

On the Complexity of Market Equilibria and Revenue Maximization

Dimitris Paparas

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2017

ProQuest Number:10257595

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10257595

Published by ProQuest LLC (2017). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

©2017

Dimitris Paparas

All Rights Reserved

ABSTRACT

On the Complexity of Market Equilibria and Revenue Maximization

Dimitris Paparas

This thesis consists of two parts. In the first part, we concentrate on the computation of Market Equilibria and settle the long-standing open problem regarding the computation of an approximate Arrow-Debreu market equilibrium in markets with CES utilities. We prove that the problem is PPAD-complete when the Constant Elasticity of Substitution parameter ρ is any constant less than -1 . Building on this result, we introduce the notion of non-monotone utilities, which covers a wide variety of utility functions in economic theory, and prove that it is PPAD-hard to compute an approximate Arrow-Debreu market equilibrium in markets with linear and non-monotone utilities.

In the second part, we study Revenue Maximization. We begin by resolving the complexity of the revenue-optimal Bayesian Unit-demand Item Pricing problem when the buyer's values for the items are independent. We show that the problem can be solved in polynomial time for distributions of support size 2; but its decision version is NP-complete for distributions of support size 3. Next, we study the optimal mechanism design problem for a single unit-demand buyer with item values drawn from independent distributions. We show that, for distributions of support-size 2 and the same high value, Item Pricing can achieve the same revenue as any menu of lotteries. On the other hand, we provide simple examples where randomization improves revenue. Finally, we show that unless the polynomial-time hierarchy collapses, namely $P^{NP} = P^{\#P}$, there is no universal efficient randomized algorithm that implements an optimal mechanism even when distributions have support size 3.

This page intentionally left blank.

Table of Contents

List of Figures	v
List of Tables	vii
1 Introduction	1
1.1 Market Equilibria	1
1.2 Revenue Maximization	4
1.3 Organization of this Thesis	6
I Market Equilibria	7
2 Preliminaries	9
2.1 Introduction	9
2.2 Related Work	10
2.3 Basic Definitions	12
2.3.1 Arrow–Debreu Markets and Market Equilibria	12
2.3.2 Polymatrix Games and Nash Equilibria	15
2.3.3 Organization of the rest of this part	16
3 The Complexity of Markets with CES utilities	17
3.1 Introduction	17
3.1.1 Elasticity of Substitution	18
3.1.2 Problem Definitions and Related Work	19
3.1.3 Main Challenges and Statement of Results	21

3.1.4	Main Ideas behind our Reduction	22
3.2	Analysing the Market of Example 3.1.1	27
3.2.1	The Excess Spending Function	27
3.2.2	Properties of the Excess Spending function	29
3.3	Markets with CES Utilities are PPAD-hard	35
3.3.1	Our Construction	35
3.3.2	Proof of Correctness	38
4	The Complexity of Non-Monotone Markets	47
4.1	Introduction	47
4.2	Definitions and Statement of Results	49
4.2.1	Examples of non-monotone Markets	51
4.3	Markets with Non-Monotone and Linear Utilities are PPAD-hard	53
4.3.1	Normalized Non-Monotone Markets	54
4.3.2	Our Construction	55
4.3.3	Proof of Correctness	59
II	Revenue Maximization	69
5	Preliminaries	71
5.1	Introduction	71
5.2	Related Work	72
5.2.1	Work on Item Pricing	73
5.2.2	Work on Lottery Pricing	74
5.3	Our Results and Organization of this part	76
6	The Complexity of Optimal Multidimensional Pricing	77
6.1	Introduction	77
6.2	Preliminaries	77
6.2.1	Problem Definition and Main Results	77
6.2.2	Tie-Breaking Rules	79

6.3	Membership in NP	83
6.4	A polynomial-time Algorithm for support size 2	87
6.4.1	An Interesting Special case.	87
6.4.2	General Case	93
6.5	NP-hardness for support size 3	96
7	On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms	103
7.1	Introduction	103
7.1.1	Ideas Behind the Proofs	106
7.2	Preliminaries	107
7.2.1	The Standard Linear Program	107
7.2.2	Properties of Optimal Solutions to $LP(\mathbf{I})$	109
7.3	Distributions with Support $\{\mathbf{a}_i, \mathbf{b}\}$	111
7.4	Two Items with Support Size 2	112
7.4.1	Upper Bounds for the Prices in L^*	114
7.4.2	Upper Bounds for the Expected Revenue	115
7.5	Small Instances where Lotteries help	118
7.6	Hardness of Optimal Mechanism Design	118
7.6.1	Linear Program Relaxation	119
7.6.2	Choices of Parameters and their Consequences	142
7.6.3	Returning to the Standard Linear Program	156
III	Conclusions	167
8	Conclusions	169
8.1	Conclusions	169
8.1.1	Summary of Results	169
8.1.2	Future Directions and Open Problems	170

IV Bibliography	175
Bibliography	177

List of Figures

3.1	A chain of markets over groups $\mathcal{R}_{i,j}$ of goods. Arrows correspond to $(n, G_i : G_j)$ -traders.	26
3.2	Market $M_{\mathbf{P}}$: Black (solid) arrows correspond to single minded-traders from H_i 's to G_j 's. Blue (dashed) arrows correspond to the chains of $\mathcal{R}_{i,j}$ markets of Figure 3.1.	27
3.3	The excess spending function f	29
3.4	The function g and the line $-\lambda x$, where $\lambda = f'(0)$	31
4.1	The excess demand function $Z_1(x)$ of Example 3.1.1.	52
7.1	Relaxed Linear Program $LP'(I)$	122

This page intentionally left blank.

List of Tables

5.1	Computational landscape of Revenue Maximization before this thesis. We use “R” to indicate that the problem is resolved in this thesis.	76
7.1	Revenue for each potentially optimal pricing (rows) and each possible valuation vector (columns).	112
7.2	An optimal menu.	113

This page intentionally left blank.

Acknowledgements

This thesis is the conclusion of my doctoral studies at Columbia University and it would have been impossible to complete without the contributions of many people.

