*-*Egalitarian House Allocation *is* NP*-complete even when the target value of the maximum envy is one.*

Proof. Let $\mathcal{I} := (G = (V, E); k)$ be an instance of Independent Set. Let $|V| = n'$ and $|E| = m'$. Since Independent Set is known to be NP-complete even on subcubic graphs, we assume that G is subcubic, i.e, that the degree of every vertex in G is at most three. We construct an instance of $[0/1]$-EHA as follows.

- We introduce a house h_v for every vertex v in V. We call these the *vertex houses.*
- We also introduce $(3m' + n') - k$ dummy houses.
- We introduce an agent a_v for every vertex $v \in V$ and three agents a_e^1, a_e^2, a_e^3 for every edge $e \in E$. We refer to these as the vertex and edge agents, respectively. We refer to the three agents corresponding to a single edge as the *cohort* around e.
- All three edge agents corresponding to the edge $e = (u, v)$ value the houses h_u and h_v.
- A vertex agent a_v values the house h_v.

Note that in this instance of EHA, there are $3m' + n'$ agents, n' vertex houses, and $3m' + n' - k$ dummy houses, and every agent approves at most two houses. We set the maximum allowed envy at one, that is, the reduced instance asks for an allocation where every agent envies at most one other agent. This completes the construction of the reduced instance. We now turn to a proof of equivalence.

The forward direction.

Let $S \subseteq V$ be an independent set of size k. Consider the allocation Φ that assigns h_v to a_v for all $v \in S$ and dummy houses to all remaining agents. Note that all vertex agents are envy-free under this allocation. If an edge agent envies two other agents, it must be two agents who received vertex houses, since recipients of dummy houses are never a cause for envy. So suppose an edge agent a_e° who approves, say, the houses h_u and h_v is envious, then it implies that both a_u and a_v belong to S while it is also true that $(u, v) \in E$; contradicting our assumption that $G[S]$ induces an independent set. This completes the argument in the forward direction.

The reverse direction.

Let Φ be an allocation with respect to which every agent envies at most one other agent. First, note that if $e = (u, v)$ is an edge in G, then note that Φ can allocate at most one of the houses h_u and h_v. Suppose not. Then notice that at least one of the agents among the cohort of edge agents corresponding to e, i.e, a_e^1, a_e^2 and a_e^3 envy the two agents who were assigned the houses h_u and h_v.

We say that Φ is nice if every vertex house is either unallocated by Φ or allocated to a vertex agent who values it. If Φ is not nice to begin with, notice that it can be converted to a nice allocation by a sequence of exchanges that does not increase the maximum envy of the allocation without changing the set of allocated houses. In particular, suppose h_v is allocated to an agent $a \neq a_v$. Then we obtain a new allocation by swapping the houses h_v and $\Phi(a_v)$ between agents a and a_v. This causes a_v to become envy-free. If a is a vertex agent, then she experiences the same envy as before. If a is an edge agent, then we have two possible scenarios. Suppose a did not value h_v: then the amount of envy experienced by a is either the same or less than, the amount of envy she had in the original allocation. On the other hand, if a did value h_v, then a envies a_v in the new allocation but nobody else, since h_u is unallocated in Φ.

Based on this, we assume without loss of generality, that Φ is a nice allocation. Now note that any nice allocation Φ is compelled to assign $3m'$ dummy houses among the $3m'$ edge agents, and this leaves us with $n' - k$ dummy houses that can be allocated among n' vertex agents. Therefore, at least k vertex agents are assigned vertex houses. Finally, observe that the corresponding vertices induce an independent set in G of size at least k. Indeed, this follows from a fact that we have already argued: any two houses corresponding to vertices that are endpoints of an edge cannot be both allocated by Φ. This concludes the argument in the reverse direction. □

Theorem 4.40. *There is a polynomial-time reduction from Multi-Colored Independent Set to $[\succ]$-Egalitarian House Allocation. This reduction shows that $[\succ]$-Egalitarian House Allocation is* NP*-complete even when the target value of the maximum envy is one.*

Proof. Let $\mathcal{I} := (G = (V, E); k)$ be an instance of Multi-Colored Independent Set with color classes $V_1, \dots, V_k$. By adding dummy global vertices if required, we assume that $|V_i| = n \geq 3$ for all $i \in [k]$, and denote the vertices in V_i by $\{u^i_1, \dots, u^i_n\}$. The global dummy vertices added in any partition V_i are adjacent to all the other vertices not in V_i.

We now construct an instance of $[\succ]$-EHA as follows.

- We introduce a house for every vertex v in V. We call these the *vertex houses* and they are denoted by:
$$H_V := \{h^1_1, \dots, h^1_n\} \uplus \dots \uplus \{h^k_1, \dots, h^k_n\}.$$
- We introduce a house h_e for every edge e in E. We call these the *edge houses* and denote this set of houses by H_E.
- For each $1 \leq i \leq k$, and for every $1 \leq p \neq q \leq n$, we introduce three houses $h^1_{[i;p,q]}$, $h^2_{[i;p,q]}$ and $h^3_{[i;p,q]}$. We call them *special houses*.
- We also introduce $k \cdot (n-1)$ additional houses, denoted by:
$$H_D := \{d^1_1, \dots, d^1_{n-1}\} \uplus \dots \uplus \{d^k_1, \dots, d^k_{n-1}\}.$$
- We introduce an agent for every vertex $v \in V$, denoted by $a_{[1,i]}, \dots, a_{[n,i]}$ for $1 \leq i \leq k$. We call them *vertex agents*.
- We introduce an agent a_e for every $e \in E$. We call them *edge agents*.
- For each $1 \leq i \leq k$, and for every $1 \leq p \neq q \leq n$, we introduce three agents $a_{[i;p,q,1]}$, $a_{[i;p,q,2]}$ and $a_{[i;p,q,3]}$. We call them *guards*.

For a set X, we use $\overline{X}$ to denote an arbitrary order on the set X. Also, for a fixed order, say σ, we use $[[\sigma]]_i$ to denote the order σ rotated i times. For example, $[[x \succ y \succ z]]_2 = z \succ x \succ y$. Note that σ is an order over n elements, then $[[\sigma]]_n = \sigma$. We also use H to denote the set of houses in the reduced instance. We are now ready to describe the preferences of the agents.

- An edge agent corresponding to an edge $e = (u^i_p, u^j_q)$ that has endpoints in V_i and V_j (with $i < j$) ranks the houses as follows:
$$\succ_e : h^i_p \succ h^j_q \succ h_e \succ \overline{H \setminus \{h^i_p, h^j_q, h_e\}}$$
- For each $1 \leq i \leq k$, for every $1 \leq p \neq q \leq n$, and $\ell \in \{1, 2, 3\}$ the guard agent $a_{[i;p,q,\ell]}$ has the following preference:

$$\succ_{[i;p,q,\ell]}: h_p^i \succ h_q^i \succ h_{[i;p,q]}^\ell \succ \overline{H \setminus \{h_p^i, h_q^i, h_{[i;p,q]}^\ell\}}.$$

- For $1 \leq i \leq k$, the vertex agents $a_{[1,i]}, \ldots, a_{[n,i]}$ rank the houses as follows

$$\begin{aligned}
\succ_{[1,i]} &: h_1^i \succ [[d_1^i \succ d_2^i \succ d_3^i \succ \cdots \succ d_{n-1}^i]]_0 && \succ \overline{H \setminus \{h_1^i, d_1^i, d_2^i, d_3^i, \ldots, d_{n-1}^i\}} \\
\succ_{[2,i]} &: h_2^i \succ [[d_1^i \succ d_2^i \succ d_3^i \succ \cdots \succ d_{n-1}^i]]_0 && \succ \overline{H \setminus \{h_2^i, d_1^i, d_2^i, d_3^i, \ldots, d_{n-1}^i\}} \\
&\vdots \\
\succ_{[j,i]} &: h_j^i \succ [[d_1^i \succ d_2^i \succ d_3^i \succ \cdots \succ d_{n-1}^i]]_{j-2} && \succ \overline{H \setminus \{h_j^i, d_1^i, d_2^i, d_3^i, \ldots, d_{n-1}^i\}} \\
&\vdots \\
\succ_{[n-1,i]} &: h_{n-1}^i \succ [[d_1^i \succ d_2^i \succ d_3^i \succ \cdots \succ d_{n-1}^i]]_{n-3} && \succ \overline{H \setminus \{h_{n-1}^i, d_1^i, d_2^i, d_3^i, \ldots, d_{n-1}^i\}} \\
\succ_{[n,i]} &: h_n^i \succ [[d_1^i \succ d_2^i \succ d_3^i \succ \cdots \succ d_{n-1}^i]]_{n-2} && \succ \overline{H \setminus \{h_n^i, d_1^i, d_2^i, d_3^i, \ldots, d_{n-1}^i\}}
\end{aligned}$$

Note that in this instance of EHA, there are $k \cdot (n-1)$ extra houses. We set the maximum allowed envy at one, that is, the reduced instance asks for an allocation where every agent envies at most one other agent. This completes the construction of the reduced instance. We now turn to a proof of equivalence.

The forward direction.

Let $S \subseteq V$ be a multicolored independent set. Let $s : [k] \to [n]$ be such that:

$$S = \{u_{s(1)}^1, u_{s(2)}^2, \ldots, u_{s(k)}^k\}.$$

We now describe an allocation Φ based on S. First, we let $\Phi(a_e) = h_e$ for all $e \in E$. Also, for each $1 \leq i \leq k$, and for every $1 \leq p \neq q \leq n$, we have $\Phi(a_{[i;p,q,1]}) = h_{[i;p,q]}^1$, $\Phi(a_{[i;p,q,2]}) = h_{[i;p,q]}^2$ and $\Phi(a_{[i;p,q,3]}) = h_{[i;p,q]}^3$.

Now, for the vertex agents corresponding to the vertices of V_i, we have the following if $s(i) \geq 2$:

$$\Phi(a_{[j,i]}) = \begin{cases} d^i_j & \text{if } 1 \leq j < s(i), \\ h^i_j & \text{if } j = s(i), \\ d^i_{j-1} & \text{if } s(i) < j \leq n, \end{cases}$$

and if $s(i) = 1$, then we proceed as follows instead:

$$\Phi(a_{[j,i]}) = \begin{cases} h^i_1 & \text{if } j = 1, \\ d^i_{j-1} & \text{if } j > 1, \end{cases}$$

Note that every house corresponding to a vertex not in S remains unallocated in Φ, implying, in particular, that exactly one vertex from each color class corresponds to an allocated vertex house in Φ. We now argue that every agent envies at most one other agent with respect to this allocation.

First, consider an edge agent a_e corresponding to an edge $e = (u^i_p, u^j_q)$ that has endpoints in V_i and V_j (with $i < j$). Recall that a_e gets her third-ranked house with respect to Φ. Since at most one of h^i_p or h^j_q is allocated with respect to Φ, we have that a_e envies at most one agent.

Similarly, every guard agent receives the special house that she ranks third, and at most one of the two top-ranked houses is allocated in Φ, since both of the top-ranked houses belong to the same color class by construction.

Now we turn to the vertex agents. It is easily verified that every vertex agent gets a house that they rank first (if they correspond to a vertex from S), second, or third. Thus, vertex agents corresponding to vertices in S are envy-free, and all other vertex agents envy at most one other agent. (If an agent $a_{[j,i]}$ receives its third-ranked house d^i_j, then notice that the first-ranked house h^i_j remains unallocated. Therefore, $a_{[j,i]}$ envies at most one other agent who might have received d^i_{j-1}, its second-ranked house.) This concludes the proof in the forward direction.

The reverse direction.

Let Φ be an allocation for the reduced instance where every agent envies at most one other agent. We make a series of claims about the allocation Φ that allows us to observe that Φ has the following properties: it allocates exactly one vertex house from the houses corresponding to vertices in a common color class, and further, it allocates such a house to a vertex agent. Such vertex houses are then easily seen to correspond to a multi-colored independent set in

G: indeed, if not, then the pair of adjacent vertices would correspond to an edge agent who is envious of at least two vertex agents, contradicting our assumption about Φ.

We first observe that we cannot allocate more than one house from among vertex houses corresponding to vertices from a common color class of G.

Claim 4.41. *Let $i \in [k]$ be arbitrary but fixed. Among the vertex houses $\{h_1^i, \ldots, h_n^i\}$, Φ leaves at least $(n-1)$ houses unallocated; in other words, Φ allocates at most one house from among these houses.*

Proof. Suppose not, and in particular, suppose Φ allocates the houses h_p^i and h_q^i for some $1 \leq p \neq q \leq n$. Then at least one of the three guard agents $a_{i;p,q,\ell}$ for $\ell \in \{1,2,3\}$ will envy the two agents who receive these two houses, which contradicts the assumption that every agent envies at most one other agent in the allocation Φ. □

The total number of houses are nk vertex houses, $(n-1)k$ dummy houses, m edge houses and $3k\binom{n}{2}$ special houses. The total number of agents are nk vertex agents, m edge agents, and $3k\binom{n}{2}$ guards. Now since at least $(n-1)k$ vertex houses remain unallocated by Claim 4.41, to ensure that every agent gets a house, at least one vertex house from each of the n vertex partitions must be allocated. This implies that exactly one house is allocated from one color class.

Since $(n-1)k$ vertex houses are unallocated, all the remaining houses must be allocated by Φ. In particular, all additional houses are allocated, and we use this fact in our next claim.

Claim 4.42. *If a vertex house is allocated in Φ, then it is assigned to a vertex agent.*

Proof. Suppose not. Let i be such that the vertex house under consideration corresponds to a vertex from V_i. Note that the additional houses $\{d_1^i, \ldots, d_{n-1}^i\}$ can be allocated among at most $(n-1)$ of the agents corresponding to the vertices in V_i. Therefore, there is at least one agent a among the agents $a_{[j,i]}, j \in [n]$ who does not receive any of the houses among her top n-ranked houses. Further, she was also not assigned her top-ranked house, by the assumption we made for the sake of contradiction. Since Φ allocates all the houses in $\{d_1^i, \ldots, d_{n-1}^i\}$, the agent a envies at least $(n-1) \geq 2$ agents, and this is the desired contradiction. □

The previous two claims imply the desired structure on Φ, and as argued earlier, the subset of vertices corresponding to allocated vertex houses induces a multicolored independent set in G, and this concludes the argument in the reverse direction. □

4.5.3 Parameterized Results for EHA

In this section, we present a linear kernel for [0/1]-EHA and discuss the tractable and the hard cases in the parameterized setting.

Theorem 4.43. *[0/1]-Egalitarian House Allocation admits a linear kernel parameterized by the number of agents. In particular, given an instance of [0/1]-Egalitarian House Allocation, there is a polynomial time algorithm that returns an equivalent instance of [0/1]-Egalitarian House Allocation with at most twice as many houses as agents.*

Proof. It suffices to prove the safety of Reduction Rule 2. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ denote an instance of HA with parameter k. Further, let $\mathcal{I}' = (H' := H \setminus X, A' := A \setminus Y, \mathcal{P}'; k)$ denote the reduced instance corresponding to $\mathcal{I}$. Recall that the parameter for the reduced instance is k as well.

If $\mathcal{I}$ is a Yes-instance of EHA, then there is an allocation $\Phi : A \rightarrow H$ with maximum envy at most k. By Claim 4.10, we may assume that Φ is a good allocation. This implies that the projection of Φ on $H' \cup A'$ is well-defined, and it is easily checked that this gives an allocation with maximum envy at most k

On the other hand, if $\mathcal{I}'$ is a Yes-instance of EHA, then there is an allocation $\Phi' : A' \rightarrow H'$ with maximum envy k. We may extend this allocation to $\Phi : A \rightarrow H$ by allocating the houses in Y to agents in X along the expansion M, that is:

$$\Phi(a) = \begin{cases} \Phi'(a) & \text{if } a \notin X, \\ M(a) & \text{if } a \in X. \end{cases}$$

Since all the newly allocated houses are not valued by any of the agents outside X and all agents in X are envy-free with respect to Φ, it is easily checked that Φ also has maximum envy k. □

The following results follow from using the algorithm described in Proposition 4.2 after guessing the allocated houses, which adds a multiplicative overhead of $\binom{m}{n} \leq 2^m$ to the running time.

Corollary 4.44. *[0/1]-Egalitarian House Allocation is fixed-parameter tractable when parameterized either by the number of houses or the number of agents. In particular, [0/1]-Egalitarian House Allocation can be solved in time $O^\star(2^m)$.*

Corollary 4.45. *[$\succ$]-Egalitarian House Allocation is fixed-parameter tractable when* *parameterized by the number of houses and can be solved in time $O^\star(2^m)$.*

The next two results follow respectively from Theorem 4.39 and Theorem 4.40 respectively.

Corollary 4.46. *$[0/1]$-Egalitarian House Allocation is para-NP-hard when parameterized by the solution size, i.e, the maximum envy, even when every agent approves at most two houses.*

Corollary 4.47. *$[\succ]$-Egalitarian House Allocation is para-NP-hard when parameterized by the solution size, i.e., the maximum envy.*

We now formulate EHA as an integer linear program, as in the case of $[0/1]$-OHA and establish the fixed-parameter traceability parameterized by the number of house types or agent types. The number of variables in our ILP will be $\mathcal{O}(n^* \cdot m^*)$, where n^* is the number of types of agents and m^* the number of types of houses. Again, by Observation 4.29, the number of variables will then be bounded separately by $2^{\mathcal{O}(m^*)}$ and $2^{\mathcal{O}(n^*)}$.

Theorem 4.48. *$[0/1]$-Egalitarian House Allocation is fixed-parameter tractable when parameterized either by the number of house types or the number of agent types.*

Given an instance $\mathcal{I} = (A, h, \mathcal{P}, k)$ of $[0/1]$-Egalitarian House Allocation, we define an ILP $\texttt{P2}(\mathcal{I})$ that encodes the instance $\mathcal{I}$. The ILP $\texttt{P2}(\mathcal{I})$ is very similar to $\texttt{P1}(\mathcal{I})$ with exactly two distinctions. (1) The ILP $\texttt{P2}(\mathcal{I})$ has all the variables of $\texttt{P1}(\mathcal{I})$. In addition, $\texttt{P2}(\mathcal{I})$ has an integer variable w that encodes the maximum envy experienced by an agent. (2) In $\texttt{P2}(\mathcal{I})$, the variable z_{ij} for $i \in [n^*], j \in [m^*]$ encodes the envy experienced by each agent of type i who receives a house of type j. Note that the envy experienced by such an agent is always either 0, or $\sum_{i' \in [n^*]} \sum_{j' \in \mathcal{P}(i)} x_{i'j'}$. Since w is the maximum envy experienced by an agent, we must also have $z_{ij} \leq w$ for every $i \in [n^*], j \in [m^*]$.

We now formally describe the ILP. Minimize w subject to the constraints in Table 4.4.

Proof Outline of Theorem 4.48. Observe that $\texttt{P2}(\mathcal{I})$differs from $\texttt{P1}(\mathcal{I})$in constraints C4.a.$i.j$, C4.b.$i.j$ and C4.c.$i.j$. These three constraints together now ensure that for any feasible solution f for $\texttt{P2}(\mathcal{I})$, we either have $f(z_{ij}) = 0$ or $f(z_{ij}) = \sum_{i' \in [n^*]} \sum_{j' \in \mathcal{P}(i)} f(x_{i'j'})$. (Also, constraint C4.c.$i.j$ subsumes the constraint C3.c$i.j$ in $\texttt{P1}(\mathcal{I})$.) The only other difference is the addition of constraints C7 and C8.$i.j$. We can show that appropriate counterparts of Claims 4.35 and 4.36 hold for $\texttt{P2}(\mathcal{I})$. So do appropriate counterparts of Claims 4.31, 4.32 and 4.34. In particular, we have $\kappa^\dagger(\mathcal{I}) = \texttt{opt}(\texttt{P2}(\mathcal{I}))$. Theorem 4.48 will then follow from Theorem 4.6. □

Remark 4.49. *By modifying $P2(\mathcal{I})$, we can formulate an integer program for $[0/1]$-UHA. We only need to remove the variable w and the constraints C7 and C8.i.j for $i \in [n^*], j \in [m^*]$ and replace*

(C1.i).	$\sum_{j\in[m^*]} x_{ij} = n_i$	for every $i \in [n^*]$
(C2.j).	$\sum_{i\in[n^*]} x_{ij} \leq m_j$	for every $j \in [m^*]$
(C3.a.$i.j$).	$x_{ij} \leq nd'_{ij}$	for every $i \in [n^*], j \in [m^*] \setminus \mathcal{P}(i)$
(C3.b.$i.j$).	$\sum_{i'\in[n^*]} \sum_{j'\in[\mathcal{P}(i)]} x_{i'j'} \leq nmz_{ij} + nm(1 - d'_{ij})$	
(C4.a.$i.j$).	$z_{ij} \leq nd_{ij}$	for every $i \in [n^*], j \in [m^*]$
(C4.b.$i.j$).	$\left(\sum_{i'\in[n^*]} \sum_{j'\in\mathcal{P}(i)} x_{i'j'}\right) - z_{ij} \leq n(1 - d_{ij})$	
(C4.c.$i.j$).	$z_{ij} \leq n \sum_{i'\in[n^*]} \sum_{j'\in\mathcal{P}(i)} x_{i'j'}$	
(C5.$i.j$).	$z_{ij} = 0$	for every $i \in [n^*], j \in \mathcal{P}(i)$
(C6.a.$i.j$).	$x_{ij} \geq 0$	for every $i \in [n^*], j \in [m^*]$
(C6.b.$i.j$).	$z_{ij} \geq 0$	
(C6.c.$i.j$).	$d_{ij} \in \{0, 1\}$	
(C6.d.$i.j$).	$d'_{ij} \in \{0, 1\}$	
(C7).	$w \geq 0$	
(C8.$i.j$).	$z_{ij} \leq w$	for every $i \in [n^*], j \in [m^*]$

Table 4.4: The constraints of the ILP $\texttt{P2}(\mathcal{I})$.

the objective function with $\sum_{i\in[n^*]}\sum_{j\in[m^*]} x_{ij}z_{ij}$. Notice that while all the constraints in this integer program are linear, the objective function is quadratic. We thus have an integer quadratic program (IQP). The value of the largest coefficient in the constraints and the objective function is nm. It is known that IQP is fixed-parameter tractable when parameterized by the number of variables plus the value of the largest coefficient (Lokshtanov, 2015). Fixed-parameter tractability results for IQP w.r.t. other parameters are also known (Eiben et al., 2019).

4.6 Utilitarian House Allocation

We now deal with the UHA problems, where the goal is to minimize total envy. We first discuss the polynomial time algorithms for Utilitarian House Allocation.

4.6.1 Polynomial Time Algorithms for UHA

Theorem 4.50. *There is a polynomial-time algorithm for $[0/1]$-Utilitarian House Allocation when the agent valuations have an extremal interval structure.*

Proof. Consider an instance $\mathcal{I} = (H, A, \mathcal{P}; k)$ of $[0/1]$-UHA. In light of Remark 4.13, assume that the valuations have a left-extremal structure. Consider the ordering on the agents such that $i < j$ if $\mathcal{P}(a_i) \subseteq \mathcal{P}(a_j)$. Our algorithm relies on the existence of an optimal allocation with some desirable properties. To that end, consider an allocation $\Phi : A \to H$. We say that an ordered pair of agents $(a_i, a_j) \in A \times A$ is rogue under Φ if $i < j$, $\Phi(a_i) \in \mathcal{P}(a_i)$ and $\Phi(a_j) \notin \mathcal{P}(a_j)$. That is, for $i < j$, (a_i, a_j) is a rogue pair if the a_i values the house that she receives and a_j does not value that she receives. Consider a rogue pair (a_i, a_j). Recall that we are in the left-extremal setting, and hence $\mathcal{P}(a_i) \subseteq \mathcal{P}(a_j)$, which implies that $\Phi(a_i) \in \mathcal{P}(a_j)$. Thus a_j is envious under Φ.

We say that Φ is rogue-free if there does not exist any rogue pair under Φ. Notice that if Φ is rogue-free, then there exists $t(\Phi) \in [n] \cup \{0\}$ such that for every $i \in [n]$ with $i > t(\Phi)$, we have $\Phi(a_i) \in \mathcal{P}(a_i)$, and hence the agent a_i is envy-free. For $i \leq t(\Phi)$, the agent a_i may or may not be envy-free.

Claim 4.51. *There exists a rogue-free optimal allocation.*

Proof. Consider an optimal allocation $\Phi : A \to H$ that minimizes the number of rogue pairs. By optimal, we mean that $\kappa^\star(\mathcal{I}) = \kappa^\star(\Phi)$. If Φ is rogue-free, then the claim trivially holds. So, assume that Φ is not rogue-free. Then there exists a rogue pair under Φ. We fix a rogue pair (a_i, a_j) as follows. Let a_i be the first agent such that (a_i, a_p) is a rogue pair for some $p \in [n]$. Then choose j such that a_j is the last agent such that (a_i, a_j) is a rouge pair. Since (a_i, a_j) is a rogue pair, we have $i < j$, $\Phi(a_i) \in \mathcal{P}(a_i)$ and $\Phi(a_j) \notin \mathcal{P}(a_j)$. Since $i < j$, we have $\mathcal{P}(a_i) \subseteq \mathcal{P}(a_j)$, which implies that $\Phi(A) \cap \mathcal{P}(a_i) \subseteq \Phi(A) \cap \mathcal{P}(a_j)$. Notice first that a_j is envious as $\Phi(a_i) \in \mathcal{P}(a_i) \subseteq \mathcal{P}(a_j)$. Also, the number of agents that a_j envies, $\mathcal{E}_\Phi(a_j) = |\Phi(A) \cap \mathcal{P}(a_j)|$.

Let Φ' be the allocation obtained from Φ by swapping the houses of a_i and a_j. That is, $\Phi'(a_i) = \Phi(a_j)$, $\Phi'(a_j) = \Phi(a_i)$ and $\Phi'(a_r) = \Phi(a_r)$ for every $r \in [n] \setminus \{i, j\}$. Then, a_i is envious under Φ' as $\Phi'(a_i) = \Phi(a_j) \notin \mathcal{P}(a_j) \supseteq \mathcal{P}(a_i)$. The number of agents a_i envies, $|\mathcal{E}_{\Phi'}(a_i)| = |\Phi'(A) \cap \mathcal{P}(a_i)|$. Now, a_j is not envious under Φ' as $\Phi'(a_j) = \Phi(a_i) \in \mathcal{P}(a_i) \subseteq \mathcal{P}(a_j)$.

Notice that $\Phi(A) = \Phi(A')$. We thus have $|\mathcal{E}_{\Phi'}(a_i)| = |\Phi'(A) \cap \mathcal{P}(a_i)| \leq |\Phi(A) \cap \mathcal{P}(a_j)| = |\mathcal{E}_\Phi(a_j)|$. Therefore, $\kappa^\star(\Phi') = \kappa^\star(\Phi) - |\mathcal{E}_\Phi(a_j)| + |\mathcal{E}_{\Phi'}(a_i)| \leq \kappa^\star(\Phi)$. Since Φ is optimal, we can conclude that Φ' is optimal as well.

Now, we claim that the number of rogue pairs under Φ' is strictly less than that under Φ, which will contradict the definition of Φ. Notice first that (a_i, a_j) is a rogue-pair under Φ but not under Φ'. Consider $p, q \in [n]$ such that (a_p, a_q) is a rogue pair under Φ', but not under Φ. Then, either $q = i$ or $p = j$. If $q = i$, then $p < q = i < j$ and (a_p, a_j) is a rogue pair under Φ, which contradicts our choice of i. If $p = j$, then $j = p < q$, then (a_i, a_q) is a rogue pair, which contradicts our choice of j. Thus the number of rogue pairs under Φ' is strictly less than that under Φ, a contradiction. Hence, we conclude that Φ is rogue-free. □

For an allocation $\Phi : A \to H$, let $S_\Phi = \{a \in A \mid \Phi(a) \notin \mathcal{P}(a)\}$ and $T_\Phi = \{a \in A \mid \Phi(a) \in \mathcal{P}(a)\}$. Note that $\{S_\Phi, T_\Phi\}$ is a partition of A (with one of the parts possibly being empty). We say that Φ is nice if no agent in S_Φ envies any other agent in S_Φ. Equivalently, Φ is nice if $\Phi(S_\Phi) \cap \mathcal{P}(S_\Phi) = \emptyset$.

Claim 4.52. *There exists an optimal rogue-free allocation that is also nice.*

Proof. Let Φ be an optimal (i.e., $\kappa^\star(\mathcal{I}) = \kappa^\star(\Phi)$) rogue-free allocation that minimizes $|S_\Phi|$. Then, there exists $t_\Phi \in [n] \cup \{0\}$ such that $S_\Phi = \left\{a_1, \ldots, a_{t(\Phi)}\right\}$ and $T_\Phi = \left\{a_{t(\Phi)+1}, \ldots, a_n\right\}$. Note that S_Φ is indeed contiguous. If not, then there exists indices i and j such that $i < j - 1$ and $a_i, a_j \in S_\Phi$ and $a_p \notin S_\Phi$ for every index p with $i < p < j$. Then (a_{j-1}, a_j) is a rogue-pair.

If $|S_\Phi| \leq 1$, then the claim trivially holds. So, assume that $|S_\Phi| \geq 2$. Suppose that there exist $a_i, a_j \in S_\Phi$ such that a_j envies a_i. Then, as $\mathcal{P}(a_j) \subseteq \mathcal{P}(a_{t(\Phi)})$, $a_{t(\Phi)}$ envies a_i as well. Let Φ' be the allocation obtained from Φ by swapping the houses of a_i and $a_{t(\Phi)}$. Then we have $S_{\Phi'} = S_\Phi \setminus \left\{a_{t(\Phi)}\right\}$ and $T_{\Phi'} = T_\Phi \cup \left\{a_{t(\Phi)}\right\}$. Thus, $|S_{\Phi'}| < |S_\Phi|$. Note that we constructed Φ' from Φ without introducing any new rogue-pairs. Additionally, we converted an envious agent under Φ (in particular, $a_{t(\Phi)}$) to an envy-free agent under Φ'. This contradicts the optimality of Φ and the fact that Φ minimizes $|S_\Phi|$. □

Claim 4.53. *Let Φ be a nice rogue-free allocation. Consider a house $h \in \Phi(T_\Phi)$. Then, (1) the number of agents who envy $\Phi^{-1}(h)$ is exactly $|\{a' \in S_\Phi \mid h \in \mathcal{P}(a')\}|$, and (2) $\kappa^\star(\Phi) = \sum_{h \in T_\Phi} |\{a' \in S_\Phi \mid h \in \mathcal{P}(a')\}|$.*

Proof. By the definition of T_Φ, no agent in T_Φ envies $\Phi^{-1}(h)$. Hence the number of agents who envy $\Phi^{-1}(h)$ is exactly equal to the number of agents in S_Φ who value h. This is precisely what assertion (1) says. Now, to compute $\kappa^\star(\Phi)$, for all $h \in \Phi(T_\Phi)$, we only need to sum the number of agents in S' value h. This is precisely what assertion (2) says. □

Informal description of our algorithm: Based on Claims 4.51-4.53, we are now ready to describe our algorithm. Informally, our algorithm works as follows. We are given an instance $\mathcal{I} = (A, H, \mathcal{P}; k)$. Suppose that Φ is the optimal allocation that we are looking for. By Claim 4.52, we can assume that Φ is rogue-free and nice. We guess t_Φ. There are at most $n+1$ guesses. For the correct guess, we correctly identify S_Φ and T_Φ. Then, Φ must allocate to each $a \in T_\Phi$ a house that a values. To each agent $a' \in S_\Phi$, Φ must allocate a house that no agent in S_Φ values. For a house $h \in H$, the envy generated by allocating h is precisely the number of agents in S_Φ who value h. We can thus reduce the problem to a minimum cost maximum matching problem, where the cost of matching each $h \in H$ to (1) an agent $a \in T_\Phi$ who values h is precisely $|\{a' \in S_\Phi \mid a' \text{ values } h\}|$; (2) an agent $a \in T_\Phi$ who does not value h is prohibitively high; (3) an agent $a' \in S_\Phi$ is 0 if no agent in S_Φ values h, and prohibitively high otherwise. We can compute a minimum cost maximum matching in polynomial time.

Algorithm: We are given an instance $\mathcal{I} = (A, H, \mathcal{P}; k)$ as input. For each fixed $t \in [n] \cup \{0\}$, we do as follows. We partition A into two sets S and T as follows: $S = \{a_1, \dots, a_t\}$ and $T = \{a_{t+1}, \dots, a_n\}$. We construct a complete bipartite graph $G_t^\star$, with vertex bipartition $A \uplus H$ and a cost function c_t on the edges defined as follows:

$$c_t((a,h)) = \begin{cases} |\{a' \in S \mid h \in \mathcal{P}(a')\}| & \text{if } a \in T \text{ and } a \text{ values } h, \\ 0 & \text{if } a \in S \text{ and no agent } a' \in S \text{ values } h, \\ k+1 & \text{otherwise.} \end{cases}$$

If $G_t^\star$ contains a matching of size n and cost at most k for any $t \in [n] \cup \{0\}$, then we return that $\mathcal{I}$ is a yes-instance of [0/1]-UHA. If $G_t^\star$ does not contain such a matching for any choice of $t \in [n] \cup \{0\}$, then we return that $\mathcal{I}$ is a no-instance of [0/1]-UHA.

Correctness: To see the correctness of our algorithm, assume first that there exists $t \in [n] \cup \{0\}$ for which $G_t^\star$ contains a matching, say M, of size n and cost at most k. Then, since $|M| = n$, M saturates A. Consider an allocation $\Phi_M : A \to H$ defined as follows: for each $(a,h) \in M$, Φ_M allocates h to a. We claim that $\kappa^\star(\Phi_M) = c_t(M) \leq k$. First, each $a \in T$ values $\Phi_M(a)$, for otherwise, $c_t((a, \Phi_M(a))) = k+1$, which is not possible. So, $a \in T$ does not envy any agent. Similarly, each $a' \in S$ does not value $\Phi_M(a'')$ for any $a'' \in S$, for otherwise, $c_t((a'', \Phi_M(a''))) = k+1$, which is not possible. So, for $a', a'' \in S$, a' does not envy a''. Also, for $(a',h) \in M$ with $a' \in S$, we have $c_t((a',h)) = 0$. Now, an agent $a' \in S$ may envy an agent $a \in T$. But note that every $h \in \Phi_M(T)$ contributes exactly $|\{a' \in S \mid h \in \mathcal{P}(a')\}| = c_t((a,h))$, where $\Phi_M(a) = h$, to $\kappa^\star(\Phi_M)$. Hence, $\kappa^\star(\Phi_M) = \sum_{\substack{h \in \Phi_M(T) \\ (a,h) \in M}} c_t((a,h)) = \sum_{\substack{(a,h) \in M \\ a \in T}} c_t((a,h)) \leq k$.

Conversely, assume that $\mathcal{I} = (A, H, \mathcal{P}; k)$ is a yes-instance. Let $\Phi : A \to H$ be an optimal allocation. By Claim 4.52, we assume without loss of generality that Φ is rogue-free and nice. Consider the iteration of our algorithm for which $t = t_\Phi$. Consider the matching M_Φ in $G_t^\star$ defined as $M_\Phi = \{(a, \Phi(a)) \mid a \in A\}$. We claim that $c_t(M_\Phi) = \kappa^\star(\Phi) \leq k$. First, consider $a \in S$. Since Φ is nice, no $a' \in S$ values $\Phi(a)$, which implies that $c_t(a, \Phi(a)) = 0$. Now, consider $a \in T$. Since Φ is rogue-free, a values $\Phi(a)$. By Claim 4.53, $\Phi(a)$ contributes exactly $|\{a' \in S \mid a' \text{ values } \Phi(a)\}| = c_t(a, \Phi(a))$ to $\kappa^\star(\Phi)$. Thus, $c_t(M_\Phi) = \sum_{(a,\Phi(a)) \in M_\Phi} c_t(a, \Phi(a)) = \sum_{\substack{(a,\Phi(a)) \in M_\Phi \\ a \in T}} c_t((a, \Phi(a))) = \sum_{a \in T} |\{a' \in S \mid a' \text{ values } \Phi(a)\}| = \kappa^\star(\Phi) \leq k$. ☐

We now turn to the restricted setting where every agent likes exactly one house. In this case, the total envy is equal to the number of envious agents, that is, $\kappa^\star(\Phi) = \kappa^\#(\Phi)$, so the following result follows from Theorem 4.17.

Corollary 4.54. *There is a polynomial-time algorithm for* $[0/1]$*-Utilitarian House Allocation* *when every agent approves exactly one house.*

4.6.2 Parameterized Results

In this section, we discuss the parameterized results for UHA. First, we design a linear kernel for [0/1]-UHA.

Theorem 4.55. *[0/1]-Utilitarian House Allocation admits a linear kernel parameterized by the number of agents. In particular, given an instance of [0/1]-Utilitarian House Allocation, there is a polynomial time algorithm that returns an equivalent instance of [0/1]-Utilitarian House Allocation with at most twice as many houses as agents.*

Proof. It suffices to prove the safety of Reduction Rule 2. Let $\mathcal{I} := (A, H, \mathcal{P}; k)$ denote an instance of HA. Further, let $\mathcal{I}' = (H' := H \setminus X, A' := A \setminus Y, \mathcal{P}'; k)$ denote the reduced instance corresponding to $\mathcal{I}$. Note that the parameter for the reduced instance is k as well.

If $\mathcal{I}$ is a Yes-instance of UHA, then there is an allocation $\Phi : A \to H$ with total envy at most k. By Claim 4.11, we may assume that Φ is a good allocation. This implies that the projection of Φ on $H' \cup A'$ is well-defined, and it is easily checked that this gives an allocation with total envy at most k.

On the other hand, if $\mathcal{I}'$ is a Yes-instance of UHA, then there is an allocation $\Phi' : A' \to H'$ with total envy at most k. We may extend this allocation to $\Phi : A \to H$ by allocating the houses in Y to agents in X along the expansion M, that is:

$$\Phi(a) = \begin{cases} \Phi'(a) & \text{if } a \notin X, \\ M(a) & \text{if } a \in X. \end{cases}$$

Since all the newly allocated houses are not valued by any of the agents outside X and all agents in X are envy-free with respect to Φ, it is easily checked that Φ also has total envy at most k. □

The following results follow from the algorithm described in Proposition 4.3 after guessing the allocated houses, which adds a multiplicative overhead of $\binom{m}{n} \leq 2^m$ to the running time.

Corollary 4.56. *[0/1]-Utilitarian House Allocation is fixed-parameter tractable when parameterized either by the number of houses or the number of agents. In particular, [0/1]-Utilitarian House Allocation can be solved in time $O^\star(2^m)$.*

$(n, m, n^\star)$	OHA			EHA			Time/Instance OHA, EHA (in sec.)
	Env. Agents ($\kappa^\star(\Phi)$)	Max Envy	Total Envy	Env. Agents	Max Envy ($\kappa^\dagger(\Phi)$)	Total Envy	
(30, 30, 1)	15.11	14.89	216.71	15.11	14.89	216.71	0.003, 0.002
(30, 30, 5)	0.95	8.78	12.33	9.07	7.56	11.29	0.19, 0.10
(30, 30, 15)	0	0	0	0	0	0	0.20, 0.33
(30, 40, 1)	10.18	19.82	191.76	20.18	9.82	188.16	0.01, 0.003
(60, 60, 1)	30.36	29.64	888.08	30.36	29.64	888.08	0.006, 0.002
(60, 60, 15)	0.01	0.31	0.31	0.21	0.21	0.21	0.57, 0.10
(60, 60, 30)	0	0	0	0	0	0	2.00, 4.21
(120, 120, 1)	59.45	60.55	3567.8	59.45	60.55	3567.8	0.01, 0.002
(120, 120, 5)	3.83	57.07	218.79	66.62	51.07	200.63	0.11, 0.20
(120, 120, 15)	0	0	0	0	0	0	1.98, 4.26
(120, 130, 5)	0	0	0	0	0	0	0.11, 0.10

Table 4.5: A summary of the results, averaged over 100 instances of each type. The OHA column corresponds to the solution from OHA ILP and the max-envy and total envy in that column shows those values when the number of envious agents is minimized. Similarly for the EHA column.

Corollary 4.57. *$[\succ]$-Utilitarian House Allocation is fixed-parameter tractable when parameterized by the number of houses and can be solved in time $O^\star(2^m)$.*

4.7 Experiments

We implemented the ILP for OHA and EHA over synthetic datasets of house allocation problems generated uniformly at random. We used Gurobi Optimizer version 9.5.1[6]. The average was taken over 100 trials for each instance. A summary is recorded in Table 4.5. For a fixed number of houses and agents, notice that as the number of agent types, $n^\star$ increases, the number of envious agents and the maximum envy decreases. Instances with identical valuations (where $n^\star = 1$) seem to admit more envy than the other extreme (where $n^\star = n$). This is due to the fact when valuations are identical, there is more contention on the specific subset of goods. On the contrary, for instances with $m^\star = 1$, envy-free allocations always exist. Indeed, when $m^\star = 1$, all houses are of the same type, which means that an agent either likes all the houses or dislikes all of them and in either case, she is envy-free no matter which house she gets. Also note that,

[6]The code can be accessed at https://github.com/anonymous1203/House-Allocation

when we increase the number of houses, for a constant number of agents and agent types, the envy decreases, which is as expected, because of the increase in the number of choices and the fact that some houses (the more contentious ones) remain unallocated.

4.8 House Allocation on Single-Peaked/Dipped Rankings

Single-peaked preferences, first formalized by Black (1948), constitute an important domain restriction in the collective decision-making problems, which not only model many real-world settings but also serve as a tractable realm for many of the otherwise hard problems. They have been extensively considered in various contexts including house allocations and matching markets (Bade, 2019), voting and electorates (Conitzer, 2007; Faliszewski et al., 2009; Sprumont, 1991) among others. A significant literature has also focused on characterizing these preferences (Ballester and Haeringer, 2011; Elkind et al., 2020; Puppe, 2018). For more details, we refer the reader to the survey of preference restrictions in social choice by Elkind et al. (2022).

For a positive integer t, we write $[t]$ to denote the set $\{1, \ldots, t\}$.

Let $\mathcal{I} = (A, H, \succ)$ be an instance of house allocation problem, with $A := \{a_1, a_2, \ldots a_n\}$ being the set of $n \in \mathbb{N}$ agents, $H := \{h_1, h_2, \ldots, h_m\}$ be the set of $m \in \mathbb{N}$ houses and $\succ$ be the set of ranking profile of all the agents. Let $\succ_a$ be the ranking of agent a over the houses H. If $h \succ_a h'$, we say that agent a strictly prefers the house h over h'. If $h \succ_a h' \; \forall \, h' \neq h$, then we say that h is a peak house for a, denoted as $peak(a)$. We say that $\text{rank}_a(h) = j$ for $j \in [m]$ if the house h appears at the j^{th} position in the ranking $\succ_a$. We denote the set of agents who prefer a house h to all other houses as $base(h)$. That is, $base(h) = \{a \mid \text{rank}_a(h) = 1\}$.

Let $\triangleright$ be an ordering on the houses. We say that $\succ_a$ is *single-peaked* with respect to the ordering $\triangleright$ if for every pair of houses $h, h' \in H$, we have that if $h \triangleright h' \triangleright \text{peak}(\succ_a)$ or $peak(\succ_a) \triangleright h' \triangleright h$, then $h' \succ_a h$. Equivalently, the definition requires that for each agent a, it holds that the favorite t houses of a with respect to $\succ_a$ form a consecutive segment within $\triangleright$. That is, as an agent moves away from his favorite house $peak(a)$ in any direction, left or right, the houses become less and less preferable for her. A preference profile $\succ$ is single-peaked with respect to $\triangleright$ if for every agent a, $\succ_a$ is single-peaked with respect to $\triangleright$ and it is single-peaked if there exists some $\triangleright$ over H such that it is single-peaked with respect to $\triangleright$ (see Figure 4.1).

A closely related domain is *single-dipped* preferences. A ranking $\succ_a$ is said to be single-dipped with respect to an ordering $\rhd$ if the reverse of $\succ_a$ is single-peaked. That is, there is a least preferred house h and the agent likes the houses better as she moves away from h. As before, a preference profile $\succ$ is single-dipped with respect to $\rhd$ if for every agent a, $\succ_a$ is single-dipped with respect to $\rhd$ and it is single-dipped if there exists some $\rhd$ over H such that it is single-dipped with respect to $\rhd$.