First of all, I am eternally grateful to my advisor, Mihalis Yannakakis, for giving me the opportunity to be his student. Under his wise guidance, I learned how to be persistent with my research and managed to overcome many obstacles. His kindness and modesty are unparalleled. He always treated me as an equal and paid attention to what I was saying (even when I was wrong) without ever showing off his brilliance. I have lost count of the number of times that, without intimidating me, he pointed out elegant ways to tackle issues in my research that I had been struggling with for weeks. Finally, I also consider him a close friend and I owe him a lot for his guidance in both professional and personal matters.

Many thanks go to Xi Chen, my secondary advisor, for all the collaboration that led to the results of this thesis. The endless hours I spent with him and Mihalis working on the whiteboard were the most valuable part of my studies at Columbia.

I also want to thank the rest of my committee members, Alex Andoni, Shuchi Chawla, and Rocco Servedio, for all the valuable comments they made on the draft of this thesis. Rocco in particular has apparently read every minute detail, or at least this is what the typos that he noticed in the proofs of a number of lemmas indicate. Shuchi was kind enough to offer me a Postdoc position at her research group, which I am glad I decided to accept and I am already enjoying as I write this paragraph.

During my studies, I was fortunate to make new friends and collaborate with many brilliant people who ended up being my coauthors. Matthew Andrews, Yigal Bejerano, Aaron Bernstein, Clément Canonne, Ilias Diakonikolas, Fernando Krell, George Matikas, Orestis Polychroniou, Igor Carboni Oliveira, Anthi Orfanou, Valia Skiani, and Xiaorui Sun, thanks for all the fun.

I was also extremely lucky to be a Summer intern at Bell Labs and IBM in 2013 and 2015 respectively, where I worked with Matthew Andrews and Nimrod Meggido and gained valuable experience. I also met some amazing friends there; Christina Christodoulakis,

Mika Göös, and Ilya Razenshteyn, I miss all our time together.

With the salary of graduate students being a running joke at every university, special thanks go to the Alexander S. Onassis Foundation for partially supporting my studies through a Fellowship that helped me not to worry about money.

The reader should also thank Olivia Winn who accepted to check the English of this thesis, despite the fact that many times I am bold enough to challenge her knowledge of her mother tongue.

Last but not least, I am grateful to my mother, Eleni, for all the sacrifices she made after the death of my father in order to raise me and my brother Themis into decent humans. Mom, I love you with all my heart and everything I have achieved is because of you.

Madison, WI

February 2nd, 2017

In memory of my father, Serafim.

This page intentionally left blank.

Chapter 1

Introduction

Over the last 15 years, both the Theoretical Computer Science and the Artificial Intelligence communities have been systematically studying problems of Economic Theory from a computational point of view, giving rise to the area of *Algorithmic Game Theory and Computational Economics*. Two of the most fundamental notions in Economic Theory are those of Revenue Maximization and of Equilibrium and, naturally, they have drawn a lot of attention from the Computational Economics community. To what extent those notions are computationally tractable is the topic of this thesis.

1.1 Market Equilibria

Perhaps the cornerstone in Economic Theory is that of Equilibrium; a state in which: 1) agents behave rationally and 2) no one has an incentive to deviate. A particular type of equilibrium is that of *Market Equilibrium*; first defined by Walras [Walras, 1874] and shown to always exist under mild assumptions in the celebrated theorem of Arrow and Debreu [Arrow and Debreu, 1954], both of whom received the Nobel Prize in economics.

Formally, an Arrow–Debreu market M consists of $n \geq 1$ traders and a set of divisible goods $\{G_1, \dots, G_m\}$ for some $m \geq 1$. Each trader comes to M with an initial endowment $\mathbf{w} \in \mathbb{R}_+^m$ of goods, where w_j denotes the amount of G_j , and also has a real-valued utility function u . Given a bundle $\mathbf{x} \in \mathbb{R}_+^m$ of goods, $u(\mathbf{x})$ is her utility from bundle \mathbf{x} . If we assign prices \mathbf{p} to the goods, where we use p_j to denote the price per unit of G_j , then each trader

will sell her endowment \mathbf{w} at prices \mathbf{p} to obtain a budget of $\mathbf{w} \cdot \mathbf{p}$ and then will spend this amount to purchase from the market a bundle of goods \mathbf{x} that maximize her utility. We say that \mathbf{p} is a market equilibrium of M if we can assign to each trader a utility maximizing bundle with respect to \mathbf{p} such that the total demand equals the total supply and the market clears. Unfortunately, the proof of existence provided by the Arrow–Debreu theorem is based on Kakutani’s fixed point theorem [Kakutani, 1941] and, hence, it is non-constructive. Therefore, although the existence of an equilibrium is usually guaranteed under some mild conditions, how to efficiently find one is a highly non-trivial, and in many cases open, problem of high importance that has received a lot of attention by the community; after all, if it is a good model for a world where rational entities interact, the equilibrium state should be reached quickly and, therefore, it should be algorithmically possible to find it quickly as well.

As it turns out, the tractability of the problem relies heavily on the type of utilities that the traders have. For example, market equilibrium prices in markets with Linear utility functions can be described as solutions to a convex program [Jain, 2007] and it is possible to recover them efficiently. This is not an isolated case; convexity plays an important role in equilibrium computation. Indeed, Codenotti and Varadarajan [Codenotti and Varadarajan, 2007] applied convex programming techniques on markets with CES (Constant Elasticity of Substitution) utilities. This is one of the most important families of utilities, parameterized by a parameter $\rho \in (-\infty, 1] \setminus \{0\}$, that is widely used by economists because it nicely incorporates the notion of elasticity of substitution: the ratio of the proportionate change in the relative demand for two goods to the proportionate change in their relative prices (see 3.1.1 for details). Their convex programming formulations apply to the problem when $\rho \in (-1, 1] \setminus \{0\}$, making this case solvable in polynomial time. Nevertheless, the complexity of the problem when $\rho < -1$ remained a major open question for more than a decade. Another property that has played a critical role in the area is WGS (Weak Gross Substitutability): increasing the price of one good while keeping all other prices fixed cannot cause a decrease in the demand of any other good. Codenotti, McCune, and Varadarajan showed [Codenotti *et al.*, 2005b] that, for markets that satisfy WGS, a discrete price-adjustment algorithm converges to an approximate equilibrium in polynomial time, nevertheless, this property

does not cover all types of markets that admit efficient algorithms. For example, the family of CES utilities with parameter $-1 \leq \rho < 0$ does not satisfy WGS.

On the other hand, it is known that the general case of the problem is hard for the class PPAD, a suitable complexity class defined by Papadimitriou [Papadimitriou, 1994]. For example, markets with Leontief utilities [Codonotti *et al.*, 2006] are PPAD-hard. Notice however that these results only hold for a few specific and isolated families of utilities and the reduction techniques developed in these proofs are all different, each fine tuned for the family of utilities under consideration. Therefore, unlike other classes of problems in complexity theory, what is the structural property that makes an equilibrium hard to find is not well-understood yet and partially addressing this question is one of the goals of this thesis.

Our Results. In more details, we take a systematic approach towards completely characterizing the complexity of finding a Market Equilibrium by providing a sufficient condition for PPAD-hardness. To this direction, we settle [Chen *et al.*, 2013] the complexity of CES markets when $\rho < -1$ by proving that the problem is PPAD-hard, thus resolving this long-standing open question. As it turns out, the main structural property that our proof of hardness exploits is the following: It is possible to use CES utilities with $\rho < -1$ to construct a market in which for some particular tuple of prices there is an over-demanded product and, furthermore, the demand for this product will increase if its price increases without changing the prices of the other products. Inspired by this observation, we introduce the notion of *non-monotonicity*. A family of utilities \mathcal{U} is non-monotone if it is possible to use utilities from \mathcal{U} and construct a market with the behavior described above. Our second result shows that, for any fixed non-monotone family \mathcal{U} of utilities, it is PPAD-hard to find a market equilibrium for markets where each trader's utility is either a linear function or belongs to \mathcal{U} . This is the first result in the literature providing a sufficient condition for PPAD-hardness that is not fine-tuned to a specific family of utility functions. It is worth mentioning, for instance, that a corollary of our result is PPAD-hardness for markets with SPLC utilities (Separable, Piece-wise Linear, and Concave). This family of markets was already known to be PPAD-hard [Vazirani and Yannakakis, 2011], nevertheless, that reduction relied on properties specific to SPLC utilities.

1.2 Revenue Maximization

One major goal in auction theory is maximizing revenue. The work of Myerson [Myerson, 1981] provides an efficient auction that maximizes the expected revenue from selling one product to many buyers. A nice property of Myerson's result is that it introduces the notion of *virtual valuations*, a closed-form characterization (and thus an efficient algorithm) of the optimal auction. A widely studied setting, complementary to the one of Myerson's, known as the *Multi-dimensional Bayesian Item Pricing* or *Deterministic* setting, is the following:

Deterministic Setting. There is one seller, one buyer, and n items. The seller has Bayesian knowledge about the buyer's valuations, namely a discrete product distribution \mathcal{D} on possible values for each item. His purpose is to assign a price p_i to each item i in a way that maximizes her expected revenue, that is, the money that the buyer spends to buy an item when her values are drawn from \mathcal{D} . Given prices for the items, the buyer will pay p_i dollars to buy the item i that maximizes her payoff $v_i - p_i$, where v_i is her actual value for item i . The computational problem here is to find revenue maximizing prices for the items when \mathcal{D} is given as input.

Although the deterministic setting favours simplicity, it turns out that it is possible to improve the seller's expected revenue [Thanassoulis, 2004] using randomization. This is achieved through a generalization known as the *Multi-dimensional Mechanism Design* or *Randomized* setting which is defined as follows:

Randomized Setting. Here, the seller offers a menu of lotteries $\{L_1, \dots, L_k\}$ where L_i consists of a tuple of probabilities (x_{i1}, \dots, x_{in}) and a price p_i , with $\sum_{j \in [n]} x_{ij} \leq 1$. To buy lottery L_i the buyer must pay p_i dollars and then she receives item j with probability x_{ij} . Given a menu of lotteries, the buyer will buy the lottery L_i that maximizes her expected payoff $\sum_j x_{ij} v_j - p_i$, where v_j is her actual value for item j . As before, the computational problem is to find a revenue maximizing menu of lotteries when \mathcal{D} is given as input.

Notice that the deterministic setting is the special case of the randomized setting where $x_{ij} \in \{0, 1\}$. These two setting can be further divided in terms of the buyer's preference: A *Unit-demand* buyer is interested in at most one item; as in the settings described above. This corresponds to items that are perfect substitutes of each other, e.g. when the seller is

a car dealer. On the other hand, an *Additive* buyer is interested in obtaining any number of items, e.g. when visiting a mall. This is equivalent to offering lotteries where $0 \leq x_{ij} \leq 1$ for all i and j in the randomized setting and offering a menu of subsets of items in the deterministic setting. In this thesis, we will concentrate on Unit-demand buyers.

Of main interest to economists is whether there is a closed form characterization of the optimal auction for the above settings, similar to Myerson's result. Relaxing this requirement, the algorithmic game theory community focused its interest on whether there exists an efficient algorithm that implements the optimal auction for the multi-item setting. A lot of work has been done towards this direction, both algorithmic and complexity theoretic. Daskalakis and Tzamos [Daskalakis *et al.*, 2014a] proved that finding a revenue maximizing menu for the *Randomized setting with an Additive buyer* is #P-hard. On the positive side, Chawla et al [Chawla *et al.*, 2007] gave a constant factor approximation algorithm for the *Deterministic setting with a Unit-demand buyer* and Cai and Daskalakis [Cai and Daskalakis, 2011] obtained a PTAS for the same setting when distributions satisfy the monotone hazard rate condition. Nevertheless, the existence of an exact efficient algorithm for both the randomized and the deterministic settings remained a major open question for the case of a Unit-demand buyer [Chawla *et al.*, 2007], [Cai and Daskalakis, 2011], [McAfee and Mc-Millan, 1988], [Manelli and Vincent, 2007]. In the second part of this thesis we resolve this question in the negative for both problems.