4.8.1 Single-Peaked Preferences

Suppose that the set of rankings $\succ$ are single-peaked with respect to the ordering $\rhd$ over the houses $(h_1 \rhd h_2 \rhd \ldots \rhd h_m)$. We say that a house h_i is a *shared peak* or a *shared house* if it is the most preferred house of more than one agent, that is, $|base(h_i)| > 1$. Otherwise, if $base(h_i) = 1$, we say it is a *individual peak* or a *individual house*. We say that a house h is non-wastefully allocated if it is allocated to an agent a such that $rank_a(h) = 1$, otherwise, it is wastefully allocated.

We define the *span* of a peak house h, denoted by span(h), as the number of houses that are identically ranked by all the agents in $base(h)$, starting from their common first ranked house. If h_i is an individual peak, then we say $span(h_i) = 0$ as there is only one agent in the set $base(h_i)$. Consider the rankings in the Figure 4.1 explicitly depicted below. Here, $span(h_7) = 0$ and $span(h_1) = 2$. If at least two agents from $\{a_2, a_3, a_4\}$ were to be envy-free in any allocation, then not only the peak house h_1 has to remain unallocated, but all the houses in the set span(h_1) must also remain unassigned under any complete allocation. The allocation of h_7, h_5, h_4, and h_3 to the four agents, respectively, makes two agents from $base(h_1)$ envy-free.

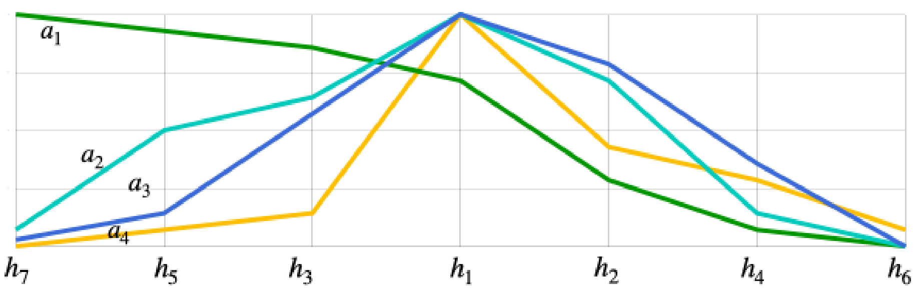

Figure 4.1: Single-peaked preferences with respect to the ordering $\rhd := h_7 \rhd h_5 \rhd h_3 \rhd h_1 \rhd h_2 \rhd h_4 \rhd h_6$. The house h_1 is a shared peak and h_7 is an individual peak. Notice that $peak(a_1) = h_7$ and $peak(a_2) = peak(a_3) = peak(a_4) = h_1$. Also, $base(h_1) = \{a_2, a_3, a_4\}$. And, span($h_1$) = 2, which contains the houses h_1 and h_2 as these are the top 2 houses identically ranked by all the agents in $base(h_1)$.

$$\begin{array}{ll} a_1: & \mathbf{h_7} \succ h_5 \succ h_3 \succ h_1 \succ h_2 \succ h_4 \succ h_6 \\ a_2: & h_1 \succ h_2 \succ h_3 \succ \mathbf{h_5} \succ h_4 \succ h_6 \succ h_7 \\ a_3: & h_1 \succ h_2 \succ \mathbf{h_4} \succ h_3 \succ h_5 \succ h_7 \succ h_6 \\ a_4: & h_1 \succ h_2 \succ \mathbf{h_3} \succ h_4 \succ h_5 \succ h_6 \succ h_7 \end{array}$$

We now present a series of structural results. If a shared peak h is assigned to one agent, it leads to envy among the other agents, with at least $|base(h) - 1|$ envious agents. The only agent who receives h is the one without envy among the $base(h)$ agents. On the other hand, we show by the following claim that even if h is not assigned, at least $|base(h) - 2|$ agents are bound to be envious under any allocation.

Lemma 4.58. *Let h be a shared peak. Then, at most 2 agents from the set $base(h)$ can be envy-free under any allocation. In other words, at least $base(h) - 2$ agents are envious under any allocation.*

Proof. Consider an allocation Φ. If house h is allocated under Φ, then it can be allocated either wastefully or non-wastefully. We show that in both cases there are at most two agents from $base(h)$ that are envy-free. If it is allocated wastefully, then all the $base(h)$ agents are envious, no matter which house they receive in Φ. If h is allocated non-wastefully to an agent, say a, then a is always envy-free in any completion of this allocation, as she receives her first ranked house. But, all other $|base(h) - 1|$ agents experience envy on account of the allocation of h to a.

If house h is not allocated under Φ, then we will prove the statement by contradiction. Suppose that at least three agents from the set $base(h)$ are envy-free under Φ. Let agents a_1, a_2, and a_3 denote three agents from $base(h)$ that are envy-free under Φ. Then, at least 2 of these envy-free agents are allocated to houses from either $[h_1, h)$ or $(h, h_m]$ where the interval of houses are from the single peak axis $\rhd$. WLOG, we assume that $\Phi(a_1) = h_i$ and $\Phi(a_2) = h_j$ such that $\{h_i, h_j\} \in [h_1, h)$. Then since h is a peak for both a_1 and a_2, by the structure of the rankings, it must be the case that both of them have the (partial) rankings as either $h_1 < h_i < h_j < h$ or $h_1 < h_j < h_i < h$. In either case, at least one of them is envious of the other, which contradicts the assumption that all three agents $\{a_1, a_2, a_3\}$ were envy-free. Therefore, at most two agents can be envy-free from the set $base(h)$, potentially, the ones that get allocated houses lying on either side of h. □

We now proceed to show another interesting structural claim that is used for allocating the individual peaks.

Lemma 4.59. *There exists an optimal allocation where all individual peaks are allocated, and* *they are allocated non-wastefully.*

Proof. Let house h be an individual peak. Suppose that h remains unallocated under an allocation $\Phi^\star$. Consider the following cases:

1. $h \in \text{span}(h_i)$ for some shared peak h_i. Then h remains unallocated and at most two agents from the set of agents $base(h_i)$, say a_1 and a_2, are envy-free under $\Phi^\star$ (this forces span(h_i) and hence, h to remain unallocated under $\Phi^\star$). Now consider the reallocation where a gets h and a_1 gets h_i (where $a_1 \in base(h_i)$). Then both a and a_1 are envy-free under this reallocation. The number of envy-free agents under this reallocation remains the same as in $\Phi^\star$. Indeed, at most two out of $\{a, a_1, a_2\}$ can be made envy-free under any allocation (by Lemma 4.58). This settles our claim in this case.

2. $h \notin \text{span}(h_i)$ for any other peak h_i. First, suppose that a is an envy-free agent under $\Phi^\star$. This implies that every house that a ranks better than $\Phi^\star(a)$ (including h) remains unallocated under $\Phi^\star$. We claim that on the re-allocation of h to a, no new envious agent is created. Suppose not. Say a' is an agent who was previously envy-free but becomes envious on the allocation of h. Since a' was envy-free previously, we can say that $\Phi^\star(a') >_{a'} \Phi^\star(a)$. If $\Phi^\star(a')$ lies to the left of $\Phi^\star(a)$, then by structure of the valuations, $\Phi^\star(a') >_{a'} h$ and therefore, a' can't be envious of the allocation of h. On the other hand, if $\Phi^\star(a')$ lies to the right of $\Phi^\star(a)$, then it must be that $\Phi^\star(a')$ also lies to the right of h. Else, a will prefer $\Phi^\star(a')$ more than $\Phi^\star(a)$, leading her to be envious. Therefore, it must be the case that $\Phi^\star(a')$ is to the right of h and hence by structure of the valuations, $\Phi^\star(a) <_{a'} h <_{a'} \Phi^\star(a')$. Therefore, a' can't be envious of the allocation of h.

   Now suppose that a was envious under $\Phi^\star$. Reallocating h to a makes her envy-free. We will first argue that there are at most two agents, say a_1 and a_2 who become newly envious of the allocation of h. Suppose the peaks of a_1 and a_2 lie to the left and right side of h respectively. Since both a_1 and a_2 have become newly envious, so $\Phi^\star(a_1)$ and $\Phi^\star(a_2)$ can not be their respective peak houses. Since a_1 and a_2 are envy-free under $\Phi^\star$, therefore, $peak(a_1) >_{a_1} h >_{a_1} \Phi^\star(a_1) >_{a_1} \Phi^\star(a)$ and $peak(a_2) >_{a_2} h >_{a_2} \Phi^\star(a_2) >_{a_2} \Phi^\star(a)$. And all houses between $\Phi^\star(a_1)$ and $peak(a_1)$ are unallocated. Similarly, all houses between $\Phi^\star(a_2)$ and $peak(a_2)$ are unallocated. Since h is a common house that lies both between $\Phi^\star(a_1)$ and $peak(a_1)$ and between $\Phi^\star(a_2)$ and $peak(a_2)$, we have that all the houses between $\Phi^\star(a_1)$ and $\Phi^\star(a_2)$ are unallocated (see Figure 4.2). Now, consider any other agent a_3 who is envy-free under

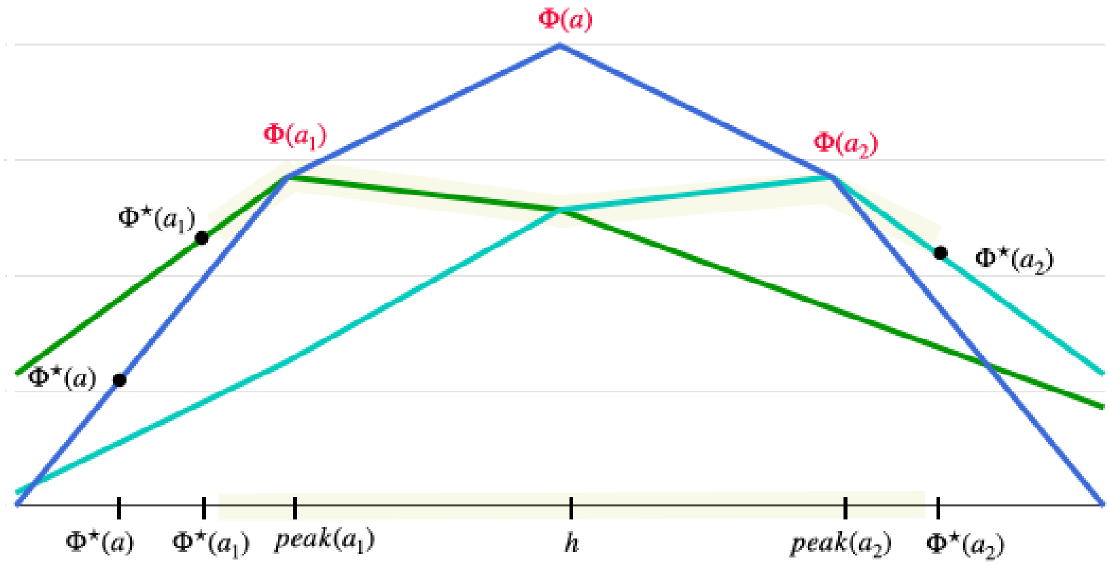

Figure 4.2: A schematic of Case 2 in the proof of Lemma 4.59.

$\Phi^\star$. We claim that she can't be envious of the allocation of h. First suppose that $\Phi^\star(a_3)$ lies to the left of $\Phi^\star(a_1)$. Then it must be the case that $peak(a_3)$ also lies to the left of $\Phi^\star(a_1)$ because a_3 is envy-free under $\Phi^\star$. This means that if a_3 is envy-free of the allocation of $\Phi^\star(a_1)$, she remains envy-free of the allocation of h as well (by the structure of the valuations). Second, if $\Phi^\star(a_3)$ lies to the right of $\Phi^\star(a_2)$. Then it must be the case that $peak(a_3)$ also lies to the right of $\Phi^\star(a_2)$ which means that if a_3 is envy-free of the allocation of a_2, she remains envy-free of the allocation of h as well (again by the structure of the valuations).

So now we have that a_1 and a_2 are the only two agents that can potentially become envious of the allocation of h. Note that $peak(a_1)$ and/or $peak(a_2)$ can not be a resolved shared peak under $\Phi^\star$. Otherwise, h must be in one of the spans, which contradicts the assumption of this case. Therefore, in the optimal allocation $\Phi^\star$, at most one agent from $base(a_1)$ and at most one agent from $base(a_2)$ is envy-free.

Based on this, we can now propose the following re-allocation: h to a, $peak(a_1)$ to a_1 and $peak(a_2)$ to a_2. It is easy to see that $\{a, a_1, a_2\}$ become newly envy-free as each of them now gets her peak house. Also, no other agent becomes newly envious of this re-allocation. Indeed, if there is such an agent, say a', then $\Phi^\star(a')$ must lie to either left of $\Phi^\star(a_1)$ or to the right of $\Phi^\star(a_2)$. Since they are envy-free under $\Phi^\star$, their respective peaks must also be on the same side as their allocated houses. This means that if they are envy-free of the allocation of $\Phi^\star(a_1)$ and $\Phi^\star(a_2)$ respectively, they do not become envious of the allocation of $peak(a_1)$, $peak(a_2)$ and h, again by the structure of the valuations. Therefore, $\{a, a_1, a_2\}$ are newly envy-free in the re-allocation, without creating any other envious agents. But under $\Phi^\star$, only $\{a_1, a_2\}$ were envy-free. This implies that $\Phi^\star$ was not optimal to begin with. This settles our claim.

Now suppose that the individual peak, say h is allocated wastefully to some agent a' under $\Phi^\star$. Let a be the unique agent in the set $base(h)$. Then, a is definitely envious. If a' is also an envious agent, then we can re-allocate h to a, which reduces the number of envious agents and contradicts the optimality of $\Phi^\star$. Therefore, a' must be an envy-free agent. Then, all the houses that a' values more than h must have remained unallocated. In particular, the peak house of a', say h', ($h' \neq h$) must have been unallocated. If h' was a shared peak, then it must have been a resolved shared peak (since h' remains unallocated) and hence at most two agents would be envy-free from $base(h')$. Notice that a' is one of them and, say a'' is the other. Then, reallocating h' to a' and h to a converts a to an envy-free agent and makes a'' envious, and does not generate any new envy. Otherwise, if h' was an individual peak, then again the re-allocation of h to a and h' to a' gives us our desired allocation. This settles our claim. □

Let the number of individual and shared peak houses be p_I and p_S respectively. Then any allocation can have at least $p_I + p_S$ many envy-free agents, just by allocating the peaks non-wastefully and completing the allocation in an arbitrary manner. Moreover, by Lemma 4.58, no allocation can have more than $2 \cdot p_S + p_I$ envy-free agents. This establishes the following result.

Lemma 4.60. *Let k be the number of envy-free agents under any allocation. Then,* $p_S + p_I \leq k \leq 2 \cdot p_S + p_I$.

The following is a generalization of Lemma 4.58.

Lemma 4.61. *Consider the set of k shared peaks $\{h_1, h_2, \ldots h_k\}$ such that $span(h_i) \cap span(h_j) \neq \emptyset$ for any $i, j \in [k]$. Then there is an allocation where at least k and at most $k+1$ agents from the set $\bigcup_{i \in k} base(h_i)$ are envy-free.*

Proof. Consider a non-wasteful allocation of the k shared peaks among k agents in the set $\bigcup_{i \in k} base(h_i)$. Clearly, these k agents are envy-free in any complete allocation, as each of them receives their favorite house. This allocation makes at least k agents envy-free. Now consider an allocation where at least one span is resolved, say span(h_1). This means that h_1 and span(h_1) remain unallocated and consequently, 2 agents from $base(h_1)$ are made envy-free by the allocation of the houses, say h_1^1 and h_1^2. Now consider any other overlapping span, say span(h_i). Then $span(h_i) \cap span(h_1) \neq \emptyset$. Then, by the structure of the rankings, either h_1^1 or h_1^2 must belong to the span(h_i). This implies that once span(h_1) is resolved (that is, h_1^1 and h_1^2 are allocated), then span(h_i) can not be resolved (at most one agent from $base(h_i)$ can be made envy-free). Since the choice of h_i was arbitrary, this holds for every other overlapping span with span(h_1). Therefore, at most $k+1$ agent can be envy-free from the set $\bigcup_{i \in k} base(h_i)$. □

We are now ready to present the main result of this section.

Theorem 4.62. *Given an instance of house allocation with single-peaked preferences, minimizing the number of envious agents admits a polynomial time algorithm.*

Proof. We first describe the algorithm. A shared peak h is said to be resolved under an allocation Φ if h and span(h) remain unallocated and exactly two agents from the set $base(h)$ are envy-free under Φ.

Algorithm 1 Minimize the number of envious agents

Require: $\{N, H, \succ, \rhd\}$
Ensure: An allocation Φ

- $\forall\, h \in p_I$:
  $\Phi(base(h)) = h$ (Allocate all the individual peaks non-wastefully).
- $\forall\, h_i, h_j$ such that $h_i \in span(h_j)$ and $h_j \in span(h_i)$
  $\Phi(a) = h_i$ for some $a \in base(h_i)$ and $\Phi(a') = h_j$ for some $a' \in base(h_j)$ (Allocate both h_i and h_j non-wastefully).
- $\forall\, h_i, h_j$ such that $h_i \in span(h_j)$ but $h_j \notin span(h_i)$
  $\Phi(a) = h_j$ for some $a \in base(h_j)$
- $p_S \leftarrow$ Set of remaining, unallocated shared peaks.
- Order the peak houses in p_S as $h_i \leq h_j$ if span(h_i) $\leq$ span(h_j). Say, $\{h_1, h_2, \ldots h_S\}$ is the ordering.
- For $i \in S$:
  m', n' = number of unallocated houses and agents under Φ
  If $m' - span(h_i) \geq n'$: Resolve h_i and $U \leftarrow span(h_i)$. Else, allocate $\{h_i, h_{i+1}, \ldots h_S\}$ non-wastefully.
- Allocate the remaining agents to any house, except U.
- Output Φ.

We now describe the correctness of Algorithm 1. Let Φ be the output of the above algorithm and let $EF(\Phi)$ denote the number of envy-free agents in Φ. Suppose $\Phi^\star$ is the optimal allocation that minimizes the number of envious agents. Clearly, $EF(\Phi^\star) \geq EF(\Phi)$. The aim is to show that $EF(\Phi^\star) = EF(\Phi)$. To this end, we will show that $|EF(\Phi^\star) \setminus EF(\Phi)| = |EF(\Phi) \setminus EF(\Phi^\star)|$.

Suppose $a \in EF(\Phi^\star) \setminus EF(\Phi)$. We will show that corresponding to a, there is a unique agent a' such that $a' \in EF(\Phi) \setminus EF(\Phi^\star)$. Note that the first ranked house by a, say h must be a shared peak, else $\Phi(a) = h$ and $a \in EF(\Phi)$, which is not the case. We now consider the following cases:

1. $\Phi^\star(a) = h$. Since $a \notin EF(\Phi)$, there is an agent a' who gets a house h' under Φ such that $\Phi(a') = h' >_a \Phi(a)$.
   (a) If $h' = h$, then, since $\Phi(a') = h$, we have that the first ranked house of a' is also h. (Since Φ always allocates the peaks non-wastefully and h is a peak for a). This implies that if a is envy-free under $\Phi^\star$, then a' must have been envious of a under $\Phi^\star$. But, as $\Phi(a') = h$, a' is envy-free under Φ. So, $a' \in EF(\Phi) \setminus EF(\Phi^\star)$.
   (b) Else, if $h' \neq h$, and h remains unallocated under Φ, then h is definitely a resolved peak under Φ. Suppose there are k overlapping spans with span(h). Then, we must have exactly $k+1$ agents, who are envy-free corresponding to the k overlapping spans (which is the only case when a shared peak is resolved). Note that $\Phi^\star$ also can have at most $k+1$ envy-free agents corresponding to the above overlapping spans (by *Lemma* 4.61). Since $a \in EF(\Phi^\star) \setminus EF(\Phi)$, we must have some agent a' among the above $k+1$ agents under Φ such that $a' \in EF(\Phi) \setminus EF(\Phi^\star)$.
2. $\Phi^\star(a) \neq h$. This implies that h remains unallocated under $\Phi^\star$. All houses that a ranks better that $\Phi^\star(a)$ remain unallocated (as $a \in EF(\Phi^\star)$). Consider all the overlapping spans with span(h). If there are k of them, then there are at most $k+1$ envy-free agents corresponding to these spans under $\Phi^\star$. Since $a \notin EF(\Phi)$, there is an agent a' who gets a house h' under Φ such that $\Phi(a') = h' >_a \Phi(a)$.
   (a) If $h' = h$, then the peak h is allocated and hence, not resolved under Φ.

   Suppose h was not resolved because it is a house that is considered in Case 2. This implies that h lies in the span(h_j) for some shared peak h_j and h_j lies in the span of h. Then, there can be at most 2 envy-free agents corresponding to the span(h) and span(h_j) in any allocation. Under Φ, there are exactly 2 envy-free agents, say $\{a_1, a_2\}$, corresponding to the spans since both h_j and h are allocated non-wastefully to say, a_1 and a_2 respectively. Since h_j is in the span(h) and h is in the span(h_j), it is easy to see that under $\Phi^\star$, at most 2 of the three agents $\{a, a_1, a_2\}$ can be envy-free. Therefore, this implies that $a_2 \in EF(\Phi) \setminus EF(\Phi^\star)$ and we are done for this case.

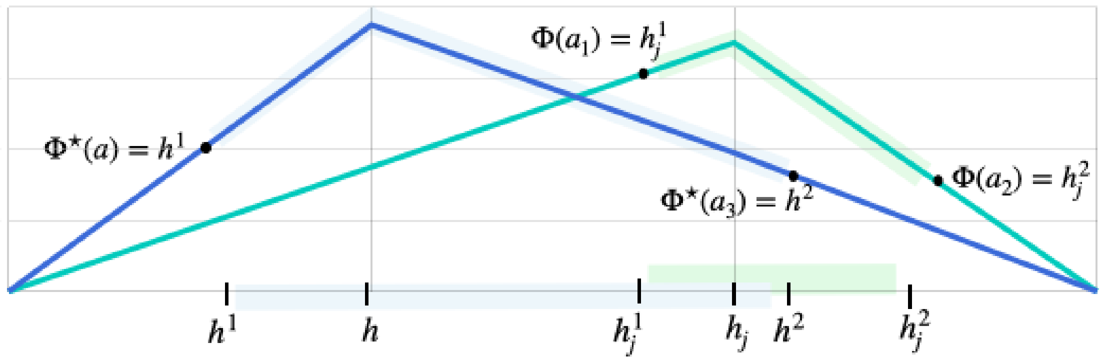

Figure 4.3: A schematic of Case 2(a) in Theorem 4.62

Now suppose h was not resolved because it is considered in Case 3. This implies that there is a shared peak h_j which lies in the span(h) (so, resolving h forces h_j to remain unallocated). Now by Lemma 4.61, we know that there can be at most 3 envy-free agents corresponding to the span(h) and span(h_j) in any allocation. In particular, two agents, say $\{a_1, a_2\}$ could be envy-free if span(h_j) was resolved and a_3 is envy-free who is the recipient of h (which is allocated non-wastefully). If h_j was resolved under Φ, then Φ has exactly 3 envy-free agents, namely $\{a_1, a_2, a_3\}$ (see Figure 4.3). Now, at most 3 agents from $\{a, a_1, a_2, a_3\}$ can be envy-free under $\Phi^\star$. But, as span(h) is resolved under $\Phi^\star$ and $h_j \in span(h)$, so h_j would not have been resolved and is also unallocated under $\Phi^\star$. This implies that only two agents, namely a and a_3 are envy-free under $\Phi^\star$. If h_j is resolved under Φ, then we have three agents $\{a_1, a_2, a_3\}$ envy-free under Φ, contradicting the optimality of $\Phi^\star$. And, if h_j is not resolved under Φ, then we have two agents envy-free under Φ, namely a_1 (recipient of h_j) and a_3 (recipient of h) (WLOG). Therefore, we get an agent a_2, such that $a_2 \in EF(\Phi) \setminus EF(\Phi^\star)$.

Otherwise, the only reason that h was not resolved under Φ was the fact that the number of unallocated houses at this point minus span(h) would have been strictly less than the number of unallocated agents, say n'. Since Φ resolves the spans in the increasing order of their sizes, let $\{h_{r_1}, h_{r_2}, \ldots h_{r_i}\}$ be the set of resolved spans and $\{h_{r_{i+1}} \ldots h \ldots h_{r_t}\}$ be the set of unresolved spans under Φ, as considered in Case 4. Suppose m' and n' are the remaining houses and agents at the beginning of Case 4. Note that we have $m' - span(h_j) - \sum_{i \in [r_i]} span(h_i) < n'$ for all $j \in [r_{i+1}, r_t]$. In particular, $m' - span(h) - \sum_{i \in [r_i]} span(h_i) < n'$.

We now argue that if span(h) was resolved under $\Phi^\star$, then there must exist at least one span in the set of resolved spans under Φ, $\{h_{r_1}, h_{r_2}, \ldots h_{r_i}\}$, which is not resolved under $\Phi^\star$. If not, and all the resolved spans under Φ are also resolved under $\Phi^\star$, then it must be that $m' - span(h) - \sum_{i \in [r_i]} span(h_i) \geq n'$. This contradicts the fact that $m' - span(h_{r_{i+1}}) - \sum_{i \in [r_i]} span(h_i) < n'$ since $span(h_{r_{i+1}}) < span(h)$ according to the ordering in Case 4. Therefore, we have a span which is resolved in Φ (say, a_1 and a_2 are two corresponding envy-free agents) but not in $\Phi^\star$. Therefore, at most one of $\{a_1, a_2\}$ is envy-free under $\Phi^\star$ and WLOG, we have $a_2 \in EF(\Phi) \setminus EF(\Phi^\star)$.

(b) Else, if $h' \neq h$, then the peak h remains unallocated under Φ. Then h is definitely a resolved peak. There are two envy-free agents a_1 and a_2 under Φ who receive the resolved peaks and span(h) remains unallocated. Now consider the set of agents $\{a, a_1, a_2\}$ under $\Phi^\star$. Since $a \in EF(\Phi^\star)$, at most one of a_1 and a_2, say a_1, can be envy-free under $\Phi^\star$ (by Lemma 4.58). Therefore, $a_2 \in EF(\Phi) \setminus EF(\Phi^\star)$.

This concludes the argument. □

4.8.2 Single-Dipped Preferences

We begin with an interesting structural claim.

Lemma 4.63. *When the preferences are single-dipped, at most two agents can be envy-free under any complete allocation.*

Proof. Suppose $\{h_1, h_2, \ldots h_m\}$ is the ordering of the houses with respect to which the preferences are single-dipped. Notice that for every agent, either h_1 or h_m is the first ranked house. If both these houses are allocated under an allocation Φ, say to agents i_1 and i_2, such that $rank_{i_1}(h_1) = 1$ and $rank_{i_2}(h_m) = 1$, then it is easy to see that both i_1 and i_2 are envy-free. Notice that these are the only envy-free agents since any other agent i is already envious of the allocation of either h_1 or h_m. This settles our claim in this case. Now suppose h_j is the first house in the ordering that is allocated (to i_1) and h_l is the last one, allocated to i_2, such that both i_1 and i_2 are envy-free. Consider any other agent i. If the dip of i lies to the left of h_j, then since $\Phi(i) \in (h_j, h_l)$, i is envious of the allocation of h_l. If the dip of i lies to between h_j and h_l, then again $\Phi(i) \in (h_j, h_l)$, and i would be envious of both h_j and h_l. Lastly, if the dip of i lies to the right of h_l, then i is envious of the allocation of h_j. This settles the claim. □

Algorithm 2 Minimize #envy for Single-Dipped Preferences

Require: $\{N, H, \succ\}$ and a single dipped axis $\rhd$
Ensure: Allocation Φ that minimizes the number of envious agents
1: $S_1 := \{h \in H | h$ is first ranked house of some agent $i\}$
2: **if** $|S_1| > 1$ **then**
3: $\quad\Phi(i_1) = h_1$ for some $h_1 \in S_1$ & $i_1 \in base(h_1)$
4: $\quad\Phi(i_2) = h_2$ for some $h_2 \in S_1$ & $i_2 \in base(h_2), h_2 \neq h_1$
5: $\quad$Order the remaining agents and let each agent choose its highest ranked house among the remaining houses.
6: **else** $|S_1| = 1$, say $S_1 = \{h\}$
7: $\quad$**if** $m - span(h) \geq n$ **then**
8: $\quad\quad S_{span(h)+1} := \{h' | h'$ is ranked $span(h) + 1$ by some agent $i\}$
9: $\quad\quad$Let $h_1, h_2 \in S_{span(h)+1}, h_1 \neq h_2$
10: $\quad\quad\Phi(i_1) = h_1$ such that $rank_{i_1}(h_1) = span(h) + 1$
11: $\quad\quad\Phi(i_2) = h_2$ such that $rank_{i_2}(h) = span(h) + 1$
12: $\quad\quad U \leftarrow span(h)$ and repeat Step 5 on the houses in $M \setminus U$
13: $\quad$**else** $m - |span(h)| < n$
14: $\quad\quad\Phi(i) = h$ for some $i \in base(h)$ and repeat Step 5
return Φ

Overview of Algorithm 2. Based on Lemma 4.63, the aim is to find two houses such that their allocation creates two envy-free agents. To that end, Algorithm 2 works as follows. Consider the set S_1 of all the houses that are ranked first by any agent. If there are two distinct houses h_1 and h_2 in S_1, then the algorithm allocates these two houses to the agents who like them the most. Notice that the agents who receive h_1 and h_2 are indeed envy-free, and all the remaining agents will be envious by Lemma 4.63. Otherwise, if $|S_1| = 1$ then everyone likes the same house, say h as their first ranked house. If $m - span(h) \geq n$, then we keep $span(h)$ unallocated and construct the set $S_{span(h)+1}$ that contains all the houses ranked at $span(h) + 1$ by any agent. By the definition of $span(h)$, this set must contain at least two distinct houses, say h_1 and h_2. We allocate these to agents i_1 and i_2 such that they rank h_1 and h_2 respectively at $span(h) + 1$. This gives us two envy-free agents (since $span(h)$ is unallocated). The remaining agents are then ordered in an arbitrary manner and allocated their best available house from $m - span(h)$ houses. Otherwise, if $m - span(h) < n$, then we claim that at most one agent can be made envy-free under any allocation and hence, we allocate h non-wastefully.

Theorem 4.64. *Given an instance $\mathcal{I} = (A, H, \succ, \rhd)$ of house allocation with single-dipped preferences, minimizing the number of envious agents admits a polynomial time algorithm.*

Proof. We show that Algorithm 2 correctly outputs an allocation Φ with minimum #envy. If

$\text{EF}(\Phi) = 2$, then we are done by Lemma 4.63. Else, $\text{EF}(\Phi) = 1$. This implies that $m - span(h) < n$. Suppose, for contradiction, there is a complete allocation Φ^* such that $\text{EF}(\Phi^*) = 2$. Then, none of the houses from $span(h)$ could have been allocated under Φ^*. Indeed, all the agents have identical ranking for the houses in $span(h)$, and allocating any of them creates $n-1$ envious agents. Therefore, $span(h)$ must remain unallocated. But then, we have that Φ^* is not a complete allocation since $m - |span(h)| < n$. This contradicts our assumption. □

When there are ties at the dip, we can have more than 2 envy-free agents but we can still minimize $\#envy$ in polynomial time.

Theorem 4.65. *Given an instance of house allocation with single-dipped preferences with ties at the dip, minimizing the number of envious agents admits a polynomial time algorithm.*

Proof. We first prove the following claim:

Lemma 4.66. *When the preferences are single-dipped with ties at the dip, either all the n agents are envy-free or at most 2 agents are envy-free.*

Proof. If there are n houses that are all tied and ranked last by all the agents, then arbitrarily allocating these n houses creates n envy-free agents. Else, if there are less than n, say $n-1$ houses that are all tied and ranked last by all the agents, then even if all of them are allocated among any $n-1$ agents, all such $n-1$ agents will envy the other agent who gets a house outside of the ties. Therefore, if there are less than n houses in the ties, then by Lemma 4.63, there can be at most 2 envy-free agents. □

If there aren't enough houses (at least n) in the ties, then we proceed similarly as in the proof of Theorem 4.64 to arrive at an allocation with two envy-free agents. This settles our claim. □

4.9 Price of Fairness

In this section, in addition to the envy-minimization, we will focus on the social welfare of an allocation, as captured by the sum of the individual agent utilities. An allocation is considered more efficient when it results in a higher level of social welfare. Minimizing the envy objectives can lead to inefficient allocations with poor social welfare. Indeed, our algorithms for OHA, EHA and UHA first check if there are enough (more than n) dummy houses and if so, allocate these dummy houses to everyone, potentially leading to an envy-free solution, but with no social welfare gain. Quantifying this welfare loss, incurred as the cost of minimizing envy is,

	$\mathbf{m = n}$	$\mathbf{m > n}$	
	Non Normalized	**Doubly Normalized**	**Normalized**
OHA / EHA / UHA	1	1	$\frac{n}{2} \leq \text{PoF} \leq n$

Table 4.6: Price of minimizing the number of envious agents, the maximum envy, and the total envy for binary valuations.

therefore, an imperative consideration. In particular, we discuss the worst-case welfare loss under different scenarios when any of the envy objectives is supposed to be minimized, and give tight bounds for the same.

We first define the Price of Fairness in the house allocation setting as follows. We use the notation PoF_{OHA} to denote the fact that the fairness notion under consideration is the minimum number of envious agents. PoF_{EHA} and PoF_{UHA} are defined analogously. When the meaning is clear from the context, we drop the subscript and simply write PoF.

Definition 4.67. *For a house allocation instance $\mathcal{I} := (A, H, \mathcal{P})$ with n agents and m houses, consider an allocation $\Phi^\star$ that maximizes the social welfare, denoted by $SW(\Phi^\star)$. Let Φ be the allocation that minimizes the number of envious agents and $SW(\Phi)$ be the social welfare of Φ. Then, the price of fairness PoF_{OHA} is defined as*

$$PoF_{OHA} = \sup_{\mathcal{I}} \frac{SW(\Phi^\star)}{SW(\Phi)} = \sup_{\mathcal{I}} \frac{\sum_{i \in A} u_i(\Phi^\star(i))}{\sum_{i \in A} u_i(\Phi(i))},$$

where the supremum is taken over all instances with n agents and m houses.

We say that an instance with binary valuations is normalized if every agent likes an equal number of houses. Moreover, if every house is also liked by an equal number of agents, then we say that the instance is doubly normalized. We now present the bounds for PoF. We show that if $m = n$, then $PoF = 1$ for all the three envy minimization objectives. When $m > n$, if the instance is doubly normalized, then $PoF = 1$ but it can be as large as n if we drop the double normalization assumption. The results are summarized in Table 4.6.

Towards proving our results for the case when $m = n$, we first present the following lemma based on the characterization that a matching M in a graph G is maximum if and only if G has no M-augmenting path. An M-augmenting path is a path in G that starts and ends at unmatched vertices (vertices not included in M) and alternates between edges in the matching M and edges that are not in M.

Proposition 4.68 (folklore)**.** *Let G be a graph. For any $A' \subseteq V(G)$, if G contains a matching that saturates A', then G contains a maximum matching that saturates A'.*

Proof. Let M' be the matching that saturates A' in G. If M' is itself a maximum matching, then we are done. Suppose not. Then there exists an M'-augmenting path P in G which starts and ends at a vertex in $A \setminus A'$. Replacing the edges of M in P by the other edges in P gives a strictly larger matching M (with exactly one more edge) which saturates A' and in addition, saturates two vertices from $A \setminus A'$. For every such augmentation, the set of saturated vertices increases by two, keeping the original set of saturated vertices intact. When no augmenting path exists, then by the characterization of maximum matchings, the resulting matching M must be a maximum matching that saturates A'. □

We now consider the case when the number of houses is equal to the number of agents, and show that $PoF = 1$ in this case. Recall that we are under the assumption that every agent values at least one house. Hence, when $m = n$ (and the valuations are binary), in any allocation, an agent either receives a house she values or she is envious; in particular, if agent a is envious, then she envies exactly $d(a)$ other agents, where $d(a)$ is the degree of a in the associated preference graph. We will crucially rely on this fact to prove that $PoF = 1$. Our arguments will also use the correspondence between matchings in the preference graph and allocations: For a matching M in the preference graph, we denote the allocation corresponding to M by Φ_M, which allocates the house $M(a)$ to the agent a and allocates houses to the remaining unmatched agents arbitrarily. Notice that the allocation Φ_M need not be unique.

Lemma 4.69. *For an instance of house allocation with $m = n$ and binary valuations, there exists an allocation that simultaneously maximizes the social welfare and minimizes the number of envious agents.*

Proof. By Proposition 4.1, we know that any allocation that minimizes the number of envious agents has exactly $|M|$ envy-free agents, where M is the maximum matching in the associated preference graph G. It is easy to see that the maximum social welfare of the instance is also exactly $|M|$, hence any allocation, say Φ_M, corresponding to a maximum matching M in G maximizes the welfare and minimizes the number of envious agents. □

Lemma 4.70. *For an instance of house allocation with $m = n$ and binary valuations, there exists an allocation that simultaneously maximizes social welfare and minimizes the maximum envy.*

Proof. Let M be a maximum matching in the associated preference graph G. Since $m = n$, all the agents that are unmatched under M are envious under Φ_M. This implies that at least $n - |M|$ agents are envious. In particular, if $|M| = n$, then under the allocation Φ_M, all agents are envy-free, and Φ_M simultaneously maximizes welfare and minimizes the maximum envy, and thus the lemma trivially holds. So, assume that $|M| < n$. As M is a maximum matching, we can conclude that there is no matching that saturates all the agents.

Let us now order the agents in the non-increasing order of their degrees, that is, $a_1, a_2, \ldots, a_n$ such that $d(a_1) \geq d(a_2) \geq \cdots \geq d(a_n)$. Let $p_1 \in [n]$ be the least index such that there is no matching that saturates all of $a_1, a_2, ..., a_{p_1}$. This implies that there is a matching that saturates $a_1, a_2, \ldots a_{p_1-1}$ but none that saturates the agents $a_1, a_2, ..., a_{p_1}$. That is, in any allocation, at least one agent among $a_1, \ldots a_{p_1}$ is envious. Thus, the maximum envy of the instance is at least $d(a_{p_1})$. But there is a matching that saturates $a_1, a_2, \ldots, a_{p_1-1}$, which can be extended to a corresponding allocation; and under such an allocation, a_{p_1} would be the first envious agent (first in the ordering $a_1, a_2, \ldots, a_n$). Therefore, we can conclude that the maximum envy of the instance is precisely equal to $d(a_{p_1})$. Now, we only have to prove that there is indeed a *maximum* matching that saturates $a_1, a_2, \ldots, a_{p_1-1}$; and Proposition 4.68 guarantees that G does contain such a maximum matching (we simply need to apply Proposition 4.68 with $A' = \{a_1, a_2, \ldots, a_{p_1-1}\}$). This implies that there is a welfare-maximizing allocation that also minimizes the maximum envy. □

Lemma 4.71. *For an instance of house allocation with $m = n$ and binary valuations, there exists an allocation that simultaneously maximizes social welfare and minimizes total envy.*

Proof. Consider an allocation, say Φ, that minimizes the total envy. We claim that the matching M corresponding to Φ in the associated preference graph G is a maximum matching. If M is not a maximum matching, then we have an augmenting path, and we get a strictly larger matching M' such that the vertices saturated by M are also saturated under M'. Now consider the allocation corresponding to M', say $\Phi_{M'}$. Since $|M| < |M'|$, the number of envious agents under $\Phi_{M'}$ is strictly less than those under Φ; also all envy-free agents under Φ remain envy-free under $\Phi_{M'}$. These arguments imply that the total envy under $\Phi_{M'}$ is less than the total envy under Φ. This is a contradiction to the fact that Φ minimizes the total envy. □

The following result now follows from *Lemma* 4.69, *Lemma* 4.70 and *Lemma* 4.71.

Theorem 4.72. *For an instance of house allocation with $m = n$ and binary valuations, $PoF = 1$ for all the three envy-minimization objectives.*

Theorem 4.72 tells us that when $m = n$, there is an allocation that simultaneously maximizes welfare and minimizes any *one* of the three measures of envy. This does raise the following question: Can we simultaneously maximize welfare and minimize *all* three measures of envy? We show that we can indeed do this; that is, there is an allocation that simultaneously minimizes the number of envious agents, the maximum and total envy while maximizing social welfare.

Theorem 4.73. *For an instance of house allocation with $m = n$ and binary valuations, there is an allocation that simultaneously minimizes the number of envious agents, the maximum and total envy, while maximizing social welfare. Moreover, we can compute such an allocation in polynomial time.*

Proof. Let M be a maximum matching in the associated preference graph G. First, the corresponding allocation Φ_M maximizes social welfare. Since $m = n$, all the agents that remain unmatched under M are envious under Φ_M. This implies that at least $n - |M|$ agents are envious under any allocation. Thus, if $|M| = n$, then every agent is envy-free under Φ_M, and thus, the theorem trivially holds. So, assume from now on that $|M| < n$.

We first order the agents in the non-increasing order of their degrees, that is, $a_1, a_2, \ldots, a_n$ such that $d(a_1) \geq d(a_2) \geq \cdots \geq d(a_n)$. Let $p_1 \in [n]$ be the least index such that G does not contain a matching that saturates $\{a_1, a_2, ..., a_{p_1}\}$. That is, there is a matching that saturates all of $a_1, a_2, \ldots a_{p_1-1}$ but none that saturates all of $a_1, a_2, ..., a_{p_1}$. Now, let $p_2 \in [n]$ be the least index such that $p_2 > p_1$ and G does not contain a matching that saturates $\{a_1, a_2, \ldots a_{p_2}\} \setminus \{a_{p_1}\}$ (if such an index p_2 exists); In general, having defined $p_1, p_2, \ldots, p_{i-1}$, we define p_i to be the least index in $[n]$ such that $p_i > p_{i-1}$ and G does not contain a matching that saturates $\{a_1, a_2, \ldots a_{p_i}\} \setminus \{a_{p_1}, a_{p_2}, \ldots, a_{p_{i-1}}\}$. Let $p_1, p_2, \ldots, p_k$ (for some $k \geq 1$), be the indices defined this way. By their definition, G contains a matching, say M', that saturates $\{a_1, a_2, \ldots a_n\} \setminus \{a_{p_1}, a_{p_2}, \ldots a_{p_k}\}$. Consider the corresponding allocation $\Phi_{M'}$; recall that $\Phi_{M'}$ allocates houses along the edges of M' and the remaining agents, i.e., $a_{p_1} \ldots a_{p_k}$, receive the remaining houses in an arbitrary manner.

Before we proceed further, let us observe that we can indeed compute M' and hence the corresponding allocation $\Phi_{M'}$ in polynomial time. The arguments we have used so far are constructive. To compute M', we only need to identify the indices $p_1, p_2, \ldots, p_k$ and find a matching that saturates $\{a_1, a_2, \ldots, a_n\} \setminus \{a_{p_1}, a_{p_2}, \ldots, a_{p_k}\}$. Notice that to find each p_i, we only need to check if G contains a matching that saturates $\{a_1, a_2, \ldots, a_j\} \setminus \{a_{p_1}, a_{p_2}, \ldots, a_{p_{i-1}}\}$ for each $j > p_{i-1}$, which we can do in polynomial time.

We will now show that the allocation $\Phi_{M'}$ satisfies all the properties required by the statement of the theorem; that is, $\Phi_{M'}$ maximizes welfare and minimizes the number of envious agents, the maximum envy and the total envy. To that end, observe first that M' does not saturate any of the agents $a_{p_1}, a_{p_2}, \ldots, a_{p_k}$. Otherwise, let p_i be the least index in $\{p_1, p_2, \ldots, p_k\}$ such that M' saturates a_{p_i}. In particular, M' is a matching that saturates $\{a_1, a_2, \ldots, a_{p_i}\} \setminus \{a_{p_1}, a_{p_2}, \ldots, a_{p_{i-1}}\}$, which contradicts the definition of p_i. By the same reasoning, we can also conclude that that M' is indeed a maximum matching. If not, then there is a larger matching, say M'', which also saturates at least one of the a_{p_i} in addition to saturating $\{a_1, a_2, \ldots a_n\} \setminus \{a_{p_1}, a_{p_2}, \ldots a_{p_k}\}$ (by Proposition 4.68), which again will lead to a contradiction.