Our Results. We begin with the deterministic unit-demand setting and initially approach it trying to upper-bound its complexity, obtaining NP membership through a partition of the pricing space into well-behaved cells. An implication of this characterization is a polynomial-time algorithm for arbitrary distributions when the number of items is constant. Building on this result, we also obtain a polynomial-time algorithm for distributions of support size 2. To complement this, we show NP-completeness for the general case of the problem, even for distributions of support size 3. Finally, we study the randomized unit-demand setting and prove that the problem does not admit polynomial-time algorithms unless $P^{NP} = P^{\#P}$. Furthermore, we prove that for distributions of support size 2 and the same high value for all items offering lotteries does not improve revenue when compared to the deterministic setting, implying a polynomial-time algorithm [Chen *et al.*, 2014]. On the other hand, we

present simple examples where randomization strictly improves revenue.

1.3 Organization of this Thesis

The rest of this thesis consists of three parts. In Part I, we study the problem of finding a Market Equilibrium. We start by providing all the necessary definitions and related work in Chapter 2. We then prove PPAD-hardness for markets with CES utilities in Chapter 3. Building on this result, we prove PPAD-hardness for markets where each trader has either a linear or a non-monotone utility function in Chapter 4. In Part II, we study the problem of revenue maximization. After introducing all concepts and related work in Chapter 5 we proceed presenting our results. In Chapters 6 and 7 we study the deterministic and randomized settings respectively, obtaining the results described above. Finally, we conclude in Part III where we summarize the results obtained in this thesis and discuss open problems and future directions.

Part I

Market Equilibria

This page intentionally left blank.

Chapter 2

Preliminaries

2.1 Introduction

General equilibrium theory [Debreu, 1959; Ellickson, 1994] is perhaps the cornerstone of Mathematical Economics. A model central to this field is that of the *Arrow–Debreu market* [Arrow and Debreu, 1954], which studies the interactions of price, demand and supply, and is established on the demand-equal-supply principle of Walras [Walras, 1874].

In an Arrow–Debreu market M , traders exchange goods to maximize their utilities¹. Formally, M consists of a $n \geq 1$ traders and a set of divisible goods $\{G_1, \dots, G_m\}$ for some $m \geq 1$. Each trader comes to the market with an initial endowment $\mathbf{w} \in \mathbb{R}_+^m$ of goods, where w_j denotes the amount of G_j , and also has a real-valued utility function u . Given a bundle $\mathbf{x} \in \mathbb{R}_+^m$ of goods, $u(\mathbf{x})$ is her utility from bundle \mathbf{x} .

Assume now that we assign prices to the goods according to a price vector $\mathbf{p} \in \mathbb{R}_+^m$, where we use p_j to denote the price of G_j . In that occasion, each trader sells her endowment \mathbf{w} at \mathbf{p} to obtain a budget of $\mathbf{w} \cdot \mathbf{p}$ and then spends this amount to purchase from the market a bundle of goods \mathbf{x} that maximize her utility. We say \mathbf{p} is a *market equilibrium* of M if we can assign to each trader a utility maximizing bundle with respect to \mathbf{p} such that the total demand equals the total supply and the market clears².

¹Arrow and Debreu also considered firms with production plans. Here we focus on the exchange setting.

²If the price of a good is 0 the total demand for that good is allowed to be less than the supply.

2.2 Related Work

Arrow and Debreu, in their celebrated theorem [Arrow and Debreu, 1954], proved that every market has an equilibrium, under some mild³ conditions. Their proof, however, uses Kakutani’s fixed point theorem [Kakutani, 1941] and, hence, is non-constructive and doesn’t suggest any algorithms; since no efficient general fixed-point algorithm is known so far.

The computational aspect of the problem was first discussed in the pioneering work of Scarf [Scarf, 1973]. Since then, it has been studied extensively and several general schemes have been proposed. Under suitable sufficient conditions, those schemes converge to an equilibrium, alas, there is no efficient general algorithm known. The difficulty of the problem is evidenced by exponential lower bounds on both the query complexity of the discrete Brouwer fixed-point problem [Hirsch *et al.*, 1989; Chen and Deng, 2008; Chen and Teng, 2007; Chen *et al.*, 2008] and the number of steps of general price adjustment schemes for market equilibria [Papadimitriou and Yannakakis, 2010], implying that any efficient algorithm for the general problem would require novel techniques.

On the complexity-theoretic aspect of the problem, Papadimitriou [Papadimitriou, 1994] defined PPAD, a complexity class that captures those problems that can be reduced in polynomial time to the problem of finding a source or a sink in a succinctly represented graph of special form. In this graph every node has both in-degree and out-degree at most one and there is also a dedicated source that cannot be accepted as a solution and its purpose is to guarantee the existence of a solution somewhere else in the graph. The graph is described succinctly by a boolean circuit \mathcal{C} which takes as input a node n and outputs a tuple of (possibly null) nodes (x, y) , where x and y are respectively the predecessor and the successor of n . The graph induced by \mathcal{C} is defined as follows: edge (s, t) exists iff $\mathcal{C}(s) = (a, t)$ and $\mathcal{C}(t) = (s, b)$ for some a and b . Papadimitriou proved PPAD-completeness for a number of archetypical total⁴ problems and conjectured that finding a Market Equilibrium is PPAD-hard. His conjecture has been confirmed the last fifteen years when, in a sequence of papers, many versions of the Market Equilibrium problem were shown to be in PPAD and/or PPAD-

³The term ‘mild’ is widely used in the literature and, hence, we adopt it.