Now, the fact that M' is a maximum matching immediately implies that the corresponding allocation $\Phi_{M'}$ maximizes social welfare and minimizes the number of envious agents. We will now use the same arguments we used in the proof of Lemma 4.70 to show that $\Phi_{M'}$ minimizes the maximum envy. Notice that the set of agents that are *not* saturated by M' is precisely $\{a_{p_1}, a_{p_2}, \ldots, a_{p_k}\}$. This fact, along with the fact that M' is a maximum matching, implies that the set of envious agents under $\Phi_{M'}$ is precisely $\{a_{p_1}, a_{p_2}, \ldots, a_{p_k}\}$. In particular, a_{p_1} is the first envious agent under $\Phi_{M'}$ (again, first in the ordering $a_1, a_2, \ldots, a_n$). Thus, the maximum envy of the allocation $\Phi_{M'}$ is $d(a_{p_i})$. By the definition of p_1, in any allocation, at least one agent among $a_1, \ldots a_{p_1}$ is envious, and hence the maximum envy in any allocation is at least $d(a_{p_i})$ (because $d(a_1) \geq d(a_2) \geq \cdots \geq d(a_n)$). From these arguments, we can conclude that $\Phi_{M'}$ minimizes the maximum envy.

To complete the proof, now we only need to argue that $\Phi_{M'}$ minimizes the total envy. To that end, we first prove the following claim.

Claim 4.74. *For every $i \in [k]$, every matching in G does* not *saturate at least i agents from the set $\{a_1, a_2, \ldots, a_{p_i}\}$. Or equivalently, for every $i \in [k]$, at least i agents from the set $\{a_1, a_2, \ldots, a_{p_i}\}$ are envious under every allocation.*

Proof. Suppose for a contradiction that G contains a matching M''' such that the number of agents from $\{a_1, a_2, \ldots, a_{p_i}\}$ that are not saturated by M''' is strictly less than i. Therefore, M''' saturates at least one agent from the set $\{a_{p_1}, a_{p_2}, \ldots, a_{p_i}\}$; let p_j be the least such index such that M''' saturates a_{p_j}. But then M''' saturates $\left\{a_1, a_2, \ldots, a_{p_j}\right\} \setminus \left\{a_{p_1}, a_{p_2}, \ldots, a_{p_{j-1}}\right\}$, which contradicts the definition of p_j. Notice that the second sentence in the statement of the claim is merely a restatement of the first, because of the correspondence between matchings and allocations: Under any allocation, the envy-free agents and the houses they received form a matching in G. □

Now, to complete the proof of the theorem, recall that the set of envious agents under $\Phi_{M'}$ is precisely $\{a_{p_1}, a_{p_2}, \ldots, a_{p_k}\}$, and thus the total envy under $\Phi_{M'}$ is precisely $\sum_{i=1}^{k} d(a_{p_i})$. Also, as we have already argued, $\Phi_{M'}$ minimizes the number of envious agents, and hence we can conclude that at least k agents are envious under *every* allocation. Assume now for a contradiction that $\Phi_{M'}$ does not minimize the total envy. Let Φ be an allocation that minimizes the total envy. Again, at least k agents are envious under Φ; let $1 \leq q_1 < q_2 < \cdots < q_k \leq n$ be such that the agents $a_{q_1}, a_{q_2}, \ldots, a_{q_k}$ are the first k envious agents under Φ (first in the ordering $a_1, a_2, \ldots, a_n$). Thus the total envy under Φ is at least $\sum_{i=1}^{k} d(a_{q_i})$. Now, by our assumption, $\Phi_{M'}$ does not minimize the total envy and Φ does, and thus $\sum_{i=1}^{k} d(a_{p_i}) > \sum_{i=1}^{k} d(a_{q_i})$; we will derive a contradiction from this. Since $\sum_{i=1}^{k} d(a_{p_i}) > \sum_{i=1}^{k} d(a_{q_i})$, there exists an index $i \in [k]$ such that $d(a_{p_i}) > d(a_{q_i})$; let $i \in [k]$ be the least index such that $d(a_{p_i}) > d(a_{q_i})$. Hence $p_i < q_i$, which implies that the set $\{a_1, a_2, \ldots, a_{p_i}\}$ contains at most $i-1$ of the agents $a_{q_1}, a_{q_2}, \ldots, a_{q_i}$. But by the definition of the q_js, the agents $a_{q_1}, a_{q_2}, \ldots, a_{q_k}$ are the first k envious agents under Φ. We can thus conclude that the set $\{a_1, a_2, \ldots, a_{p_i}\}$ contains at most $i-1$ agents who are envious under Φ, which contradicts Claim 4.74.

We have thus shown that $\Phi_{M'}$ minimizes the total envy, and this completes the proof of the theorem. □

We now consider the case when $m > n$ and the setting of the doubly normalized valuations. The following result follows from the proof of Theorem 5 in Bhaskar et al. (2023), where they show that for such structured valuations, every good can be assigned non-wastefully such that every agent derives the value of at least one. In our setting, we can find an envy-free allocation that is non-wasteful and hence, welfare maximizing as well, proving the following result.

Corollary 4.75. *For $m > n$ and doubly normalized binary valuations, PoF $= 1$ for all three envy-minimization objectives.*

It is easy to see that when there are dummy houses, then there can be instances where the fair allocation can be highly inefficient. Consider the case when number of dummy houses is at least the number of agents. Then, irrespective of the individual valuations, every agent can get a dummy house and be envy-free, leading to no social welfare at all. But, the following result suggests that even when there aren't any dummy houses, there are instances where the price of fairness can be high.

Theorem 4.76. *For $m > n$ and binary normalized valuations, PoF $= \Theta(n)$, for all the three* *envy optimization objectives, even when there are no dummy houses to begin with.*

Proof. We first show that $PoF \geq n/2$. We say that a house is allocated non-wastefully if the receiving agent values it at 1. Consider an instance with $2n$ agents and $3n$ houses, such that $n > 2$. Agent 1 likes the first n houses, 2 likes the next n houses, and all of the remaining $2n - 2$ agents collectively like the last n houses. The instance does not have any dummy houses. Note that any welfare-maximizing allocation $\Phi^\star$ allocates $n + 2$ houses non-wastefully. So $SW(\Phi^\star) \geq n + 2$. But in this instance, there exists an envy-free allocation Φ that allocates exactly 2 houses non-wastefully to the first two agents. The last n houses remain unallocated under Φ, as they create Hall's violator for the set of last $2n - 2$ agents. Therefore, the last $2n - 2$ agents receive the houses they don't like, and the set of their liked houses remains unallocated. This implies that $SW(\Phi) \leq 2$. The PoF for this particular instance is, therefore, $(n+2)/2 \approx n$. Since the instance had $2n$ agents to begin with, we have that in general, $PoF \geq n/2$.

We now show that $PoF \leq n$. Under any arbitrary instance with n agents and binary valuations, the maximum possible social welfare under any allocation is n. Consider an envy-free allocation Φ, if it exists. We claim that there is at least one house that is allocated non-wastefully under Φ. Suppose not. Then every house is allocated wastefully, and hence, $SW(\Phi) = 0$. But since there are no dummy houses, there is at least one agent a who likes (a subset of) the allocated houses. But since welfare is zero, a is an envious agent, which contradicts the fact that we started with an envy-free allocation. Therefore, $SW(\Phi) \geq 1$. This implies that $PoF \leq n$.

If an envy-free allocation does not exist, consider Φ to be an allocation that minimizes the number of envious agents. Suppose the number of envious agents under Φ is k and let a be one such envious agent. If $SW(\Phi) = 0$, then everyone, including a receives a house that they do not like. Now a is envious because one of his liked houses, say h is allocated to some agent Φ' and that too wastefully. Consider the re-allocation of the house h to a and the house $\Phi(a)$ to a'. Then a becomes envy-free, and no agent (including a') becomes newly envious of this re-allocation, since the set of allocated houses is exactly the same and $SW(\Phi) = 0$. This re-allocation has $k - 1$ envious agents, which contradicts the fact that we started with an optimal allocation. Therefore, $SW(\Phi) \geq 1$. This implies that $PoF \leq n$. The argument for EHA and UHA is analogous. This settles our claim. □

In the proof of *Theorem* 4.76, although the number of agent types in the lower bound instance is 3, the PoF is linear in the number of agents. This implies that we do not expect better bounds for instances with few agent types.

4.10 Conclusion

We studied three kinds of natural quantifications of envy to be minimized in the setting of house allocation: OHA, EHA, and UHA. Most of our results are summarized in Table 4.1. We also show that both OHA and EHA are FPT when parameterized by the total number of house types or agent types. We also look at the price of fairness in the context of house allocation and give tight bounds for the same.

We leave several questions related to UHA open. The complexity of EHA and UHA is open for single-peaked preferences. It will be interesting to see how ties play a role in the complexity of OHA and settle the complexity of OHA on single-plateaued preferences.

Chapter 5

Price of Equitability

"For everything you have missed, you have gained something else, and for everything you gain, you lose something else."

- Ralph Waldo Emerson, *Compensation (Essays, First Series)*

5.1 Introduction

Tradeoffs are inevitable when we pursue multiple optimization objectives that are typically not simultaneously achievable. Quantifying such tradeoffs is a fundamental problem in computation, game theory, and economics. Our focus in this work is on the "price of fairness" in the context of fair division problems, which is a notion that captures tradeoffs between *fairness* and *welfare*.

Recall that a fair division instance in the discrete setting involves a set of n *agents* $N = \{1, 2, \ldots, n\}$, m indivisible *goods* $M = \{g_1, \ldots, g_m\}$, and $\mathcal{V} := \{v_1, v_2, \ldots, v_n\}$, a *valuation profile* consisting of each agent's valuation of the goods. For any agent $i \in N$, her valuation function $v_i : 2^M \to \mathbb{N} \cup \{0\}$ specifies her numerical value (or *utility*) for every subset of goods in M. We will assume that the valuations are normalized, that is, for all $i \in N$, $v_i(M) = W$, where W is the normalization constant. Our goal is to devise an *allocation* of goods to agents; defined as an ordered partition[1] of the m goods into n "bundles", where the bundles are (possibly empty) subsets of M, and the convention is that the i^{th} bundle in the partition is the set of goods assigned to the agent i.

The *welfare* of an allocation is a measure of the utility that the agents derive from the allocation. For additive valuations, the individual utility that an agent i derives from their bundle Φ_i is simply the sum of the values that they have for the goods in Φ_i. The overall welfare of an allocation Φ is typically defined by aggregating individual utilities in various ways. Not surprisingly, there are several notions of welfare corresponding to different approaches to consolidating the individual utilities. For instance, the *utilitarian social welfare* is the sum of utilities of agents under Φ; the *egalitarian social welfare* is the lowest utility achieved by any agent with respect to Φ; and the *Nash social welfare* is the geometric mean of utilities of agents under Φ. One may view all of these welfare notions as special cases of the p-mean welfare (where $p \in (-\infty, 0) \cup (0, 1]$), which is defined as the generalized p-mean of utilities of agents under Φ, i.e., $W_p(\Phi) := \left(\frac{1}{n}\sum_{i \in N}\left(v_i(\Phi_i)\right)^p\right)^{1/p}$. Note that for $p > 1$, the p-mean optimal allocation tends to concentrate the distribution among fewer agents (consider the simple case of two identical agents with additive valuations who value each of two goods at 1), which is contrary to the spirit of fairness. Hence we focus on $p \leqslant 1$.

A natural goal for a fair division problem is to obtain an allocation that maximizes the overall welfare. However, observe that optimizing exclusively for welfare can lead to undesirable allocations. To see this, consider an instance with additive valuations where all the valuation functions are the same, i.e., the utility of any good g is the same for all agents in N. In this case, *every* allocation has the same utilitarian welfare. So, when we only optimize for—in this example, utilitarian—welfare, we have no way of distinguishing between, say, the allocation that allocates all goods to one agent and one that distributes the goods more evenly among the agents. To remedy this, one is typically interested in allocations that not only maximize welfare, but are also "fair".

The price of fairness is informally the cost of achieving a specific fairness notion, where the cost is viewed through the lens of a particular welfare concept. For a fairness notion $\mathcal{F}$ (such as EQ1 or EF1) and a welfare notion $\mathcal{W}$ (such as egalitarian or utilitarian welfare), the price of fairness is the "worst-case ratio" of the maximum welfare (measured by $\mathcal{W}$) that can be obtained by *any* allocation, to the maximum welfare that can be obtained among allocations that are fair according to $\mathcal{F}$. For example, it is known from the work of Caragiannis et al. (2019b) that under

[1] Unless otherwise specified, we implicitly assume that allocations are *complete*, i.e., every good is assigned to some agent.

additive valuations, any allocation that maximizes the Nash social welfare satisfies EF1. Thus, the price of fairness of EF1 with respect to Nash social welfare is 1. Further, Barman et al. (2020b) show that the price of EF1 with respect to utilitarian welfare is $O(\sqrt{n})$ for normalized subadditive valuations.

In this contribution, we focus on bounds for the price of fairness in the context of EQ1, a notion that we will henceforth refer to as the *price of equity* (PoE) when there is no ambiguity. Much of the existing literature on the price of fairness analysis focuses on *specific* welfare measures (e.g., utilitarian, egalitarian, and Nash social welfare). Our work deviates from this trend by analyzing the *entire* family of generalized p-mean welfare measures (i.e., for *all* $p \leqslant 1$); recall that this captures the notions of egalitarian, utilitarian, and Nash welfare as special cases. Our results therefore address the behavior of the price of equity for a wide spectrum of welfare notions.

Further, we obtain bounds in terms of the *number of agent types* — which we denote by r — rather than the total number of agents. The number of agent types of a fair division instance is the largest number of agents whose valuations are mutually distinct: in other words, it is the number of distinct valuation functions in the instance. Note that the number of agent types is potentially *much* smaller than the total number of agents. The notion of agent types has been popular in the fair division literature for the reason that it is a natural quantification of the "simplicity" of the structure of the instance as given by the valuations. Note that the well-studied special case of identical valuations is equivalent to the class of instances for which $r = 1$, and therefore one might interpret parameterizing by r as a smooth generalization of the case of identical valuations. For a representative selection of studies that focus on instances with a bounded number of agent types, we refer the reader to (Bliem et al., 2016; Bouveret et al., 2017; Garg et al., 2021; Brânzei et al., 2016).

We restrict ourselves to the setting of *binary submodular* (also known as matroid rank) valuations. A valuation function v_i is submodular if for any subsets of goods $S, S' \subseteq M$ such that $S \subseteq S'$, and for any good $g \notin S'$, $v_i(S \cup g) - v_i(S) \geqslant v_i(S' \cup g) - v_i(S')$. That is, the marginal value of adding g to S is at least that of adding g to a superset of S. Valuation v_i is binary submodular if for any subset of goods $S \subseteq M$ and any good g, the marginal value $v_i(S \cup g) - v_i(S) \in \{0, 1\}$. Binary submodular valuations are frequently studied in fair division and are considered to be a useful special case such as in allocating items under a budget, or with exogenous quotas (Benabbou et al., 2021; Viswanathan and Zick, 2023). It also provides algorithmic leverage: many computational questions of interest that are hard in general turn out to be tractable once we restrict our attention to binary submodular valuations. As an example, while it is NP-hard to compute a Nash social welfare maximizing allocation even for identical additive valuations (Roos and Rothe, 2010), such an allocation can be computed in polynomial time under binary submodular valuations in conjunction with other desirable properties such as strategyproofness, envy-freeness up to any good, and ex-ante envy-freeness (Babaioff et al., 2021).

A strict subset of binary submodular valuations is the class of *binary additive* valuations—this is a subclass of additive valuations wherein each value $v_i(g)$ is either 0 or 1. Binary additive valuations provide a simple way for agents to express their preferences as they naturally align with the idea of agents "approving" or "rejecting" goods. These are also widely studied in the literature on fair division, for example, see (Ortega, 2020; Kyropoulou et al., 2020; Babaioff et al., 2021; Amanatidis et al., 2021; Aleksandrov and Walsh, 2020; Aziz and Rey, 2021). In the case of voting too, binary additive valuations play a role. Darmann and Schauer (2015) consider the complexity of maximizing Nash social welfare when scores inherent in classical voting procedures are used to associate utilities with the agents' preferences, and find that the case of approval ballots — which happen to lead to binary additive valuations — are a tractable subclass.

Our Contributions and Techniques

We now turn to a discussion of our findings (see Table 5.1 for a summary of our results for binary additive valuations). Given an instance of fair division with binary submodular valuations, let $\Phi^\star$ be an allocation that maximizes the Nash social welfare. It is implicit from the results of Benabbou et al. (2021) that $\Phi^\star$ also has maximum p-mean welfare for all $p \leqslant 1$ (for details, refer to Section 5.3). We show an analogous result for EQ1 allocations, by demonstrating that there exists an EQ1 allocation (which we call Φ^T, or the *truncated allocation*) that maximizes the p-mean welfare for all p. To this end, in allocation $\Phi^\star$, let i be an agent with minimum value, and let $\ell = v_i(\Phi_i^\star)$. If the allocation is not already EQ1, then we reallocate "excess" goods from the bundles of agents who value their bundles at more than $\ell + 1$ and give them to agent i. Notice that agent i must have marginal value 0 for all these excess goods, otherwise this reallocation would improve the Nash welfare. It turns out that this allocation Φ^T is EQ1 and also has — among EQ1 allocations — the highest p-mean welfare.

Theorem 5.1. *For any $p \in \mathbb{R} \cup \{-\infty\}$ and binary submodular valuations, the p-mean welfare of the truncated allocation Φ^T is at least that of any other* EQ1 *allocation.*

Notice that together with the result of Benabbou et al. (2021), Theorem 5.1 allows us to focus only on the maximum Nash social welfare allocation $\Phi^\star$ and the truncated allocation Φ^T to

PoE	Agent types (r)	
	Lower bound	Upper bound
Utilitarian welfare ($p = 1$)	$r - 1$	r
Nash welfare ($p = 0$)	$\frac{(r-1)/e}{\ln(r-1)}$	$\frac{(r-1)}{\ln(r-1)/e}$
Egalitarian welfare ($p \to -\infty$)	1	1 (Sun et al., 2023b)
$p \in (0, 1)$	$p(r-1)/e$	$2r - 1$
$p \in (-1, 0)$	$2^{1/p}(r-1)^{1/(1-p)}$	$2^{-1/p}(-p)^{1/p(1-p)}(r-1)^{1/(1-p)}$
$p \leqslant -1$	$2^{1/p}(r-1)^{1/(1-p)}$	$2(r-1)^{1/(1-p)}$

Table 5.1: Summary of results for the price of equity (PoE). Each cell indicates either the lower or the upper bound (columns) on PoE for a specific welfare measure (rows) as a function of the number of *agent types* r. Our contributions are highlighted by shaded boxes. The lower bounds are from Theorem 5.2, while the upper bounds are shown in Theorem 5.3 and Theorem 5.4. Section 5.8 presents the upper and lower bounds graphically as a function of r, for $p = 1$, $p = 0$, $p = -1$, and $p = -10$.

obtain upper bounds on the PoE for all $p \leqslant 1$ simultaneously.

We now describe our bounds on the PoE for binary additive valuations. Our lower bounds are based on varying parameters in a single basic instance. The parameters are r, the number of agent types, and W, the normalization constant for the agents. Given r and W, the instance has $m = rW$ goods, divided into r groups of W goods each. The groups are $M_1, M_2, \ldots, M_r$. There are $W + 1$ agents who value all the goods in M_1 at 1 each and everything else at 0. Further, for each $2 \leqslant i \leqslant r$, we have exactly one agent who values the goods in M_i and nothing else.

To summarize, we have $W + 1$ agents of the first type, who have a common interest in W goods. Any allocation must leave one of these agents with zero value. Beyond these coveted goods, each of the remaining goods is valued by exactly one agent. A welfare maximizing allocation will allocate each good in $M_2 \cup \cdots \cup M_r$ to the unique agent who values it; however, an EQ1 allocation is constrained by the fact that an agent of the first type must get value 0.[2] It turns

[2]For $p \leqslant 0$, we use the standard convention that allocation Φ is a p-mean optimal allocation if (a) it maximizes number of agents with positive value, and (b) among all allocations that satisfy (a), maximizes the p-mean welfare when restricted to agents with positive value.

out that using this family of instances, we can obtain the following bounds. We note here that we rely on normalization as a crucial assumption for deriving both the lower bounds and the upper bounds. This puts equitability in contrast with envy-freeness—the compatibility between EF1 and p-mean welfares hold even without normalization for binary submodular valuations (Barman et al., 2020a).

Theorem 5.2 (PoE lower bounds). *Let $s := r - 1$. The price of equity for binary additive valuations is at least:*

1. *s, for $p = 1$,*
2. *$\frac{p}{e}s$, for $p \in (0, 1)$,*
3. *$\frac{s}{e \ln s}$, for $p = 0$,*
4. *$2^{1/p}s^{1/(1-p)}$, for $p < 0$.*

We now turn to the upper bounds for binary additive valuations. It turns out that the PoE for utilitarian welfare is bounded by the *rank* of the instance, where the rank is simply the rank of the $n \times m$ matrix $\{v_i(g_j)\}_{1 \leqslant i \leqslant n; 1 \leqslant j \leqslant m}$. Observe that the rank is a lower bound for the number of agent types, so this result also implies an upper bound of r on the PoE. In fact, the rank could be logarithmic in the number of agent types, and hence this is a significantly tighter bound than the number of agent types.

To obtain this upper bound, in allocation Φ^T (which, as shown in Theorem 5.1, maximizes the utilitarian welfare among EQ1 allocations) we show that the number of wasted goods is at most $m(1 - \frac{1}{k})$, where k is the rank of the instance. This implies the theorem.

Theorem 5.3 (Utilitarian PoE upper bound). *Under binary additive valuations and utilitarian welfare as the objective, the price of equity is at most the rank of the instance.*

For other values of p, we obtain the following upper bounds.

Theorem 5.4 (PoE upper bounds). *Let $s := r - 1$. The price of equity for binary additive valuations is at most*

1. *$1 + s$ for $p = 1$*
2. *$1 + 2s$ for $p \in (0, 1)$*
3. *$\frac{s}{\ln(s/e)}$ for $p = 0$ (i.e., the Nash social welfare)*
4. *$s^{1/(1-p)}2^{-1/p}(-1/p)^{1/p(p-1)}$ for $p \in (-1, 0)$*
5. *$2s^{1/(1-p)}$ for $p \leqslant -1$*

We note that for any fixed p, the lower bounds (Theorem 5.2) and upper bounds (Theorem 5.4) are within a constant factor of each other.

Conceptually, for the proof of the upper bounds, we show that the worst case for the PoE is in fact the family of instances used for showing our lower bounds in Theorem 5.2. In particular, any instance can be transformed into one belonging to the lower bound family, without improving the PoE. Note that for the PoE, we can focus on the allocations $\Phi^\star$ and Φ^T irrespective of the p-mean welfare measure, since these maximize the p-mean welfare for all $p \leqslant 1$ simultaneously. For a given instance, let l be the minimum value of any agent in $\Phi^\star$. We divide the agent types into two groups: types for which every agent has value at most $l+1$ in $\Phi^\star$, and types for which an agent has value $> l+1$. Note that for a type in the first group, each agent of this type retains her value in Φ^T, while for a type in the second group, the value of each agent of this type is truncated to $l+1$. Our proof shows that agents in the first group must have total value at least W, as in the lower bound example. We also use W as an upper bound for the total value of each agent type in the second group. Then letting α be the fraction of agents in the first group, and optimizing over α, gives us the required upper bounds.

We then consider the PoE for binary additive valuations with the additional structure that both the rows and the columns are normalized. That is, each agent values exactly W goods, and each good is valued by exactly W_c agents. For such *doubly normalized* instances, we show the PoE is 1.

Theorem 5.5. *For doubly normalized instances under binary additive valuations, the PoE for the p-mean welfare is 1 for all $p \leqslant 1$.*

Finally, we obtain bounds on the PoE for binary submodular valuations. For identical valuations, it follows from similar results for EF1 that the PoE is 1.

Proposition 5.6. *When all agents have identical binary submodular valuations, the PoE is 1 for p-mean welfare measure for all $p \leqslant 1$.*

However, this is the best that can be obtained, in the sense that even with just *two* agent types, the PoE for utilitarian welfare is at least $n/6$, where n is the number of agents. Hence we cannot obtain bounds on the PoE that depend on the number of agent types, as we did for binary additive valuations.

Theorem 5.7. *The PoE for utilitarian welfare when agents have binary submodular valuations is at least $n/6$ (where n is the number of agents), even when there are just two types of agents.*

Nonetheless, we do obtain an upper bound of $2n$ on the PoE for binary submodular valuations.

Theorem 5.8. *For binary submodular valuations and any $p \leqslant 1$, the PoE for p-mean welfare is at most $2n$.*

Related Work

The notion of *price of fairness* was proposed in the works of Bertsimas et al. (2011) and Caragiannis et al. (2012). These formulations were inspired from similar notions in game theory—specifically, *price of stability* and *price of anarchy*—that capture the loss in social welfare due to strategic behavior.[3] Caragiannis et al. (2012) studied the price of fairness for divisible and indivisible resources under three fairness notions: *proportionality* (Steinhaus, 1948), *envy-freeness* (Gamow and Stern, 1958; Foley, 1967), and *equitability* (Dubins and Spanier, 1961). For indivisible resources, they defined price of fairness only with respect to those instances that admit some allocation satisfying the fairness criterion.

Recently, Bei et al. (2021) studied price of fairness for indivisible goods for fairness notions whose existence is guaranteed; in particular, they studied *envy-freeness up to one good* (EF1) and *maximum Nash welfare* allocations.[4] In a similar vein, Sun et al. (2023a) studied price of fairness for allocating indivisible *chores* for different relaxations of envy-freeness and maximin share. Perhaps closest to our work is a recent paper by Sun et al. (2023b). This work studies price of equitability for EQX for indivisible goods as well as indivisible chores under utilitarian and egalitarian welfare. The valuations are assumed to be additive but not necessarily binary. For indivisible goods, the price of equity is shown to be between $n-1$ and $3n$, where n is the number of agents, while for egalitarian welfare, a tight bound of 1 is provided.

5.2 Preliminaries

In this chapter, we restrict our attention to equitability and its approximations as our fairness criteria, coupled with p-mean welfare notions. We refer the reader to Section 1.2 for the definitions of the input instance, fairness, and welfare notions.

[3] Price of anarchy was defined by Koutsoupias and Papadimitriou (2009) and subsequently studied in the notable work of Roughgarden and Tardos (2002), while price of stability was defined by Anshelevich et al. (2008).

[4] The EF1 notion was formulated by Budish (2011) although subsequently it was observed that an algorithm of Lipton et al. (2004) achieves this guarantee for monotone valuations. The Nash social welfare function was originally proposed in the context of the bargaining problem (Nash Jr, 1950) and subsequently studied for resource allocation problems by Eisenberg and Gale (1959).

We will primarily focus on binary submodular valuations in Section 5.3 and 5.7, and on binary additive valuations in Section 5.4, 5.5, and 5.6.

Agents i and j are said to be of the same *type* if their valuation functions are identical, i.e., if for every subset of goods $S \subseteq M$, $v_i(s) = v_j(S)$. We will use r to denote the number of distinct agent types in an instance. Further, an instance is *normalized* if for some constant W, $v_i(M) = W$ for all agents i. Our work focuses on instances with normalized valuations since there are trivial instances where the price of equity for any p-mean welfare for $p \in \mathbb{R}$ is large without this assumption (e.g., the simple instance with 2 agents and k goods, where agent 1 has value 1 for the first good and zero for the others, and agent 2 has value 1 for all goods, has price of equity $k/3$ for the utilitarian welfare).

Price of fairness.

Given a fairness notion $\mathcal{F}$ (e.g., EQ1) and a p-mean welfare measure, the price of fairness of $\mathcal{F}$ with respect to a welfare measure $\mathcal{W}_p$ is the supremum over all fair division instances with n agents and m goods of the ratio of the maximum welfare (according to $\mathcal{W}_p$) of any allocation and the maximum welfare of any allocation that satisfies $\mathcal{F}$.

Formally, let $\mathcal{I}_{n,m}$ denote the set of all fair division instances with n agents and m items. Let $\mathcal{A}(I)$ denote the set of all allocations in the instance I, and further let $\mathcal{A}_{\mathcal{F}}(I)$ denote the set of all allocations in the instance I that satisfy the fairness notion $\mathcal{F}$.

Then, the price of fairness (PoF) of the fairness notion $\mathcal{F}$ with respect to the welfare measure $\mathcal{W}_p$ is defined as:

$$\mathrm{PoF}(\mathcal{F}, \mathcal{W}_p) := \sup_{I \in \mathcal{I}_{n,m}} \frac{\max_{\Phi^* \in \mathcal{A}(I)} \mathcal{W}_p(\Phi^*)}{\max_{\Phi \in \mathcal{A}_{\mathcal{F}}(I)} \mathcal{W}_p(\Phi)}.$$

As indicated earlier, throughout this chapter we will focus on equitability up to one good (EQ1) as the fairness notion of choice (i.e., $\mathcal{F}$ is EQ1). For notational simplicity, we will just write PoF instead of $\mathrm{PoF}(\mathcal{F}, \mathcal{W})$ whenever the welfare measure $\mathcal{W}$ is clear from context, and we will refer to this ratio as the price of equity (PoE) whenever the fairness notion in question is EQ1.

Some properties of p-mean welfare

We state here some basic properties of the p-mean welfare that will be useful in due course.

Claim 5.9. *For all $p < 1$, the p-mean welfare is a concave function of the agent valuations.*

This proof was shown by Ahle (2020). We reproduce it here for completeness.

Proof. Consider $f(x) = \left(\Sigma_i x_i^p\right)^{1/p}$. The Hessian matrix H is then given by:

$$H_{ij} = (1-p)f^{1-2p}A \quad \text{where } A_{ij} = \begin{cases} -x_i^{p-2}\Sigma_{k\neq i}x_k^p & \text{if } i = j \\ x_i^{p-1}x_j^{p-1} & \text{if } i \neq j \end{cases}$$

For $p \leqslant 1$, the matrix H is negative semidefinite, since the initial coefficient $(1-p)f^{1-2p} \geqslant 0$, and for any vector v,

$$v^T Av \;=\; \left(\sum_{i=1}^{n} v_i x_i^{p-1}\right)^2 - \sum_{i=1}^{n} v_i^2 x_i^{p-2} \sum_j x_i^p \;\leqslant\; 0$$

where the last inequality follows by applying the Cauchy-Schwarz inequality[5] to $\left(v_i x^{p/2-1}\right) \cdot \left(x_i^{p/2}\right)$. Hence, the function f is concave. □

Corollary 5.10. *Given a vector of values for n agents $x \in \mathbb{R}_+^n$ and a subset $S \subseteq N$ of agents, let x' be the vector where $x_i' = x_i$ if $i \notin S$, and $x_i' = \sum_{j\in S} x_j/|S|$ if $i \in S$. Then for all $p \leqslant 1$,*

$$\left(\frac{1}{n}\sum_{i=1}^{n}(x_i)^p\right)^{1/p} \leqslant \left(\frac{1}{n}\sum_{i=1}^{n}(x_i')^p\right)^{1/p},$$

i.e., averaging out the value for a subset of agents weakly increases the p-mean welfare.

Claim 5.11. *Given $l \in \mathbb{N}$, and a vector $(x_1, \dots, x_l) \in \mathbb{R}_+^l$, for $p \in [0,1]$,*

$$\frac{1}{l}\sum_{i=1}^{l} x_i^{1-p} \leqslant \left(\frac{1}{l}\sum_{i=1}^{l} x_i\right)^{1-p},$$

while for $p < 0$, the opposite inequality holds.

Proof. For $p \in \{0,1\}$, the claim can be seen by simply substituting these values. For $p \in (0,1)$, the function $f(x) = x^{1-p}$ is concave, hence an application of Jensen's inequality gives us the claim. For $p < 0$, the function $f(x) = x^{1-p}$ is convex, hence again, Jensen's inequality gives us the claim. □

[5] For any vectors a, b, $(\sum_i \Phi_i b_i)^2 \leqslant (\sum_i \Phi_i^2) \times (\sum_i b_i^2)$

5.3 Optimal Allocations for Binary Submodular Valuations

We first show that for obtaining bounds on the price of equity for the class of binary submodular valuations (and hence, for binary additive valuations), we can focus on two allocations: the first is the Nash welfare optimal allocation $\Phi^\star$, which obtains the optimal p-mean welfare for all $p \leqslant 1$, and the second is the truncated allocation Φ^T, which obtains the optimal p-mean welfare among all EQ1 allocations for all $p \in \mathbb{R} \cup \{-\infty\}$.

Benabbou et al. (2021) show the following results.

Proposition 5.12 (Benabbou et al., 2021, Theorem 3.14)**.** *Let $\Lambda : \mathbb{Z}^n \to \mathbb{R}$ be a symmetric strictly convex function, and let $\Psi : \mathbb{Z}^n \to \mathbb{R}$ be a symmetric strictly concave function. Let Φ be some allocation. For binary submodular valuations, the following statements are equivalent:*

1. *Φ is a minimizer of Λ over all the utilitarian optimal allocations,*
2. *Φ is a maximizer of Ψ over all the utilitarian optimal allocations,*
3. *Φ is a leximin allocation, and*
4. *Φ maximizes Nash social welfare.*

Proposition 5.13 (Benabbou et al., 2021, Theorem 3.11)**.** *For binary submodular valuations, any Pareto optimal allocation is utilitarian optimal.*

For $p \leqslant 1$, if the p-mean welfare function was strictly concave, then it would follow immediately that the Nash welfare optimal allocation $\Phi^\star$ in fact simultaneously maximizes the p-mean welfare for all $p \leqslant 1$. However, in general the p-mean welfare is concave (Claim 5.9), but not strictly concave. E.g., for any $p \leqslant 1$ and any vector of values $v = (v_1, \ldots, v_n)$ with $v_i > 0$ for all agents i, let us overload notation slightly and define $\mathcal{W}_p(v) = \left(\frac{1}{n}\sum_{i=1}^n v_i^p\right)^{1/p}$. Then $\mathcal{W}_p(2v) = (\mathcal{W}_p(v) + \mathcal{W}_p(3v))/2$, violating strict concavity. However, we can slightly modify the proof of Theorem 3.14 from Benabbou et al. (2021), to obtain the following result[6]

Proposition 5.14. *For binary submodular valuations, any Nash welfare maximizing allocation (and hence, leximin allocation) simultaneously maximizes the p-mean welfare for all $p \leqslant 1$.*

Proof. The following property of leximin allocations is useful in the proof.

[6] Babaioff et al. (2021) also show that when agents are truthful, their prioritized egalitarian mechanism guarantees an allocation that is Lorenz-dominating, leximin and also maximizes Nash welfare.

Proposition 5.15. *For agents with binary submodular valuations, let Φ be a utilitarian optimal allocation so that $\max_i v_i(\Phi_i) \leqslant \min_i v_i(\Phi_i) + 1$. Then Φ is a leximin allocation.*

Proof. (of Proposition 5.15) Assume without loss of generality that $v_i(\Phi_i) \leqslant v_{i+1}(\Phi_{i+1})$ for $i \in [n-1]$. If Φ is not a leximin allocation, let Φ' be a leximin allocation. Let permutation $\pi \in S_n$ be such that $v_{\pi(i)}(\Phi'_{\pi(i)}) \leqslant v_{\pi(i+1)}(\Phi'_{\pi(i+1)})$ for $i \in [n-1]$. Then there exists $k \in [n]$, so that $v_i(\Phi_i) = v_{\pi(i)}(\Phi'_{\pi(i)})$ for $i < k$, and $v_k(\Phi_k) < v_{\pi(k)}(\Phi'_{\pi(k)})$. Note that for $i > k$,

$$v_i(\Phi_i) \leqslant v_k(\Phi_k) + 1 \leqslant v_{\pi(k)}(\Phi'_{\pi(k)}) \leqslant v_{\pi(i)}(\Phi'_{\pi(i)}) .$$

But then $\sum_{i=1}^{n} v_{\pi(i)}(\Phi'_{\pi(i)}) > \sum_{i=1}^{n} v_i(\Phi_i)$, and allocation Φ cannot be utilitarian optimal. □

The following results are shown by Benabbou et al. (2021).

Proposition 5.16 (Benabbou et al., 2021, Lemma 3.12)**.** *For agents with binary submodular valuations, let Φ be a utilitarian optimal allocation that is not a leximin allocation. Let agents i, j be such that $v_j(\Phi_j) \geqslant v_i(\Phi_i) + 2$.*[7] *Then there is another allocation Φ' that is utilitarian optimal and satisfies (i) $v_j(\Phi'_j) = v_j(\Phi_j) - 1$, (ii) $v_i(\Phi'_i) = v_i(\Phi_i) + 1$, and (iii) the values for other agents are unchanged.*

Note that in the above proposition, the allocation Φ' is a lexicographic improvement on Φ.

Proposition 5.17 (Benabbou et al., 2021, Lemma 3.13)**.** *Let Ψ be a symmetric concave function, and Φ be a utilitarian optimal allocation with agents i, j such that $v_j(\Phi_j) \geqslant v_i(\Phi_i) + 2$. Let Φ' be another utilitarian optimal allocation that satisfies (i) $v_j(\Phi'_j) = v_j(\Phi_j) - 1$, (ii) $v_i(\Phi'_i) = v_i(\Phi_i) + 1$, and (iii) the values for other agents are unchanged. Then $\Psi(\Phi') \geqslant \Psi(\Phi)$.*

We can now prove Proposition 5.14. Firstly, note that the Nash welfare maximizing allocation is also leximin from Proposition 5.12, and hence maximizes the egalitarian welfare (in other words, maximizes p-mean welfare for $p \to -\infty$). For any fixed $p \leqslant 1$, let Φ be an allocation that maximizes the p-mean welfare. Since the p-mean welfare is strictly increasing, allocation Φ is Pareto optimal, and hence from Proposition 5.13 is also utilitarian optimal. We will show that there exists an allocation Φ^T so that $\mathcal{W}_p(B) = \mathcal{W}_p(\Phi)$, and Φ^T is a leximin allocation. If Φ is not leximin, then by Propositions 5.16 and 5.17, there is an allocation Φ' so that $\mathcal{W}_p(\Phi') \geqslant \mathcal{W}_p(\Phi)$, and Φ' lexicographically dominates Φ. Then either Φ' is a leximin allocation, or we can continue in this manner until we get a leximin allocation Φ^T with $\mathcal{W}_p(B) \geqslant \mathcal{W}_p(\Phi)$, as

[7] By Proposition 5.15, such agents must exist.

required. Finally, by Proposition 5.13, an allocation is leximin if and only if it maximizes the Nash social welfare, hence if allocation Φ^* maximizes the Nash social welfare, it also maximizes the p-mean welfare for all $p \leqslant 1$. □

We now show that similarly, there exists an EQ1 allocation Φ^T that maximizes the p-mean welfare for all p. Given $\Phi^\star$, allocation Φ^T is obtained as follows, which we call the *truncated allocation*. Let $l = \min_i v_i(\Phi_i^\star)$ be the smallest value that any agent obtains in $\Phi^\star$, and let i_l be an agent that has this minimum value. Note that for any agent i, if $v_i(\Phi_i^\star) \geqslant l+2$, then all goods allocated to i must have marginal value 0 for the agent i_l, i.e., for all $g \in \Phi_i^\star$, $v_{i_l}(\Phi_{i_l}^\star \cup \{g\}) = v_{i_l}(\Phi_{i_l}^\star)$ (else we can increase the Nash social welfare by re-allocating any good that violates this to agent i_l).

For the EQ1 allocation that we would like to construct, for any agent i with $v_i(\Phi_i^\star) \geqslant l+2$, we remove goods from $\Phi_i^\star$ until i's value for the remaining bundle is $l+1$. We allocate the removed goods to agent i_l (that has marginal value 0 for these goods). Let Φ^T be the resulting allocation. Then clearly, if $v_i(\Phi_i^\star) \in \{l, l+1\}$, then $v_i(\Phi_i^T) = v_i(\Phi_i^\star)$, else $v_i(\Phi_i^T) = l+1$. Thus, allocation Φ^T, our truncated NSW allocation, is EQ1.

Theorem 5.1. *For any $p \in \mathbb{R} \cup \{-\infty\}$ and binary submodular valuations, the p-mean welfare of the truncated allocation Φ^T is at least that of any other* EQ1 *allocation.*

Proof. Let n_1 be the number of agents that have value l in allocation Φ^T, and n_2 be the number of agents with value $l+1$. Clearly, $n = n_1 + n_2$. Consider any other allocation C. We will show that the following statement is true: either (i) there exists an agent i with $v_i(C_i) \leqslant l-1$, or (ii) if all agents have value $v_i(C_i) \geqslant l$, then at most n_2 agents have value $\geqslant l+1$ (and hence at least n_1 agents have value $\leqslant l$).

Assuming the statement is true, if C is an EQ1 allocation, either (i) every agent has value $\leqslant l$, or (ii) at most n_2 agents have value $l+1$, and at least n_1 agents have value $\leqslant l$. It follows that allocation Φ^T maximizes any symmetric non-decreasing function of agent valuations in the set of EQ1 allocations, and hence Φ^T maximizes the p-mean welfare among all EQ1 allocations for all $p \in \mathbb{R}$. Since the minimum agent valuation in Φ^T is the same as in $\Phi^\star$, which by Proposition 5.12 also maximizes the egalitarian welfare, allocation Φ^T maximizes the p-mean welfare for $p = -\infty$ as well.

Lastly, to prove the statement, by the truncation procedure that yields allocation Φ^T, the number of agents $|\{i : v_i(\Phi_i^\star) \geqslant l+1\}|$ that have value at least $l+1$ in allocation $\Phi^\star$ is also n_2. Further, by Proposition 5.12, $\Phi^\star$ is also a leximin allocation, and hence no allocation in which every agent has value at least l, can have more than n_2 agents with value at least $l+1$. The statement follows. □

5.4 Lower Bounds on the PoE for Binary Additive Valuations

Theorem 5.2 (PoE lower bounds). *Let $s := r-1$. The price of equity for binary additive valuations is at least:*

1. *s, for $p=1$,*
2. *$\frac{p}{e}s$, for $p \in (0,1)$,*
3. *$\frac{s}{e\ln s}$, for $p=0$,*
4. *$2^{1/p}s^{1/(1-p)}$, for $p<0$.*

Note that as $p \to -\infty$, $2^{1/p}s^{1/(1-p)} \to 1$.

Proof. All our lower bounds are based on varying parameters in a single instance. The parameters are r, the number of agent types, and W, the normalization constant for the agents. Given r, W, the instance has $m = rW$ goods, divided into r groups of W goods each. The groups are M_1, M_2, …, M_r. There are $W+1$ agents of the first agent type, and 1 agent each of the remaining $r-1$ types (thus, $n = W+r$). Agents of type t have value 1 for the goods in group M_t, and value 0 for all other goods. The instance is thus *disjoint*; no good has positive value for agents of two different types.

We note the following properties of our lower-bound instance:

1. For any $p \leqslant 1$, an optimal p-mean welfare allocation has value 1 for W agents of the first type, and value W for each of the remaining $r-1$ agents.
2. For any $p \leqslant 1$, the EQ1 allocation with maximum p-mean welfare gives value 1 to all agents except for one agent of the first type (since there are $W+1$ agents of the first type, and only W goods for which they have positive value).

We use Λ_p to denote the PoE for this instance. Then Λ_p is exactly

$$\Lambda_p = \left(\frac{\frac{1}{W+r-1}\left(W \times 1^p + (r-1) \times W^p\right)}{\frac{1}{W+r-1}\left(W \times 1^p + (r-1) \times 1^p\right)} \right)^{1/p} = \left(\frac{W + s \times W^p}{W+s} \right)^{1/p}.$$

Note that although there are $W+r$ agents, in any allocation one agent must have value 0, hence the p-mean average is taken over $W+r-1$ agents. For each of the cases in the theorem, we will now show how to choose W, s to obtain the bound claimed.

For $p=1$, choose $W=s^2$. Then

$$\Lambda_p \geqslant \frac{W+sW}{W+s} = \frac{s^2+s^3}{s^2+s} = s,$$

giving the required bound.

For $p \in (0,1)$, choose $W = ps$. Then

$$\begin{aligned}\Lambda_p &\geqslant \left(\frac{W + s \times W^p}{W+s} \right)^{1/p} \\ &= \left(\frac{ps + s \times (ps)^p}{ps+s} \right)^{1/p} = \left(\frac{p + (ps)^p}{p+1} \right)^{1/p} \\ &\geqslant ps(p+1)^{-1/p} \geqslant ps/e, \qquad \text{since } 1+x \leqslant e^x.\end{aligned}$$

For $p=0$, the p-mean welfare is the Nash social welfare. Note that in the EQ1 allocation, each of $W+s$ agents has value 1, hence the NSW is 1. In the optimal Nash social welfare allocation, W agents have value 1, and s agents have value W, hence the NSW is $W^{s/W+s}$, which is also the PoE for this instance. Now choose $W = s/\ln s$. Then

$$\begin{aligned}\Lambda_p &\geqslant \exp\frac{s\ln W}{s+W} = \exp\frac{s(\ln s - \ln\ln s)}{s + s/\ln s} = \exp\frac{\ln s - \ln\ln s}{1 + 1/\ln s} \\ &\geqslant \exp\frac{\ln s - \ln\ln s}{1 + 1/(\ln s - \ln\ln s - 1)} \\ &= \exp\left(\ln s - \ln\ln s - 1\right) = \frac{s}{e\ln s}.\end{aligned}$$

Lastly, for $p < 0$, choose W so that $W = sW^p$, or $W = s^{1/(1-p)}$. Then

$$\begin{aligned}
\Lambda_p = & \left(\frac{W + s \times W^p}{W + s}\right)^{1/p} \\
= & \left(\frac{2W}{W + s}\right)^{1/p} \\
= & 2^{1/p}\left(\frac{W^p}{1 + W^p}\right)^{1/p} \\
\geqslant & 2^{1/p} s^{1/(1-p)},
\end{aligned}$$

where the last inequality is because $p < 0$. □

5.5 Upper Bounds on the PoE for Binary Additive Valuations

We first consider the case of utilitarian welfare, and then present our results for $p < 1$.

5.5.1 Upper bounds on the PoE for p=1

We assume that each good has value 1 for at least one agent, else the good can be removed without consequence. Given an instance with binary additive valuations for the agents, for an agent i, we overload notation and let $v_i := (v_i(g))_{g \in M}$ denote the vector of values for the individual goods. Define V to be the matrix whose ith row is given by v_i.

We say that an instance has rank k if the matrix V has rank k (equivalently, there are k linearly independent valuation vectors among the agents). Note that the rank is a lower bound on both the number of agent types, as well as the number of good types. Finally, since the rank is k, we assume the agents are ordered so that the vectors $v_1, \ldots, v_k$ are linearly independent; the corresponding agents are called basis agents.

Theorem 5.3 (Utilitarian PoE upper bound). *Under binary additive valuations and utilitarian welfare as the objective, the price of equity is at most the rank of the instance.*

Proof. Let k denote the rank of the instance, and consider allocation Φ^T that maximizes the utilitarian welfare among all EQ1 allocations. Recall that a good g is *wasted* if it is assigned to agent i such that $v_i(g) = 0$. We will show that the number of wasted goods is at most $m(1 - \frac{1}{k})$. Thus, allocation Φ has social welfare at least m/k. Since the optimal social welfare is at most m, this would be sufficient to prove the theorem.

Since allocation Φ is EQ1, there exists a utility level ℓ such that for each agent i, $v_i(\Phi_i) \in \{\ell, \ell + 1\}$. We say an agent i is *poor* if $v_i(\Phi_i) = \ell$, else agent i is *rich*. If $v_i(\Phi_i) = \ell$ for all agents, then all agents are poor.

Suppose for a contradiction that strictly more than $m(1 - \frac{1}{k})$ goods are wasted. Consider a wasted good g and a poor agent i. It must be true that $v_i(g) = 0$, else we could assign g to i and increase the utilitarian welfare while maintaining EQ1. Hence if agent i is poor, then $v_i(g) = 0$ for each wasted good g. Hence, $v_i(g) = 1$ for strictly less than m/k goods. Then due to normalization, every agent has value 1 for strictly less than m/k goods. In particular, the k basis agents have value 1 for strictly less than m/k goods each. Thus, there is a good — say g^* — for which each basis agent has value 0.

By definition, the value of each agent for g^* is a linear combination of the values of the basis agents for g^*. Since the basis agents have value 0 for g^*, it follows that every agent must have value 0 for g^*, yielding the required contradiction. □

It follows immediately from the theorem that the price of equity is also bounded by the number of agent types.

Corollary 5.18. *Under binary additive valuations and utilitarian welfare as the objective, the price of equity is at most r, the number of agent types.*

5.5.2 Upper bounds on the PoE for p < 1

From Proposition 5.12 and Theorem 5.1, to bound the PoE for any $p < 1$, it suffices to obtain an upper bound on the ratio of the p-mean welfare for the two allocations $\Phi^\star$ (which maximizes the Nash welfare) and Φ^T (the truncated allocation).

We will use various properties of the allocations $\Phi^\star$ and Φ^T in the following proofs. To state these, define T_k as the set of agents of type k, and let S_k be the set of goods allocated to agents in T_k by $\Phi^\star$. That is, $S_k := \cup_{i \in T_k} \Phi_i^\star$. Let $m_k := |S_k|$, and $n_k := |T_k|$. Then note that for each agent $i \in T_k$,

$$v_i(\Phi_i^\star) = |\Phi_i^\star| \in \left\{ \left\lfloor \frac{m_k}{n_k} \right\rfloor, \left\lceil \frac{m_k}{n_k} \right\rceil \right\}.$$

We reindex the types in increasing order of the averaged number of goods assigned by $\Phi^\star$, so that $m_i/n_i \leqslant m_{i+1}/n_{i+1}$. Now define

$$\lambda := \begin{cases} \lceil \frac{m_1}{n_1} \rceil & \text{if } m_1/n_1 \text{ is fractional} \\ 1 + \frac{m_1}{n_1} & \text{if } m_1/n_1 \text{ is integral}\,. \end{cases}$$

Thus λ is integral, $\lambda > m_1/n_1$, and $\lambda \geqslant 2$ (since the p-mean welfare is only taken over agents with positive valuation, $m_1 \geqslant n_1$). Note that in $\Phi^\star$, the smallest value of any agent is $\lfloor m_1/n_1 \rfloor$, and $\lambda \leqslant 1 + \lfloor m_1/n_1 \rfloor$. Hence agents with value at most λ in $\Phi^\star$ will retain their value in allocation Φ^T, by definition of Φ^T, while other agents will have their values truncated to λ.

Now let ρ be the highest index so that $\lambda \geqslant m_\rho/n_\rho$. Thus,

$$\lambda \geqslant \frac{\sum_{i=1}^{\rho} m_i}{\sum_{i=1}^{\rho} n_i}\,. \tag{5.1}$$

As stated above, any agent of type $k \leqslant \rho$ will retain their value, i.e., $v_i(\Phi_i^T) = v_i(\Phi_i^\star)$ for an agent i of type $k \leqslant \rho$.

We claim that agents of the first ρ types must have at least W goods assigned to them in $\Phi^\star$.

Claim 5.19. $\sum_{i=1}^{\rho} m_i \geqslant W$.

Proof. For a contradiction, let $\sum_{i=1}^{\rho} m_i < W$. Since $\lambda > m_1/n_1$, there is an agent $i^\star$ of type 1 with value $v_i(\Phi_i^\star) = \lambda - 1$. Since $\sum_{i=1}^{\rho} m_i < W$, a good g that has value 1 for agents of type 1 is allocated in $\Phi^\star$ to an agent i' of type $k > \rho$. Since $m_k/n_k > \lambda$ by definition of ρ, there is an agent i'' of type k with value $v_{i''}(\Phi_{i''}^\star) \geqslant \lambda + 1$. Since $\Phi^\star$ maximizes the Nash social welfare, any good $h \in \Phi_{i''}^\star$ has value 1 for both agents i'' and i'. Then it is easy to see that transferring any good from i'' to i', and then transferring good g from i' to $i^\star$, will increase the Nash social welfare. Since $\Phi^\star$ maximizes the Nash social welfare, we have a contradiction. □

Then from (5.1) and Claim 5.19, we obtain

$$\lambda \geqslant W / \sum_{i=1}^{\rho} n_i\,. \tag{5.2}$$

We now obtain a general expression for bounding the PoE for all $p \leqslant 1$. We will then optimize this expression for different ranges of p, to obtain upper bounds on the PoE.

Lemma 5.20. *The price of equity for p-mean welfare for instances with r types is at most*

1. $\sup_{\alpha\in[0,1]} \left(\alpha + \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p}$ *for* $p < 0$
2. $\sup_{\alpha\in[0,1]} \left(\frac{s\alpha}{1-\alpha}\right)^{(1-\alpha)}$ *for* $p = 0$,
3. $\sup_{\alpha\in[0,1]} \left(\alpha + 2^p \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p}$ *for* $p \in (0,1)$.

where as before, $s = r - 1$.

Proof. The p-mean welfare for the NSW optimal allocation $\Phi^\star$ is

$$\mathcal{W}_p(\Phi^\star) = \left(\frac{1}{n}\sum_{i=1}^{n} v_i(\Phi_i^\star)^p\right)^{1/p} = \left(\frac{1}{n}\sum_{k=1}^{r}\sum_{i\in T_k} v_i(\Phi_i^\star)^p\right)^{1/p},$$

where in the last expression, we partition the agents by their respective types.

We now consider the agent types $k \leqslant \rho$ and $k > \rho$ separately. For agents of type $k > \rho$, we average out the values and replace their individual values by the average value m_k/n_k, and use Corollary 5.10 to obtain

$$\mathcal{W}_p(\Phi^\star) \leqslant \left(\frac{1}{n}\left(\sum_{k=1}^{\rho}\sum_{i\in T_k} v_i(\Phi_i^\star)^p + \sum_{k=\rho+1}^{r} n_k\left(\frac{m_k}{n_k}\right)^p\right)\right)^{1/p}.$$

The truncated allocation Φ^T is an EQ1 allocation, and we will consider the ratio $\mathcal{W}_p(\Phi^\star)/\mathcal{W}_p(B)$. This is clearly an upper bound on the price of equity. For allocation Φ^T, recall that for agents i of type $k \leqslant \rho$, $v_i(\Phi_i^T) = v_i(\Phi_i^\star)$ since these are not truncated, while for agents i of type $k > \rho$, $v_i(\Phi_i^T) = \lambda$. Hence the PoE is

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\frac{\sum_{k=1}^{\rho}\sum_{i\in T_k} v_i(\Phi_i^\star)^p + \sum_{k=\rho+1}^{r} n_k\left(\frac{m_k}{n_k}\right)^p}{\sum_{k=1}^{\rho}\sum_{i\in T_k} v_i(\Phi_i^\star)^p + \lambda^p\sum_{k=\rho+1}^{r} n_k}\right)^{1/p}. \tag{5.3}$$

We will split the rest of the analysis into three cases: (1) $p < 0$, (2) $p > 0$, and (3) $p = 0$.

<u>Case I</u>: $p < 0$

Noting that the first term in the numerator and the denominator in (5.3) is the same, to simplify this further, we will use Proposition 5.21. The proposition is easily verified, and we skip a formal proof.

Proposition 5.21. *Consider non-negative real numbers x, y, a, b such that $x \geqslant y$, $b \geqslant a$, and $y + a > 0$. Then for any fixed $p < 0$,*

$$\left(\frac{x+a}{x+b}\right)^{1/p} \leqslant \left(\frac{y+a}{y+b}\right)^{1/p}.$$

In (5.3) we then let $x = \sum_{k=1}^{\rho} \sum_{i \in T_k} v_i(\Phi_i^\star)^p$, $y = \lambda^p \sum_{k=1}^{\rho} n_k$, $a = \sum_{k=\rho+1}^{r} n_k \left(\frac{m_k}{n_k}\right)^p$, and $\Phi^T = \lambda^p \sum_{k=\rho+1}^{r} n_k$. Then since $x \geqslant y$, $b \geqslant a$, and $y + a > 0$, from Proposition 5.21 we get

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\frac{\lambda^p \sum_{k=1}^{\rho} n_k + \sum_{k=\rho+1}^{r} n_k \left(\frac{m_k}{n_k}\right)^p}{n\lambda^p}\right)^{1/p} = \left(\frac{\sum_{k=1}^{\rho} n_k}{n} + \frac{\sum_{k=\rho+1}^{r} n_k \left(\frac{m_k}{n_k}\right)^p}{n\lambda^p}\right)^{1/p}.$$

We define $\alpha := \sum_{k=1}^{\rho} n_k / n$, i.e., the ratio of number of types that retain their values in Φ^T. Replacing in the above expression, and using that $\lambda \geqslant W / \sum_{i=1}^{\rho} n_i$ from (5.2),

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\alpha + \frac{\sum_{k=\rho+1}^{r} n_k \left(\frac{m_k}{n_k}\right)^p}{nW^p / \left(\sum_{i=1}^{\rho} n_i\right)^p}\right)^{1/p}.$$

For each type k, $m_k \leqslant W$, since for agents of each type at most W goods have positive value. Hence

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\alpha + \frac{\sum_{k=\rho+1}^{r} n_k \left(\frac{W}{n_k}\right)^p}{nW^p / \left(\sum_{i=1}^{\rho} n_i\right)^p}\right)^{1/p} = \left(\alpha + \frac{\sum_{k=\rho+1}^{r} n_k^{1-p}}{n / \left(\sum_{i=1}^{\rho} n_i\right)^p}\right)^{1/p}$$
$$= \left(\alpha + \alpha^p \sum_{k=\rho+1}^{r} \left(\frac{n_k}{n}\right)^{1-p}\right)^{1/p}.$$

We now use Claim 5.11, choosing $x_k = n_k/n$, which gives us

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\alpha + \alpha^p(r-\rho)\left(\frac{\sum_{k=\rho+1}^{r} n_k}{n(r-\rho)}\right)^{1-p}\right)^{1/p}$$

$$= \left(\alpha + \alpha^p(r-\rho)^p\left(\frac{n-\sum_{k=1}^{\rho} n_k}{n}\right)^{1-p}\right)^{1/p} = \left(\alpha + \alpha^p(r-\rho)^p(1-\alpha)^{1-p}\right)^{1/p}.$$

Finally, since $\rho \geqslant 1$, $r-\rho \leqslant s$ (where we defined $s = r-1$), hence we get the claim.

Case II: $p > 0$

Noting that the first term in the numerator and the denominator in (5.3) is the same, to simplify this further, we will use Proposition 5.22. The proposition is easily verified, and we skip a formal proof.

Proposition 5.22. *Consider non-negative real numbers x, y, a, b such that $x \geqslant y$, $a \geqslant b$ and $y + b > 0$. Then for any fixed $p > 0$,*

$$\left(\frac{x+a}{x+b}\right)^{1/p} \leqslant \left(\frac{y+a}{y+b}\right)^{1/p}$$

Observe that for any agent $i \in [n]$ such that $v_i(\Phi_i^*) > 0$, we have that $\lambda \leqslant 2 \cdot v_i(\Phi_i^*)$. Indeed, if $\lambda > 2 \cdot v_i(\Phi_i^*)$, then from the discussion in Section 5.5.2, it follows that $2 \cdot v_i(\Phi_i^*) < 1 + v_i(\Phi_i^*)$, which, for integral valuations, implies that $v_i(\Phi_i^*) = 0$.

In (5.3) we then let $x = 2^p \sum_{k=1}^{\rho} \sum_{i \in T_k} v_i(\Phi_i^\star)^p$, $y = \lambda^p \sum_{k=1}^{\rho} n_k$, $a = 2^p \sum_{k=\rho+1}^{r} n_k \left(\frac{m_k}{n_k}\right)^p$, and $\Phi^T = 2^p \lambda^p \sum_{k=\rho+1}^{r} n_k$. Then since $x \geqslant y$, $a \geqslant b$, and $y + b > 0$, from Proposition 5.22 we get

$$\left(\frac{\sum_{k=1}^{\rho}\sum_{i\in T_k}2^p v_i(\Phi_i^\star)^p+2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{\sum_{k=1}^{\rho}\sum_{i\in T_k}2^p v_i(\Phi_i^\star)^p+2^p\lambda^p\sum_{k=\rho+1}^{r}n_k}\right)^{1/p} \leqslant \left(\frac{\lambda^p\sum_{k=1}^{\rho}n_k+2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{\lambda^p\sum_{k=1}^{\rho}n_k+2^p\lambda^p\sum_{k=\rho+1}^{r}n_k}\right)^{1/p}$$
$$\leqslant \left(\frac{\lambda^p\sum_{k=1}^{\rho}n_k+2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{\lambda^p\sum_{k=1}^{\rho}n_k+\lambda^p\sum_{k=\rho+1}^{r}n_k}\right)^{1/p}$$
$$\leqslant \left(\frac{\lambda^p\sum_{k=1}^{\rho}n_k+2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{n\lambda^p}\right)^{1/p}. \tag{5.4}$$

The LHS in (5.4) is equal to the RHS in (5.3). Thus, we get that

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\frac{\lambda^p\sum_{k=1}^{\rho}n_k+2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{n\lambda^p}\right)^{1/p} = \left(\frac{\sum_{k=1}^{\rho}n_k}{n}+\frac{2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{n\lambda^p}\right)^{1/p}.$$

We define $\alpha := \sum_{k=1}^{\rho} n_k/n$, i.e., the ratio of number of types that retain their values in Φ^T. Replacing in the above expression, and using that $\lambda \geqslant W/\sum_{i=1}^{\rho} n_i$ from (5.2)

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\alpha+\frac{2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{m_k}{n_k}\right)^p}{nW^p/\left(\sum_{i=1}^{\rho}n_i\right)^p}\right)^{1/p}$$

For each type k, $m_k \leqslant W$, since for agents of each type at most W goods have positive value. Hence

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\alpha+\frac{2^p\sum_{k=\rho+1}^{r}n_k\left(\frac{W}{n_k}\right)^p}{nW^p/\left(\sum_{i=1}^{\rho}n_i\right)^p}\right)^{1/p} = \left(\alpha+\frac{2^p\sum_{k=\rho+1}^{r}n_k^{1-p}}{n/\left(\sum_{i=1}^{\rho}n_i\right)^p}\right)^{1/p}$$
$$= \left(\alpha+2^p\alpha^p\sum_{k=\rho+1}^{r}\left(\frac{n_k}{n}\right)^{1-p}\right)^{1/p}$$

We now use Claim 5.11, choosing $x_k = n_k/n$, which gives us

$$\frac{\mathcal{W}_p(\Phi^\star)}{\mathcal{W}_p(B)} \leqslant \left(\alpha + 2^p \alpha^p (r-\rho) \left(\frac{\sum_{k=\rho+1}^{r} n_k}{n(r-\rho)}\right)^{1-p}\right)^{1/p}$$

$$= \left(\alpha + 2^p \alpha^p (r-\rho)^p \left(\frac{n - \sum_{k=1}^{\rho} n_k}{n}\right)^{1-p}\right)^{1/p} = \left(\alpha + 2^p \alpha^p (r-\rho)^p (1-\alpha)^{1-p}\right)^{1/p}$$

Finally, since $\rho \geqslant 1$, $r - \rho \leqslant s$ (where we defined $s = r - 1$), hence we get the claim.

<u>Case III</u>: $p = 0$. In this case, the Nash welfare of an allocation is the geometric mean of the values of the agents. By definition of the truncated allocation Φ^T, agents of the first ρ types have the same value in Φ^* and Φ^T, hence

$$\frac{\mathcal{W}_0(\Phi^\star)}{\mathcal{W}_0(B)} = \left(\frac{\Pi_{i=1}^{n} v_i(\Phi_i^*)}{\Pi_{i=1}^{n} v_i(\Phi_i^T)}\right)^{1/n}$$

$$= \left(\frac{\Pi_{k=1}^{\rho} \Pi_{i\in T_k} v_i(\Phi_i^*) \cdot \Pi_{k=\rho+1}^{r} \Pi_{i\in T_k} v_i(A^s{*}_i)}{\Pi_{k=1}^{\rho} \Pi_{i\in T_k} v_i(\Phi_i^T) \cdot \Pi_{k=\rho+1}^{r} \Pi_{i\in T_k} v_i(\Phi_i^T)}\right)^{1/n}$$

$$= \left(\frac{\Pi_{k=\rho+1}^{r} \Pi_{i\in T_k} v_i(\Phi_i^*)}{\Pi_{k=\rho+1}^{r} \Pi_{i\in T_k} v_i(\Phi_i^T)}\right)^{1/n}$$

In Φ^*, by Corollary 5.10, for a fixed type k, we can bound $\Pi_{i\in T_k} v_i(\Phi_i^*)$ from above by $(m_k/n_k)^{n_k} \leqslant (W/n_k)^{n_k}$. Further, each agent of type $> \rho$ has $v_i(\Phi_i^T) = \lambda$, and from (5.2), $\lambda \geq W / \sum_{i=1}^{\rho} n_i$.

Let $n' := \sum_{i=1}^{\rho} n_i$ be the number of agents of the first ρ types. Then substituting these values, we get

$$\frac{\mathcal{W}_0(\Phi^\star)}{\mathcal{W}_0(B)} \leqslant \left(\frac{\Pi_{k=\rho+1}^{r}(W/n_k)^{n_k}}{\Pi_{k=\rho+1}^{r}(W/n')^{n_k}} \right)^{1/n}$$
$$= \left(\frac{(n')^{n-n'}}{\Pi_{k=\rho+1}^{r}(n_k)^{n_k}} \right)^{1/n}$$

Noting that $\sum_{k=\rho+1}^{r} n_k = n - n'$, and each $n_k \geqslant 1$, the product $\Pi_{k=\rho+1}^{r}(n_k)^{n_k}$ is maximized when the n_k's are equal, hence each $n_k = (n - n')/(r - \rho)$. With this substitution,

$$\frac{\mathcal{W}_0(\Phi^\star)}{\mathcal{W}_0(B)} \leqslant \left(\frac{n'}{(n-n')/(r-\rho)} \right)^{(n-n')/n}$$

Recalling that $\alpha = n'/n$, and further $s = r - 1 \geqslant r - \rho$,

$$\frac{\mathcal{W}_0(\Phi^\star)}{\mathcal{W}_0(B)} \leqslant \left(\frac{s\,\alpha}{(1-\alpha)} \right)^{1-\alpha}$$

which is the required expression. □

We are now ready to present our upper bounds.

Theorem 5.4 (PoE upper bounds). *Let $s := r - 1$. The price of equity for binary additive valuations is at most*

1. $1 + s$ *for* $p = 1$
2. $1 + 2s$ *for* $p \in (0, 1)$
3. $\frac{s}{\ln(s/e)}$ *for* $p = 0$ *(i.e., the Nash social welfare)*
4. $s^{1/(1-p)} 2^{-1/p} (-1/p)^{1/p(p-1)}$ *for* $p \in (-1, 0)$
5. $2s^{1/(1-p)}$ *for* $p \leqslant -1$

Proof. Our starting point is Lemma 5.20. For $p \to 0$, the PoE is at most $\sup_{\alpha \in [0,1]} (s\alpha/(1-\alpha))^{(1-\alpha)}$. Let $\beta := \alpha/(1-\alpha)$, then $1-\alpha = 1/(1+\beta)$, and hence the upper bound on the PoE is $\sup_{\beta \geqslant 0} (s\beta)^{1/(\beta+1)}$.

Some calculus shows that the maximum value of this function is $\exp(W(s/e))$, where $W(\cdot)$ is the Lambert W function, which is the inverse of the function $f(x) = xe^x$. Further, $W(x) \leqslant \ln x - \ln\ln x + \frac{e \ln\ln x}{(e-1)\ln x}$ for $x \geqslant e$. The last term $\frac{e \ln\ln x}{(e-1)\ln x} \leqslant 1$ for $x \geqslant e$. Hence for $s \geqslant e^2$, we get that the PoE is bounded by

$$\exp(W(s/e)) \leqslant \exp(\ln(s/e) - \ln\ln(s/e) + 1) = \frac{s}{\ln(s/e)}$$

giving the required bound on the PoE.

For $p \in (0,1)$, again from Lemma 5.20, the upper bound on the PoE can be written as

$$\sup_{\alpha \in [0,1]} \left(\alpha \times 1 + (1-\alpha) \times (2s\alpha/(1-\alpha))^p\right)^{1/p}$$

Since $p \in (0,1)$, $f(x) = x^{1/p}$ is a convex function, and hence by Jensen's inequality this is at most

$$\sup_{\alpha \in [0,1]} \left(\alpha \times 1^{1/p} + (1-\alpha) \times (2s\alpha/(1-\alpha))\right) = \sup_{\alpha \in [0,1]} (\alpha + 2s\alpha) = 1 + 2s$$

which is the upper bound claimed, for $p \in (0,1)$.

For $p < 0$, we separate the two cases $\alpha \geqslant 1/2$ and $\alpha \leqslant 1/2$. If $\alpha \geqslant 1/2$, then the expression from Lemma 5.20 evaluates to

$$\left(\alpha + \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p} \leqslant \alpha^{1/p} \leqslant (1/2)^{1/p} . \tag{5.5}$$

If $\alpha \leqslant 1/2$, then $(1-\alpha) \geqslant 1/2$, and hence,

$$\left(\alpha + \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p} \leqslant \left(\alpha + \alpha^p s^p 2^{-1+p}\right)^{1/p} .$$

Let $z := \alpha + \alpha^p(2s)^p/2$ be the parenthesized expression; our goal is to minimize this (since the exponent $1/p$ is negative, this will give us an upper bound on the PoE). Differentiating w.r.t. α gives us

$$\frac{dz}{d\alpha} = 1 + \frac{p}{2}(2s)^p\alpha^{p-1}\,.$$

Since p is negative, this increases with α, and hence the derivative is a convex function with a unique minima, obtained at

$$\alpha\Phi^* = \frac{(-2/p)^{1/(p-1)}}{(2s)^{p/(p-1)}}$$

or $(\alpha\Phi^* 2s)^p = -2\alpha\Phi^*/p$. Replacing this value gives us

$$\left(\alpha + \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p} \leqslant {\alpha^*}^{1/p}(1-1/p)^{1/p}$$

For $p < 0$, $1 - 1/p \geqslant 1$, and hence $(1-1/p)^{1/p} \leqslant 1$. Hence

$$\begin{aligned}\left(\alpha + \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p} \leqslant {\alpha^*}^{1/p} &= (2s)^{1/(1-p)}(-2/p)^{1/p(p-1)} \\ &= s^{1/(1-p)}2^{-1/p}(-1/p)^{1/p(p-1)}\,.\end{aligned}$$

This is greater than $2^{-1/p}$, the expression we obtain for $\alpha \geqslant 1/2$ in Equation (5.5), and hence this is a bound on the PoE for $p < 0$.

Finally for $p \leqslant -1$, let us consider the coefficient of $s^{1/(1-p)}$ obtained previously, namely $2^{-1/p}(-1/p)^{1/p(p-1)}$. This is an increasing function of p, and hence the maximum value obtained is 2, at $p = -1$. Hence for $p \leqslant -1$, the PoE is at most $2s^{1/(1-p)}$. □

5.6 PoE Bounds for Doubly Normalized Instances

So far, we have considered instances with binary additive normalized valuations, where each agent values the same number W of goods. In this case, for the utilitarian welfare, we have seen that the PoE can be as bad as r, the number of types of agents. In this section, we consider instances with further structure. In *doubly normalized* instances, each good g is valued by the same number W_c of agents. Thus, $v_i(M) = W$ for all $i \in N$, and additionally, $\sum_{i \in N} v_i(g) = W_c$ for every good $g \in M$. The valuation matrix V is thus both row and column normalized. Such instances are intuitively "balanced," and we ask if this balance is reflected in the PoE for such instances. This indeed turns out to be the case.

Theorem 5.5. *For doubly normalized instances under binary additive valuations, the PoE for the p-mean welfare is 1 for all $p \leqslant 1$.*

For an undirected graph, the edge-incident matrix X has entry $X_{i,e} = 1$ if edge e is incident on vertex i, and $X_{i,e} = 0$ otherwise. We will use the following well-known property of edge-incidence matrices for bipartite graphs.

Proposition 5.23 (e.g., Schrijver, 1998)**.** *If G is a bipartite graph, then the edge-incidence matrix of G is totally unimodular.*

Proof of Theorem 5.5. Let V be the valuation matrix for a doubly normalized instance, where each row sums to W and each column sums to W_c. Divide each entry by W_c. Let V^f be the resulting matrix. Then V^f satisfies: (i) each entry is either 0 or $1/W_c$, (ii) each column sums to 1, and (iii) each row sums to W/W_c. We will show that the matrix V^f can be represented as the convex combination of nonnegative integer matrices $X^1, \ldots, X^t$ so that for any matrix X^k in this decomposition, each column sums to 1 and each row sums to either $\lceil W/W_c \rceil$ or $\lfloor W/W_c \rfloor$. Assuming such a decomposition, fix any such matrix X^k in this decomposition. Clearly, due to (ii) and nonnegativity, each entry of X^k is either 1 or 0. Further if the entry $X^k_{i,g} = 1$, then $V^f_{i,g} = 1/W_c$ since V^f is a convex combination of the M-matrices, and hence $V_{i,g} = 1$,

Consider then the allocation Φ that assigns good g to agent i if $X^k_{i,g} = 1$. In this allocation, following the properties of X^k, each good is assigned to an agent that has value 1 for it, and each agent is assigned either $\lceil W/W_c \rceil$ or $\lfloor W/W_c \rfloor$ goods. The allocation is thus EQ1 and maximizes the utilitarian welfare. Further by Proposition 5.15 this is also a leximin allocation, and hence by Proposition 5.14 and Proposition 5.12 this maximizes the p-mean welfare for all $p \leqslant 1$, proving the theorem.

It remains to show that V can be decomposed as stated. To see this, consider a complete bipartite graph $G = (A \cup B, E)$ with $|A| = n$ and $|B| = m$. To each edge $\{i, g\}$ with $i \in A$, $g \in B$, we associate a variable x_{ig}. Consider now the set of linear constraints:

$$\begin{aligned} \forall i \in A, \quad & \textstyle\sum_{g\in B} x_{ig} && \geqslant \lfloor W/W_c \rfloor \\ \forall i \in A, \quad & \textstyle\sum_{g\in B} x_{ig} && \leqslant \lceil W/W_c \rceil \\ \forall g \in B, \quad & \textstyle\sum_{i\in A} x_{ig} && = 1 \\ \forall i \in A, g \in B, \quad & x_{ig} && \geqslant 0 \end{aligned}$$

Together, these linear constraints ask for a fractional set of edges that have degree 1 for each vertex in Φ^T and degree between $\lfloor W/W_c \rfloor$ and $\lceil W/W_c \rceil$ for each vertex in Φ.

Consider the polytope obtained by these inequalities. Taking $x_{ig} = V^f_{ig}$ satisfies these constraints. Further, it can be seen that the constraint matrix is equal to the edge-incidence matrix for the bipartite graph G (with the rows corresponding to vertices $i \in A$ repeated, and the identity matrix appended for nonnegativity of the variables). Hence, the constraint matrix is totally unimodular by Proposition 5.23, and thus the extreme points of the polytope are integral. Since V^f is a point in the polytope, V^f can be represented as the convex combination of nonnegative integral matrices $X^1, \ldots, X^t$ corresponding to the vertices of the polytope, as required. □

We make two remarks. Firstly, note that since each matrix X^k in the convex decomposition of V^f gives us an EQ1 allocation with maximum utilitarian welfare, the convex combination gives us a randomized allocation that is ex ante EQ, and ex post EQ1 and welfare optimal. Secondly, the doubly normalized constraint is sufficient, but not necessary, for the price of equity to be 1. Consider an instance with 3 agents $\{a_1, a_2, a_3\}$ and 4 goods $\{g_1, g_2, g_3, g_4\}$ such that a_1 values $\{g_1, g_2\}$ while a_2 and a_3 both value $\{g_3, g_4\}$. This instance is not column normalized, but admits an EQ1 allocation with optimal welfare.

We now turn to an alternate proof of Theorem 5.5, based on a so-called "eating argument" and an extension of Hall's theorem. We first state some results that we will use. Consider a bipartite graph G with bipartition A and B. A set of edges $T \subseteq E(G)$ is said to be a *q-expansion* from A to B if every vertex of A is incident to exactly q edges in T and exactly $q|A|$ vertices in B are incident on T. A perfect matching, for instance, is a 1-expansion, and a star with q leaves is a

q-expansion.

Lemma 5.24 (Cygan et al., 2015, Lemma 2.17)**.** *Let G be a bipartite graph with bipartition A and B. There is a q-expansion from A to B if and only if $|N(X)| \geq q|X|$ for every $X \subseteq A$. Furthermore, if there is no q-expansion from A to B, then a set $X \subseteq A$ such that $|N(X)| < q|X|$ can be found in polynomial time.*

A non-negative, square ($m \times m$) matrix Y is said to be *doubly stochastic* if the sum of entries in each row and each column is 1, that is, $\sum_{i \in [m]} Y[i][j] = 1 \; \forall j \in [m]$ and $\sum_{j \in [m]} Y[i][j] = 1 \; \forall i \in [m]$. A *permutation matrix* is a doubly stochastic matrix such that all the entries are either 0 or 1. The following result, known as the Birkhoff-von Neumann Theorem, states that a doubly stochastic matrix can be represented as a convex combination of permutation matrices.

Theorem 5.25 (Birkhoff, 1946; von Neumann, 1953)**.** *Let Y be a doubly stochastic matrix. Then, there exist positive weights $w_1, w_2, \ldots w_k$ and permutation matrices $P_1, P_2, \ldots P_k$ such that $\sum_{i \in [k]} w_i = 1$ and $Y = \sum_{i \in [k]} w_i P_i$.*

In other words, the convex hull of the set of all permutation matrices is the set of doubly-stochastic matrices.

We are now ready to prove that doubly normalized instances have PoE 1.

Proof. (of Theorem 5.5) Let $\mathcal{I} = \langle N, M, \mathcal{V} \rangle$ be a doubly normalized instance. Let $G = (N, M)$ be the corresponding (W, W_c)-regular bipartite graph with agents and goods as bi-partitions and $(i, g) \in E(G)$ if and only if agent i values the good g. Note that the number of edges in G is nW, as exactly W edges are incident on each agent. Likewise, as exactly W_c edges are incident on each of the m goods, therefore, $|E(G)| = nW = mW_c$.

We first consider the case when $n/m = W/W_c = p$ for some integer p. That is, the number of agents is an integer multiple of the number of goods, and show the existence of a non-wasteful EQ allocation that allocates a utility of p to every agent.

Consider any subset $S \subseteq N$. The number of edges from S to its neighborhood $N(S)$ is exactly $W|S|$. The number of edges incident on $N(S)$ in G is exactly $W_c|N(S)|$. Then $W_c|N(S)| \geqslant W|S|$, and hence $|N(S)| \geqslant p|S|$.

By Lemma 5.24, G must have a p-expansion, say T, from N to M. Now consider the allocation Φ that allocates along the edges of T. Precisely, if $(i, g) \in T$, then good g is allocated to the agent i under Φ. Then by definition of p-expansion, Φ is an indeed an EQ allocation as it allocates exactly p goods to every agent. Since Φ is non-wasteful, it achieves the optimal utilitarian welfare m.

Now suppose W is not an integer multiple of W_c. We propose the following version of the probabilistic serial algorithm that constructs a non-wasteful EQ1 allocation.[8]

Let $W = pW_c + q$ for some constant p and $q \neq 0$. We create $p+1$ copies of every agent, say $\{a_i^1, a_i^2, \ldots a_i^{p+1}\}$. We also add $t = (p+1)\,n - m$ many dummy goods to the instance which are valued at zero by everyone. Note that the new instance has an equal number of agents and goods, precisely $(p+1)n$. Now each good is represented as food, and the agent copies start eating away all the available goods that they like, all at once. By the structure of the instance, exactly W_c agent copies eat the same W goods at the same time, and the same speed – at the rate of one good per unit time. In particular, at the t^{th} timestep, $1/(W-(t-1)W_c)$ fraction of the remaining good is consumed by the t^{th} copy of the agent. This gives us a square matrix Y with $(p+1)\,n$ columns as goods and $(p+1)\,n$ rows as copies of the agents. The entry $Y[a_i^j][g]$ corresponds to the fraction of good g eaten by j^{th} copy of agent i at timestep j.

In particular, in the first timestep, W goods are consumed by all the first copies (a_i^1) of W_c agents (who like them) simultaneously, each of whom eats $1/W$ fraction of W goods each. At second timestep, $(1 - W_c/W)$ fraction of these r goods remain, out of which $1/(W-W_c)$ fraction is consumed by the second copy of all the W_c agents and so on. That is, assuming N_g be the set of W_c agents who like g, we have:

$$\sum_{i \in N_g} v_i(\Phi_i^1) = \frac{W_c}{W}$$

$$\sum_{i \in N_g} v_i(\Phi_i^2) = W_c \left(\frac{1}{W - W_c} \left(1 - \frac{W_c}{W} \right) \right) = \frac{W_c}{W}$$

$$\sum_{i \in N_g} v_i(\Phi_i^3) = W_c \left(\frac{1}{W - 2W_c} \left(1 - \frac{2W_c}{W} \right) \right) = \frac{W_c}{W}$$

$$\vdots$$

$$\sum_{i \in N_g} v_i(\Phi_i^p) = W_c \left(\frac{1}{W - (p-1)W_c} \left(1 - \frac{(p-1)W_c}{W} \right) \right) = \frac{W_c}{W}$$

[8]The probabilistic serial algorithm was proposed by Bogomolnaia and Moulin (2001) in the context of the assignment problem where the number of goods and agents is the same. Subsequently, Aziz et al. (2023a) used this algorithm (in combination with the Birkhoff-von Neumann decomposition) to study fair allocation with an unequal number of goods and agents. The work of Aziz et al. (2023a) focuses on computing a randomized allocation with desirable ex-ante and ex-post *envy-freeness* guarantees. By contrast, our work uses the technique of Aziz et al. (2023a) to achieve an *equitability* guarantee.

For the last agent copy a_i^{p+1}, the fraction of each of the W goods that remain is $1 - {}^{pW_c}/_{W} = {}^{(W-pW_c)}/_{W} = {}^{q}/_{W}$. This is divided equally among the last copy of all the W_c agents, each of them getting ${}^{q}/_{WW_c}$ fraction.

Also, a_i^{p+1} eats ${}^{q}/_{WW_c}$ of W goods, thereby summing to ${}^{q}/_{W_c}$. Now ${}^{q}/_{W_c} < 1$, and we have t dummy goods remaining to be consumed. Therefore, at this timestep, a_i^{p+1} for $i \in [n]$, start eating $\left(\frac{1-q/W_c}{t}\right)$ fraction of each of the dummy goods, thereby consuming one unit of good, in aggregate.

We now claim that Y is a doubly stochastic matrix. To this end, we first show that the fractions in every column of Y adds up to 1. For a column c corresponding to an original good g, summing over the $p+1$ copies of s agents who like g, we get:

$$\sum_{j=1}^{p+1} \sum_{i \in N_g} v_i(\Phi_i^j) = \sum_{j=1}^{p} \sum_{i \in N_g} v_i(\Phi_i^j) + v_i(a_i^{p+1}) = p\left(\frac{W_c}{W}\right) + W_c\left(\frac{q}{WW_c}\right) = \frac{pW_c + q}{W} = 1$$

Now for a column c corresponding to a dummy good g, each of the agent copies eat $\left(\frac{1-q/W_c}{t}\right)$ fraction of the dummy goods. Since there are n such agents, the sum of the fractions in column c is

$$n\left(\frac{1-q/W_c}{t}\right) = \frac{n(W_c - q)}{W_c(p+1)n - mW_c} = \frac{n(W_c - q)}{npW_c + nW_c - nW} = \frac{W_c - q}{W_c - q} = 1$$

Also, it is easy to see that rows in Y add up to 1. For $j \in [1,p]$, a_i^j starts eating at j^{th} timestep, when $(1 - {}^{(j-1)W_c}/_{W})$ fraction of any good g remains. She eats ${}^{1}/_{(W-(j-1)W_c)}$ fraction of W such goods, that adds to $W(\frac{1}{W-(j-1)W_c})(1 - \frac{(j-1)W_c}{W}) = 1$. As for the row corresponding to a_i^{p+1}, it adds up to 1 by construction.

This establishes the following claim.

Claim 5.26. *Y is a doubly stochastic matrix.*

By Theorem 5.25, Y can be represented as a convex combination of permutation matrices. An illustration of this is shown in Example 5.28. For the final allocation, one of these permutation matrices, say P, is selected with probability equal to the corresponding weight. A good is allocated to a_i^j if and only if $P[a_i^j][g] = 1$. Finally, all the goods allocated to the copies of agent i are said to be allocated to agent i.

We now claim that the resulting allocation is EQ1 with optimal utilitarian welfare.

Claim 5.27. *Every integral allocation returned by the above algorithm satisfies* EQ1*.*

Proof. (of Claim 5.27) Since the matrices in the decomposition are permutation matrices and the number of goods is equal to the number of agents, each of the agent-copies gets exactly one good. This implies that all the agents end up with an equal number of goods, precisely, $p+1$. Since the dummy goods are consumed by only the last agent copy, therefore, every agent gets at most one dummy good in the final allocation. Also, all the original goods are allocated non-wastefully, as except for the dummy goods, agents eat only the goods that they like. Therefore, whoever ends up with a dummy good has a utility of p and the remaining agents have a utility of $p+1$, resulting in an EQ1 allocation with optimal utilitarian welfare. □

Therefore the price of equity for doubly normalized instances is 1. This finishes the proof of Theorem 5.5. □

Example 5.28. *Consider an instance with* 4 *agents,* $\{a_1, \ldots a_4\}$ *and* 6 *goods,* $\{1, 2, \ldots 6\}$*. We set* $W = 3$ *and* $W_c = 2$*.* a_1 *likes* $\{1, 2, 3\}$*,* a_2 *likes* $\{4, 5, 6\}$*,* a_3 *likes* $\{2, 3, 4\}$ *and* a_4 *likes* $\{1, 5, 6\}$*. Here, since* $p = 1$*, we create* $p + 1 = 2$ *copies of every agent, and introduce* $(p+1)n - m = 2 \cdot 4 - 6 = 2$ *dummy goods. The corresponding matrix* Y *and its decomposition is as follows.*

$$
\begin{pmatrix}
1/3 & 1/3 & 1/3 & & & & 0 & 0 \\
1/6 & 1/6 & 1/6 & & & & 1/4 & 1/4 \\
& & & 1/3 & 1/3 & 1/3 & 0 & 0 \\
& & & 1/6 & 1/6 & 1/6 & 1/4 & 1/4 \\
& 1/3 & 1/3 & 1/3 & & & 0 & 0 \\
& 1/6 & 1/6 & 1/6 & & & 1/4 & 1/4 \\
1/3 & & & & 1/3 & 1/3 & 0 & 0 \\
1/6 & & & & 1/6 & 1/6 & 1/4 & 1/4
\end{pmatrix}
= 0.16
\begin{pmatrix}
1 & & & & & & & \\
& 1 & & & & & & \\
& & & 1 & & & & \\
& & & & 1 & & & \\
& & 1 & & & & & \\
& & & & & & 1 & \\
& & & & & 1 & & \\
& & & & & & & 1
\end{pmatrix}
+ 0.08
\begin{pmatrix}
1 & & & & & & & \\
& & 1 & & & & & \\
& & & 1 & & & & \\
& & & & & 1 & & \\
& 1 & & & & & & \\
& & & & & & 1 & \\
& & & & 1 & & & \\
& & & & & & & 1
\end{pmatrix}
$$

$$
+0.08
\begin{pmatrix}
1 & & & & & & & \\
& & 1 & & & & & \\
& & & 1 & & & & \\
& & & & & 1 & & \\
& 1 & & & & & & \\
& & & & & & & 1 \\
& & & & 1 & & & \\
& & & & & & 1 &
\end{pmatrix}
+ 0.08
\begin{pmatrix}
& 1 & & & & & & \\
1 & & & & & & & \\
& & & & 1 & & & \\
& & & 1 & & & & \\
& & 1 & & & & & \\
& & & & & & & 1 \\
& & & & & 1 & & \\
& & & & & & 1 &
\end{pmatrix}
+ 0.16
\begin{pmatrix}
& 1 & & & & & & \\
& & & & & & 1 & \\
& & & & 1 & & & \\
& & & & & & & 1 \\
& & & 1 & & & & \\
& & 1 & & & & & \\
1 & & & & & & & \\
& & & & & 1 & &
\end{pmatrix}
$$

$$
+0.16
\begin{pmatrix}
& & 1 & & & & & \\
& & & & & & 1 & \\
& & & & & 1 & & \\
& & & & & & & 1 \\
& 1 & & & & & & \\
& & & 1 & & & & \\
1 & & & & & & & \\
& & & & 1 & & &
\end{pmatrix}
+ 0.08
\begin{pmatrix}
& & 1 & & & & & \\
& & & & & & & 1 \\
& & & & & 1 & & \\
& & & & & & 1 & \\
& 1 & & & & & & \\
& & & 1 & & & & \\
1 & & & & & & & \\
& & & & 1 & & &
\end{pmatrix}
+ 0.16
\begin{pmatrix}
& & 1 & & & & & \\
& & & & & & & 1 \\
& & & & & 1 & & \\
& & & & & & 1 & \\
& & & 1 & & & & \\
& 1 & & & & & & \\
& & & & 1 & & & \\
1 & & & & & & &
\end{pmatrix}
$$

5.7 PoE Bounds for Binary Submodular Valuations

We now consider the more general case of binary submodular valuations. Here we focus on the utilitarian welfare, and show that our results for binary additive valuations that bound the PoE by the number of types of agents do not extend to binary submodular valuations. We first show that from prior work (see Proposition 5.29 below), it follows that if the agents have identical valuations, then PoE is 1 for the p-mean welfare objective for all $p \leqslant 1$.