⁴A problem is total if it always accepts a solution.

hard [Codenotti *et al.*, 2006; Huang and Teng, 2007; Deng and Du, 2008; Chen *et al.*, 2009a; Vazirani and Yannakakis, 2011; Papadimitriou and Wilkens, 2011; Chen and Teng, 2009; Chen and Teng, 2011]. Nevertheless, the Market Equilibrium problem in its general case, even when the necessary conditions for existence of equilibrium are satisfied, is unlikely to lie in PPAD due to the algebraic nature of its solutions. To address this, Etessami and Yannakakis [Etessami and Yannakakis, 2010] defined FIXP, a more general class capturing all problems that can be reduced in polynomial time to the problem of finding a fixed point of a Brouwer function that is provided in the form of an algebraic circuit over operators $\{+, -, *, /, \max, \sqrt{\cdot}\}$, and proved that the Market Equilibrium problem with algebraic demand functions as well as a number of other important problems are FIXP-complete. Their work has recently been extended by more FIXP completeness and membership results [Chen *et al.*, 2013; Garg *et al.*, 2014; Garg *et al.*, 2016]. The above results imply that, unless standard complexity assumptions are refuted, the problem of finding a Market Equilibrium does not admit an efficient algorithm in its general case. Notice however, that these results only hold for a few specific and isolated families of utilities and the reduction techniques developed in these proofs are all different, each fine tuned for the family of utilities under consideration. Therefore, unlike other classes of problems in complexity theory, what is the structural property that makes an equilibrium hard to find is not well-understood yet and partially addressing this question is one of the goals of this thesis.

On the algorithmic aspect of the problem, Deng, Papadimitriou, and Safra [Deng *et al.*, 2003] initiated a line of work on the computation and approximation of equilibria that, during the last fifteen years, has led to numerous results and a lot of progress for various market models. This includes efficient algorithms for the market equilibrium problem for various utility functions and notions of approximation [Jain *et al.*, 2003; Devanur and Vazirani, 2003; Chen *et al.*, 2004; Devanur and Vazirani, 2004; Garg and Kapoor, 2004; Garg *et al.*, 2004; Codenotti *et al.*, 2005b; Codenotti *et al.*, 2005a; Codenotti *et al.*, 2005c; Jain *et al.*, 2005; Jain and Mahdian, 2005; Jain and Varadarajan, 2006; Chen *et al.*, 2006; Jain, 2007; Ye, 2007; Devanur and Kannan, 2008; Devanur *et al.*, 2008; Ye, 2008], many of which are based on the convex-programming approach of [Eisenberg and Gale, 1959; Nenakov and Primak, 1983].

A property that has played a critical role in the above work is WGS (Weak Gross Substitutability). A family \mathcal{U} of utilities satisfies WGS if for any market consisting of traders with utilities from \mathcal{U} , increasing the price of one good while keeping all other prices fixed cannot cause a decrease in the demand of any other good. WGS implies that the set of equilibria is convex. Arrow, Block, and Hurwicz [Arrow *et al.*, 1959] showed that the continuous tatonnement process [Walras, 1874; Samuelson, 1947] converges for any market satisfying WGS. Recently, Codenotti, McCune, and Varadarajan [Codenotti *et al.*, 2005b] showed that a discrete tatonnement algorithm converges to an approximate equilibrium in polynomial time, if equipped with an excess demand oracle; an algorithm that takes as input a set of prices and outputs the total demand for each of the goods under those prices. Another general property that implies convexity of equilibria is WARP (Weak Axiom of Revealed Preference [Mas-Colell *et al.*, 1995]). While many families of utilities satisfy WGS or WARP, they do not seem to cover all the efficiently solvable market problems, e.g., the family of CES utilities with parameter $-1 \leq \rho < 0$ (see chapter 3 for definition) does not satisfy WGS or WARP but has a convex formulation [Codenotti *et al.*, 2005a].

2.3 Basic Definitions

Notation. We use \mathbb{R}_+ and \mathbb{Q}_+ to denote the sets of nonnegative reals and nonnegative rationals respectively. Given an integer $n > 0$, we use $[n]$ to denote the set $\{1, \dots, n\}$. Given two integers m and n , where $m \leq n$, we use $[m : n]$ to denote the set $\{m, m + 1, \dots, n\}$. Given a vector $\mathbf{y} \in \mathbb{R}^m$ and $c > 0$, we use $B(\mathbf{y}, c)$ to denote the set of \mathbf{x} with $\|\mathbf{x} - \mathbf{y}\|_\infty \leq c$.

2.3.1 Arrow–Debreu Markets and Market Equilibria

An Arrow–Debreu exchange market M consists of a finite set $\{T_1, \dots, T_n\}$ of $n > 0$ traders and a finite set $\{G_1, \dots, G_m\}$ of $m > 0$ divisible goods. Each trader T_i owns an initial endowment $\mathbf{w}_i \in \mathbb{R}_+^m$, where $w_{i,j}$ denotes her initial amount of good G_j . Each trader T_i also has a utility function $u_i : \mathbf{R}_+^m \rightarrow \mathbf{R}_+$, where $u_i(x_{i,1}, \dots, x_{i,m})$ represents the utility she derives if for each $j \in [m]$ she owns $x_{i,j}$ units of G_j . In the rest of the paper, we will refer to an Arrow–Debreu exchange market simply as a market for convenience.