Proposition 5.6. *When all agents have identical binary submodular valuations, the PoE is 1 for p-mean welfare measure for all $p \leqslant 1$.*

As earlier, an allocation $\Phi = (\Phi_1, \ldots, \Phi_n)$ is *clean* if for all agents i, $v_i(\Phi_i) = |A_i|$, that is, no good is wastefully allocated. We note that, given any allocation Φ, we can obtain a clean (possibly partial) allocation $\hat{\Phi}$ so that $v_i(\Phi_i) = v_i(\hat{\Phi}_i)$ for all agents i by repeatedly removing wasted items from the allocation Φ. We will use the following result due to Benabbou et al. (2021).

Proposition 5.29 (Benabbou et al., 2021, Corollary 3.8)**.** *For binary submodular valuations, any clean, utilitarian optimal (partial) allocation Φ that minimizes $\sum_i v_i(\Phi_i)^2$ among all utilitarian optimal allocations is* EF1*.*

Proof of Proposition 5.6. Let $\Phi^\star$ be a Nash welfare maximizing allocation for the given instance. We will show that under identical binary submodular valuations, Φ^* can be transformed into an EQ1 allocation without any change in the Nash welfare objective, thus implying that PoE is 1 for Nash welfare. Furthermore, from Proposition 5.14, we know that any allocation that maximizes Nash welfare also simultaneously maximizes p-mean welfare for all $p \leqslant 1$. This would imply that PoE is 1 for p-mean welfare objective for all $p \leqslant 1$.

First, we will transform Φ^* into a clean partial allocation via the following procedure: For each agent i with $v_i(\Phi_i^*) > |\Phi_i^*|$, there must be a wasted good in Φ_i^*; we simply remove such wasted goods until we get a clean partial allocation $\hat{\Phi}$. Next, we will add back the removed goods *arbitrarily* to obtain a complete allocation Φ (in particular, adding back the removed goods may get back the original allocation Φ^*).

Note that $\hat{\Phi}$ is a partial allocation with $v_i(\hat{\Phi}_i) = v_i(\Phi_i^*)$ for each agent i; in other words, Φ^* and $\hat{\Phi}$ have the same p-mean welfare for all $p \leqslant 1$. We know from Proposition 5.14 that, for all $p \leqslant 1$, Φ^* maximizes the p-mean welfare. The same holds true for $\hat{\Phi}$.

By adding the removed goods back to $\hat{\Phi}$, the utility of any agent cannot decrease; that is, for every agent i, $v_i(\Phi_i) = v_i(\hat{\Phi}_i)$. This means that Φ is a complete allocation that simultaneously maximizes the p-mean welfare for all $p \leqslant 1$.

By Proposition 5.12, allocation Φ minimizes the strictly convex function $\sum_i v_i(\Phi_i)^2$ among all utilitarian allocations, and the same holds for the partial allocation $\hat{\Phi}$. Then, by Proposition 5.29, we get that $\hat{\Phi}$ is EF1. By the identical valuations assumption, $\hat{\Phi}$ is also EQ1.

In going from $\hat{\Phi}$ to Φ, each good that is added back has zero marginal value for the agent it is assigned to under Φ. Thus, the allocation Φ is also EQ1, which readily implies that for all $p \leqslant 1$, the PoE for p-mean welfare is 1, as desired. □

The bound in Proposition 5.6 is, in a certain sense, the best that can be obtained. We will now show that with more than one type of agent under binary submodular valuations, the PoE is at least $n/6$ for utilitarian welfare. Hence we cannot obtain bounds on the PoE that depend on the number of agent types for all $p \leqslant 1$, as we did for binary additive valuations.

Theorem 5.7. *The PoE for utilitarian welfare when agents have binary submodular valuations is at least $n/6$ (where n is the number of agents), even when there are just two types of agents.*

Proof. In our example for the lower bound, we represent goods as vectors (i.e., elements of a linear matroid). Then the value of an agent for a bundle is just the number of linearly independent vectors in the bundle. Fix $k \in \mathbb{N}$. Our example will have $2k$ agents and $k^2 + k$ goods.

Goods: There are $k(k+1)$ goods, consisting of $k+1$ groups of k goods each. The groups are $G_1, G_2, \ldots, G_{k+1}$.

Agents: There are $2k$ agents, with k agents of type 1 and k agents of type 2. Agents of type 1 see goods in G_1 as the standard basis vectors for $\mathbb{R}^k$, and goods in G_j for $j \neq 1$ as zero vectors. Thus, for an agent i of type 1, $v_i(G_1) = v_i(M) = k$, and $v_i(G_j) = 0$ for $j > 1$.

Agents of type 2 see the goods in each group G_i as the standard basis vectors for $\mathbb{R}^k$, and hence for an agent i of type 2, $v_i(G_j) = v_i(M) = k$, for all $j \in [k+1]$. Thus, the valuations are normalized.

In an EQ1 allocation, each agent of type 1 has value at most 1, and hence the social welfare is at most $3k$. In the optimal allocation, each agent of type 1 gets a single vector from G_1. Each agent of group 2 gets assigned an entire group G_j of vectors, and hence has value k. The optimal social welfare is thus $k + k^2$, and hence the PoE is at least $k/3$, or $n/6$, where n is the number of agents. □

Note that for the example in the proof of Theorem 5.7, for any $p \in (0,1]$, the PoE is

$$\Lambda_p = \left(\frac{\frac{1}{2k}(k \times 1 + k \times k^p)}{\frac{1}{2k}(k \times 1 + k \times 2^p)} \right)^{1/p} = \left(\frac{1+k^p}{1+2^p} \right)^{1/p} \geqslant \frac{k}{3^{1/p}},$$

and hence the PoE depends on the number of agents, even with two types. Similarly, for the Nash social welfare, one obtains the PoE as $\sqrt{k/2} = \sqrt{n/4}$.

For $p < 0$, for this example, the PoE is a constant that depends on p (for example, for $p = -1$, the PoE for this example is 1.5). It is possible that for $p < 0$ the PoE may depend on the agent types, rather than number of agents. We leave this as an open question.

Despite this, we show that $2n$ is an upper bound on the PoE for all $p \leqslant 1$. For an allocation $A = (A_1, \ldots, A_n)$ of the goods, we say good g is *valuable* for i if $v_i(\Phi_i \cup g) > v_i(\Phi_i)$ (and i *values* g in this case).

Theorem 5.8. *For binary submodular valuations and any $p \leqslant 1$, the PoE for p-mean welfare is at most $2n$.*

Proof. As before, let $\Phi^\star$ be an allocation with optimal Nash welfare. If $\Phi^\star$ is an EQ1 allocation, we are done, since from Proposition 5.14, $\Phi^\star$ simultaneously maximizes the p-mean welfare for all $p \leqslant 1$. Otherwise, we construct the truncated allocation Φ^T as described in Section 5.3. We will show that for every agent i with non-zero value in Φ^T, $v_i(\Phi_i^T) \geqslant W/(2n)$, where W is the normalization constant. It follows that the PoE is bounded by $2n$ for all $p \leqslant 1$.

Consider the allocation $\Phi^\star$. Let i_l be a minimum positive value agent in $\Phi^\star$. Note that $\Phi^*_{i_l} = B_{i_l}$. Let ν be the value of agent i_l under $\Phi^\star$. Since $v_{i_l}(M) = W$, there are $W - \nu$ goods that i_l values that are allocated to other agents. Further, any agent $i \neq i_l$ is allocated at most $\nu + 1$ goods that i_l values, since otherwise, we can transfer a good that i_l values from i to i_l and increase the Nash social welfare of allocation Φ^*. Hence, $W - \nu \leqslant (n-1)(\nu+1)$, or $W \leqslant n\nu + n - 1 \leqslant 2n\nu$ for $\nu \geqslant 1$. Thus for any agent i, $v_i(\Phi_i^T) \geqslant v_i(B_{i_l}) = \nu \geqslant W/(2n)$, as required. □

5.8 Visualizing the PoE Bounds

To help in visualising our upper and lower bounds on the PoE (as presented in Table 5.1), we also present the PoE bounds graphically below (see Figures 5.1 to 5.4). Each graph shows the lower and upper bounds obtained as a function of r, the number of agent types. In each figure, the upper line in blue represents the upper bound, and the lower line in orange represents the lower bound. We provide the plots for four values of p, namely $p = 1$ (the utilitarian welfare), $p = 0$ (the Nash social welfare), $p = -1$, and $p = -10$ (recall that for $p \to -\infty$, the egalitarian welfare, the PoE is 1).

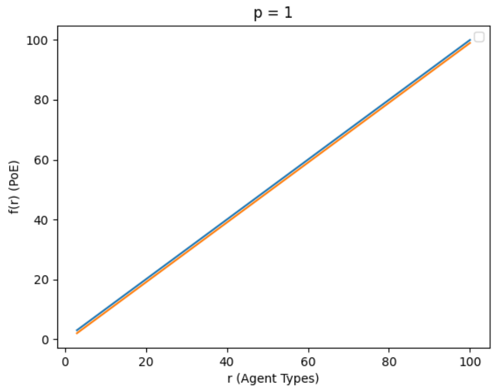

Figure 5.1: PoE as a function of r for $p = 1$

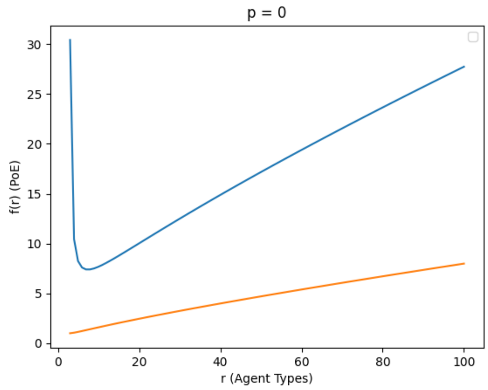

Figure 5.2: PoE as a function of r for $p = 0$

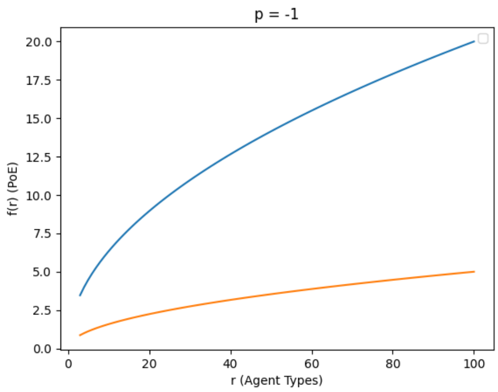

Figure 5.3: PoE as a function of r for $p = -1$

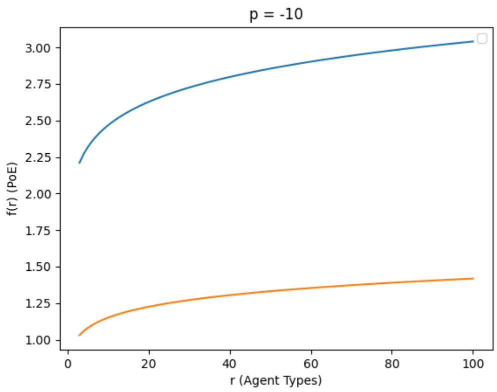

Figure 5.4: PoE as a function of r for $p = -10$

5.9 Some Concluding Remarks on Chores

Our focus in the paper has been on goods, where agents have non-negative marginal utility for all items. We briefly remark on the case of bads or chores, where all marginal utilities are non-positive. Consider any instance with binary additive valuations, i.e., the value of each item is either 0 or -1. It is not hard to see that in these instances, there is always a utilitarian optimal

EQ1 allocation: if chore c has value 0 for an agent i, assign c to i. The remaining chores have value -1 for all agents, and can be assigned using the round robin procedure. This allocation is clearly EQ1 and also achieves the best possible utilitarian welfare.

For more general additive instances with chores, we now show that the PoE is unbounded, even in very simple cases.[9] To this end, consider the following example involving $2n$ items and $n+1$ agents. The first n agents mildly dislike the first n chores and severely dislike the last n, while it is the opposite for the $(n+1)$th agent, who strongly dislikes the first n items and mildly dislikes the last n.

	$c_1, \cdots, c_n$	$c_{n+1}, \cdots, c_{2n}$
$a_1, \ldots, a_n$	$-\epsilon$	-1
a_{n+1}	-1	$-\epsilon$

In this example, the maximum utility is $-2n\epsilon$: assign the first n chores to the first agent and the last n chores to the last agent. On the other hand, in any EQ1 allocation, the last agent can get at most 2 chores, and hence some agent gets a chore that they value at -1. The PoE is thus at least $1/(2n\epsilon)$, which can be made arbitrarily large by choosing ϵ appropriately. Note that this instance has two item types, two agent types, and only two distinct entries in the valuation matrix. Relaxing any of these conditions implies identical valuations, where the PoE is 1; so, in some sense, this is a "minimally complex" example that already exhibits unbounded PoE. There is thus a sharp change in the PoE between instances where the values are in $\{0, -1\}$ and those where the values are in $\{-\epsilon, -1\}$. While the PoE is unbounded as ϵ approaches 0, it "snaps back" to 1 at $\epsilon = 0$.

To conclude, we obtain nearly tight bounds on the price of equity in terms of agent types for the p-mean welfare spectrum. This captures, as special cases, the notions of utilitarian, egalitarian, and Nash welfare. Our bounds are in terms of agent types (r) rather than the number of agents. Overall, our results provide a fine-grained perspective on the behavior of the price of equity parameterized by p and r.

In future work, it would be interesting to extend the insights that we obtain in this work beyond

[9] For chores, we adopt the natural definition of PoE: the ratio of the utilitarian welfare of the best EQ1 allocation, to the maximum utilitarian welfare obtainable in any allocation. Note that if the denominator is 0, then so is the numerator (and this can be identified in polynomial time).

the domain of binary valuations. We also propose obtaining bounds on the PoE parameterized by other structural parameters, such as the number of item types. We note that for additive valuations, the rank of the valuation matrix is a lower bound on the number of item types, and hence Theorem 5.3 bounds the PoE in this case by the number of item types as well.

Chapter 6

Equitable and Efficient Allocations for Mixtures of Goods and Bads

"Why are numbers beautiful? It's like asking why is Beethoven's Ninth Symphony beautiful. If you don't see why, someone can't tell you. I know numbers are beautiful. If they aren't beautiful, nothing is."

- Paul Erdos, *Quoted in The Man Who Loved Only Numbers by Paul Hoffman*

6.1 Introduction

Consider fairly allocating courses to students where some of them enjoy courses with mathematical rigor while others find them daunting. Consider a fair assignment of research papers among the reviewers who may have subjective opinions about the papers depending on the domain. Further, consider a group of friends narrowing down upon a common activity to pursue. Some of them may find window shopping enjoyable while it may be a chore for others. Likewise, a host of real-world problems have this flavor where a single item/event could be a 'good' for one agent (valued positively) but a 'chore' for another (valued negatively). All these scenarios can be modeled as allocation problems of indivisible items among rational and opinionated agents and in this chapter, we study the computational landscape of finding fair (equitable) allocations for the mixtures of goods and chores.

Envy-freeness has been a prominent and well-studied criterion and has also been looked at for the mixture of goods and chores (see Aziz et al. (2022); Aleksandrov and Walsh (2020); Aziz and Rey (2021); Bérczi et al. (2020)). Equitability in the context of this mixed setting is relatively less studied. In this work, we try to reduce this gap and present a fair and comprehensive understanding of equitability for mixed items.

The choice of fairness notion also depends upon the context and both these notions have their own mutually exclusive desirable properties. As far as practicality, empirical relevance, and perceived fairness are considered, experiments suggest that equitability (or inequality aversion) is a preferable criterion over envy-freeness (Herreiner and Puppe, 2009, 2010; Gal et al., 2016).

Gourvès et al. (2014) showed that when all items are goods (additive valuations), an EQX (hence, EQ1) allocation always exists and is efficiently computable. Freeman et al. (2019) and Freeman et al. (2020) studied EQ1 in conjunction with Pareto optimality and envy-freeness when all items are goods and chores respectively. Recently, Barman et al. (2024a) considered EQX for the mixed instances with both goods and chores and showed that computing an EQX allocation (even without any efficiency requirement) in the mixed setting is weakly NP-Hard even for two agents and strongly NP-hard for more agents. They showed that some restricted scenarios are tractable, such as when the items are objective goods and chores and the number of agents is two. The hardness for EQX motivates to look at the weaker requirement of EQ1 and in this context, we answer both the existential and computational questions for finding approximately equitable allocations in the setting of mixed items, standalone and coupled with efficiency guarantees as well.

Our Contributions

Our work provide a deep dive into the existence and computational boundaries of EQ1 allocations for mixtures of goods and chores with and without economic efficiency notions (e.g. Pareto optimality and utilitarian/egalitarian social welfare). We first show that unlike the goods-only and chores-only settings, an EQ1 allocation may not exist for mixed items (Example 6.1). Moreover, deciding the existence of an EQ1 allocation is weakly NP-Hard (Theorem 6.2).

Upon further scrutiny, we observe that the non-existence of EQ1 is primarily due to the sharp contrast between how agents subjectively evaluate the entire set of items. This observation motivates the study of *normalized valuations* where the value of the entire set of items is constant (positive, negative, or zero) for all agents. Normalization is a common assumption in the cake-cutting literature[1], and has been studied in the context of online fair division (Beynier et al., 2019b) or studying the price of fairness to derive positive results (Bhaskar et al., 2023). Lange et al. (2020) showed that normalization is crucial for tractability. Table 6.1 provides a summary of our algorithmic results. We highlight the following results:

EQ1 for objective valuations.

For objective valuations, but not necessarily normalized, we show that an EQ1 allocation always exists and can be computed efficiently (Theorem 6.3). Our analysis gives rise to a lemma (Lemma 6.5) which enables the design of new algorithms along with efficiency. It states that if there exists a partial EQ1 allocation of subjective items, it can be completed by allocating the remaining objective items in an EQ1 manner.

EQ1 for normalized valuations.

Under normalized valuations, when instances are trivalued in the form of $\{-w, 0, w\}$, we show that an EQ1 allocation always exists and can be computed in polynomial time for any number of agents (Theorem 6.9). These instances (i.e. $\{-w, 0, w\}$ valuations) generalize both binary and bivalued preferences and capture realistic scenarios involving approvals/dis-approvals/neutrality. Our algorithmic techniques involve carefully transferring items among 'rich' and 'poor' agents.

Furthermore, we show when valuations are in addition *type-normalized* (i.e. the sum of chores and the sum of goods are independently constant), an EQ1 allocation can be computed efficiently for two agents (Theorem 6.10). Along the way, we highlight several challenges in achieving EQ1 for any normalized valuation.

EQ1+PO for normalized valuations.

In Section 6.4, we show an EQ1+PO allocation may fail to exist even for $\{-1, 1\}$ normalized valuations. The corresponding computational problem remains intractable even for the special class of type-normalized valuations (Theorem 6.16). Nonetheless, we develop a polynomial-time algorithm for $\{-w, 0, w\}$ normalized valuations that computes an EQ1+PO allocation, when one exists (Theorem 6.18). We note here that the presence of zeros has been a source of

[1] In fair cake-cutting, the common assumption is that each agent values the entire cake as 1 (or -1 for burnt cakes)

computational challenge for EQ1+PO allocations even in the only goods setting Freeman et al. (2019); Garg and Murhekar (2024). Our algorithmic technique involves reducing the instance to a partial instance such that every item is non-negatively valued by at least one agent, computing a Nash optimal solution, and invoking Lemma 6.5 to assign any remaining chore.

In Section 6.6, we look into the compatibility of EQ with EF and show that every EQ+PO allocation is EF for $\{1, 0, -1\}$ valuations (Proposition 6.21). Also, for such valuations, an EF1+EQ1+PO allocation can be computed in polynomial time, whenever it exists (Corollary 6.24). On the other hand, for general type-normalized valuations, deciding the existence of EF (or EF1) + EQ1 + PO is strongly NP-Hard (Corollary 6.23), while deciding the existence of EF (or EF1)+EQ1 allocations is weakly NP-Hard (Corollary 6.22).

EQX+Welfare.

In Section 6.5, we develop a pseudo-polynomial time algorithms (for a fixed number of agents) for finding EQX allocations that maximize utilitarian or egalitarian welfare (Theorem 6.20). Note that maximizing welfare along with approximate equitability is already known to be weakly NP-Hard, even when all the items are goods Sun et al. (2023b). Table 6.1 gives a partial summary of our results.

Additional Related Work

Equitability as a fairness notion was first studied in the context of divisible items, where it is possible to assign a fraction of an item to any agent (Dubins and Spanier, 1961). It is known that equitable allocations always exist for divisible items (Cechlárová et al., 2013; Chèze, 2017). There is a computational limitation to finding such an allocation Procaccia and Wang (2017) but approximate equitable allocations admit efficient algorithms (Cechlárová and Pillárová, 2012a,b). In the setting of indivisible items, Freeman et al. (2019, 2020) studied equitability along with efficiency guarantees but in the context of only goods and only chores. Perhaps closest to our work is that of Barman et al. (2024a), who study EQX allocations for mixtures. Recently, the price of equitability in terms of welfare loss has also been looked at (Caragiannis et al., 2012; Aumann and Dombb, 2015; Sun et al., 2023b; Bhaskar et al., 2023).

The mixtures and chores-only setting differs substantially from the goods counterpart. Many of the existing algorithmic techniques from the latter fail to work in the former (Aziz et al., 2022; Bhaskar et al., 2020; Bogomolnaia et al., 2017) and hence, the mixed setting demands independent analysis and scrutiny. For the goods-only case, maximizing Nash welfare gives

	Non-Normalized	Normalized	Type-Normalized	Objective Items
EQ1	× (6.1) Pseudo-Poly (6.20) (weakly) NP-Complete (6.2)	✓, $\{-w, 0, w\}$, $\mathbb{P}$ (6.9) Pseudo-Poly (6.20)	✓, $n = 2$, $\mathbb{P}$ (6.10) ✓, $n > 2$, $\{-w, 0, w\}$, $\mathbb{P}$, (6.9) Pseudo-Poly (6.20)	✓, $\mathbb{P}$ (6.3)
EQ1+PO	×, $\{1, 0, -1\}$, $\mathbb{P}$ (6.18) (strongly) NP-Hard (6.16)	×, $\{1, 0, -1\}$, $\mathbb{P}$ (6.18) (strongly) NP-Hard (6.16)	✓, $n = 2$, $\{-w, 0, w\}$, $\mathbb{P}$, (6.17) ×, $\{1, 0, -1\}$, $\mathbb{P}$, (6.18) (strongly) NP-Hard (6.16)	× (strongly) NP-Hard*
SW×EQX SW/EQX	× Pseudo-Poly (6.20) (weakly) NP-Hard†			

Table 6.1: A partial summary of our results. Each cell contains existence/computation results with ✓ implying existence, and × implying non-existence. The table entry, say '✓, $\{-w, 0, w\}$, $\mathbb{P}$' conveys that the corresponding fair and efficient allocation always exists for $\{-w, 0, w\}$ valuations and admits a polynomial time algorithm. The results marked '$\star$' and '$\dagger$' follow from Freeman et al. (2019) and Sun et al. (2023b) respectively.

both fairness (EF1) and efficiency (PO) guarantees (Caragiannis et al., 2019b; Barman et al., 2018a) but due to the presence of negative utilities, Nash looses its guarantees in the mixed setting. In addition, for mixtures, EF1+PO is known to exist only for 2 agents and the case of arbitrary agents stands as a major open problem (Aziz et al., 2022). The existence of EFX for only goods and only chores is known to exist only for certain special cases (Chaudhury et al., 2020a; Plaut and Roughgarden, 2020) but is open in general. On the other hand, for mixtures, Bérczi et al. (2020) shows the non-existence of EFX under non-monotone non-additive valuations while Hosseini et al. (2023d) establishes the non-existence of EFX allocations under additive valuations. Hosseini et al. (2023b) identify a class of lexicographic preferences where EFX + PO allocations exist. Other fairness notions like maximin share and PO (Kulkarni et al., 2021a,b), and competitive equilibrium (Bogomolnaia et al., 2017) have also been looked at for mixtures.

6.2 Preliminaries

Setting.

A fair division $([n], [m], \mathcal{V})$ instance consists of $n \in \mathbb{N}$ agents, $m \in \mathbb{N}$ items and valuations $V = \{v_1, v_2, \ldots v_n\}$ where $v_i : 2^{[m]} \to \mathbb{Z}$, that captures the value that an agent i derives from a subset of items. We restrict our attention to *additive valuations* where for any subset $S \subseteq [m]$, we have $v_i(S) = \sum_{o \in S} v_i\{o\}$. Also, $\{-w, w\}$ valuations denote that every item is either valued at a constant w or $-w$ by every agent.

Allocations and Equitability.

An allocation $\Phi = \{\Phi_1, \Phi_2, \ldots \Phi_n\}$ is a partition of m items into n bundles, one for each agent. The utility that an agent i derives from her bundle is $v_i(\Phi_i) = \sum_{o \in A_i} v_i\{o\}$. An allocation Φ is said to be *equitable* (EQ) if all the agents derive equal utility from their respective bundles, that is, for every pair of agents i and j, we have $v_i(\Phi_i) = v_j(\Phi_i)$. An allocation Φ is said to be *equitable up to one item* (EQ1) if for any pair of agents i and j such that $v_i(\Phi_i) < v_j(\Phi_i)$, either there exists some good g in A_j such that $v_i(\Phi_i) \geq v_j(\Phi_j \setminus \{g\})$ or there exists some chore c in Φ_i such that $v_i(\Phi_i \setminus \{c\}) \geq v_j(\Phi_i)$. Further, an allocation Φ is said to be *equitable up to any item* (EQX) if for any pair of agents i and j such that $v_i(\Phi_i) < v_j(\Phi_i)$, we have $v_i(\Phi_i) \geq v_j(\Phi_j \setminus \{g\})$ for all goods g in Φ_j, and $v_i(\Phi_i \setminus \{c\}) \geq v_j(\Phi_i)$ for all chores c in Φ_i.

Efficiency.

An allocation Φ is said to be Pareto dominated by another allocation Φ' if $v_i(\Phi'_i) \geq v_i(\Phi_i)$ for all agents i and there exists at least one agent j such that $v_j(\Phi'_j) > v_j(\Phi_i)$. If an allocation is not Pareto dominated by any other allocation, then it is said to be Pareto optimal (PO). The *welfare* of an allocation is an aggregation of individual utilities in various ways. The *utilitarian welfare* of an allocation Φ is the arithmetic mean of individual utilities under Φ, denoted by $UW(\Phi) = 1/n \sum_{i \in [n]} v_i(\Phi_i)$. The *egalitarian welfare* is the minimum utility among all the individual utilities, denoted by $EW(\Phi) = \min v_i(\Phi_i)$. Further, we say that an agent i is poor under an allocation Φ if $v_i(\Phi_i) \leq v_j(\Phi_i) \; \forall \, j \in [n]$ and it is rich if $v_i(\Phi_i) > v_j(\Phi_i) \; \forall \, j \in [n]$.

Normalization(s).

Let O be the set of all items. We say that an item is an *objective good* if it is valued non-negatively by all the agents. That is, $o \in O$, such that $v_i(o) \geq 0 \; \forall \, i \in N$ and denote the set of all objective goods as O^+. Similarly, all items $o \in O$ such that $v_i(o) \leq 0$ are called *objective chores*, denoted by O^-. The remaining objects are called *subjective items*, denoted by $O^{\pm}$. We now discuss the two types of normalizations that we consider. If in an instance, all agents value the entire bundle at a constant, that is, $v_i(O)$ is a constant for all i, then we say it is *normalized instance*. If for all agents i, all the goods sum up to a constant, say g', and all the chores sum up to a constant, say c', that is, $\sum_{o \in G} v_i(o) = g'$ and $\sum_{o \in C} v_i(o) = c'$, then we call it *type-normalized instance*. An instance is not normalized unless otherwise stated. Note that type-normalization implies normalization (but not converse).

6.3 EQ1 Allocations

When the input instance has subjective items, an EQ1 allocation may not exist, even for 2 agents. The following example illustrates this.

Example 6.1. *Consider an instance with two agents Alice and Bob such that Alice values all three items at* -1 *while Bob values all the three items at* 1. *In any allocation, Alice always values her bundle less than Bob's valuation of her own bundle, even after the hypothetical removal of some item from any of the bundles. Therefore, no allocation in this instance is EQ1.*

Moreover, we show in the following result that deciding if there is an EQ1 allocation is weakly NP-Hard. This is in contrast to the only goods or only chores setting where even for non-normalized instances, computing an EQ1 allocation admits efficient algorithms.

Theorem 6.2. *For any allocation instance with mixed items, deciding the existence of an EQ1 allocation is NP-complete.*

Proof. We exhibit a reduction from 2-PARTITION, where given a multiset $U = \{b_1, b_2 \dots b_m\}$ of integers with sum $2T$, the task is to decide if there is a partition into two subsets S and $U \setminus S$ such that sum of the numbers in both the partitions equals T. We construct the reduced allocation instance as follows. We create 4 agents, m set-items $\{o_1, \dots, o_m\}$, and 4 dummy items $\{d_1, \dots d_4\}$. The first two agents value the set-items at $\{b_1, \dots b_m\}$ and all the dummy items at $-3T$. The last two agents value the set items at 0 and the dummy items at T. This completes the construction. We now argue the equivalence, where we show that if there is an EQ1 allocation in the reduced instance, then 2-PARTITION is a yes instance and vice-versa.

Forward Direction.

Suppose that the 2-PARTITION instance is a yes-instance. Let S and $U \setminus S$ be the said partitions. Then the allocation where agent Φ_1 gets S, Φ_2 gets $U \setminus S$, a_3 gets $\{d_1, d_2\}$ and a_4 gets $\{d_3, d_4\}$ is clearly an EQ1 allocation.

	o_1	o_2	$\dots$	o_m	d_1	d_2	d_3	d_4
a_1	b_1	b_2	$\dots$	b_m	$-3T$	$-3T$	$-3T$	$-3T$
a_2	b_1	b_2	$\dots$	b_m	$-3T$	$-3T$	$-3T$	$-3T$
a_3	0	0	$\dots$	0	T	T	T	T
a_4	0	0	$\dots$	0	T	T	T	T

Table 6.2: Reduced instance as in the proof of Theorem 6.2.

Reverse Direction.

Suppose there is an EQ1 allocation, say Φ, under the reduced instance. Then, note that Φ assigns at most one dummy item to the agents a_1 and a_2, each. If not, say WLOG, a_1 receives $\{d_1, d_2\}$ under Φ, then the maximum utility a_1 can derive is $-4T$ (where it gets all the set items as well). But, a_1 violates EQ1 with respect to a_3 and a_4, whose minimum utility is 0 each and

a_1 derives negative utility even if it removes the item d_1 from its bundle. Therefore, a_1 and a_2 can get at most one dummy item each. We now consider the following cases:

- a_1 and a_2 do not receive any dummy item. Note that since Φ is EQ1, anyone from a_3 and a_4 cannot get three dummy items, else either one of a_1 or a_2 violates EQ1. Therefore, a_3 and a_4 both get two dummy items each. This forces a_1 and a_2 to receive a utility of at least T, thereby corresponding to a partition.
- a_1 receives d_1 and a_2 receives d_2. Note that if the remaining two dummy items are allocated to any one agent, say a_3, then EQ1 is violated. Indeed, the maximum utility of a_1 and a_2 is negative, and even if they choose to ignore any item, they fall short of the utility derived by a_3 (which is $2T$). Therefore, a_3 receives d_3 and a_4 receives d_4 under Φ. This forces a_1 and a_2 to receive a utility of T each from the set items, thereby forcing a partition.
- a_1 receives d_1 and a_2 does not receive any dummy item. Then, WLOG, a_3 gets one dummy item and a_4 gets two dummy items. To be consistent with EQ1 against a_4, agent a_1 must get all the set items, thereby deriving a utility of $-3T + 2T$. This leaves a_2 empty-handed and hence, it violates EQ1 with respect to a_4. Therefore, since Φ is EQ1, this case does not arise.

This concludes the argument. □

This non-existence and hardness motivate us to explore the tractable scenarios in this context. Below we show some positive results for some restricted settings. We first present a few tractable cases when an EQ1 allocation always exists and can be found efficiently, without any assumption on normalization(s).

Theorem 6.3. *For any allocation instance with mixed items containing objective goods and chores, an EQ1 allocation always exists and can be computed in polynomial time.*

Proof. The algorithm iteratively picks the least happy agent (agent with the least utility)—breaking ties arbitrarily—and allocates her most valuable good among the remaining items in O^+ according to her preference. Once O^+ is exhausted, it picks the happiest agent who then receives her most disliked chore from O^-. The correctness of the above algorithm is as follows. In the beginning, when no one is allocated any item, EQ1 is satisfied vacuously. We argue that if the allocation is EQ1 before the assignment of an item, it remains EQ1 after that assignment as well. Let i be the least happy agent ($v_i(\Phi_i) \leq v_j(\Phi_j)\ \forall\ j \neq i$) at iteration t. Let g be the good most valued by i among the remaining goods in O^+, which is added to Φ_i

in the iteration $t+1$. Now either i continues to be the least happy agent in which case, the allocation continues to be EQ1. Otherwise, consider an agent j such that $v_i(\Phi_i \cup \{g\}) > v_j(\Phi_j)$. Then, since g was the last added good in Φ_i, it is also the least favorite item of i in her entire bundle Φ_i (note that nothing has been allocated from O^- till this iteration). This implies that $v_i(\Phi_i \cup \{g\} \setminus g) = v_i(\Phi_i) < v_j(\Phi_j)$. Hence, any other agent violates EQ with respect to agent i only up to the recent addition of g into Φ_i. Since the empty allocation in the beginning is vacuously EQ1, this settles the claim that the allocation is EQ1, until O^+ is exhausted. Now consider when everything from O^+ has been allocated. The instance now reduces to the one with only chores. Suppose agent i is the happiest agent at this point, that is, $v_i(\Phi_i) \geq v_j(\Phi_j)\ \forall\, j \in [n]$. The algorithm picks i and allocates it a chore c that it dislikes the most. If $v_i(\Phi_i \cup c) \geq v_j(\Phi_j)\ \forall\, j \in [n]$, then the allocation continues to be EQ1. Else, $v_i(\Phi_i \cup c) < v_j(\Phi_j)$ for some agent j. Since i was the happiest agent previously and hence became a potential recipient of c, therefore, i can choose to hypothetically remove c from his bundle in order to value her own bundle more than j. That is, $v_i(\Phi_i \cup c \setminus \{c\}) = v_i(\Phi_i) \geq v_j(\Phi_j)$. This implies that the allocation remains EQ1 after the allocation of c. This settles the claim. □

We note here that Theorem 6.3 does not guarantee the existence of EQX allocations even for identical objective valuations, as illustrated in the following example.

Example 6.4. *Consider the following instance with identical objective valuations. The output of the algorithm in Theorem 6.3 is highlighted, which fails EQX. Indeed, to satisfy EQX, the poor agent Alice should have at least as much utility as Bob upon the hypothetical removal of any good from Bob's bundle or any chore from her own bundle. But if Alice chooses to ignore the item o_2 from Bob's bundle and reduce her utility from 1 to -1, then also it remains a poor agent. But this allocation is EQ1, since it can ignore the chore o_7 from her bundle, thereby deriving equal utility as Bob, that is, 1.*

	o_1	o_2	o_3	o_4	o_5	o_6	o_7
Alice	(2)	2	(2)	2	(−3)	−3	(−3)
Bob	2	(2)	2	(2)	−3	(−3)	−3

We now show that if there exists a partial EQ1 allocation that allocates $O^{\pm}$ entirely, then it can always be extended to a complete EQ1 allocation.

Lemma 6.5 (Completion Lemma). *Consider a partial EQ1 allocation Φ that allocates a subset of items, say $S \subset [m]$. Then, Φ can be completed while preserving EQ1 if $O^{\pm} \subseteq S$.*

Proof. Let Φ be the partial EQ1 allocation that allocates S such that $O^{\pm} \subseteq S \subset [m]$. Since the remaining items in the instance are all objective items, we allocate a poor agent under Φ its most liked item from O^{+}. Once O^{+} is exhausted, we allocate a rich agent its most disliked item from O^{-}. The correctness follows by a similar argument as in Theorem 6.3. □

The above algorithm(s) output an EQ1 allocation even for instances that are not normalized. We now turn our attention to the normalized and type-normalized instances.

Theorem 6.6. *For any allocation instance with $\{w, -w\}$ normalized valuations, an EQ1 allocation always exists and can be computed in polynomial time.*

Proof. We show that Algorithm 3 returns an EQ1 allocation when the instance has $\{-w, w\}$ normalized valuations. First note that since the instance is normalized, the reduced instance restricted to items in only $O^{\pm}$ is also normalized. That is, every agent assigns a value of w to exactly k_1 items from $O^{\pm}$ and a value of $-w$ to exactly k_2 items from $O^{\pm}$, where k_1 and k_2 are constants. The idea is to first allocate all the items in $O^{\pm}$ in an almost equitable way and then to extend the partial EQ1 allocation to a complete allocation using Lemma 6.5. To that end, we first show that there exists a partial EQ1 allocation that allocates all items in $O^{\pm}$.

Claim 6.7. *There exists a partial EQ1 allocation restricted to the items in $O^{\pm}$ such that it exhausts $O^{\pm}$. Also, such an allocation can be computed in polynomial time.*

Proof. We iteratively pick a poor agent and allocate it an item it values at w (a good) from $O^{\pm}$. If there are no such items in $O^{\pm}$, then we iteratively pick one of the rich agents and allocate it an item it values at $-w$ (a chore) from $O^{\pm}$. At this point, it is easy to see that the partial allocation is EQ1. We denote the set of agents with minimum utility as poor agents (P) and those with maximum utility as rich agents (R). Now suppose that the remaining items in $O^{\pm}$ are all chores for the poor agents and goods for the rich agents and if they are allocated to any of them, they increase the amount of inequity and hence may violate EQ1. Therefore, in order to move towards an EQ1 allocation that allocates all the items in $O^{\pm}$, we now aim to convert a poor (rich) agent to rich (poor) by re-allocating one of the allocated items, so that the converted agent can now be a potential owner of one of the remaining unallocated items while maintaining EQ1. We do so by executing one of the following transfers at a time.

Algorithm 3 n agents, $\{w, -w\}$ Normalized Valuations

Require: n agents, m items, $\{-w, w\}$ Normalized Valuations
Ensure: An EQ1 allocation Φ

1: $\Phi \leftarrow$ An empty allocation
2: $P = \{p_1, p_2, \ldots p_s\} \leftarrow$ Set of all poor agents under Φ
3: $R = \{r_1, r_2, \ldots r_t\} \leftarrow$ Set of all rich agents under Φ
4: **while** $O^{\pm} \neq \emptyset$ **do**
5:     **while** Any p values an unallocated item o in $O^{\pm}$ at w **do**
6:         $\Phi_p = \Phi_p \cup \{o\}$
7:     **while** Any r values an unallocated item o in $O^{\pm}$ at $-w$ **do**
8:         $\Phi_r = \Phi_r \cup \{o\}$
9:     **if** Any p values a good o in any Φ_r at w **then**
10:         $\Phi_r = \Phi_r \setminus \{o\}; \Phi_p = \Phi_p \cup \{o\}$ (**Rich-Poor Transfer**)
11:         Repeat from Step 5
12:     **if** Any r values a chore o in any Φ_p at $-w$ **then**
13:         $\Phi_p = \Phi_p \setminus \{o\}; \Phi_r = \Phi_r \cup \{o\}$ (**Poor-Rich Transfer**)
14:         Repeat from Step 5
15:     **if** Any r' values a good o in any Φ_r at $-w$ **then**
16:         $\Phi_r = \Phi_r \setminus \{o\}; \Phi_{r'} = \Phi_{r'} \cup \{o\}$ (**Rich-Rich Transfer**)
17:         Repeat from Step 5
18:     **if** Any p' values a chore o in any Φ_p at w **then**
19:         $\Phi_p = \Phi_p \setminus \{o\}; \Phi_{p'} = \Phi_{p'} \cup \{o\}$ (**Poor-Poor Transfer**)
20:         Repeat from Step 5
21:     **if** $|P| > |R|$ **then**
22:         **for** $j \in [t]$ **do**
23:             $\Phi_{r_j} = \Phi_{r_j} \setminus \{o : v_{r_j}(o) = w\}$
24:             $\Phi_{p_j} = \Phi_{p_j} \cup \{o\}$
25:         Repeat from Step 5
26: Use **Completion Algorithm** of Lemma 6.5
27: **return** Φ

1. **Transfers.** Consider a pair of agents $p \in P$ and $r \in R$ such that there is an item o in Φ_r and $v_r(o) = w = v_p(o)$. Then, we re-allocate o to p, consequently converting p to a rich

agent and r to a poor agent. We call it a *rich-poor transfer*. Likewise, if there is a pair of agents $p \in P$ and $r \in R$ such that there is an item o in Φ_r and $v_r(o) = -w = v_p(o)$, we re-allocate o to r, again converting p to a rich agent and r to a poor agent (see Table 6.3). We call it a *poor-rich transfer*. If there is a pair of rich agents, say r and r' such that there is a good o in $\Phi_{r'}$ which is valued at $-w$ by r, then we can re-allocate o to r, thereby decreasing the utility of both r and r' by $-w$ (*rich-rich transfer*). Likewise, we do a *poor-poor transfer* for a pair of poor agents p, p' such that there is a $-w$-valued item o in $\Phi_{p'}$, $v_p(o) = w$ and re-allocate o to p.

We execute one of the transfers at a time, thereby converting at least one poor (rich) agent to rich (poor). Now, the remaining item o which was earlier a chore for all the poor agents and a good for all the rich agents, has a potential owner. We continue this until we find a poor agent who can be converted to a rich one by such a transfer and can be allocated one of the remaining items, or until $O^{\pm}$ is exhausted.

	Φ_r	Φ_p
r	w w	$-w$
p	w	w w $-w$

Table 6.3: Rich-poor and poor-rich transfers. In one step, only one of these transfers is executed, not both.

2. Suppose there is no such transfer feasible. Then, we consider the following cases.

   (2a) Firstly, suppose $|P| > |R|$. Then we take a w-valued item from Φ_r and allocate it to p who necessarily values it at $-w$ (else, we would have executed a rich-poor transfer). We do this for $|R|$ disjoint pairs of a rich and a poor agent, thereby decreasing the utility of all these pairs by $-w$. But since $|P| > |R|$, we have a poor agent p' whose value remains intact, which in turn, makes him one of the rich agents after the above transfers. Now, the remaining unallocated item which was a chore for p' can be allocated to it without violating EQ1.

   (2b) Second, suppose $|P| \leq |R|$. In this case, we have that all the rich agents value all the goods in every other rich bundle at w (else, we would have executed a transfer). Now suppose $v_r(\Phi_r) = (k+1)w$ and $v_p(\Phi_p) = kw$ where k is some constant. Then there are $|R|(k+1)$ many allocated items valued at w by a rich agent r. Also, it values all the remaining unallocated items $O^{\pm}$ at w, so there are at least $|R|(k+1)+1$ many valuable items for r. Since all the remaining unallocated items in $O^{\pm}$ are chores for p, there must be at least $|R|(k+1)+1$ many items among the allocated ones which are valued at w by p (by normalization). Even if p values every good in every other poor agent's bundle at w, then also only $|P|k$ out of $|R|(k+1)+1$ such items are accounted for. Also, note that p cannot value any of the $|R|(k+1)$ items allocated to the rich agents at w, otherwise, we are done by appropriate transfers). This implies that there must be $(|R|-|P|)k+|R|+1$ many extra items valued at w by p, outside of the $|R|(k+1)$ and $|P|k$ items allocated as goods to the rich and poor agents respectively. Let's call this set of extra goods for p as E_g. Note that no item from E_g adds any value to any of the agents it is allocated to. (If it is allocated to a rich agent and adds value to its bundle, then its bundle will be at $(k+2)w$. If it is allocated to a poor agent and adds value to her, then its bundle will be at $(k+1)w$). Therefore, they must have been allocated in pairs with chores. Hence there are $|E_g|$ more extra items valued at $-w$ by p. We call them E_c (see Table 4 in the appendix). Now if any of the rich agents values any $o \in E_c$ at w, then p again falls short of good items and it violates normalization. Otherwise, if r values all E_c at $-w$, then $E_c \notin O^{\pm}$, which is a contradiction to the fact that the partial allocation only allocates items from $O^{\pm}$ at this point. Therefore, this case does not arise and this settles our claim.

□

Once $O^{\pm}$ is exhausted, we end up with a partial EQ1 allocation Φ. Now we can allocate the items from O^{+} to a poor agent iteratively and once O^{+} is exhausted, we allocate the item from O^{-} to a rich agent iteratively (Lemma 6.5). This ensures that Φ is EQ1 at every step of the algorithm. Consequently, the complete allocation Φ is EQ1. □

If the instances are allowed to have 0-valued items, then the approach of assigning $O^{\pm}$ first and then completing the allocation fails even for type-normalized valuations, as illustrated in Example 6.8.

Example 6.8. *In the following instance, there is no (partial) allocation Φ of $O^{\pm} = \{o_1, o_2, o_3\}$ such that Φ is EQ1.*

	$O^{\pm}$													
	Allocated Items												Unallocated Items	
	o_1	o_2	o_3	o_4	o_5	o_6	o_7	o_8	o_9	o_{10}	E_g	E_c	o	o
r_1	w	w	w	w	w	w						$-w$ ... $-w$	w	w
r_2	w	w	w	w	w	w						$-w$ ... $-w$	w	w
p_1	$-w$	$-w$	$-w$	$-w$	$-w$	$-w$	w	w	w	w	w ... w	$-w$ ... $-w$	$-w$	$-w$
p_2	$-w$	$-w$	$-w$	$-w$	$-w$	$-w$	w	w	w	w	w ... w	$-w$ ... $-w$	$-w$	$-w$

Table 6.4: A schematic of Case 2(a) of the proof of Claim 6.7. r_1 and r_2 are two rich agents with utility $3c$ each and p_1 and p_2 are two poor agents with utility $2c$ each. The values in gray are forced, otherwise, we are done by Case 1. The values in red depict the contradiction that $E_c \in O^-$.

	o_1	o_2	o_3	o_4	o_5	o_6	o_7	o_8	o_9
Alice	1	1	1	0	0	0	−1	−1	−1
Bob	−1	−1	−1	1	1	1	0	0	0

Therefore, when we have $\{w, 0, -w\}$ valuations, we need a slightly different approach of allocating items from the entire bundle rather than exhausting a subset of items first. Using this approach, we give an efficient algorithm for this case as follows.

Theorem 6.9. *For any allocation instance with $\{-w, 0, w\}$ normalized valuations, an EQ1 allocation always exists and can be computed in polynomial time.*

Proof. The idea is similar to Theorem 6.6 except here, we do not aim to exhaust $O^{\pm}$ first, instead, allocate the items from the entire bundle. This is so because a partial allocation restricted to $O^{\pm}$ may not satisfy EQ1 (see Example 6.8).

We iteratively pick a poor agent p and allocate it an item valued at w or 0. If there is no such item for p, then we iteratively pick one of the rich agents and allocate it an item valued at $-w$ or 0. At this point, it is easy to see that the partial allocation Φ is EQ1. We denote the set of agents with minimum utility under Φ as poor agents (P) and those with maximum utility under Φ as rich agents (R). Either we arrive at a complete EQ1 allocation by continuing in the above manner or there are unallocated items which are chores for agents in P and goods for agents in R. Allocating any of the remaining items only increases the inequity and hence violates EQ1. Therefore, in order to move towards a complete EQ1 allocation, we now aim to convert a poor (rich) agent to rich (poor) by re-allocating one of the allocated items, so that the converted agent can now be a potential owner of one of the remaining unallocated items while maintaining EQ1. We do so by executing one of the following transfers at a time.

Transfers.

Consider a pair of agents $p \in P$ and $r \in R$ such that there is a w-valued good g in Φ_r and $v_p(g) \in \{0, w\}$. Then, we re-allocate g to p. We call it a *rich-poor transfer*. Likewise, if there is a pair of agents $p \in P$ and $r \in R$ such that there is a $-w$-valued chore w in Φ_p and $v_r(c) \in \{-w, 0\}$, we re-allocate w to r (*poor-rich transfer*). If there is a pair of two rich agents r and r' such that there is a w-valued good g in $\Phi_{r'}$ and $v_r(g) \in \{-w, 0\}$, we do a *rich-rich transfer* and allocate g to r. Likewise, we do a *poor-poor transfer* for a pair of poor agents p, p' such that there is a $-w$-valued chore w in $\Phi_{p'}$ and $v_p(c) \in \{0, w\}$ and re-allocate w to p.

Note that each of the above transfers convert at least one poor (rich) agent to rich (poor). Suppose none of the above transfers is possible. Then we consider the following two cases. If $|P| > |R|$. Then we take a w-valued item from Φ_r and allocate it to p who necessarily values it at $-w$ (else, we would have executed a rich-poor transfer). We do this for $|R|$ disjoint pairs of a rich and a poor agent, thereby decreasing the utility of all the agents involved by $-w$. But since $|P| > |R|$, we have a poor agent p' whose value remains intact, which in turn, makes him one of the new rich agents after the above transfers. Now, the remaining unallocated item which was a chore for p' can be allocated to it without violating EQ1.

We now argue the remaining case of $|P| \leq |R|$. If there are no transfers feasible, then we have that all the rich agents value the w-valued goods in every other rich agent bundle at w. WLOG, we assume the normalization constant to be positive. Suppose $v_r(\Phi_r) = (k+1)w$ and $v_p(\Phi_p) = kw$, where k is some constant integer. Then there are at least $|R|(k+1)$ many allocated items valued at w by a rich agent r. Also, it values all the remaining unallocated items at w, so there are at least $|R|(k+1) + 1$ many valuable items for r. Also, note that p cannot value any of the $|R|(k+1)$ many w-valued items allocated to the rich agents at w or 0, else we have our desired transfer. This implies that every p values the $|R|(k+1)$ allocated goods to the rich agents as well as the remaining unallocated items at $-w$. Because p has a significant number of $-w$-valued items, normalization then forces p to value at least $|R|(k+1)+1$ allocated items at w (as the normalization constant is assumed to be positive). Even if it values all the $|P|k$ many w-valued items allocated to all the poor agents, $|R|(k+1)+1 > |P|k$, and hence, there are at least $|R|(k+1)+1-|P|k$ extra items in the allocated set that are valued at w by p. We call this set E_g. Note that no item from E_g adds any value to any of the agents it is allocated to. (If it is allocated to a rich agent and adds value to her bundle, then her bundle will be at $(k+2)w$. If it is allocated to a poor agent and adds value to her, then her bundle will be at $(k+1)w$). Therefore, they must have been allocated in pairs with chores. Hence there are $|E_g|$ more extra items valued at $-w$ by p. We call them E_c. Consider such a pair (o_g, o_c) is allocated to p such that $o_g \in E_g$ and $o_c \in E_c$. Now if r values o_c at 0 or $-w$, we have a desired transfer (re-allocate o_c to r). Otherwise, it is the case that every r values the items in E_c at w. But this implies p again falls short of the w-valued items and hence violates normalization. Consider the other case when a pair (o_g, o_c) is allocated to r such that exactly one of them is w-valued for r. This again implies that p again falls short of the w-valued items and hence violates normalization. This settles our claim. □

We now consider general valuation beyond $\{w, 0, -w\}$ valuations and present some tractable cases for two agents and type-normalized valuations as follows.

Theorem 6.10. *For any allocation instance with type-normalized valuations and two agents, an EQ1 allocation always exists and can be computed in polynomial time.*

Proof. The algorithm starts by allocating items from $O^{\pm}$. For every $o \in O^{\pm}$, o is allocated to the agent who values it positively. (Since $o \in O^{\pm}$ and $n = 2$, exactly one of the agents values it positively). Once $O^{\pm}$ is allocated entirely, let v_1 and v_2 be the utilities derived by the two agents respectively, and say $v_1 > v_2$. Then by type-normalization, there must be enough items in O^+ such that agent 2 derives at least $v_1 - v_2$ utility. The algorithm then allocates all such items to agent 2, starting from her favorite good from the remaining items, until her utility becomes at least as much as agent 1. Indeed, this partial allocation is EQ1. The remaining instance is the one that contains objective goods and chores. Using Lemma 6.5, we get a complete EQ1 allocation, as desired. □

Theorem 6.11. *For any allocation instance with type-normalized valuations, two agents, and only subjective items, every PO allocation is EQ. Hence, an EQ allocation always exists and can be computed in polynomial time.*

Proof. Since $O^+ = O^- = \emptyset$ and $n = 2$, an item that is positively valued by one agent is negatively valued by the other. Consider a PO allocation $\Phi = \{\Phi_1, \Phi_2\}$. Under Φ, agent 1 gets all the items that it values positively, and agent 2 gets the remaining items (which are all positively valued by her, by the structure of the instance). Because of type-normalization ($\sum_{o \in G} v_i(o) = g$), we have $v_1(\Phi_1) = g = v_2(\Phi_2)$, hence Φ is EQ. □

We note here that in the presence of both subjective and objective items with general valuations, even with type-normalized instances, the intuitive idea of allocating the least happy agent her most favorite item, if it exists, else allocating the happiest agent her most disliked chore does not work. Consider the following instance.

	o_1	o_2	o_3	o_4	o_5	o_6	o_7	o_8
Alice	(10)	3	2	2	−2	−2	−2	−1
Bob	17	−1	−1	−1	−1	−1	−1	−1

The above instance is a type-normalized instant with the normalization constant 10. If Alice is chosen as the least happy agent (the choice is arbitrary in the first step), it gets o_1. Now it is easy to verify that no completion of this partial allocation can be EQ1 as now, the poor agent, Bob does not have any item that it values positively.

In addition, a leximin++ allocation (defined below), which is known to satisfy not only EQ1 but stronger EQX property for certain restricted instances, fails to achieve EQ1 for mixed instances, that too normalized and with few agents and items. A leximin allocation is the one that maximizes the minimum utility of an agent, subject to that maximizes the second-minimum utility, and so forth. Although there can be many leximin optimal allocations, they all induce a unique utility vector. Leximin++ allocation is an allocation that maximizes the minimum utility, and then maximizes the size of the bundle of the agent with minimum utility and so on. It is known that a leximin++ allocation is also EQX, for objective instances with identical chores (Barman et al., 2024a). When we go beyond objective instances to look for EQ1 allocations, the Example 6.12 suggests that leximin++ may not be EQ1, even for a normalized instance with two agents and two items.

Example 6.12. *Leximin++ is not EQ1 even for a normalized instance with two agent* $\{\Phi_1, \Phi_2\}$ *and two items* $\{o_1, o_2\}$. Φ_1 *values the items at* $\{10, -15\}$ *while* Φ_2'*s valuation is* $\{-2, -3\}$. *There are at most 4 possible allocations in this instance with utility vectors* $\{0, -5\}, \{-5, 0\}, \{10, -3\}, \{-15, -2\}$ *for* $\{\Phi_1, \Phi_2\}$ *respectively. The only leximin++*

allocation is the one with utility vector $\{10, -3\}$ but it is not EQ1. Even if we restrict the valuations to $\{1, -1\}$, the Example 6.13 suggests that leximin++ may not be EQ1.

Example 6.13. *Leximin++ is not EQ1 even with $\{1, -1\}$-normalized valuations. The example in Table 6.5 illustrates this. The highlighted allocation is leximin++ with the utility vector as $\{3, 1, 1, 1, 1\}$. This is not EQ1 since even if the last 5 agents choose to hypothetically ignore a good from $\Phi_1's$ bundle, they still fall short of the equitable utility.*

	o_1	o_2	o_3	o_4	o_5	o_6	o_7	o_8	o_9	o_{10}	o_{12}	o_{13}
Φ_1	-1	(1)	(1)	(1)	(1)	(1)	(-1)	(-1)	-1	-1	-1	-1
Φ_2	(1)	-1	-1	-1	-1	-1	-1	-1	1	1	1	1
a_3	1	-1	-1	-1	-1	-1	-1	-1	(1)	1	1	1
a_4	1	-1	-1	-1	-1	-1	-1	-1	1	(1)	1	1
a_5	1	-1	-1	-1	-1	-1	-1	-1	1	1	(1)	1
a_6	1	-1	-1	-1	-1	-1	-1	-1	1	1	1	(1)

Table 6.5: Leximin++ is not EQ1 even with {1, -1} normalized valuation. The highlighted allocation is leximin++ with the utility vector as $\{3, 1, 1, 1, 1\}$. This is not EQ1 since even if the last 5 agents choose to hypothetically ignore a good from $\Phi_1's$ bundle, they still fall short of the equitable utility.

6.4 EQ1+PO Allocations

Theorem 6.14. *For any allocation instance with identical (even non-normalized) valuations, an EQ1+PO allocation always exists and admits a polynomial time algorithm.*

Proof. Notice that for identical valuations, any complete allocation is PO. Consider any complete allocation Φ' which allocates a good g allocated to an agent i. Under any re-allocation of g, i becomes strictly worse off. On the other hand, consider any chore c allocated to some agent. Then under any re-allocation of c to any agent j, the receiving agent j becomes strictly worse off. Therefore, every complete allocation is PO. Now to compute an EQ1+PO allocation, since identical valuations are a subset of objective instances, we can use Theorem 6.3 to compute an EQ1 allocation, Φ. Then, because Φ is a complete allocation, it is also PO. This settles our claim. □

Consider a non-normalized instance with two agents, Alice, and Bob, and four items. Alice values the first 3 items at -1 and the last item at 1 while Bob values the first 3 items at 1 and the last item at -1. Here, the only PO allocation is the one under which Bob receives the first 3 items but this unique allocation is easily verified to be not EQ1. Further, the following instance illustrates that EQ1+PO may not exist even with type-normalized valuations.

Example 6.15. *Consider the following instance. The highlighted allocation is PO (unique) but not EQ1. Therefore, an EQ1+PO allocation may not exists even with type-normalized* $\{1, -1\}$ *valuations.*

	o_1	o_2	o_3	o_4	o_5	o_6
Φ_1	①	①	①	-1	-1	-1
Φ_2	-1	-1	-1	①	①	1
a_3	-1	-1	-1	1	1	①

We show below that deciding the existence of EQ1+PO allocations is hard even when the valuations are type-normalized.

Theorem 6.16. *For any allocation instance with type-normalized valuations, deciding the existence of an EQ1+PO allocation is strongly NP-Hard.*

Proof. We present a reduction from 3-Partition, known to be strongly NP-hard, where the problem is to decide if a multiset of integers can be partitioned into triplets such that all of them add up to a constant. Formally, the input is a set $S = \{b_1, b_2, \dots b_{3r}\}$; $r \in \mathbb{N}$; and the output is a partition of S into r subsets $\{S_1, S_2, \dots S_r\}$ such that $\sum_{b_i \in S_i} b_i = T = \frac{1}{r} \sum_{b_i \in S} b_i$. Given any instance of 3-Partition, we construct an instance of allocation problem as follows. We create r set-agents namely $\{a_1, a_2, \dots a_r\}$ and one dummy agent, a_{r+1}. We create $3r$ set-items namely $\{o_1, o_2, \dots o_{3r}\}$ and two dummy items $\{o_{3r+1}, o_{3r+2}\}$. All the set agents value the set items identically at b_i and the two dummy items at $-T$ and $-(r-3)T$ respectively. The dummy agent values the set-items $\{o_1, \dots o_{3r-1}\}$ at 0, o_{3r} at $-(r-2)T$ and the two dummy items at T and $(r-1)T$ respectively. Note that this is a type-normalized instance such that every agent values all the items together at $2T$. This completes the construction (see Table 6.6). We now argue the equivalence.

Forward Direction.

Suppose 3-Partition is a yes-instance and $\{S_1, S_2, \ldots S_r\}$ is the desired solution. Then, the corresponding EQ1+PO allocation, Φ, can be constructed as follows. For $a_i : i \in [r]$, we set $\Phi(a_i) = \{o_j : o_j \in S_i\}$. Finally, $\Phi(a_{r+1}) = \{o_{3r+1}, o_{3r+2}\}$. It is easy to see that Φ is an EQ1 allocation since $v_i(\Phi_i)) = T$ for all $a_i : i \in [r]$ and $v_i(\Phi_i) \geq v_{r+1}(\Phi(a_{r+1}) \setminus \{o_{3r+2}\}) = T$. Also, Φ is PO since it is a welfare-maximizing allocation where each item is assigned to an agent who values it the most.

Reverse Direction.

Suppose that Φ is an EQ1+PO allocation. Then, both the dummy items $\{o_{3r+1}, o_{3r+2}\}$ must be allocated to the dummy agent a_{r+1}, who is the only agent who values both the items positively. (Since any set-agent derives a negative utility from these items, allocating them to her violates PO). Since Φ is also EQ1, this forces every set-agent to derive the utility of at least T under Φ so that when it removes the item $\{o_{3r+2}\}$ from $\Phi(a_{r+1})$, EQ1 is preserved. ☐

	$o_1 \ldots o_{3r-1}$	o_{3r}	o_{3r+1}	o_{3r+2}
a_1	$b_1 \ldots b_{3r-1}$	b_{3r}	$-T$	$-(r-3)T$
a_2	$b_1 \ldots b_{3r-1}$	b_{3r}	$-T$	$-(r-3)T$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
a_r	$b_1 \ldots b_{3r-1}$	b_{3r}	$-T$	$-(r-3)T$
a_{r+1}	$0 \ldots 0$	$-r+2T$	T	$(r-1)T$

Table 6.6: Reduced instance as in the proof of Theorem 6.16

Given the non-existence and hardness even for the type-normalized valuations, we now present some tractable cases.

Theorem 6.17. *For $n = 2$ and type-normalized $\{w, 0, -w\}$ valuations, an EQ1+PO always exists and can be computed efficiently.*

Proof. We will show that the allocation Φ as returned by the Algorithm 4 is EQ1+PO. Notice that Φ allocates all items from $O^{\pm} \cup O^{+}$ non-wastefully. Also, $o \in O^{-}$ is allocated to an agent who values it at 0, otherwise, if both the agents value it at $-w$, then PO is satisfied irrespective of which agent ends up receiving that item. Therefore, Φ satisfies PO. We now argue that it also satisfies EQ1. Suppose i is the least happy agent by the end of Step 6. Say, the utility of i at this point is $k_1 \cdot c$ and that of agent j is $k_2 \cdot c$. Since i is the poor agent, therefore, $k_1 < k_2$. Then, there must be at least $k_2 - k_1$ items in O^{+} valued at $\{w, 0\}$ by $\{i, j\}$ (by type normalization). By construction, all these $k_2 - k_1$ items are allocated to i under Φ. This compensates for the inequity experienced by i so far. Now the allocation is extended by iteratively allocating the least happy agent her most valuable item, which maintains EQ1 till O^{+} is exhausted. For the items in O^{-} which are valued at $-w$ by both the agents, the happiest agent gets that item, which thereby maintains EQ1. For the remaining chores in O^{-}, at least one of the agents values them at 0 and ends up receiving the same. This does not violate EQ1 as it does not change the utility of the agents. This settles the claim. □

When we increase the number of agents from 2 to 3, an EQ1+PO allocation may not exist even with binary valuations. Consider 3 agents and 6 items such that two of them value the items at $\{1, 1, 1, 0, 0, 0\}$ and the last agents values at $\{0, 0, 0, 1, 1, 1\}$. Then any PO allocation gives a utility of 3 to the last agent, but at least one of the first two agents gets a maximum utility of 1, violating EQ1. Given this non-existence, we present the following algorithm that computes an EQ1+PO allocation for $\{1, 0, -1\}$ instances, whenever it exists.

Theorem 6.18. *For $\{1, 0, -1\}$ valuations (even non-normalized), an EQ1+PO allocation can be found efficiently if it exists.*

Proof. We show that Algorithm 5 returns an EQ1+PO allocation, whenever it exists. Towards correctness, we first establish the following result.

Lemma 6.19. *There is an EQ1+PO allocation for the partial instance $\mathcal{I} := O^{\pm} \cup O^{+}$ if and only if the Nash optimal allocation Φ' for the instance $\mathcal{I}^{\mathcal{G}}$ is EQ1.*

Proof. Suppose the Nash optimal allocation Φ' for instance $\mathcal{I}^{\mathcal{G}}$ is EQ1. We will show that Φ' is EQ1+PO for $\mathcal{I}$ as well. Note that Φ' is EQ1+PO for $\mathcal{I}^{\mathcal{G}}$ (since Nash satisfies PO). Since Φ' is PO for $\mathcal{I}^{\mathcal{G}}$, it never allocates an item to an agent who values it at 0. Therefore, in the instance, $\mathcal{I}$, Φ' never allocates any item from $O^{\pm} \cup O^{+}$ to an agent who values it at -1 or 0. Therefore, all items are allocated to agents who value them the most, at 1, and hence Φ' is PO (and EQ1) for $\mathcal{I}$.

Algorithm 4 EQ1+PO, n=2 and $\{w, 0, -w\}$ Type-Normalized Valuations

Input: An instance with 2 agents, m items and $\{1, 0, -1\}$-type-normalized valuation function.
Output: An EQ1+PO allocation Φ.

$O^+ : \{o : v_i(o) \geq 0 \text{ for } i \in [1,2]\}$
$O^- : \{o : v_i(o) \leq 0 \text{ for } i \in [1,2]\}$
$O^\pm : O \setminus \{O^+ \cup O^-\}$
$\Phi \leftarrow$ An empty allocation
while $O^\pm \neq \emptyset$ **do**
    $\Phi_i = \Phi_i \cup \{o\}$, where $o \in O^\pm$ such that $v_i(o) \geq 0$
if $\exists\, o \in O^+$ such that $v_i(o) = 0$ for some i **then**
    $\Phi(j) = \Phi(j) \cup \{o\}$ such that $j \neq i$
while $O^+ \neq \emptyset$ **do**
    $i \leftarrow$ least happy agent
    $\Phi_i = \Phi_i \cup \{o\}$ where o is the most valuable good from O^+ for i
if $\exists\, o \in O^-$ such that $v_i(o) = 0$ for some i **then**
    $\Phi_i = \Phi_i \cup \{o\}$
while $O^- \neq \emptyset$ **do**
    $i \leftarrow$ happiest agent
    $\Phi_i = \Phi_i \cup \{o\}$ where o is the most disliked chore from O^- for i
return Φ

On the other hand, suppose there is an EQ1+PO allocation Φ for instance $\mathcal{I}$. If the Nash optimal allocation for I^G is EQ1, we are done. Otherwise, suppose that the Nash optimal allocation for I^G is not EQ1. We will now argue by contradiction that this is not possible.

Since Φ is PO for the instance $\mathcal{I}$, any item $o \in O^\pm$ must have been allocated to an agent a who values it at 1. Indeed if not, then there is a Pareto improvement by allocating o to a which makes a strictly better off without making any other agent worse off. Likewise, all items $o \in O^+$ are allocated to respective agents who value them at 1. So, we have $\sum_{i \in N} v_i(\Phi_i) = m'$ where $m' = |O^\pm \cup O^+|$. Therefore, Φ is not only PO but also achieves the maximum utilitarian welfare in $\mathcal{I}$. Consider the same allocation Φ in the instance I^G. Since the only modification from $\mathcal{I}$ to $\mathcal{I}^{\mathcal{G}}$ is that for the agents who valued items in $O^\pm$ at -1 in $\mathcal{I}$, now value them at 0 in $\mathcal{I}^{\mathcal{G}}$ and everything else remains the same. So, we have $\sum_{i \in N} v_i(\Phi_i) = m'$ in $\mathcal{I}^{\mathcal{G}}$ as well. Hence, Φ maximizes the utilitarian welfare in I^G and hence, is PO. Also, Φ is EQ1 by assumption. therefore, for the instance $\mathcal{I}^{\mathcal{G}}$, Φ is an EQ1+PO allocation. But, Lemma 22 of Freeman et al.

(2019) shows that if an instance with binary valuations admits some EQ1 and PO allocation, then every Nash optimal allocation must satisfy EQ1. Therefore, the Nash optimal allocation in $\mathcal{I}^{\mathcal{G}}$ satisfies EQ1, which is a contradiction to our assumption. $\square$

If Φ is EQ1+PO for $\mathcal{I}$, then we show that the allocation of O^- as mentioned above maintains EQ1+PO. If Φ is not EQ1+PO, then we show that if the EQ1 violators under Φ can not be resolved by using the remaining chores, then there is no EQ1+PO allocation for the original instance.

We now show that if Φ is not EQ1, then there is no EQ1+PO allocation for the instance. If the EQ1 violators in Φ' (the Nash allocation on the reduced instance) cannot be resolved by using the remaining chores, then there is no EQ1+PO allocation. We argue this claim by contradiction. Suppose there is a complete EQ1+PO allocation, say $\Phi^\star$. Let t be the number of items in O^- that are valued at -1 by all the agents (they are 'universal' chores). Under the Nash optimal partial allocation Φ', let $v_p = \min_i v_i(\Phi'_i)$ be the utility of the poorest agent, $v_r = v_p + 1$ be the utility of rich agents, and the remaining agents are EQ1 violators with utility strictly greater than v_r. We denote the set of violators as S. It is easy to see that if $\sum_{s \in S}(v_s - v_r) \leq t$, then the t many -1's from O^- could have been used to bring down the utility of all the violators to v_r, ensuring that the completion of Φ' into Φ is EQ1. Therefore, we have $\sum_{s \in S}(v_s - v_r) > t$.

Now consider the EQ1+PO allocation $\Phi^\star$. Let $\bar{\Phi}^\star$ denote the restriction of $\Phi^\star$ to $O^\pm \cup O^+$. Then, $\bar{\Phi}^\star$ must be a PO allocation but is not EQ1, otherwise the Nash optimal allocation must have been EQ1. Let $v^\star_p,\ v^\star_r,\ v^\star_s$ denote the utilities of poor, rich, and violators in $\bar{\Phi}^\star$ and $S^\star$ be the set of violators. It must be the case that $\sum_{v \in V^\star}(v^\star_v - v^\star_r) \leq t$, since the completion $\Phi^\star$ is an EQ1 allocation. Also, notice that since both $\bar{\Phi}^\star$ and Φ' are Pareto optimal, they both allocate the items in $O^\pm \cup O^+$ non-wastefully. Therefore, the sum of the utilities under both the partial allocations is exactly m', where $m' := |O^\pm \cup O^+|$. We now consider the following cases.

- $v_p \geq v^\star_p$. Then,

$$\begin{aligned}
\sum_{p,\,r,\,s}(v_p + v_r + v_s) &> \sum_{p,\,r}(v_p + v_r) + (t + v_r|S|) \\
&\geq \sum_{p,\,r}(v^\star_p + v^\star_r) + (t + v^\star_r|S|) \\
&\geq \sum_{p,\,r,\,v}(v^\star_p + v^\star_r + v^\star_v)
\end{aligned}$$

  which is a contradiction as the first and the last term is equal to m'.

- $v_p < v_p^\star$. Then we will show that A' is not a Nash optimal allocation. We have $v_r = v_p + 1 < v_p^\star + 1 = v_r^\star$. Consider the nash welfare

$$\begin{aligned} NW(A') &= (\prod_{p,r,s} v_p v_r v_s)^{\frac{1}{n}} < (\prod_{v,r,s} v_p^\star v_r^\star v_s)^{\frac{1}{n}} \\ &< (\prod_{p,r,s} v_p^\star v_r^\star v_s^\star)^{\frac{1}{n}} \\ &= NW(\bar{A}^\star) \end{aligned}$$

  The last but one inequality follows from the fact that $\sum_{p,\, r,\, s}(v_p + v_r + v_s) = \sum_{p,\, r,\, s}(v_p^\star + v_r^\star + v_s^\star) = m'$ and the product of a set of numbers with a constant sum has the highest outcome if they are closer/equal to each other rather than being further away ($\sum_{s \in V^\star}(v_s^\star - v_r^\star) \le t < \sum_{v \in V}(v_s - v_r)$ ensures that $v_r^\star$ and $v_s^\star$ are closer to each other than v_s and v_r).

This settles the claim. □

Algorithm 5 EQ1+PO

Input: n agents, m items and $\{1, 0, -1\}$ valuations

Output: An EQ1+PO allocation Φ, if it exists, else, the most equitable allocation Φ within the set of PO allocations.

$O^+ : \{o : v_i(o) \ge 0 \,\forall\, i \in [n]\}$
$O^- : \{o : v_i(o) \le 0 \,\forall\, i \in [n]\}$
$O^\pm : O \setminus \{O^+ \cup O^-\}$
$\Phi \leftarrow$ An empty allocation
while $\exists o \in O^- : v_i(o) = 0$ for some agent i **do**
    $\Phi_i \leftarrow \Phi_i \cup \{o\}$
For $o \in O^\pm$ such that $v_i(o) = -1$, set $v_i(o) = 0$.
Let Φ' be the Nash optimal allocation on the set of goods $O^\pm \cup O^+$.
$\Phi_i \leftarrow \Phi_i \cup A'_i$
while $\exists o \in O^-$ **do**
    Let i be the agent with maximum utility under Φ_i
    $\Phi_i \leftarrow \Phi_i \cup \{o\}$
return Φ

6.5 EQX and Social Welfare

In this section, we discuss the computational complexity of finding allocations that maximize the utilitarian welfare (UW) or egalitarian welfare (EW) within the set of EQX allocations for mixtures. Even in the case when all the underlying items are goods, finding a UW or EW allocation within the set of EQ1 allocations (UW/EQ1 or EW/EQ1), for a fixed number of agents, is weakly NP-Hard (Sun et al., 2023b). This rules out the possibility of a polynomial-time algorithm for mixed instances. Nonetheless, we present pseudo-polynomial time algorithms for finding such allocations in mixed instances with a constant number of agents. Our algorithmic technique extends those developed by Aziz et al. (2023b) to the mixed setting, and relies on dynamic programming.

Theorem 6.20. *For any mixed fair division instance with a constant number of agents, computing a Utilitarian or Egalitarian maximal allocation within the set of EQX allocations admits a pseudo-polynomial time algorithm.*

Proof. We present a dynamic programming algorithm that keeps a set of states representing the set of possible allocations. At each state, it considers allocating the item o_k to one of the n agents. Finally, it chooses the state that optimizes for social welfare and respects EQX. We denote $\sum_{o \in O} \max_i v_i(o) = V_g$ and $\sum_{o \in O} \min_i v_i(o) = -V_c$. The states are of the form $(k, \mathbf{v}, \mathbf{g}, \mathbf{c})$ where $k \in [m], v_i \in [-V_c, V_g],\ g_i, c_i \in [m]\ \forall i \in [n]$. The items g_i and c_i refer to the least valuable good and the least disliked chore in the bundle of agent i. The state $(k, \mathbf{v}, \mathbf{g}, \mathbf{c}) :=$ True if and only if there is an allocation of objects $\{o_1, \dots o_k\}$ the value of agent i is at least v_i and the bundle of i contains g_i and c_i as the least valuable good and the least disliked chore respectively. The initial state $(0; 0, 0 \dots 0; \mathbf{0}, \mathbf{0})$ refers to the empty allocation and is vacuously true. Consider the case when o_1, which is a good for agent i is allocated to i. Then, the state $(1, 0, \dots v_i(o_1), \dots 0; \emptyset, \dots o_1, \dots \emptyset; \emptyset \dots \emptyset)$ is True and every other state is False. The state that corresponds to the allocation of o_k to some agent, say i is given as follows: If $g_i \neq o_k$ and $c_i \neq o_k\ \forall i \in [n]$, then

$$\begin{aligned}(k, v_1, \dots, v_i, \dots v_n; g_1, \dots, g_n; c_1, \dots, c_n) = \\ \vee_{i \in [n]}(k-1, v_1, \dots, v_i - v_i(o_k), \dots v_n; \\ g_1, \dots g_n; c_1, \dots, c_n)\end{aligned} \tag{6.1}$$

where if o_k is a good for agent i then $v_i(g_i) \leq v_i(o_k)$ else, $v_i(c_i) \geq v_i(o_k)$. Else, if $g_i = o_k$, then,

$$(k, v_1, \dots v_n; g_1, \dots, g_i = o_k, \dots g_n; c_1, \dots, c_n) = \\ \vee_{g \in [o_1, o_{k-1}]} (k-1, v_1, \dots, v_i - v_i(o_k), \dots v_n; \\ g_1, \dots, g, \dots g_n; c_1, \dots, c_n) \tag{6.2}$$

where $g \in \cup_{r \in [k-1]} o_r$ and $v_i(o_k) \leq v_i(g)$.

Else, if $c_i = o_k$, then,

$$(k, v_1, \dots, v_i, \dots v_n; g_1, \dots, g_n; c_1, \dots c_i = o_k, \dots c_n) = \\ \vee_{c \in [o_1, o_{k-1}]} (k-1, v_1, \dots, v_i + v_i(o_k), \dots v_n; \\ g_1, \dots, g_i, \dots g_n; c_1, \dots c, \dots c_n) \tag{6.3}$$

where $c \in \cup_{r \in [k-1]} o_r$ and $v_i(o_k) \geq v_i(c)$. The states $(m, \mathbf{v}, \mathbf{g}, \mathbf{c})$ correspond to the final allocation. An allocation corresponding to one of the final states is EQX if and only if $v_i \geq v_j - v_j(g_j)$ and $v_i - v_i(c_i) \geq v_j$ for every pair of agents $i, j \in [n]$.

We now argue that every entry in the DP table is indeed computed correctly. To that end, we do an induction on the number of items allocated. For the base case, when one item is allocated, to say agent i, then only i derives the value of $v_i(o_1)$, and the rest of the agents get a value 0. Depending on whether o_1 is a good or a chore for agent i, either $g_i = o_1$ or $c_i = o_1$ and everything else is $\emptyset$. This is correctly captured in the base case. By induction hypothesis, suppose all the table entries that allocate the first $k-1$ items are computed correctly. Consider the allocation of k^{th} item as captured by the table entry $(k, \mathbf{v}, \mathbf{g}, \mathbf{c})$.

First consider the case when $g_i \neq o_k$ and $c_i \neq o_k$ $\forall i \in [n]$ in the entry $(k, \mathbf{v}, \mathbf{g}, \mathbf{c})$. Suppose RHS of Equation (6.1) is True. Suppose o_k is a good for some agent i such that $v_i(g_i) \leq v_i(o_k)$ (the case of chores can be argued similarly) and RHS of Equation (6.1) is True for the index i. That is, $(k-1, v_1, \dots, v_i - v_i(o_k), \dots v_n; g_1, \dots g_i, \dots g_n; c_1, \dots c_i, \dots c_n)$ is True. It means that there is allocation of $k-1$ items such that everyone gets a utility of $(v_1 \dots v_i - v_i(o_k), \dots v_n)$ and $(g_1 \dots g_n)$ and $(c_1, \dots c_n)$ are the least valued goods and the least disliked chores in the respective bundles. Then allocating o_k to i^{th} agent gives an allocation with utilities $(v_1, \dots v_i, \dots v_n)$ such that the set of the least valued goods and least disliked chores remain the same for all the agents (because $v_i(g_i) \leq v_i(o_k)$). Therefore, LHS of Equation (6.1) is True.

On the other hand, suppose the LHS of Equation (6.1) is True. Suppose o_k is a good and belongs to agent $i's$ bundle. Then since $g_i \neq o_k$, it means that g_i is the least valued item in agent $i's$ bundle. Consider the allocation after removing o_k from $i's$ bundle. Then, it corresponds to an allocation of $k-1$ items such that each agent gets a utility of $(v_1, \ldots v_i - v_i(o_k), \ldots v_n)$ such that $(g_1, \ldots g_n)$ and $(c_1, \ldots c_n)$ are the least valued goods and the least valued chores in the respective bundles. Therefore, $(k, v_1, \ldots v_i - v_i(o_k), \ldots v_n, g_1, \ldots g_n, c_1, \ldots c_n)$ is True and hence RHS of Equation (6.1) is True.

If o_k is a good for agent i such that $v_i(g_i) > v_i(o_k)$, RHS of Equation (6.1) is False by definition, and LHS of Equation (6.1) is False since g_i was the least valuable item in $i's$ bundle but $g_i \neq o_k$.

Now consider the case when $g_i = o_k$ for some $i \in [n]$ in the table entry $(k, \mathbf{v}, \mathbf{g}, \mathbf{c})$. This means that $(k, \mathbf{v}, \mathbf{g}, \mathbf{c})$ corresponds to a state where o_k is allocated to i who considers it to be a good. Now suppose RHS of Equation (6.2) is True. Then, adding o_k to $i's$ bundle increases its utility to v_i from $v_i - v_i(o_k)$ and since $v_i(o_k) < v_i(g)$, o_k is the new least valued item in $i's$ bundle, which is captured in the LHS and hence LHS is True. On the other hand, if LHS of Equation (6.2) is True, then consider the allocation post removing the item o_k from $i's$ bundle. Then, RHS of Equation (6.2) is True for a least valued item g in $i's$ bundle from $\{o_1, \ldots o_{k-1}\}$. The case when $c_i = o_k$ for some $i \in [n]$ is argued similarly. This settles the claim that the table entries are correctly computed at every step.

The states $(m, \mathbf{v}, \mathbf{g}, \mathbf{c})$ correspond to the final allocation. An allocation corresponding to one of the final states is EQX if and only if $v_i \geq v_j - v_j(g_j)$ and $v_i - v_i(c_i) \geq v_j$ for every pair of agents $i, j \in [n]$.

Among the states that correspond to EQX allocations, the algorithm selects the one that maximizes UW, that is, $\sum_{i \in [n]} v_i$ or EW, that is, $\min_{i \in [n]} v_i$. The total number of possible states is $O(m^{2n+1} \cdot V^n)$, where $V = V_g + |V_c|$. Computing one state requires look-ups of at most n previously computed states. Therefore, the time required to compute all the states is $O(n \cdot m^{2n+1} \cdot V^n)$. Finding the UW of an EQX allocation (corresponding to every final state that returns True) takes $O(n)$ time and computing EW takes atmost $O(n^2)$ time. The final step is to compute the maxima of these values. This takes a quadratic overhead in the number of final states that corresponds to True. Hence, the total runtime is bounded by $O(n^3 \cdot m^{4n+2} \cdot V^{2n})$.

Note that when valuations are $\{1, 0, -1\}$, we have $V \leq 2m$ and, hence, the algorithm runs in polynomial time for a fixed number of agents. □

6.6 EF+EQ+PO Allocations

Proposition 6.21. *For $\{w, 0, -w\}$-valuations, an EQ+PO allocation is also envy-free (EF).*

Proof. Suppose Φ is an EQ+PO allocation for the given instance. EQ implies that we have $v_i(\Phi_i) = v_i(\Phi_j) = k \cdot c$ and PO ensures that if an agent receives an item o with value $-w$, then everyone else values o at $-w$. (Else, if there is an agent i such that $v_i(o) = 0$ *or* c, then allocating o to i is a Pareto improvement). Likewise, if an agent receives an item that it values at 0, then everyone else values that item at either 0 or $-w$, again for the same reason. Now suppose Φ is not EF. Then, there is a pair of agents i and j such that $v_i(\Phi_i) = k \cdot c$ but $v_i(\Phi_j) > k \cdot c = v_j(\Phi_j)$. This implies that there is an item o in $j's$ bundle that is valued at 0 (or $-w$) by j but valued at w (or 0) by i. Allocating o to i is a Pareto improvement, which is a contradiction. Therefore, Φ is EF. □

Since the allocations constructed in the proof of Theorem 6.2 and Theorem 6.16 is EF, therefore, we get the following results.

Corollary 6.22. *Deciding whether an instance admits an allocation that is simultaneously EF+EQ1 or EF1+EQ1 is (weakly) NP-complete.*

Corollary 6.23. *Deciding whether a type-normalized instance admits an allocation that is simultaneously EF+EQ1+PO or EF1+EQ1+PO is (strongly) NP-hard.*

Notice that for $\{1, 0, -1\}$ valuations, a Pareto optimal allocation is EF1 if and only if it is EFx, and is EQ1 if and only if it is EQX. Therefore, the above two results hold for all combinations of X + Y + PO, where X ∈ {EFx, EF1} and Y ∈ {EQX, EQ1}. The allocation constructed in Theorem 6.18 can be easily verified to be EF1, thereby confirming the following result.

Corollary 6.24. *For $\{1, 0, -1\}$ valuations, an EF1+EQ1+PO allocation can be computed in polynomial time, whenever such an allocation exists.*

6.7 Concluding Remarks

We present a comprehensive picture of the existence and computational complexity of approximate equitable allocations for mixed items, coupled with efficiency and welfare notions. We show that finding an EQ1 allocation for instances with mixed items is NP-Hard, unlike the 'only goods' and 'only chores' settings where the problem admits efficient algorithms. We also present several tractable cases that can be solved efficiently by careful transfers of items among agents. Further, we show that deciding the existence of EQ1+PO allocation is strongly NP-Hard even for type-normalized valuations. We settle the complexity of finding welfare-maximizing EQ1 allocations by presenting a pseudo-polynomial time algorithm. A polynomial time algorithm is ruled out because of the known (weak) hardness result for the same. The question of deciding the existence of EQ1 allocations under normalized valuations stands open.

Chapter 7

Generalized Consesus Allocations- Valuing the Perspectives of Others

> *"If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is."*
>
> - John von Neumann, *Archaeology of Computers–Reminiscences, 1945-1947*

7.1 Introduction

The question at the heart of resource allocation is to find a *good* allocation — and as the reader may anticipate, this further prompts the issue of what makes one allocation better than another, and also if there is an absolute sense in which an allocation can be thought of as a "good" allocation.

In the context of cake cutting, the existence and complexity of "exact divisions"(where all agents agree on the value of the division) has been well-studied, usually referred to as consensus halving (Simmons and Su, 2003; Deligkas et al., 2021b; Filos-Ratsikas and Goldberg, 2019). In this problem, a homogenous resource is to be divided into two parts such that every agent values the two parts equally. Formally, there are n agents with valuation functions over the interval say $\mathcal{I} = [0, 1]$. The goal is to divide the interval into pieces using at most n cuts and assign a label from $\{+, -\}$ to each piece, such that every agent values the total amount of $\mathcal{I}$ labeled $'+'$ and the total amount of $\mathcal{I}$ labeled $'-'$ equally. Prior results have shown that consensus halving and its approximate versions (in which there is a small discrepancy between the values of the two portions), are hard even for agents with simple restricted valuation functions (Filos-Ratsikas et al., 2020; Filos-Ratsikas and Goldberg, 2018, 2019). Generalizations of consensus halving where instead of two parts, the homogenous resource is to be divided into k parts of equal value to all the agents have also been explored (Simmons and Su, 2003; Filos-Ratsikas et al., 2020).

In a similar vein, in the context of indivisible items, we pose the "exact-and-equitable" division question. In contrast to the consensus halving scenario, in our setting, the items can not be fractionally assigned and there is no underlying geometry. In particular, suppose agent i values item j at u_{ij}: we would like to divide the m items into n bundles in such a way that *every* agent values *every* bundle at v, for some common value v. Note that this is already weakly NP-hard to determine between two agents and m items with identical valuations[1]. Our first result is to show that the problem is NP-complete even in the setting of additive binary valuations (which is to say that all utilities u_{ij} are either 0 or 1, and the value that an agent has for a bundle is the number of items in it that it values at 1). In particular, we show hardness when the problem is to determine if m items can be divided into n bundles so that every agent values every bundle at 1, a constant independent of the number of items.

Technical Motivation.

In practice, a perfectly equitable consensus may, in general, be too much to ask for. A natural relaxation to ask for an approximate consensus: where all agents agree that all bundles have a value in some specified range, say $[p, q]$. We show that even in the setting of additive binary valuations, the problem of dividing m items into a collection of k bundles so that all k bundles are valued at either 0 or 1 is NP-complete, even when each agent values two items and each item is valued by at most four agents. From an algorithmic standpoint, note that iterating over all possible ways of dividing m items into any number of bundles requires $2^{O(m \log m)} \cdot n^{O(1)}$ time. This can be used to determine the existence of an allocation with any desired set of properties with an overhead in running time proportional to the time required to validate the properties

[1] This can be shown by a standard reduction from the Partition problem involving k numbers $\{n_1, \ldots, n_k\}$: introduce k items and have both agents value the j^{th} item at n_j: this instance admits an exact and equitable allocation if and only if the k numbers can be partitioned into two subsets whose sums are equal.