Assume now that we assign prices to the goods according to a nonnegative price vector $\mathbf{p} = (\pi_1, \dots, \pi_m) \neq \mathbf{0}$, where we use π_j to denote the price of G_j . Each trader T_i will then sell her initial endowment \mathbf{w}_i at prices \mathbf{p} and obtain a budget $\sum_{j \in [m]} w_{i,j} \cdot \pi_j$. She will then spend this amount to purchase from the market a bundle of goods $\mathbf{x}_i \in \mathbb{R}_+^m$ that maximize her utility. We say \mathbf{p} is a *market equilibrium* of M if we can assign to each trader an optimal (a utility maximizing) bundle with respect to \mathbf{p} such that the total demand of each good equals the total supply (or is at most the total supply for goods priced at 0) and the market clears. Formally, given \mathbf{p} we let $\text{OPT}_i(\mathbf{p})$ be the set of optimal bundles of T_i with respect to \mathbf{p} , that is, $\text{OPT}_i(\mathbf{p})$ is the set of solutions to the program

$$\begin{aligned} & \max u_i(\mathbf{x}) \\ \text{s.t. } & \sum_{j \in [m]} x_j \cdot \pi_j \leq \sum_{j \in [m]} w_{i,j} \cdot \pi_j \end{aligned}$$

Next we define the (aggregate) excess demand of a good with respect to a price vector \mathbf{p} :

Definition 1 (Excess Demand). *Given \mathbf{p} , the excess demand $Z(\mathbf{p})$ consists of all vectors \mathbf{z} of the form $\mathbf{z} = \mathbf{x}_1 + \dots + \mathbf{x}_m - (\mathbf{w}_1 + \dots + \mathbf{w}_m)$, where \mathbf{x}_i is an optimal bundle in $\text{OPT}_i(\mathbf{p})$ for each $i \in [n]$. For each good G_j we also use $Z_j(\mathbf{p})$ to denote the projection of $Z(\mathbf{p})$ on the j th coordinate.*

In general, $Z(\mathbf{p})$ is a set and Z is a correspondence. We usually refer to a subset of traders in a market as a submarket, and sometimes we are interested in the excess demand of a submarket, for which the sums of \mathbf{x}_i 's and \mathbf{w}_i 's are only taken over traders in the subset. Finally we define market equilibria:

Definition 2 (Market Equilibria). *We say \mathbf{p} is a market equilibrium of M if $Z(\mathbf{p})$ contains a vector \mathbf{z} such that $z_j \leq 0$ for all $j \in [m]$ and $z_j < 0$ implies that $\pi_j = 0$.*

Notice that if $z_j > 0$, then the traders request more than the total available amount of G_j and if $z_j \leq 0$ then they request at most as much amount of it as is available in the market. As $\text{OPT}_i(\mathbf{p})$ is invariant under scaling of \mathbf{p} (by a positive factor), it is easy to see that the set of market equilibria is closed under scaling.

In general, a market equilibrium may not exist. The pioneering existence theorem of Arrow and Debreu [Arrow and Debreu, 1954] states that if all the utility functions are quasi-concave, then under certain mild conditions a market always has an equilibrium. Here, we use the weaker sufficient condition of Maxfield [Maxfield, 1997].

Definition 3 (Local Non-Satiation). *We say a utility function $u : \mathbb{R}_+^m \rightarrow \mathbb{R}_+$ is locally non-satiated if for any $\mathbf{x} \in \mathbb{R}_+^m$ and any $\epsilon > 0$, there exists a $\mathbf{y} \in B(\mathbf{x}, \epsilon) \cap \mathbb{R}_+^m$ such that $u(\mathbf{y}) > u(\mathbf{x})$. We say u is non-satiated with respect to the k th good, if for any $\mathbf{x} \in \mathbb{R}_+^m$, there exists a $\mathbf{y} \in \mathbb{R}_+^m$ such that $u(\mathbf{y}) > u(\mathbf{x})$ and $y_j = x_j$ for all $j \neq k$.*

If the utility of a trader is locally non-satiated, then her optimal bundle must exhaust her budget. Therefore, if every trader in M has a non-satiated utility then Walras' law holds: $\mathbf{z} \cdot \mathbf{p} = 0$ for all $\mathbf{z} \in Z(\mathbf{p})$.

Definition 4 (Economy Graphs). *Given a market M , we define a directed graph as follows. Each vertex of the graph corresponds to a good G_i in M . For two goods G_i and G_j in M , we add an edge from G_i to G_j if there is a trader T_k such that $w_{k,i} > 0$ and u_k is non-satiated with respect to G_j , i.e., T_k owns a positive amount of G_i and is interested in G_j . We call this graph the economy graph of M [Maxfield, 1997].⁵*

We then say a market M is strongly connected if its economy graph is strongly connected. A simplified version of Maxfield's existence theorem [Maxfield, 1997] is the following:

Theorem 1 (Maxfield [Maxfield, 1997]). *If the following two conditions hold, then M has a market equilibrium: 1) Every utility function is continuous, quasi-concave, and locally non-satiated; and 2) M is strongly connected. Moreover, the price of every good is positive in a market equilibrium.*

Clearly, when a market satisfies the conditions of Theorem 1, \mathbf{p} is an equilibrium if and only if $\mathbf{0} \in Z(\mathbf{p})$. Here, we are also interested in the problem of finding an approximate equilibrium in a market that satisfies the conditions of Theorem 1. For this we define two notions of approximate equilibria:

⁵ Maxfield defines this as a graph between the traders instead of the goods, but it is easy to see that the sufficient condition of strong connectivity is equivalent between the two versions.

Definition 5 (ϵ -Approximate Market Equilibria). *We call \mathbf{p} an ϵ -approximate market equilibrium of M for some $\epsilon > 0$ if there exists a vector $\mathbf{z} \in Z(\mathbf{p})$ such that $z_j \leq \epsilon \sum_{i \in [n]} w_{i,j}$ for all $j \in [m]$.*

Definition 6 (ϵ -Tight Approximate Market Equilibria). *We say \mathbf{p} is an ϵ -tight approximate market equilibrium of M for some $\epsilon > 0$ if there exists $\mathbf{z} \in Z(\mathbf{p})$ such that $|z_j| \leq \epsilon \sum_{i \in [n]} w_{i,j}$ for all $j \in [m]$.*

Both notions of approximate equilibria have been used in the literature. Although the two-sided notion of tight approximate market equilibria is more commonly used, in our work [Chen *et al.*, 2013] we present an unexpected market for which any $(1/2)$ -tight approximate market equilibrium \mathbf{p} must have a doubly exponentially small entry when $\sum_j \pi_j = 1$.