<table>
<tr><th></th><th>Strong Security</th><th>Weak Security</th><th>Strong Abundance</th><th>Weak Abundance</th></tr>
<tr><td>Capacitated</td><td colspan="2">NP-complete
(Theorem 7.4 and Corollary 7.6)</td><td rowspan="3">NP-complete
(Theorem 7.16 and Corollary 7.17)</td><td rowspan="3">Open</td></tr>
<tr><td>Egalitarian</td><td colspan="2">NP-complete
(Theorem 7.7 and Corollary 7.8)</td></tr>
<tr><td>Utilitarian</td><td colspan="2">NP-Complete
(Corollary 7.9)</td></tr>
<tr><td>(Left) Extremal Valuations</td><td>P</td><td>P</td><td>P</td><td>P</td></tr>
</table>

Table 7.1: A partial summary of our results. The egalitarian version of the problem asks if every agent's utility is at least a given target, while the utilitarian version asks if the total utility derived by all agents meets a given target. For (left) extremal instances, we give efficient algorithms for all the capacitated problems.

sought. Assuming ETH, we show that there is no $2^{o(m \log m)} \cdot n^{O(1)}$ algorithm for the following closely related question: given n agents valuations of m items and a number k, and parameters p, q and c, is it possible to divide m items into k bundles so that every agent values all but c of the k bundles in the range $[p, q]$? We show this hardness result for $p = 1, q = 5$, and $c = 1$. We note that in our setting of approximate division, it is possible that $k < n$, and the final allocation of bundles to agents associated with this division would implicitly allocate empty bundles to some agents.

Given the hardness of finding even "almost" exact equitable divisions, we consider other ways of relaxing the demands we make from a perfectly equitable consensus allocation. Here, we treat a common valuation as a target lower or upper bound, instead of an exact goal. Given an allocation $(\Phi_1, ..., \Phi_n)$, we think of the base value of a bundle Φ_i as being $v_i(\Phi_i)$: the value that the agent the bundle has been assigned to has for it. Now we consider allocations where all external valuations of any bundle are: (a) at least the base value; and (b) at most the base value. These concepts relate also to the quality of the valuations from a "user experience" perspective, and we motivate this with the following concrete examples.

Real-World Motivation.

Our first notion aims to find an allocation where the agents feel validated/secure by the opinions of others. An agent feels secure if the other agents value the items she owns highly – the sense of security for an individual is linked to how much other agents appreciate her possessions. If an agent owns a bundle that she values highly but that everyone else deems worthless, then it is natural for the agent to be unhappy about the situation, or at any rate, be suspect about their own judgement. For instance, consider that a piece of land may be valued very highly by an agent, whom we call a, based on its virtues as explained by a real estate agent. However, it is possible that the said piece of land is not resourceful in the traditional sense: so the other agents would not value it as much. The valuations of the other agents may alert a to the possibility that their choice is sub-optimal. Overall, in any situation where goods or resources are being allocated, people often want to be recognized or appreciated for what they own or contribute.

The opposite situation is also something that one might want to avoid: when the agents feel that their bundle is useless, even as other agents disagree. That is, the agents underestimate the worth of their bundles. For example, say a company has announced a collection of goods as bonuses for their employees. If an employee, say e, ends up with a good they already have, she may not value it very much, while another employee may value it much more. Note that e may not be envious of any other employees in this situation, and may not even trigger envy among other agents. It is plausible that e values all goods similarly, and all other agents value their bundles at least as much as they value e's bundle. However, we propose that e may still be unhappy because of their bundle being *perceived as being valuable by other employees* while they feel they know better. The dissatisfaction is from a sense of being misunderstood for having more than what (they feel) they actually have. In a housing market, a tenant may be given a smaller unit than she expected (contrary to others who consider the unit rich enough), leading her to underestimate the value of her living space and potentially feel dissatisfied with the living situation.

In recent work, Hosseini (2023) argues that the current fairness criteria fall short of encompassing the intricacies of human decision judgment and psyche. The author says that human value judgment is rarely self-reflective and is often influenced by various individual, social, and cognitive factors. The author explicitly asks the question *"What (and how) cognitive and behavioral factors influence individuals' perception of fairness, and how should these human judgments form new fairness concepts and inform the design of new algorithms?"*. We believe that our work takes a step forward in answering the above question as we introduce novel concepts that capture elements of human psychology molded and influenced by the perspectives of others.

We defer the formal definitions to Section 7.2, but introduce the main ideas here, which are easily stated. Fix an allocation Φ of a set of m goods among n agents, and let a and b be an arbitrary but fixed pair of agents. Let's say that a's value for his bundle happens to be u, while

b values that bundle at w. The key notions we explore are the following. The agent b makes a *insecure* if b does not value $a's$ possessions as much, that is, $w < u$. And the agent a is *modest* with respect to b if a herself does not value her bundle as much, that is, $w > u$. If an allocation has no insecure agents, we call it a strongly secure allocation, and if an allocation has no modest agents, we call it a strongly abundant allocation. We also suggest that it is *irrational* for agent a to be rattled by agents who it perceives as being not as well-off as herself. Therefore, if we insist that a's value for b's bundle being at least u is a pre-requisite for triggering the emotions above, we say that the agent's behavior is *rational*. An allocation that satisfies the weaker requirement of not having rationally insecure agents is called a weakly secure allocation, and analogously, an allocation that does not have rationally modest agents is called a weakly abundant allocation.

Allocations that are both strongly secure and strongly abundant is called an *exact division* or a *consensus allocation*. In particular, a complete allocation Φ is said to be an exact division with ratios $(w_1, \ldots, w_m)$ if all agents "agree" that the value of bundle i is w_i for all $i \in [m]$. Our notion of an exact-and-equitable allocation extends the notion of a consensus allocation further by requiring that all the ratios are the same.

Despite the desire for traditional fairness objectives such as envy-freeness (where each agent prefers their own bundle over others') and proportionality (ensuring each agent receives a proportional share), it's important to note that these goals may not always be guaranteed in a complete allocation scenario. The definitions we propose overcome this non-existence. Strongly secure (respectively, abundant) allocations always exist. They can be obtained, for example, by giving all goods to the agent who values the entire bundle the least (respectively, the most). However, these allocations are evidently not desirable, so to pose a more reasonable question, we impose cardinality and welfare constraints on the bundles, requiring that no agent receives more than k items. Our main results are summarized in Table 7.1. We show that finding good allocations in the sense of security and abundance is intractable in conjunction with cardinality constraints and welfare goals in general. To complement the hardness, we give efficient algorithms for the cases when agent preferences are "extremal" (Elkind and Lackner, 2015), demonstrating a non-trivial tractable subclass of instances. We note that the practical applicability of the extremal preference model extends to situations with location constraints. For example, in scenarios related to land allocation, areas in close proximity to amenities may emerge as the preferred choice among agents.

We note here that the goals we present here may be at loggerheads with the traditional fairness goals. For instance, consider an example with two goods g_a and g_b and two agents a and b, where a approves only g_a and b approves only g_b. It is easy to verify that all secure allocations have an envious agent and the only envy-free allocation, which allocates g_a to a and g_b to b is indeed insecure for both the agents! This is precisely because the goal here is to shift the focus from envy (which an agent feels towards others) to validation (which an agent expects from others).

Related Work.

One of the widely used notions of judging the quality of an allocation is the extent of envy. Given an allocation, an agent envies another if she perceives the bundle of the other agent to be more valuable than her own. An allocation is *envy-free* if no agent envies any other agent (Gamow and Stern, 1958; Foley, 1967; Budish, 2011; Lipton et al., 2004). Another compelling notion is *equitability*, in which agents derive equal utilities from their assigned shares (Dubins and Spanier, 1961; Gourvès et al., 2014; Freeman et al., 2019). Note that the trivial allocation that leaves every agent empty-handed is vacuously envy-free and equitable. Therefore, in the interest of a non-trivial pursuit, one is typically interested in fair allocations that also satisfy some criteria of economic *efficiency*. Completeness (every good should be allocated to some agent), Non-wastefulness (no agent receives a good that is worth nothing to her and worth something to another agent), and Pareto-efficiency (there is no other allocation that would make at least one agent strictly better off while not making any of the others worse off) are a few standard notions of efficiency in the literature. Allocations are also judged by *welfare*, which is a measure of the utility that the agents derive from their respective bundles. There are several notions of welfare corresponding to the way individual utilities are aggregated – egalitarian welfare is the lowest utility achieved by any agent; utilitarian and Nash welfare are the sum and the geometric mean of individual utilities respectively (Caragiannis et al., 2019b; Freeman et al., 2019; Barman et al., 2018a; Aziz et al., 2023b).

Recent works have also considered consensus halving when the resource is a set of homogenous items but without a linear ordering (Goldberg et al., 2022). As a discrete counterpart to consensus halving, a well-studied problem is Necklace Splitting (Filos-Ratsikas and Goldberg, 2019). Here, the input is an open necklace with ka_i beads of color i, for $1 \leq i \leq n$. The task is to cut the necklace in $(k-1) \cdot n$ places and partition the resulting substrings into k collections, each containing precisely a_i beads of color $i, 1 \leq i \leq n$. An "open necklace" means that the beads do not form a cycle, but a string. Here, it is implicitly assumed that the beads are valued identically by everyone.

Towards incorporating the perspectives of others, Shams et al. (2022) proposed the notion of approval envy, where an agent a experiences approval envy towards b if she is envious of b, and sufficiently many agents agree that this should be the case, from their own perspectives. Our work resembles this line of thought that takes into account the others' perspectives but demands that the other agents should agree on a certain valuation of the bundle in question.

Unlike envy-freeness, our notions involve an interpersonal utility comparison, which is well substantiated by Herreiner and Puppe (2009). Here, the author asserts, with empirical evidence, that human perceptions of fairness are seldom aligned with theoretical properties like envy-freeness and, instead, are often shaped by interpersonal comparisons, a facet deeply rooted in human behavior. The author further argues that interpersonal comparisons are dominant and envy-freeness plays a secondary role in situations in which Pareto optimality and inequality aversion (again, an interpersonal utility comparison) are not sufficient to determine a fair allocation.

7.2 Preliminaries

In this chapter, we focus entirely on *binary valuations*, which is the special case when $v_{i,j} \in \{0,1\}$ for all $i \in [n], j \in [m]$. We note here that binary valuations are a crucial subclass with simple elicitation and several works in computational social choice literature (Brams and Fishburn, 1978; Lackner and Skowron, 2023; Halpern et al., 2020; Barman et al., 2018b) have paid special attention to the binary case. An instance $\mathcal{I} = \{A, O, \mathbf{v}\}$ of the allocation problem is said to have an *extremal interval structure with respect to goods* if there exists an ordering σ of the goods such that the goods liked (valued at 1) by any agent a, denoted by $\mathcal{P}(a)$, forms a prefix or suffix of σ (Elkind and Lackner, 2015). Further, we say that $\mathcal{I}$ has a ***left** (respectively, **right**) extremal interval structure with respect to goods* if there exists an ordering σ of the goods such that for every agent a, the set of goods $\mathcal{P}(a)$ forms a prefix (respectively, suffix) of σ. We note here that the practical applicability of the extremal preference model extends to situations with location constraints, for instance, in scenarios related to land allocation, areas near amenities may emerge as the preferred choice among agents.

We now turn to the definitions that, to the best of our knowledge, are introduced in this work. First, we say that an allocation is an exact and equitable allocation if every agent values every bundle at a common value; in other words, there exists a value ν such that for all $i, j \in [n]$ $u_i(\Phi_j) = \nu$. An (k, p, q, c)-approximate exact and equitable allocation is a division of m items into k bundles where every agent has an "almost identical" valuation of all but c bundles; that is, for all $i \in [n]$ and for at least $k - c$ values of $j \in [k]$, it is true that $u_i(\Phi_j) \in [p, q]$. For the remaining definitions, we let a and b be an arbitrary but fixed pair of agents, Φ be a fixed allocation. Say a values her bundle at u.

Strongly Secure Allocation.

Agent a is *insecure* with respect to b if b's valuation of a's bundle is less than u. That is, $u_b(\Phi_a) < u_a(\Phi_a)$. An allocation is *strongly secure* if no agent is insecure with respect to another. That is, $\forall a, b : u_b(\Phi_a) \geq u_a(\Phi_a)$.

Strongly Abundant Allocation.

Agent a is *modest* with respect to b if b's valuation of a's bundle is more than u. That is, $u_b(\Phi_a) > u_a(\Phi_a)$. An allocation is *strongly abundant* if no agent is modest with respect to another. That is, $\forall a, b : u_b(\Phi_a) \leq u_a(\Phi_a)$.

Weakly Secure/Abundant Allocation.

Agent a *cares for* b if a's valuation of b's bundle is at least u. That is, $u_a(\Phi_b) \geq u_a(\Phi_a)$. Else, a *does not care* for b. An allocation is *weakly secure* if for every pair of agents a, b; it holds that if a is insecure with respect to b, then a does not care for b. That is, if $\exists a, b : u_b(\Phi_a) < u_a(\Phi_a)$, then $u_a(\Phi_b) < u_a(\Phi_a)$. Likewise, an allocation is *weakly abundant* if for every pair of agents a, b; it holds that if a is modest with respect to b, then a does not care for b. That is, if $\exists a, b : u_b(\Phi_a) > u_a(\Phi_a)$, then $u_a(\Phi_b) < u_a(\Phi_a)$

Note that an allocation is a consensus allocation if it is both strongly secure and strongly abundant. Also, observe that every strongly secure (abundant) allocation is also weakly secure (abundant), but the converse is not necessarily true.

Existence and Computational Questions.

Consider the following allocations: Φ assigns the grand bundle to the agent who values it the least and assigns the empty bundle to all other agents; while Φ' assigns every good to the agent who values it the *least*. It is easy to verify that both Φ and Φ' are strongly secure allocations. This implies that a strongly (and hence, weakly) secure allocation always exists, contrary to the fact that their fairness counterparts, i.e. envy-free or equitable allocations may not exist. However, note that neither of the above allocations are satisfying, in an intuitive sense of the phrase. The first one is problematic since it is rather skewed, and the second one seems to be optimizing for the worst possible welfare. It is also straightforward to come up with analogously trivial allocations that are strongly abundant, and equally unremarkable. We therefore propose and study the following computational questions, where we add capacity constraints so as to obtain secure (and abundant) allocations with better welfare guarantees than Φ and Φ'.

Computational Questions.

We conclude this section with the computational questions that we address in this paper. We focus on allocations with *cardinality constraints*, wherein we require the number of goods in all bundles to be at most k, and k is a part of the input. Specifically, we consider the following variations of consensus-based division and one-sided consensus problems:

Exact Equitable Allocations
Input: A set A of n agents, a set O of m goods, a valuation matrix $\mathbf{v} \in \{\mathbb{N}\}^{m \times n}$, and $p \in \mathbb{N}$.
Question: Does there exist a partition of the m goods into n bundles such that each agent values each bundle at p?

Approximate Exact Equitable Allocations
Input: A set A of n agents, a set O of m goods, a valuation matrix $\mathbf{v} \in \{\mathbb{N}\}^{m \times n}$, and $k, p, q, c \in \mathbb{N}$.
Question: Does there exist a (k, p, q, c)-approximate exact and equitable allocation?

Strongly Secure Capacitated Allocation
Input: A set A of n agents, a set O of m goods, a valuation matrix $\mathbf{v} \in \{0,1\}^{m \times n}$, and $k \in \mathbb{N}$.
Question: Does there exist a strongly secure allocation where each bundle has at most k goods?

Strongly Secure Capacitated Allocation
(Egalitarian)
Input: A set A of n agents, a set O of m goods, a valuation matrix $\mathbf{v} \in \{0,1\}^{m \times n}$, and $k, \ell \in \mathbb{N}, \ell \neq 0$.
Question: Does there exist a strongly secure allocation where each bundle has at most k goods and each agent has a utility of at least ℓ?

Strongly Secure Capacitated Allocation
(Utilitarian)
Input: A set A of n agents, a set O of m goods, a valuation matrix $\mathbf{v} \in \{0,1\}^{m \times n}$, and $k, \ell \in \mathbb{N}, \ell \neq 0$.
Question: Does there exist a strongly secure allocation where each bundle has at most k goods and total utility is at least ℓ?

The problems for weakly secure, strongly abundant, and weakly abundant allocations are defined analogously.

7.3 Exact Equitable Allocations

We begin by showing that finding an exact allocation that is equitable is hard even in the setting of additive binary valuations, for a constant target valuation.

Theorem 7.1. *The Exact Equitable Allocation problem is NP-complete by a reduction from 3-Coloring.*

Proof. Let $G = (V, E)$ be an instance of 3-coloring, where $V = \{v_1, \ldots, v_p\}$ and $E = \{e_1, \ldots, e_q\}$. Recall that G is a YES-instance of 3-coloring if and only if V can be partitioned into three parts so that every edge has at most one endpoint in each part. We assume WLOG that G is connected. We describe now an instance of Exact Equitable Allocation based on G as follows. We have $m := p + 2q + 3$ items overall, one corresponding to each vertex and two corresponding to each edge, and three "special" items. The vertex items are denoted g_i^v for $1 \leqslant i \leqslant n$ and are called "vertex items". We use g_i^e for $1 \leqslant i \leqslant n$ to denote one set of items corresponding to edges, and these are called "ID" items, and we use g_i^d for $1 \leqslant i \leqslant n$ to denote the other set of items corresponding to edges, and these are called "dummy" items. Finally, the special items are denoted $\{\alpha, \beta, \gamma\}$. We also have $n := q + 3$ agents, denoted $\{a_1, \ldots, a_q\}$, one for each edge, and $\{b_1, b_2, b_3\}$, who we refer to as the *forcing* agents. The agent corresponding to the edge $e_j = (v_a, v_b)$ likes the vertex items g_a^v and g_b^v, and the ID item g_j^e corresponding to the edge e_j, and all the dummy goods. The three forcing agents approve all the dummy goods and the special items. Overall, now we want a partition of $p + 2q$ items into $q + 3$ parts so that each agent values each part at 1. We now argue the equivalence of these instances.

Forward direction.

Let $A, B,$ and C be a partition of V into three independent sets.. We can then construct an allocation as follows. For the set A, create a bundle that contains all the vertex items g_i^v corresponding to vertices in A, along with the ID items g_i^e corresponding to edges whose endpoints lie in B and C, and the special item α. Let this bundle be denoted by $\mathcal{B}_A$. Define the bundles $\mathcal{B}_B$ and $\mathcal{B}_C$ analogously (with the special items β and γ respectively). Further, assign every dummy item g_j^d to its own individual bundle, denoted $\mathcal{B}_j$. With this allocation, it is easy to verify that every agent values all parts at 1.

Reverse direction.

Suppose we have an exact equitable allocation for the instance as described above, then every agent values each of the n bundles in the allocation at exactly 1. Indeed, the common value cannot be less than one since there are items that are valued non-trivially, and it cannot be more than 1 since all agents value the set of all items at n. Observe that all dummy goods must be assigned to separate singleton parts, since the dummy items are universally valued and all other items are valued by at least one agent — in particular, every vertex good is valued by at least one agent corresponding to an edge because of our assumption that G is connected. The remaining vertex and ID items must then be split up into three sets in such a way that no edge agent assigns a value of more than 1 to each set. Given these constraints, the only possible way to achieve this is if the vertex items in each of these remaining sets form an independent set in G. Thus, any exact equitable allocation of our transformed instance implies a valid 3-coloring of the original graph G. □

Theorem 7.2. *The Approximate Exact Equitable Allocation problem is NP-complete when $p = 0, q = 1, c = 0,$ and $k = 3$ by a reduction from 3-Coloring.*

Proof Sketch. Let $G = (V, E)$ be an instance of 3-coloring, where $V = \{v_1, \ldots, v_p\}$ and $E = \{e_1, \ldots, e_q\}$. We describe now an instance of Approximate Exact Equitable Allocation based on G as follows. We have p items, one corresponding to each vertex. The vertex items are denoted g_i^v for $1 \leqslant i \leqslant n$. We also have q agents, denoted $\{a_1, \ldots, a_q\}$, one for each edge. The agent corresponding to the edge $e_j = (v_a, v_b)$ approves exactly two items: g_a^v and g_b^v. Set $p = 0$, $q = 1$, $c = 0$, and $k = 3$. The proof of equivalence is self-evident. This result holds even when each agent values two items and each item is valued by at most four agents since the 3-coloring problem is hard even for graphs where the maximum degree is four (Garey et al., 1976). □

Theorem 7.3. *Assuming ETH, Approximate Exact Equitable Allocation cannot be solved in time $2^{o(m \log m)} \cdot n^{O(1)}$, even when $p = 2$, $q = 3$, and $c = 1$ by a reduction from $(\ell \times \ell)$-Permutation Clique.*

Proof. The $(\ell \times \ell)$-Permutation Clique problem is the following. The input is an integer ℓ and a graph G with a vertex set $[\ell] \times [\ell]$, and the question is to determine if there exists a set $X \subseteq V(G)$ that is a clique in G and that induces a permutation of $[\ell]$. It is known that unless ETH fails, $(\ell \times \ell)$-Permutation Clique cannot be solved in time $\mathcal{O}^*\left(2^{o(\ell \log \ell)}\right)$ Lokshtanov et al. (2018).

We assume, without loss of generality, that there are no edges between any pair of vertices that have a common value for the first coordinate (i.e, vertices of the form $(i, \star)$ for any $1 \leqslant i \leqslant k$); and no edges between any pair of vertices that have a common value for the second coordinate (i.e, vertices of the form $(\star, j)$ for any $1 \leqslant j \leqslant k$). We show a polynomial-time reduction from $(\ell \times \ell)$-Permutation Clique to Approximate Exact Equitable Allocation where the reduced instance has $p = 2, q = 3, c = 1, k = \ell$. We note that in the reduced instance, we will have that $m = O(\ell)$.

First, introduce 2ℓ items: denote these by $\{p_1, \ldots, p_\ell\} \cup \{q_1, \ldots, q_\ell\}$. We have two agents P and Q. The agent P values the items $\{p_1, \ldots, p_\ell\}$ at 2 and $\{q_1, \ldots, q_\ell\}$ at 0, while Q values the items $\{p_1, \ldots, p_\ell\}$ at 0 and $\{q_1, \ldots, q_\ell\}$ at 2. For every $1 \leqslant i, j, a, b \leqslant k$ such that (i, a) and (j, b) are *not* adjacent, we introduce an agent $C_{i,j,a,b}$ who values the items p_i and q_b at 1.5, q_a and p_j at 2.5, and all other items at 1. This completes the construction: we now show the equivalence of the instances.

Forward Direction.

Let σ be a permutation of k such that $(1, \sigma(1)), \cdots (k, \sigma(k))$ forms a clique in G. Then consider the following k bundles: $\{p_1, q_{\sigma(1)}\}, \cdots, \{p_k, q_{\sigma(k)}\}$. Note that P and Q value each bundle at 2. For an agent $C_{i,j,a,b}$, each bundle has value either 2 (if both items have value 1), 3 (if both items have value 1.5), 4 (if one item has value 1.5 and the other has value 2.5), or 5 (if both items have value 2.5). Note that there cannot be two bundles that have value 4, since the only way for this to happen is if $\{p_i, q_a\}$ and $\{p_j, q_b\}$ are bundles, but this will not be the case since this would imply that the vertices (i, a) and (j, b) chosen for the clique are non-adjacent — a contradiction. Based on this, it is easy to verify that there is at most one bundle whose value exceeds 3 for the agent $C_{i,j,a,b}$, and all bundles have a value of at least two since all bundles have exactly two items.

Reverse Direction

Let $B_1, \ldots, B_k$ be k bundles such that every agent values all but at most one bundle at either 2 or 3. Note that $B_i \cap \{p_1, \ldots, p_\ell\} = 1$ and $B_i \cap \{q_1, \ldots, q_\ell\} = 1,$ for all $1 \leqslant i \leqslant k$. Indeed, if this is not the case, then some bundle has more than one p-item (or q-item) and therefore also some bundle that has no p-items (respectively, q-items), leading to more than one bundle for the agent P (respectively, Q) that they value differently from 2 or 3, which is a contradiction. By renaming the bundles, we can ensure that $p_i \in B_i$ for all $1 \leqslant i \leqslant k$. Denoting by $\sigma(i)$ the index j for which we have $q_j \in B_i$, we claim that $\{(i, \sigma(i)) \mid 1 \leqslant i \leqslant k\}$ forms a clique in G. Indeed, suppose not: let us say that there is no edge between (i, a) and (j, b), where $a := \sigma(i)$ and $b := \sigma(j)$. Then the agent $C_{i,j,a,b}$ values the bundles $\{p_i, q_a\}$ and $\{p_j, q_b\}$ at $1.5 + 2.5 = 4$, leading to two bundles for which the agent's value is neither 2 nor 3, again a contradiction. Because of this, also observe that σ is a permutation since no two vertices that share a coordinate in common have an edge between them. This concludes the proof. ☐

7.4 Secure Allocations

7.4.1 Hardness Results for Secure Allocations

Theorem 7.4. *The Weakly Secure Capacitated Allocation problem is NP-complete by a reduction from Equitable 3-Coloring.*

Proof. Let $\mathcal{I} := (G = (V, E))$ be an instance of Equitable 3-Coloring, where G is a connected graph on $n = 3k$ vertices and m edges. $\mathcal{I}$ is a yes instance only if G can be partitioned into 3 independent sets of size k each. We construct the reduced instance $\mathcal{I}' := (A, O, \mathbf{v}; k + 1)$ of the allocation problem as follows:

We introduce an *edge agent* a_e for every edge $e \in E$ and three special agents x, y, and z. We introduce a *vertex good* g_v for every vertex $v \in V$. We also introduce $(k + 1)m$ copies of an ordinary $\star$ good and three special goods, namely, $\star_x, \star_y$ and $\star_z$.

An edge agent a_e corresponding to the edge $e = (uv)$ likes all the special star goods and all the vertex goods except those corresponding to its points. That is, a_e likes

	$\star$	…	$\star$	$\star_x$	$\star_y$	$\star_z$	g_1	g_2	g_3	g_4	g_5	g_6
a_{12}	0	…	0	1	1	1	0	0	1	1	1	1
a_{23}	0	…	0	1	1	1	1	0	0	1	1	1
a_{34}	0	…	0	1	1	1	1	1	0	0	1	1
a_{45}	0	…	0	1	1	1	1	1	1	0	0	1
a_{56}	0	…	0	1	1	1	1	1	1	1	0	0
a_{61}	0	…	0	1	1	1	0	1	1	1	1	0
x	1	…	1	$\boxed{0}$	1	1	$\boxed{1}$	1	$\boxed{1}$	1	1	1
y	1	…	1	1	$\boxed{0}$	1	1	$\boxed{1}$	1	1	$\boxed{1}$	1
z	1	…	1	1	1	$\boxed{0}$	1	1	1	$\boxed{1}$	1	$\boxed{1}$

Table 7.2: Reduced Allocation Instance from Equitable 3-Coloring (Theorem 7.4), where G in the original instance is a cycle on 6 vertices with the edge set $\{(12),(23),(34),(45),(56),(61)\}$

$\{\star_x, \star_y, \star_z, g_1, \ldots g_n\} \setminus \{g_u, g_v\}$. The special agents x, y and z like all the goods except $\star_x$, $\star_y$ and $\star_z$ respectively.

Note that there are $m+3$ agents and $(k+1)m + 3 + 3k = (m+3)(k+1)$ goods. This completes the construction. Table 7.2 shows an example of the above construction when the graph G in the original instance is a cycle in 6 vertices. We now argue the equivalence.

The forward direction.

Suppose G has a 3-coloring given by vertex partitions X, Y and Z, such that $|X| = |Y| = |Z| = k$. Then consider the allocation Φ under which an edge agent a_e gets $k+1$ ordinary $\star$ goods, x gets $\star_x \cup X$, y gets $\star_y \cup Y$ and z gets $\star_z \cup Z$. Notice that every agent gets at most $k+1$ goods. We claim that every agent is strongly (and hence, weakly) secure with respect to the allocation Φ. Indeed, note that an edge agent values his bundle at 0, and hence is secure trivially. The special agents x, y, and z are secure with respect to each other, as they value each others' bundle at the maximum value $k+1$. Further, an edge agent a value all but at most one good from any of the special agent's bundle. This is true as the special agents get an independent set, any edge agent dislikes only her endpoints and at most one of the endpoints of an edge can appear in any independent set. Therefore $u_a(\Phi_x) \geq k = u_x(\Phi_x)$, and so, all the special agents are secure with respect to the edge agents.

The reverse direction.

Suppose there exists weakly secure capacitated allocation Φ for the reduced instance $\mathcal{I}'$.

Claim 7.5. *Under the allocation Φ, the special agents x, y and z must get the special goods $\star_x$, $\star_y$ and $\star_z$ respectively.*

Proof. We prove the above claim for the agent x. It follows for y and z analogously. Notice that by pigeonholing and cardinality constraint, every agent receives exactly $(k+1)$ goods. Suppose x does not get $\star_x$ under Φ, which is the only good she dislikes. Then she derives the utility of $k+1$ from her bundle Φ_x. Now for x to be weakly secure, everyone else (except for the one who gets $\star_x$) must value Φ_x at no less than $k+1$. Observe that x can not get any ordinary $\star$ good. Otherwise, as $\star$ goods are not liked by any edge agent, therefore, x is made insecure by the edge agent a who did not receive $\star_x$. Indeed, x cares about $a's$ valuation of her bundle, but $u_a(\Phi_x) < k+1$. So, $\Phi_x \subseteq \{\star_y, \star_z, g_{v:v \in V(G)}\}$. Consider the following cases.

- x gets at least 2 vertex items. Choose a good $g_v \in \Phi_x$ such that the edge agent a_v corresponding to v did not receive $\star_x$. Such an agent exists as the graph is connected. Note that x cares for a_v but is made insecure by her as $u_{a_v}(\Phi_x) < k+1$.
- x gets at most one vertex good g_v. If x gets g_v and there is a corresponding edge agent a_v who did not get $\star_x$, then x is insecure and we are done. Suppose the only edge agent a_v corresponding to g_v gets $\star_x$. Then x does not care for a_v and hence is secure at least with respect to a_v. But, as everyone gets exactly $k+1 \geq 2$ goods, so x must receive at least one of the $\star_y$ or $\star_z$, say $\star_y$. Then, consider the agent y who now values Φ_y at $k+1$. Observe that y must get at least one vertex item, say g_u again because of the cardinality constraint. Then y cares for the edge agent a_u incident on the vertex u but is made insecure by a_u as $u_{a_u}(\Phi_y) < k+1$. We get a similar contradiction when x gets none of the vertex goods.

This establishes the claim. □

Now notice that since x gets $\star_x$, she values her bundle at exactly k and everyone else's bundle at $k+1$, hence cares for everyone else's opinion. Notice that the remaining k goods in $x's$ bundle can not be all $\star$ goods, as then edge agents, who do not like them, will make x insecure. So, x must get some vertex goods. Also, note that x can not get any of the $\star$ goods. Else, consider the vertex good g_v from $x's$ bundle and the corresponding edge agent a_v. Then $u_{a_v}(\Phi_x) \leq k-1$ (as a_v dislikes $\star$ good, and g_v). Therefore, x must get k vertex goods, say $X \subseteq V$. Moreover, X must form an independent set in G. If not, then say $e \in E(X)$. Consider the agent a_e who now dislikes 2 goods from $x's$ bundle, corresponding to her two end-points and hence values Φ_x at at most $k-1$, again making x insecure. A similar argument shows that y and z must also get $\star_y \cup Y$ and $\star_z \cup Z$ respectively, where Y and Z are vertex goods corresponding to the k-sized independent sets in G. This shows the existence of 3 independent sets of equal size in G and hence, concludes the argument in the reverse direction. □

Corollary 7.6. *The Strongly Secure Capacitated Allocation problem is NP-complete.*

Proof. Notice that in the proof of Theorem 7.4, the allocation Φ in the forward direction is a strongly secure allocation, and the reverse direction relies on the strictly weaker notions of security. Therefore, the equivalence of the reduction from Equitable 3-Coloring also holds for Strongly Secure Capacitated Allocation. □

For our next result, we observe that deciding whether there exists a vertex cover of size $\frac{n}{2}$ in a graph on n vertices is NP-complete. This can be seen by a reduction from vertex cover. Consider an instance $\mathcal{I} := (G, k)$ of vertex cover. Suppose G has n vertices. Then we construct a graph G' from G by adding a clique of size $(n-k+1)$ and an independent set of size $(k-1)$ to G. The number of vertices in G' is $n + (n-k+1) + (k-1) = 2n$. We claim that (G', n) is a yes-instance if and only if (G, k) is a yes-instance. In the forward direction, if there is a k-sized vertex cover S in G then, S along with the $(n-k)$ vertices from the clique form a vertex cover of size n in G'. In the reverse direction, any vertex cover S' of G' must have $(n-k)$ vertices from the clique. Then, $S' \cap V(G)$ is a vertex cover of size k in G.

Theorem 7.7. *The Weakly Secure Capacitated Allocation (Egalitarian) problem is NP-complete by a reduction from Vertex Cover.*

Proof. Let $\mathcal{I} := (G = (V, E); k)$ be an instance of Vertex Cover where G is a simple graph on n vertices and m edges and $k = \frac{n}{2}$. We construct the reduced instance $\mathcal{I}' := (A, O, \mathbf{v}; k+1, 1)$ of the allocation problem as follows:

- We introduce an *edge agent* a_e for every edge $e \in E$ and two *special agents* s and t. We introduce an *edge good* g_e for every edge $e \in E$ and a *vertex good* g_v for every vertex $v \in V$. We also introduce two *special goods* g_s and g_t.
- An edge agent a_e corresponding to the edge $e = (uv)$ likes all the edge goods and $\{g_t, g_u, g_v\}$. The special agents s and t like all the edge goods and $\{g_s, g_t\}$.

Note that there are $m+2$ agents and $m+n+2$ goods. All the edge goods and g_t are universal goods, in the sense that they are liked by all the agents. This completes the

	g_s	g_t	g_{12}	g_{23}	g_{34}	g_{14}	g_1	g_2	g_3	g_4
a_{12}	0	1	$\boxed{1}$	1	1	1	1	1	0	0
a_{23}	0	1	1	$\boxed{1}$	1	1	0	1	1	0
a_{34}	0	1	1	1	$\boxed{1}$	1	0	0	1	1
a_{41}	0	1	1	1	1	$\boxed{1}$	1	0	0	1
s	$\boxed{1}$	1	1	1	1	1	$\boxed{0}$	0	$\boxed{0}$	0
t	1	$\boxed{1}$	1	1	1	1	0	$\boxed{0}$	0	$\boxed{0}$

Table 7.3: Reduced Allocation Instance from Vertex Cover (Theorem 7.7), where G in the original instance is a cycle on 4 vertices with the edge set $\{(12), (23), (34), (41)\}$.

construction. Table 7.3 shows an example of the construction when the graph G in the original instance is a cycle on 4 vertices. We now argue the equivalence.

The forward direction.

Suppose $S \subseteq V$ be a vertex cover of size $\frac{n}{2}$. Then consider the allocation Φ under which the edge agent a_e gets corresponding edge good g_e, s gets $\{g_s, g_{u:u \in S}\}$ and t gets $\{g_t, g_{u:u \notin S}\}$. Notice that the above allocation Φ gives the utility of 1 to every agent, and the number of goods in any bundle is at most $\frac{n}{2} + 1$. Also, Φ is strongly (and hence, weakly) secure. Indeed,

- s is secure with respect to any edge agent. As s gets the goods corresponding to vertex cover S, every edge agent values $s's$ bundle at at least 1, which is at least $u_s(\Phi_s)$. s and t are secure with respect to each other, as they value each others bundle identically.
- Any edge agent a_e and t are secure with respect to any other agent as they get a good that is liked by everyone.

The reverse direction.

Let Φ be the capacitated allocation which is weakly secure and gives the utility of at least 1 to every agent. Since the agents s and t like the goods $\{g_s, g_t, g_{e:e \in E}\}$, they must get one of them under Φ in order to derive the utility of at least 1. Suppose

- Φ allocates g_s and g_t to s and t respectively. Then, any edge agent a_e *cares* for t and therefore must get a good liked by t. This implies that every edge agent a must get some edge good g. This forces s to *care* for all the edge agents and hence get a vertex cover of size at most $\frac{n}{2}$ in order to be secure. The case when Φ allocates g_s to t and g_t to s is analogous.

- Φ allocates some edge goods g_e and $g_{e'}$ to s and t respectively. Then, if there is any edge agent who derives utility from vertex goods only, then she is weakly insecure with respect to s and t, so every edge agent must get exactly one good from $\{g_s, g_t, g_{e:e\in E}\} \setminus \{g_e, g_{e'}\}$. Consider the edge agent a who gets g_s. Since a does not like g_s, she must get a vertex cover of size at most $\frac{n}{2}$ in order to derive the minimum utility of 1 and be secure against every other edge agent. (Note that a *cares* for every other edge agent). The case when Φ allocates g_t to s and some edge good g_e to t is analogous.
- Φ allocates g_s to s and some edge good g_e to t. By a similar argument as above, every edge agent must get exactly one good from $\{g_t, g_{e:e\in E}\} \setminus \{g_e\}$. Then, for s to be secure with respect to all the edge agents, she must get a vertex cover of size at most $\frac{n}{2}$. The case when Φ allocates g_s to t and some edge good g_e to s is analogous.

In all the cases, at least one agent must get a vertex cover S of size at most $\frac{n}{2}$ in order to be either secure with respect to other agents or to derive the minimum utility of 1. Since any super-set of S is also a vertex cover, therefore $\mathcal{I}$ is a yes-instance and this completes the argument in the reverse direction. □

Corollary 7.8. *The* Strongly Secure Capacitated Allocation (Egalitarian) *problem is NP-complete.*

Proof. Notice that in the proof of Theorem 7.7, the allocation Φ in the forward direction is a strongly secure allocation, and the reverse direction relies on the strictly weaker notions of security. Therefore, the equivalence of the reduction from Vertex Cover also holds for Strongly Secure Capacitated Allocation (Egalitarian). □

Corollary 7.9. *The* Strongly Secure Capacitated Allocation (Utilitarian) *and* Weakly Secure Capacitated Allocation (Utilitarian) *problems are NP-complete.*

Proof. Notice that in the proof of Theorem 7.4, the allocation Φ constructed in the reduced instance has a total utility of $3k$. Therefore, setting $l = 3k$ in both the Strongly and Weakly Secure Capacitated Allocation (Utilitarian) problems establishes the hardness. □

7.4.2 Algorithms for Secure Allocations

We say that a good is universal if it is liked by all the agents. A good is said to be wastefully allocated if it is allocated to an agent who does not derive any utility from that good.

Theorem 7.10. *For left extremal instances, the Strongly Secure Capacitated Allocation problem admits a polynomial time algorithm.*

Proof. Consider a left-extremal instance $\mathcal{I} := (A, O, \mathbf{v}; k)$ of Strongly Secure Capacitated Allocation problem. The right-extremal setting is analogous. We first arrange the agents $\{a_1, a_2, \ldots a_n\}$ in the increasing order of the number of goods liked by them (that is, in the increasing order of the length of their intervals). Note that a_i values all the goods valued by $a_1, \ldots a_{i-1}$ and all the goods valued by a_1 are universal goods. We say that the goods liked by the agent a_i form the interval I_i. We now make the following claim.

Claim 7.11. *For left-extremal instances, under any secure allocation, all the non-universal goods must be allocated wastefully.*

Proof. Suppose a non-universal good g is allocated non-wastefully to an agent a. By assumption, since g is not a universal good, therefore, $a \neq a_1$. Also, a_1 does not like g and so, can be a potential cause of insecurity to a. To compensate, the agent a must be allocated a good that is not liked by him but valued by a_1. But due to the left extremal structure of the valuations, there is no such good – all the goods liked by a_1 are liked by a. Therefore, g must be allocated wastefully. □

The algorithm works as follows. Consider a bipartite graph $G = (A, O, E)$ where the left and right bi-partitions consist of goods and agents respectively. There is an edge (with capacity 1) between an agent a and a good g is a does not like g. Add a source vertex s adjacent to all the goods with a capacity of 1. Add a sink node t adjacent to all the agents with a capacity of k. If G has a flow of value at least w, where w is the number of non-universal goods in the instance, then the algorithm returns that $\mathcal{I}$ is a yes instance with the following allocation Φ. It allocates all the non-universal goods corresponding to the flow edges. Now for the remaining universal goods, they are allocated to agents, respecting the capacity constraints. That is, they are allocated to agents whose bundle size is less than k. If there is no such agent then there were more than nk goods to begin with, in which case, no allocation can respect the cardinality constraint. Otherwise, at the end, all the goods are allocated and every agent gets at most k

of them. If there is no flow of value at least w in G, then the algorithm returns that $\mathcal{I}$ is a no instance.

To see the correctness of the above algorithm, we will show that it returns yes is and only if $\mathcal{I}$ is a yes instance. Indeed if it returns yes, then it returns Φ, and G must have had the flow value of at least w. Since there are only w non-wasteful goods, this implies that all the non-universal goods are allocated wastefully under Φ along the flow edges, which does not cause any agent to be insecure. Also note that every agent derives utility only from the universal goods, and therefore any other agent also value their bundle at exactly the number of universal goods in the said bundle. Therefore, Φ is the required capacitated secure allocation, and hence, $\mathcal{I}$ is a yes instance.

On the other hand, suppose $\mathcal{I}$ is a yes-instance, with a witness allocation Φ' Then by Claim 7.11, all non-universal goods must be allocated wastefully under Φ'. This implies at least w goods are allocated wastefully and therefore there is a flow corresponding to these goods has a value of at least w. Then, the algorithm would have constructed the allocation Φ according to the flow edges and returned yes based on the number of remaining goods and the capacities. If it returned no, then the number of goods to begin with must have been greater than nk, which would contradict the existence of Φ'. This settles the claim. □

Theorem 7.12. *For extremal instances, the Strongly Secure Capacitated Allocation problem admits a polynomial time algorithm.*

Proof. Consider an instance $\mathcal{I} := (A, O, \mathbf{v}, k)$ of Strongly Secure Capacitated Allocation problem. Let $A_L := (l_1, l_2, \ldots l_p)$ be the agents who prefer the left extremal goods and $A_R := (r_1, r_2, \ldots r_q)$ be the ones who prefer the right extremal goods, arranged in the increasing order of their respective interval lengths. Let U_L and U_R be the set of items liked by everyone in A_L and A_R respectively. We first make the following claim.

Claim 7.13. *Under any secure allocation on extremal instances, any good g such that $g \notin U_L \cup U_R$ must be allocated wastefully.*

Proof. Consider a good $g \notin U_L \cup U_R$. Then agents l_1 and r_1 do not value g. Suppose under some secure allocation Φ, g is allocated non-wastefully to a left agent l_i. Note that $l_i \neq l_1$ as $g \notin U_L$ and l_1 only values goods in U_L. Since l_1 does not value g, it can be a potential cause of insecurity to l_i. To compensate, l_i must receive a good valued by l_1 but not valued by herself. But there is no such good due to the interval structure – all the goods liked by a_1 are liked by a. Similarly, if g is allocated non-wastefully to a right agent r_i, then r_1, who does not value g, will be a potential cause of insecurity to the agent r_i. Therefore, g must be allocated wastefully under any secure allocation. $\square$

Claim 7.14. *Under any secure allocation, if any left (right) agent receives t goods from U_L (U_R), she must also receive at least t goods from U_R (U_L).*

Proof. Suppose not. Consider an agent $a \in A_L$ who gets t goods from U_L, but at most $t-1$ goods from U_R under a secure allocation Φ. Then $v_{r_1}(\Phi(a)) \leq t-1 < t \leq v_a(\Phi(a))$. This implies that a is insecure with respect to r_1. $\square$

The algorithm works as follows. Suppose without loss of generality $|U_L| < |U_R|$. Let $P := \{(g, g') : g \in U_L \text{ and } g' \in U_R\}$. Then, P has exactly $|U_L|$ many pairs of goods. It first guesses the number of agents in A_L and A_R who receive the pairs from P. Say, p_1 agents from A_L and p_2 agents from A_R receive these pairs. Then it chooses the last p_1 agents from A_L $(l_p, l_{p-1}, \dots)$ and the last p_2 agents from A_R $(r_q, r_{q-1}, \dots)$ and exhaust the pairs from P by allocating them to the said agents. Call this partial allocation Φ. Now for all the remaining goods, we construct a bipartite graph $G = (A, O', E)$ where the left and right bi-partitions consist of unallocated goods and agents respectively. There is an edge (with capacity 1) between an agent a and a good g if a does not like g. Add a source vertex s adjacent to all the goods with a capacity of 1. Add a sink node t adjacent to all the agents. The capacity of the edge from an agent a to the sink t is $k - 2P_a$, where P_a is the number of pairs that a gets from P. If G admits a flow of value at least $m - 2|U_L|$, then the algorithm outputs Φ by allocating the remaining $m - 2|U_L|$ goods wastefully according to the flow edges. Else, if G does not admit a flow of value at least $m - 2|U_L|$, the algorithm returns that $\mathcal{I}$ is a no instance.