Finally, because we have to deal with real numbers, we need the following definition to make formal statements:

Definition 7. *We say a real number β is moderately computable if there is an algorithm that, given a rational number $\gamma > 0$, outputs a γ -rational approximation β' of β : $|\beta' - \beta| \leq \gamma$, in time polynomial in $1/\gamma$.*

2.3.2 Polymatrix Games and Nash Equilibria

To obtain our results in Chapters 3 and 4, we provide polynomial-time reductions from the problem of computing an approximate Nash equilibrium in a polymatrix game [Janovskaya, 1968] with two pure strategies for each player. Such a game with n players can be described by a $2n \times 2n$ rational matrix \mathbf{P} , with all entries between 0 and 1, $\sum_{i \in [2n]} P_{i,j} = \xi$, and $P_{2i-1,2i-1} = P_{2i-1,2i} = P_{2i,2i-1} = P_{2i,2i} = \psi$ for all j and some constants $\xi > \psi > 0$. An ϵ -well-supported Nash equilibrium is a vector $\mathbf{x} \in \mathbb{R}_+^{2n}$ such that for all $i \in [n]$, we have $x_{2i-1} + x_{2i} = 1$ and

$$\begin{aligned} \mathbf{x}^T \cdot \mathbf{P}_{2i-1} > \mathbf{x}^T \cdot \mathbf{P}_{2i} + \epsilon &\Rightarrow x_{2i} = 0 \\ \mathbf{x}^T \cdot \mathbf{P}_{2i} > \mathbf{x}^T \cdot \mathbf{P}_{2i-1} + \epsilon &\Rightarrow x_{2i-1} = 0, \end{aligned}$$

where \mathbf{P}_{2i-1} and \mathbf{P}_{2i} denote the $(2i-1)$ th and $(2i)$ th column vectors of \mathbf{P} , respectively.

Normalization. For convenience, we normalize \mathbf{P} into a $2n \times 2n$ matrix \mathbf{P}' by setting

$$P'_{i,2j-1} = 1/2 + (P_{i,2j-1} - P_{i,2j})/2 \quad \text{and} \quad P'_{i,2j} = 1/2 - (P_{i,2j-1} - P_{i,2j})/2$$

for all $i \in [2n], j \in [n]$. It is clear that \mathbf{P}' is also a rational matrix with entries between 0 and 1 and in addition

$$P'_{i,2j-1} + P'_{i,2j} = 1, \quad \text{for all } i \in [2n] \text{ and } j \in [n]. \quad (2.1)$$

From the definition of ϵ -well-supported Nash equilibria, it is easy to show that \mathbf{P} and \mathbf{P}' have the same set of ϵ -well-supported equilibria for any $\epsilon \geq 0$.

Let **POLYMATRIX** denote the following problem:

Definition 8. Given a normalized polymatrix game \mathbf{P} (i.e. the entries of \mathbf{P} satisfy (2.1)), find an ϵ -well-supported Nash equilibrium with $\epsilon = 1/n$.

It was shown in [Daskalakis *et al.*, 2009] that finding an exact Nash equilibrium of a polymatrix game with two pure strategies for each player is PPAD-hard (it is not stated explicitly there but follows from the proof of Lemma 6.3). It turns out that **POLYMATRIX** is PPAD-hard as well. The proof uses techniques developed in previous work on Nash equilibria [Daskalakis *et al.*, 2009; Chen *et al.*, 2009b]. While its PPAD-hardness is used here as a bridge to establish Theorem 3 and Theorem 6, we think the result on **POLYMATRIX** is interesting for its own right and refer the interested reader to the original paper [Chen *et al.*, 2013] for its proof.

Theorem 2. **POLYMATRIX** is PPAD-complete.⁶

2.3.3 Organization of the rest of this part

In the remaining of this part we prove hardness of finding a market equilibrium for some interesting types of exchange markets. In Chapter 3 we study markets with CES utilities and prove that, for any constant $\rho < -1$, finding an ϵ -approximate Market Equilibrium is PPAD-hard for a suitable chosen ϵ . In Chapter 4, we extend this result to all families of utilities that satisfy a property that we call *non-monotonicity*.

⁶Rubinstein [Rubinstein, 2015] has subsequently shown the PPAD-hardness of the ϵ -approximate equilibrium problem for a polymatrix game even for some constant $\epsilon > 0$.

Chapter 3

The Complexity of Markets with CES utilities

3.1 Introduction

In this chapter, we study the complexity of approximating market equilibria in markets with CES (constant elasticity of substitution) utility functions [Mas-Colell *et al.*, 1995].

Definition 9 (CES utility function). *We call $u : \mathbb{R}_+^m \rightarrow \mathbb{R}_+$ a CES function with parameter $\rho < 1$, $\rho \neq 0$, if it is of the form*

$$u(x_1, \dots, x_m) = \left(\sum_{j \in [m]} \alpha_j \cdot x_j^\rho \right)^{\frac{1}{\rho}}$$

where the coefficients $\alpha_1, \dots, \alpha_m \in \mathbb{R}_+$.

The family of CES utility functions was first introduced in [Solow, 1956; Dickinson, 1954]. It was then used in [Arrow *et al.*, 1961] to model production functions and predict economic growth. It has been one of the most widely used families of utility functions in economics literature [Shoven and Whalley, 1992; de La Grandville, 2009] due to its versatility and flexibility in economic modeling. For example, the popular modeling language MPSGE [Rutherford, 1999] for equilibrium analysis uses CES functions (and their generalization to nested CES functions) to model consumption and production.

Consider a trader T with a CES utility function u in which $\alpha_j > 0$ iff $j \in S \subseteq [m]$. Let \mathbf{w} denote the initial endowment of T and let \mathbf{p} denote a price vector with $\pi_j > 0$ for all $j \in [m]$. Then using the KKT conditions (on the optimization problem of maximizing T 's utility subject to the budget constraint), we have the following folklore formula for the *unique* optimal bundle of T :

$$x_j = \left(\frac{\alpha_j}{\pi_j} \right)^{1/(1-\rho)} \times \frac{\mathbf{w} \cdot \mathbf{p}}{\sum_{k \in S} \alpha_k^{1/(1-\rho)} \cdot \pi_k^{-\rho/(1-\rho)}}, \quad \text{for each } j \in S \quad (3.1)$$

It is also clear that if $\pi_j = 0$ for some $j \in S$ then T would demand an infinite amount of G_j . This implies that when a CES market is strongly connected π_j must be positive for all $j \in [m]$ in any (exact or approximate) market equilibrium of M .