To see the correctness of the above algorithm, we will show that it returns yes if and only if $\mathcal{I}$ is a yes instance. Indeed, if it returns yes, it returns the above-described allocation Φ. It is easy to see that Φ is a capacitated allocation by construction. Also, every agent derives utility only from U_L (or U_R) and not both. Say agent a derives utility from U_L. Then she is secure with respect to all the left agents as U_L is valued by all of them. Also, by Claim 7.14, a must also receive goods from U_R, which would imply that all the right agents also value her bundle sufficiently. Hence, a is secure with respect to everyone. This concludes that Φ is a capacitated secure allocation.

On the other hand, suppose $\mathcal{I}$ is a yes instance. Let Φ' be the capacitated secure allocation. Then by Claim 7.13, under Φ', all the goods $g \notin U_L \cup U_R$ must be allocated wastefully. Suppose p'_1 many agents from A_L received goods from P and p'_2 from A_R received these pairs under Φ'. We modify Φ' to an allocation Φ, wherein, the pairs from P are allocated only to the last contiguous chunk of agents, that is, to $\{l_p, l_{p-1} \ldots\}$ from A_L and to $\{r_q, r_{q-1}, \ldots\}$ from A_R. We will show that Φ is also capacitated secure allocation. Suppose a pair from P, say (g, g'), is allocated to l_i. Suppose l_j such that $j > i$, does not get any such pair. Then l_j must have received wasteful goods that she does not value. By extremal structure, l_i also does not value any of the $\Phi'(l_i)$. We can then swap a pair of goods from $\Phi'(l_j)$ and (g, g'). It is easy to see that this does not cause either l_i or l_j to be insecure. Indeed, l_j receives an equal number of goods from U_L and U_R, so she is secure with respect to both left and right agents. Also, l_i receives a pair of goods she does not value, and so is secure. We repeat the swaps for every such pair of agents and call the final allocation Φ. Now when the algorithm guesses $p_1 = p_1'$, and $p_2 = p_2'$, it outputs the partial allocation that overlaps with Φ when restricted to the pairs in P. Now the existence of Φ itself guarantees that the remaining goods can be allocated wastefully and hence there must be a flow of value at least $m - 2|U_L|$ in G. This implies that the algorithm finally returns yes and this settles our claim. □

Since every strongly secure allocation is also weakly secure, therefore, from Theorem 7.12, we immediately have the following corollary.

Corollary 7.15. *For extremal instances, the* Weakly Secure Capacitated Allocation *problem admits a polynomial time algorithm.*

7.5 Abundant Allocations

7.5.1 Hardness Results for Abundant Allocations

Theorem 7.16. *The* Strongly Abundant Capacitated Allocation *problem is NP-complete by a reduction from Vertex Cover.*

Proof. Let $\mathcal{I} := (G = (V, E); k)$ be an instance of Vertex Cover, where G is a simple graph on n vertices and m edges and $k = \frac{n}{2}$. We construct the reduced instance $\mathcal{I}' := (A, O, \mathbf{v}; k + 1)$ of the allocation problem as follows:

- We introduce an *edge agent* a_e for every edge $e \in E$ and two *special agents* s and t. We introduce a *vertex good* g_v for every vertex $v \in V$. We also introduce $m(k + 1) + 1$ many copies of a special $\star$ good.
- An edge agent a_e corresponding to the edge $e = (uv)$ likes all the goods except those

corresponding to her endpoints. That is, a_e likes $\{\star, g_1, g_2, \dots g_n\} \setminus \{g_u, g_v\}$. The special agents s and t like all vertex goods.

Note that there are $m+2$ agents and $m(k+1)+1+n$ goods. This completes the construction. Table 7.4 shows an example of the construction when the graph G in the original instance is a cycle on 4 vertices. We now argue the equivalence of the reduction.

	$\star$	$\dots$	$\star$	g_1	g_2	g_3	g_4
a_{12}	1	$\dots$	1	0	0	1	1
a_{23}	1	$\dots$	1	1	0	0	1
a_{34}	1	$\dots$	1	1	1	0	0
a_{41}	1	$\dots$	1	0	1	1	0
s	$\boxed{0}$	$\dots$	0	$\boxed{1}$	1	$\boxed{1}$	1
t	0	$\dots$	0	1	$\boxed{1}$	1	$\boxed{1}$

Table 7.4: Reduced Allocation Instance from Vertex Cover (Theorem 7.16), where G in the original instance is a cycle on 4 vertices with the edge set $\{(12), (23), (34), (41)\}$.

The forward direction.

Suppose there exist a vertex cover S of size k in the original instance $\mathcal{I}$. Then, consider the allocation Φ that allocates $(k+1)$ copies of the $\star$ good to every edge agent, $\star \cup S$ to the agent s and $V \setminus S$ to t. Note that as $k = \frac{n}{2}$, every agent gets at most $k+1$ goods under Φ. Also, Φ is a strongly abundant allocation, as described below:

- s and t value their bundle at $|S| = k$. They both are abundant with respect to each other, as they both value the same goods. Also, any edge agent a, dislikes at least 1 but at most 2 goods (corresponding to her endpoints) from S and hence value s's bundle at either $(k-2)+1$ or $(k-1)+1$. Therefore $u_a(\Phi_s) \leq k \leq u_s(\Phi_s)$, making s abundant with respect to the edge agents. The argument for abundance of t with respect to a_e is similar.
- Any edge agent values her bundle at the maximum possible value, which is the bundle size, $k+1$, and hence is always abundant with respect to any agent.

The reverse direction.

Suppose there exist a strongly abundant capacitated allocation Φ for the reduced instance $\mathcal{I}'$. Notice that under Φ, because of the cardinality constraint, either s or t must get at least 1

$\star$ good, which she does not like. Say, s gets $\star$. Since every edge agent a_e likes the $\star$ good, therefore, for s to be abundant with respect to all of them, she must have at least 1 good in her bundle disliked by all the edge agents. This forces s to get a vertex cover S of size at most k. Indeed, if there is an edge agent a who is not covered by the vertices corresponding to what s gets, then a likes the entire z_s. That is, $u_a(\Phi_s) = k+1 > k = u_s(\Phi_s)$, thereby making s modest. Since any superset of S is also a vertex cover, therefore, $\mathcal{I}$ is a yes-instance. ☐

Corollary 7.17. *Strongly Abundant Capacitated Allocation (Egalitarian) and Strongly Abundant Capacitated Allocation (Utilitarian) are NP-complete.*

Proof. Notice that in the proof of Theorem 7.16, the allocation Φ constructed in the forward direction, allocates a utility of at least k to every agent. Therefore, setting $l = k$ and $l = 3k+1$ establishes the hardness of the Egalitarian and Utilitarian versions of the problem respectively. ☐

7.5.2 Algorithms for Abundant Allocations

Theorem 7.18. *For left extremal instances, the Strongly Abundant Capacitated Allocation problem admits a polynomial time algorithm.*

Proof. Consider a left-extremal instance $\mathcal{I} := (A, O, \mathbf{v}; k)$ of Strongly Abundant Capacitated Allocation problem. We first arrange the agents $\{a_1, a_2, \ldots a_n\}$ in the increasing order of the number of goods liked by them (that is, in the increasing order of the length of their intervals). Note that a_i values all the goods valued by $a_1, \ldots a_{i-1}$ and all the goods valued by a_1 are universal goods. We first make the following claim.

Claim 7.19. *For left-extremal instances, under any abundant allocation, no good can be allocated wastefully.*

Proof. Let Φ be an abundant allocation under which g is allocated wastefully to an agent a_i. Note that a_i can not be the last agent in the ordering, as every good is liked by at least one agent. Then, the last agent a_n who values g and the entire bundle of a_i as well (by the extremal structure), violates the abundance of a_i. Hence, no good can be wastefully allocated under any abundant allocation. ☐

The algorithm works as follows. Consider a bipartite graph $G = (A, O, E)$ where the left and right bi-partitions consist of goods and agents respectively. There is an edge (with capacity 1)

between an agent a and a good g if a likes g. Add a source vertex s adjacent to all the goods with a capacity of 1. Add a sink node t adjacent to all the agents with a capacity of k. If G admits a flow of value at least m, where m is the number of goods, then the algorithm returns yes and the allocation Φ that allocates the goods according to the flow edges. Else, if G does not admit a flow of value at least m, then the algorithm returns that $\mathcal{I}$ is a no instance.
To see the correctness of the algorithm, we will show that it returns yes if and only if $\mathcal{I}$ is a yes instance. Suppose it returns yes and the allocation Φ. Note that Φ is an abundant allocation, as it is non-wasteful and every agent values their bundle at her size which is the maximum possible utility and hence they are abundant with respect to any other agent. It is easy to see that Φ is capacitated by construction, hence $\mathcal{I}$ is a yes instance. On the other hand, suppose $\mathcal{I}$ is a yes instance. Then there is a capacitated abundant allocation Φ'. By Claim 7.19, all goods must have been allocated non-wastefully under Φ'. Therefore, the flow in G corresponding to the allocation Φ' has a value of at least m. Hence the algorithm returns yes. □

Since strongly abundant allocations are also weakly abundant, so, Theorem 7.18 gives us the following corollary.

Corollary 7.20. *For left-extremal instances, the Weakly Abundant Capacitated Allocation problem admits a polynomial time algorithm.*

7.6 Experiments

We randomly generated 100 instances of fair division with binary valuations and conducted 10 iterations to explore the presence of exact equitable, approximate exact equitable, capacitated strongly secure/abundant allocations. The instances generated had a small number of agents $(2-4)$ and items $(5-8)$. In every iteration, less than 5% of the instances had exact equitable allocations, while approximately 55% − 65% exhibited approximate exact equitable allocations. Notably, capacitated secure and abundant allocations were present in over 80% of the instances.

7.7 Concluding Remarks

We introduced the notions of exact equitable allocations, secure allocations, and abundant allocations. After establishing the hardness of finding exact equitable allocations and approximate variations, we further studied the problems of finding secure and abundant allocations in conjunction with cardinality constraints and utility goals from a computational perspective. The additional constraints are well-motivated, since it is easy to come up with "lazy" allocations that meets the said criteria but are not interesting. The goals we present here may be at loggerheads with the traditional fairness goals. For instance, consider an example with two goods g_a and g_b and two agents a and b, where a approves only g_a and b approves only g_b. It is easy to verify that all secure allocations have an envious agent and the only envy-free allocation also happens to be insecure! This motivates the question of determining the "price of" one of these goals relative to another, after fixing appropriate quantifications of these criteria.

We leave open the question of the complexity of finding strongly and weakly secure or abundant allocations without cardinality constraints in the presence of utility targets, either egalitarian or utilitarian. We were also inconclusive about the problem of finding weakly abundant allocations with cardinality constraints. It seems natural to use a round-robin approach to allocate goods in such a way that achieves weak abundance, but it turns out that most variations of a round-robin theme fail to produce weakly abundant allocations. Extending the realm of tractability from extremal intervals is another open direction in this

Bibliography

Abdulkadiroğlu, Atila and Sönmez, Tayfun. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999. ISSN 0022-0531. 68

Abraham, David J; Cechlárová, Katarína; Manlove, David F, and Mehlhorn, Kurt. Pareto optimality in house allocation problems. In *International Symposium on Algorithms and Computation*, pages 3–15. Springer, 2004. 63

Ahle, Thomas. How Can We Prove the Harmonic Mean is a Concave Function? Mathematics Stack Exchange, 2020. (Version: 2020-08-06). 147

Aigner-Horev, Elad and Segal-Halevi, Erel. Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, Nov 2021. ISSN 0020-0255. 69

Akrami, Hannaneh and Rathi, Nidhi. Epistemic efx allocations exist for monotone valuations, 2024. 11

Akrami, Hannaneh; Chaudhury, Bhaskar Ray; Garg, Jugal; Mehlhorn, Kurt, and Mehta, Ruta. EFX allocations: Simplifications and improvements. *CoRR*, abs/2205.07638, 2022. 10

Akrami, Hannaneh; Alon, Noga; Chaudhury, Bhaskar Ray; Garg, Jugal; Mehlhorn, Kurt, and Mehta, Ruta. Efx: A simpler approach and an (almost) optimal guarantee via rainbow cycle number. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, page 61, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047. 20

Aleksandrov, Martin and Walsh, Toby. Two algorithms for additive and fair division of mixed manna. In Schmid, Ute; Klügl, Franziska, and Wolter, Diedrich, editors, *KI 2020: Advances in Artificial Intelligence*, pages 3–17, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58285-2. 142, 178

Amanatidis, Georgios; Markakis, Evangelos, and Ntokos, Apostolos. Multiple birds with one stone: Beating 1/2 for efx and gmms via envy cycle elimination. *Theoretical Computer Science*, 841:94–109, 2020. 10

Amanatidis, Georgios; Birmpas, Georgios; Filos-Ratsikas, Aris; Hollender, Alexandros, and Voudouris, Alexandros A. Maximum Nash Welfare and Other Stories about EFX. *Theoretical Computer Science*, 863:69–85, 2021. 142

Anshelevich, Elliot; Dasgupta, Anirban; Kleinberg, Jon; Tardos, Éva; Wexler, Tom, and Roughgarden, Tim. The Price of Stability for Network Design with Fair Cost Allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008. 146

Aumann, Yonatan and Dombb, Yair. The efficiency of fair division with connected pieces. *ACM Trans. Econ. Comput.*, 3(4), jul 2015. ISSN 2167-8375. 20, 180

Aziz, Haris and Rey, Simon. Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, IJCAI'20, 2021. ISBN 9780999241165. 142, 178

Aziz, Haris; Gaspers, Serge; Mackenzie, Simon, and Walsh, Toby. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015. 9, 19

Aziz, Haris; Caragiannis, Ioannis; Igarashi, Ayumi, and Walsh, Toby. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1), apr 2022. ISSN 1387-2532. 178, 180, 182

Aziz, Haris; Freeman, Rupert; Shah, Nisarg, and Vaish, Rohit. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. *Operations Research*, 2023a. 168

Aziz, Haris; Huang, Xin; Mattei, Nicholas, and Segal-Halevi, Erel. Computing welfare-maximizing fair allocations of indivisible goods. *European Journal of Operational Research*, 307(2):773–784, 2023b. ISSN 0377-2217. 14, 202, 212

Babaioff, Moshe; Ezra, Tomer, and Feige, Uriel. Fair and Truthful Mechanisms for Dichotomous Valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, volume 35, pages 5119–5126, 2021. 10, 142, 149

Bade, Sophie. Matching with single-peaked preferences. *Journal of Economic Theory*, 180:81–99, 2019. ISSN 0022-0531. 119

Ballester, Miguel A. and Haeringer, Guillaume. A characterization of the single-peaked domain. *Social Choice and Welfare*, 36(2):305–322, 2011. ISSN 01761714, 1432217X. 119

Barman, Siddharth; Krishnamurthy, Sanath Kumar, and Vaish, Rohit. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, page 557–574, New York, NY, USA, 2018a. Association for Computing Machinery. ISBN 9781450358293. 14, 60, 182, 212

Barman, Siddharth; Krishnamurthy, Sanath Kumar, and Vaish, Rohit. Greedy algorithms for maximizing nash social welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, page 7–13, Richland, SC, 2018b. International Foundation for Autonomous Agents and Multiagent Systems. 71, 213

Barman, Siddharth; Bhaskar, Umang; Krishna, Anand, and Sundaram, Ranjani G. Tight Approximation Algorithms for p-Mean Welfare Under Subadditive Valuations. In *Proceedings of the 28th Annual European Symposium on Algorithms*, volume 173, page 11, 2020a. 144

Barman, Siddharth; Bhaskar, Umang, and Shah, Nisarg. Optimal Bounds on the Price of Fairness for Indivisible Goods. In *Proceedings of the 16th International Conference on Web and Internet Economics*, pages 356–369, 2020b. 9, 14, 66, 69, 141

Barman, Siddharth; Bhaskar, Umang; Pandit, Yeshwant, and Pyne, Soumyajit. Nearly equitable allocations beyond additivity and monotonicity. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'24/IAAI'24/EAAI'24. AAAI Press, 2024a. ISBN 978-1-57735-887-9. 13, 178, 180, 194

Barman, Siddharth; Kar, Debajyoti, and Pathak, Shraddha. Parameterized guarantees for almost envy-free allocations. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, page 151–159, Richland, SC, 2024b. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864. 10

Bei, Xiaohui; Lu, Xinhang; Manurangsi, Pasin, and Suksompong, Warut. The Price of Fairness for Indivisible Goods. *Theory of Computing Systems*, 65(7):1069–1093, 2021. 14, 20, 66, 69, 146

Benabbou, Nawal; Chakraborty, Mithun; Igarashi, Ayumi, and Zick, Yair. Finding Fair and Efficient Allocations for Matroid Rank Valuations. *ACM Transactions on Economics and Computation*, 9(4):1–41, 2021. 141, 142, 149, 150, 171

Bérczi, Kristóf; Bérczi-Kovács, Erika R.; Boros, Endre; Gedefa, Fekadu Tolessa; Kamiyama, Naoyuki; Kavitha, Telikepalli; Kobayashi, Yusuke, and Makino, Kazuhisa. Envy-free relaxations for goods, chores, and mixed items. *CoRR*, abs/2006.04428, 2020. 178, 182

Berendsohn, Benjamin Aram; Boyadzhiyska, Simona, and Kozma, László. Fixed-Point Cycles and Approximate EFX Allocations. In Szeider, Stefan; Ganian, Robert, and Silva, Alexandra, editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-256-3. doi: 10.4230/LIPIcs.MFCS.2022.17. 10

Berger, Ben; Cohen, Avi; Feldman, Michal, and Fiat, Amos. Almost full efx exists for four agents. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):4826–4833, Jun. 2022. 10, 20

Bertsimas, Dimitris; Farias, Vivek F, and Trichakis, Nikolaos. The Price of Fairness. *Operations Research*, 59(1):17–31, 2011. 14, 20, 66, 69, 146

Bertsimas, Dimitris; Ivanhoe, Joshua; Jacquillat, Alexandre; Li, Michael; Previero, Alessandro; Lami, Omar Skali, and Bouardi, Hamza Tazi. Optimizing vaccine allocation to combat the covid-19 pandemic. *MedRxiv*, pages 2020–11, 2020. 2

Beynier, Aurélie; Chevaleyre, Yann; Gourvès, Laurent; Harutyunyan, Ararat; Lesca, Julien; Maudet, Nicolas, and Wilczynski, Anaëlle. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agent Systems*, 33(5):591–627, 2019a. 68

Beynier, Aurélie; Chevaleyre, Yann; Gourvès, Laurent; Harutyunyan, Ararat; Lesca, Julien; Maudet, Nicolas, and Wilczynski, Anaëlle. Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agent Systems*, 33, 09 2019b. 179

Bhaskar, Umang; Sricharan, A. R., and Vaish, Rohit. On approximate envy-freeness for indivisible chores and mixed resources. In *International Workshop and International*

Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 2020. 180

Bhaskar, Umang; Misra, Neeldhara; Sethia, Aditi, and Vaish, Rohit. The price of equity with binary valuations and few agent types. In Deligkas, Argyrios and Filos-Ratsikas, Aris, editors, *Algorithmic Game Theory*, pages 271–289, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-43254-5. 18, 20, 66, 69, 136, 179, 180

Bilò, Vittorio; Caragiannis, Ioannis; Flammini, Michele; Igarashi, Ayumi; Monaco, Gianpiero; Peters, Dominik; Vinci, Cosimo, and Zwicker, William S. Almost envy-free allocations with connected bundles. *Games and Economic Behavior*, 131:197–221, 2022. ISSN 0899-8256. 20

Birkhoff, G. Three observations on linear algebra. *Univ. Nac. Tacuman, Rev. Ser. A*, 5:147–151, 1946. 167

Bismuth, Samuel; Bliznets, Ivan, and Segal-Halevi, Erel. Fair division with bounded sharing: Binary and non-degenerate valuations, 2024. 12, 37

Black, Duncan. On the rationale of group decision-making. *Journal of Political Economy*, 56(1): 23–34, 1948. ISSN 00223808, 1537534X. 119

Bliem, Bernhard; Bredereck, Robert, and Niedermeier, Rolf. Complexity of Efficient and Envy-Free Resource Allocation: Few Agents, Resources, or Utility Levels. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 102–108, 2016. 141

Bliznets, Ivan; Bukov, Anton, and Sagunov, Danil. Fair division with minimal withheld information in social networks. *Theoretical Computer Science*, 991:114446, 2024. ISSN 0304-3975. 13

Bogomolnaia, Anna and Moulin, Hervé. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100(2):295–328, 2001. 168

Bogomolnaia, Anna; Moulin, Hervé; Sandomirskiy, Fedor, and Yanovskaya, Elena. Competitive division of a mixed manna. *Econometrica*, 85(6):1847–1871, 2017. 180, 182

Bouveret, Sylvain; Cechlárová, Katarína; Elkind, Edith; Igarashi, Ayumi, and Peters, Dominik. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, page 135–141. AAAI Press, 2017. ISBN 9780999241103. 20, 141

Brams, Steven J and Fishburn, Peter C. Approval voting. *American Political Science Review*, 72 (3):831–847, 1978. 213

Brams, Steven J and Taylor, Alan D. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996. 12

Brânzei, Simina; Lv, Yuezhou, and Mehta, Ruta. To Give or Not to Give: Fair Division for Single Minded Valuations. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 123–129, 2016. 141

Brustle, Johannes; Dippel, Jack; Narayan, Vishnu V.; Suzuki, Mashbat, and Vetta, Adrian. One dollar each eliminates envy. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 23–39, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379755. 12, 37

Bu, Xiaolin; Song, Jiaxin, and Yu, Ziqi. Efx allocations exist for binary valuations. In Li, Minming; Sun, Xiaoming, and Wu, Xiaowei, editors, *Frontiers of Algorithmics*, pages 252–262, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-39344-0. 10, 20

Budish, Eric. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011. 5, 11, 37, 146, 212

Caragiannis, Ioannis; Kaklamanis, Christos; Kanellopoulos, Panagiotis, and Kyropoulou, Maria. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012. 14, 20, 66, 69, 146, 180

Caragiannis, Ioannis; Gravin, Nick, and Huang, Xin. Envy-freeness up to any item with high nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 527–545, New York, NY, USA, 2019a. Association for Computing Machinery. ISBN 9781450367929. 10, 20

Caragiannis, Ioannis; Kurokawa, David; Moulin, Hervé; Procaccia, Ariel D; Shah, Nisarg, and Wang, Junxing. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation*, 7(3):1–32, 2019b. 5, 9, 14, 37, 140, 182, 212

Caragiannis, Ioannis; Garg, Jugal; Rathi, Nidhi; Sharma, Eklavya, and Varricchio, Giovanna. New fairness concepts for allocating indivisible items. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, IJCAI '23, 2023. ISBN 978-1-956792-03-4. 11

Cechlárová, Katarína and Pillárová, Eva. On the computability of equitable divisions. *Discrete Optimization*, 9(4):249–257, 2012a. ISSN 1572-5286. 180

Cechlárová, Katarína and Pillárová, Eva. A near equitable 2-person cake cutting algorithm. *Optimization*, 61(11):1321–1330, 2012b. 180

Cechlárová, Katarı´na; Doboš, Jozef, and Pillárová, Eva. On the existence of equitable cake divisions. *Information Sciences*, 228:239–245, 2013. ISSN 0020-0255. 180

Chan, Hau; Chen, Jing; Li, Bo, and Wu, Xiaowei. Maximin-aware allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, IJCAI'19, page 137–143. AAAI Press, 2019. ISBN 9780999241141. 10

Chaudhury, Bhaskar Ray; Garg, Jugal, and Mehlhorn, Kurt. Efx exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 1–19, New York, NY, USA, 2020a. Association for Computing Machinery. ISBN 9781450379755. 10, 20, 182

Chaudhury, Bhaskar Ray; Kavitha, Telikepalli; Mehlhorn, Kurt, and Sgouritsa, Alkmini. A little charity guarantees almost envy-freeness. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*, pages 2658–2672. SIAM, 2020b. 10, 37

Chaudhury, Bhaskar Ray; Garg, Jugal; Mehlhorn, Kurt; Mehta, Ruta, and Misra, Pranabendu. Improving efx guarantees through rainbow cycle number. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, EC '21, page 310–311, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450385541. 10

Chevaleyre, Yann; Endriss, Ulle; Estivie, Sylvia, and Maudet, Nicolas. Reaching envy-free states in distributed negotiation settings. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, page 1239–1244, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc. 11

Choo, Davin; Ling, Yan Hao; Suksompong, Warut; Teh, Nicholas, and Zhang, Jian. Envy-free house allocation with minimum subsidy. *Operations Research Letters*, 54:107103, 2024. ISSN 0167-6377. 12, 69

Christodoulou, George; Fiat, Amos; Koutsoupias, Elias, and Sgouritsa, Alkmini. Fair allocation in graphs. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC

'23, page 473–488, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047. 20, 21, 22

Chèze, Guillaume. Existence of a simple and equitable fair division: A short proof. *Mathematical Social Sciences*, 87:92–93, 2017. ISSN 0165-4896. 180

Conitzer, Vincent. Eliciting single-peaked preferences using comparison queries. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8, 2007. 119

Cygan, Marek; Fomin, Fedor V.; Kowalik, Lukasz; Lokshtanov, Daniel; Marx, Dániel; Pilipczuk, Marcin; Pilipczuk, Michal, and Saurabh, Saket. *Parameterized Algorithms*. Springer, 2015. 76, 78, 167

Darmann, Andreas and Schauer, Joachim. Maximizing nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2):548–559, 2015. ISSN 0377-2217. 35, 142

Dayal, Pratyush and Misra, Neeldhara. Deleting to structured trees. In *Proceedings of the 25th International Computing and Combinatorics Conference*, pages 128–139. Springer, 2019. 42

Deligkas, Argyrios; Eiben, Eduard; Ganian, Robert; Hamm, Thekla, and Ordyniak, Sebastian. The parameterized complexity of connected fair division. In *IJCAI*, pages 139–145, 2021a. 20

Deligkas, Argyrios; Fearnley, John; Melissourgos, Themistoklis, and Spirakis, Paul G. Computing exact solutions of consensus halving and the borsuk-ulam theorem. *J. Comput. Syst. Sci.*, 117:75–98, 2021b. 207

Desirable Allocations, Image, Accessed: 30 July 2024. URL `https://dailyevergreen.com/115316/opinion/covid-19-vaccination-\more-important-than-personal-liberty/`. xvii, 3

Downey, Rodney G. and Fellows, Michael R. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. ISBN 978-1-4471-5558-4. 76

Dubins, Lester E and Spanier, Edwin H. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1P1):1–17, 1961. 6, 146, 180, 212

Edward Su, Francis. Rental harmony: Sperner's lemma in fair division. *The American mathematical monthly*, 106(10):930–942, 1999. 1, 12

Eiben, Eduard; Ganian, Robert; Knop, Dusan, and Ordyniak, Sebastian. Solving integer quadratic programming via explicit and structural restrictions. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 1477–1484. AAAI Press, 2019. 113

Eisenberg, Edmund and Gale, David. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959. 7, 146

Elkind, Edith and Lackner, Martin. Structure in dichotomous preferences. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, page 2019–2025. AAAI Press, 2015. ISBN 9781577357384. 16, 67, 72, 211, 213

Elkind, Edith; Faliszewski, Piotr, and Skowron, Piotr. A characterization of the single-peaked single-crossing domain. *Social Choice and Welfare*, 54(1):167–181, Jan 2020. ISSN 1432-217X. 119

Elkind, Edith; Lackner, Martin, and Peters, Dominik. Preference restrictions in computational social choice: A survey, 2022. 119

Faliszewski, Piotr; Hemaspaandra, Edith; Hemaspaandra, Lane A., and Rothe, Jörg. The shield that never was: societies with single-peaked preferences are more open to manipulation and control. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK '09, page 118–127, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585604. 119

Fellows, Michael R.; Hermelin, Danny; Rosamond, Frances, and Vialette, Stéphane. On the parameterized complexity of multiple-interval graph problems. *Theoretical Computer Science*, 410(1):53–61, 2009. ISSN 0304-3975. 26, 77

Filos-Ratsikas, Aris and Goldberg, Paul W. Consensus halving is ppa-complete. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, page 51–64, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355599. 208

Filos-Ratsikas, Aris and Goldberg, Paul W. The complexity of splitting necklaces and bisecting ham sandwiches. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 638–649, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367059. 207, 208, 212

Filos-Ratsikas, Aris; Hollender, Alexandros; Sotiraki, Katerina, and Zampetakis, Manolis. Consensus-halving: Does it ever get easier? In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC '20, page 381–399, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379755. 208

Foley, Duncan. Resource allocation and the public sector. *Yale Economic Essays*, pages 45–98, 1967. 5, 146, 212

Food Bill Division, Image, Accessed: 30 July 2024. URL `http://surl.li/wcpxpa`. xvii, 2

Freedman, Rachel; Borg, Jana Schaich; Sinnott-Armstrong, Walter; Dickerson, John P., and Conitzer, Vincent. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence*, 283:103261, 2020. ISSN 0004-3702. 2

Freeman, Rupert; Sikdar, Sujoy; Vaish, Rohit, and Xia, Lirong. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, IJCAI'19, page 280–286. AAAI Press, 2019. ISBN 9780999241141. xx, 6, 13, 16, 178, 180, 181, 199, 212

Freeman, Rupert; Sikdar, Sujoy; Vaish, Rohit, and Xia, Lirong. Equitable allocations of indivisible chores. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '20, page 384–392, Richland, SC, 2020. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450375184. 13, 16, 178, 180

Gal, Ya'akov (Kobi); Mash, Moshe; Procaccia, Ariel D., and Zick, Yair. Which is the fairest (rent division) of them all? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, page 67–84, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450339360. 1, 13, 178

Gamow, George and Stern, Marvin. Puzzle-Math, 1958. 5, 146, 212

Gan, Jiarui; Suksompong, Warut, and Voudouris, Alexandros A. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019. 64, 69, 73, 103

Garey, Michael R. and Johnson, David S. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA, 1990. ISBN 0716710455. 77

Garey, M.R.; Johnson, D.S., and Stockmeyer, L. Some simplified np-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976. ISSN 0304-3975. 217

Garg, Jugal and Murhekar, Aniket. Computing pareto-optimal and almost envy-free allocations of indivisible goods. *J. Artif. Intell. Res.*, 80:1–25, 2024. URL `https://doi.org/10.1613/jair.1.15414`. 180

Garg, Jugal; Kulkarni, Pooja, and Murhekar, Aniket. On Fair and Efficient Allocations of Indivisible Public Goods. In *Proceedings of the 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 213 of *LIPIcs*, pages 22:1–22:19, 2021. 141

Goldberg, Paul W.; Hollender, Alexandros; Igarashi, Ayumi; Manurangsi, Pasin, and Suksompong, Warut. Consensus halving for sets of items. *Math. Oper. Res.*, 47(4):3357–3379, nov 2022. ISSN 0364-765X. 12, 212

Gourvès, Laurent; Monnot, Jérôme, and Tlilane, Lydia. Near fairness in matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence*, pages 393–398, 2014. 5, 6, 13, 178, 212

Halpern, Daniel and Shah, Nisarg. Fair division with subsidy. In Fotakis, Dimitris and Markakis, Evangelos, editors, *Algorithmic Game Theory*, pages 374–389, Cham, 2019. Springer International Publishing. ISBN 978-3-030-30473-7. 12

Halpern, Daniel; Procaccia, Ariel D; Psomas, Alexandros, and Shah, Nisarg. Fair division with binary valuations: One rule to rule them all. In *Web and Internet Economics: 16th International Conference, WINE 2020, Beijing, China, December 7–11, 2020, Proceedings 16*, pages 370–383. Springer, 2020. 35, 71, 213

Herreiner, Dorothea K. and Puppe, Clemens. Inequality aversion and efficiency with ordinal and cardinal social preferences—an experimental study. *Journal of Economic Behavior & Organization*, 76(2):238–253, 2010. ISSN 0167-2681. 13, 178

Herreiner, Dorothea K. and Puppe, Clemens D. Envy freeness in experimental fair division problems. *Theory and Decision*, 67(1):65–100, Jul 2009. ISSN 1573-7187. 13, 178, 213

Hosseini, Hadi. The fairness fair: Bringing human perception into collective decision-making. Technical report, arXiv. org, 2023. 210

Hosseini, Hadi; Sikdar, Sujoy; Vaish, Rohit; Wang, Jun, and Xia, Lirong. Fair division through information withholding. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 12, 37

Hosseini, Hadi; Kavner, Joshua; Sikdar, Sujoy; Vaish, Rohit, and Xia, Lirong. Hide, not seek: Perceived fairness in envy-free allocations of indivisible goods, 2023a. 12

Hosseini, Hadi; Mammadov, Aghaheybat, and Wąs, Tomasz. Fairly allocating goods and (terrible) chores. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 2738–2746, 2023b. 182

Hosseini, Hadi; Payan, Justin; Sengupta, Rik; Vaish, Rohit, and Viswanathan, Vignesh. Graphical house allocation. In *22nd International Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 2023c. 64, 69

Hosseini, Hadi; Sikdar, Sujoy; Vaish, Rohit, and Xia, Lirong. Fairly dividing mixtures of goods and chores under lexicographic preferences. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 152–160, 2023d. 182

Hosseini, Hadi; Kumar, Medha, and Roy, Sanjukta. The degree of fairness in efficient house allocation. *CoRR*, abs/2407.04664, 2024a. 69

Hosseini, Hadi; McGregor, Andrew; Sengupta, Rik; Vaish, Rohit, and Viswanathan, Vignesh. Tight approximations for graphical house allocation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, page 825–833, Richland, SC, 2024b. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864. 64, 69

House Allocation, Image, 2017. URL `https://www.shutterstock.com/image-photo/\structural-engineers-commenting-allowed-build-strong-702548389`. Online; accessed 30 July 2024. xvii, 2

Hylland, Aanund and Zeckhauser, Richard. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979. 63, 68

Inheritance, Image, Accessed: 30 July 2024. URL `https://www.pixtastock.com/\illustration/75192875`. xvii, 2

Irving, Robert W.; Kavitha, Telikepalli; Mehlhorn, Kurt; Michail, Dimitrios, and Paluch, Katarzyna E. Rank-maximal matchings. *ACM Trans. Algorithms*, 2(4):602–610, oct 2006. ISSN 1549-6325. 63

Kamiyama, Naoyuki; Manurangsi, Pasin, and Suksompong, Warut. On the complexity of fair house allocation. *Operations Research Letters*, 49(4):572–577, 2021. 64, 69, 92

Klijn, Flip. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare*, 17(2):201–215, 2000. 12

Koutsoupias, Elias and Papadimitriou, Christos. Worst-Case Equilibria. *Computer Science Review*, 3(2):65–69, 2009. 146

Krysta, Piotr; Manlove, David F.; Rastegari, Baharak, and Zhang, Jinshan. Size versus truthfulness in the house allocation problem. *Algorithmica*, 81(9):3422–3463, 2019. 63

Kulkarni, Rucha; Mehta, Ruta, and Taki, Setareh. Indivisible mixed manna: On the computability of mms+po allocations. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, EC '21, page 683–684, New York, NY, USA, 2021a. Association for Computing Machinery. ISBN 9781450385541. 182

Kulkarni, Rucha; Mehta, Ruta, and Taki, Setareh. On the PTAS for maximin shares in an indivisible mixed manna. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 5523–5530. AAAI Press, 2021b. 182

Kurz, Sascha. The price of fairness for a small number of indivisible items. In Lübbecke, Marco; Koster, Arie; Letmathe, Peter; Madlener, Reinhard; Peis, Britta, and Walther, Grit, editors, *Operations Research Proceedings 2014*, pages 335–340, Cham, 2016. Springer International Publishing. ISBN 978-3-319-28697-6. 14

Kyropoulou, Maria; Suksompong, Warut, and Voudouris, Alexandros A. Almost Envy-Freeness in Group Resource Allocation. *Theoretical Computer Science*, 841:110–123, 2020. 142

Lackner, Martin and Skowron, Piotr. *Multi-winner voting with approval preferences*. Springer Nature, 2023. 71, 213

Lange, Pascal; Nguyen, Nhan-Tam, and Rothe, Jörg. The price to pay for forgoing normalization in fair division of indivisible goods. *Annals of Mathematics and Artificial Intelligence*, 88(7): 817–832, Jul 2020. ISSN 1573-7470. 179

Lee, Euiwoong. Apx-hardness of maximizing nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017. ISSN 0020-0190. 14

Lenstra, Hendrik W. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983. 27, 77

Lipton, Richard J; Markakis, Evangelos; Mossel, Elchanan, and Saberi, Amin. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004. 5, 9, 11, 19, 37, 146, 212

Lokshtanov, Daniel. Parameterized integer quadratic programming: Variables and coefficients. *CoRR*, abs/1511.00310, 2015. 113

Lokshtanov, Daniel; Marx, Dániel, and Saurabh, Saket. Slightly superexponential parameterized problems. *SIAM J. Comput.*, 47(3):675–702, 2018. 218

Madathil, Jayakrishnan; Misra, Neeldhara, and Sethia, Aditi. The complexity of minimizing envy in house allocation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, page 2673–2675, Richland, SC, 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321. 18

Manlove, David F. *Algorithmics of Matching Under Preferences*, volume 2 of *Series on Theoretical Computer Science*. WorldScientific, 2013. ISBN 978-981-4425-24-7. 63

Manurangsi, Pasin. Inapproximability of maximum biclique problems, minimum k-cut and densest at-least-k-subgraph from the small set expansion hypothesis. *Algorithms*, 11(1), 2018. ISSN 1999-4893. doi: 10.3390/a11010010. 92, 94

Mas-Colell, Andreu; Whinston, Michael D., and Green, Jerry R. *Microeconomic Theory*. Number 9780195102680 in OUP Catalogue. Oxford University Press, 1995. ISBN ARRAY(0x937da240). 60

Maskin, Eric and Feiwel, G. *On the Fair Allocation of Indivisible Goods*, pages 341–349. MacMillan, 1987. 12

Mathieson, Luke and Szeider, Stefan. Editing graphs to satisfy degree constraints: A parameterized approach. *Journal of Computer and System Sciences*, 78(1):179–191, 2012. ISSN 0022-0000. JCSS Knowledge Representation and Reasoning. 91

Misra, Neeldhara and Sethia, Aditi. Fair division is hard even for amicable agents. In *SOFSEM 2021: Theory and Practice of Computer Science: 47th International Conference on Current Trends in Theory and Practice of Computer Science, SOFSEM 2021, Bolzano-Bozen, Italy, January 25–29, 2021, Proceedings*, page 421–430, Berlin, Heidelberg, 2021. Springer-Verlag. ISBN 978-3-030-67730-5. 18

Moulin, Hervé. *Fair Division and Collective Welfare.* MIT press, 2004. 8

Nash Jr, John F. The Bargaining Problem. *Econometrica*, pages 155–162, 1950. 7, 146

Nguyen, Trung Thanh and Rothe, Jörg. How to decrease the degree of envy in allocations of indivisible goods. In *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, 2013, Proceedings*, page 271–284, Berlin, Heidelberg, 2013. Springer-Verlag. ISBN 978-3-642-41574-6. 12, 64

Nguyen, Trung Thanh and Rothe, Jörg. Minimizing Envy and Maximizing Average Nash Social Welfare in the Allocation of Indivisible Goods. *Discrete Applied Mathematics*, 179:54–68, 2014. 14

Nicosia, Gaia; Pacifici, Andrea, and Pferschy, Ulrich. Price of fairness for allocating a bounded resource. *European Journal of Operational Research*, 257(3):933–943, 2017. ISSN 0377-2217. 14

Organ Donation, Image, Accessed: 30 July 2024. URL https://currentaffairs.adda247.com/11th-national-organ-donation-day/. xvii, 2

Ortega, Josué. Multi-Unit Assignment under Dichotomous Preferences. *Mathematical Social Sciences*, 103:15–24, 2020. 142

Payan, Justin; Sengupta, Rik, and Viswanathan, Vignesh. Relaxations of envy-freeness over graphs. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, page 2652–2654, Richland, SC, 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321. 20

Plaut, Benjamin and Roughgarden, Tim. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020. 10, 20, 182

Procaccia, Ariel D. and Wang, Junxing. A lower bound for equitable cake cutting. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 479–495, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450345279. 180

Puppe, Clemens. The single-peaked domain revisited: A simple global characterization. *Journal of Economic Theory*, 176:55–80, 2018. ISSN 0022-0531. 119

Raghavendra, Prasad and Steurer, David. Graph expansion and the unique games conjecture. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, page 755–764, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781450300506. 92

Roos, Magnus and Rothe, Jörg. Complexity of Social Welfare Optimization in Multiagent Resource Allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 641–648, 2010. 142

Roth, Alvin E. and Peranson, Elliott. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89 (4):748–780, September 1999. 2

Roughgarden, Tim and Tardos, Éva. How Bad is Selfish Routing? *Journal of the ACM*, 49(2): 236–259, 2002. 146

Sandomirskiy, Fedor and Segal-Halevi, Erel. Fair division with minimal sharing. *CoRR*, abs/1908.01669, 2019. 12, 37, 38, 39, 41, 42, 57, 58, 59, 61, 62

Schrijver, Alexander. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1998. 165

Shams, Parham; Beynier, Aurélie; Bouveret, Sylvain, and Maudet, Nicolas. Minimizing and balancing envy among agents using Ordered Weighted Average. In *7th International Conference on Algorithmic Decision Theory*, Toulouse, France, November 2021. 12, 64

Shams, Parham; Beynier, Aurélie; Bouveret, Sylvain, and Maudet, Nicolas. Fair in the eyes of others. *Journal of Artificial Intelligence Research*, 75:913–951, 2022. 212

Shapley, Lloyd and Scarf, Herbert. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974. ISSN 0304-4068. 68

Shende, Priyanka and Purohit, Manish. Strategy-proof and envy-free mechanisms for house allocation. *J. Econ. Theory*, 213:105712, 2020. 69

Simmons, Forest W. and Su, Francis Edward. Consensus-halving via theorems of borsuk-ulam and tucker. *Mathematical Social Sciences*, 45(1):15–25, 2003. ISSN 0165-4896. ix, 15, 17, 207, 208

Sprumont, Yves. The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica: Journal of the Econometric Society*, pages 509–519, 1991. 119

Steinhaus, Hugo. The Problem of Fair Division. *Econometrica*, 16(1):101–104, 1948. 146

Sun, Ankang; Chen, Bo, and Doan, Xuan Vinh. Fairness criteria for allocating indivisible chores: connections and efficiencies. *Autonomous Agents and Multi-Agent Systems*, 37(2), aug 2023a. ISSN 1387-2532. 14, 20, 146

Sun, Ankang; Chen, Bo, and Doan, Xuan Vinh. Equitability and welfare maximization for allocating indivisible items. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, page 561–563, Richland, SC, 2023b. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321. xx, 13, 14, 20, 66, 69, 143, 146, 180, 181, 202

Svensson, Lars-Gunnar. Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4):939–54, 1983. 12

Tadenuma, Koichi and Thomson, William. The fair allocation of an indivisible good when monetary compensations are possible. *Mathematical Social Sciences*, 25(2):117–132, February 1993. 12

Thomassé, Stéphan. A 4k2 kernel for feedback vertex set. *ACM Trans. Algorithms*, 6(2), April 2010. ISSN 1549-6325. 78

Velez, Rodrigo A. Equitable rent division. *ACM Trans. Econ. Comput.*, 6(2), sep 2018. ISSN 2167-8375. 1

Viswanathan, Vignesh and Zick, Yair. A General Framework for Fair Allocation with Matroid Rank Valuations. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 1129–1152, 2023. 141

von Neumann, John. *1. A Certain Zero-sum Two-person Game Equivalent to the Optimal Assignment Problem*, pages 5–12. Princeton University Press, Princeton, 1953. 167

Zeng, Jinghan A and Mehta, Ruta. On the structure of envy-free orientations on graphs, 2024. 20

Zhou, Lin. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990. 63