3.1.1 Elasticity of Substitution

The parameter ρ of a CES utility function is related to the elasticity of substitution σ , which measures the percentage change in the ratio of the demands of two goods in response to a percentage change in their prices, or intuitively, how easy it is to substitute different goods or resources [Hicks, 1932; Robinson, 1933] (namely $\rho = (\sigma - 1)/\sigma$). Selecting specific values for ρ between 1 and $-\infty$ yields various basic utility functions and models different points in the substitutes-complements spectrum. This ranges from the perfect substitutes case when $\rho = 1$, which corresponds to linear utilities

$$u(x_1, \dots, x_m) = \sum_{j \in [m]} \alpha_j x_j,$$

to the intermediate case when $\rho \rightarrow 0$, which corresponds to Cobb-Douglas utilities

$$u(x_1, \dots, x_m) = \prod_{j \in [m]} x_j^{\alpha_j},$$

in which case the trader spends a fixed proportion of her budget on each of the goods, to the perfect complements case when $\rho \rightarrow -\infty$, which corresponds to Leontief utilities. A Leontief utility function has the form

$$u(x_1, \dots, x_m) = \min_{j \in S} \{x_j/c_j\},$$

for some subset $S \subseteq [m]$ of goods and positive constants $c_j > 0$ for all $j \in S$. This represents the utility of a trader who wants to acquire goods in S in quantities proportional to the c_j . This function is the limit of the functions $(\sum_{j \in S} (x_j/c_j)^\rho)^{1/\rho}$ as $\rho \rightarrow -\infty$; that is, the Leontief function is the limit of CES functions with coefficients $\alpha_j = 1/c_j^\rho$ for $j \in S$ and $\alpha_j = 0$ for $j \notin S$.

3.1.2 Problem Definitions and Related Work

Markets with CES utilities may not have any rational equilibria in general, even when ρ and all the coefficients are rational, therefore we need to study the approximation of market equilibria. For this purpose we define the following three computational problems:

1. **CES**: The input of the problem is a pair (k, M) , where k is a positive integer encoded in unary (k represents the desired number of bits of precision), and M is a strongly connected market in which all utilities are CES, with the parameter $\rho_i < 1$ of each trader T_i being rational and given in unary (because ρ appears in the exponent of the utility and demand functions). The ρ_i parameters may be the same or different for different traders, and there may be a mixture of positive and negative parameters. The endowments $w_{i,j}$ and coefficients $\alpha_{i,j}$ are rational and encoded in binary. The goal is to find a price vector \mathbf{p} that is within $1/2^k$ of some equilibrium in every coordinate, i.e., such that there exists an (exact) equilibrium \mathbf{p}^* of M with $\|\mathbf{p} - \mathbf{p}^*\|_\infty \leq 1/2^k$.
2. **CES-APPROX**: The input of the problem is the same as **CES**. The goal is to find an ϵ -approximate market equilibrium of M , where $\epsilon = 1/2^k$.
3. **ρ -CES-APPROX** for any fixed rational number $\rho < -1$: The input is the same as **CES**, except that the utilities of all the traders have the same fixed parameter ρ , which is considered as a constant, not part of the input. The goal is to find an ϵ -approximate market equilibrium of M , where $\epsilon = 1/k$.

The output of the first problem (**CES**) is usually referred to in the literature as a strongly approximate equilibrium. It is also possible to define **CES** under a model of real

computation and ask for an exact equilibrium. Finally, to justify the use of ϵ -approximate market equilibria instead of ϵ -tight approximate market equilibria, in both **CES-APPROX** and ρ -**CES-APPROX**, we refer the interested reader to [Chen *et al.*, 2013] for an example of a market for which an exponential number of bits is required to represent any of its $(1/2)$ -tight approximate equilibria. By contrast, in the same paper we show that for the one-sided notion there is always an ϵ -approximate equilibrium with a polynomial number of bits and, in fact, **CES-APPROX** is in PPAD and **CES** is in FIXP.

Nenakov and Primak [Nenakov and Primak, 1983] gave a convex program that characterizes the set of equilibria when $\rho = 1$, i.e. all utilities are linear. Jain [Jain, 2007] discovered the same convex program independently and used the ellipsoid algorithm to give a polynomial-time exact algorithm for the linear case. It turns out that this convex program can also be applied to characterize the set of equilibria in CES markets with $\rho > 0$. Codenotti, McCune, Penumatcha, and Varadarajan [Codenotti *et al.*, 2005a] gave a different convex formulation for the set of equilibria in CES markets with $\rho \in [-1, 0)$. The range of $\rho < -1$ however has remained an intriguing open problem. For this range, it is known that the set of equilibria can be disconnected, and thus one cannot hope for a direct convex formulation. For example, the following market has three disconnected equilibria.

Example 3.1.1. Consider the following market M with two goods G_1, G_2 and two traders T_1, T_2 . T_1 has 1 unit of G_1 , T_2 has 1 unit of G_2 , and the utilities are

$$u_1(x_1, x_2) = (\alpha \cdot x_1^\rho + x_2^\rho)^{1/\rho} \quad \text{and} \quad u_2(x_1, x_2) = (x_1^\rho + \alpha \cdot x_2^\rho)^{1/\rho}$$

respectively. When $\rho < -1$ and α is large enough, [Gjerstad, 1996] shows that M has three disconnected equilibria $(1, 1)$, $(1-\theta, 1+\theta)$, and $(1+\theta, 1-\theta)$ for some $\theta > 0$, and furthermore, the excess demand function of both goods is increasing at $(1, 1)$. We will carefully study this market in Section 3.2 since it plays an important role in the proof of Theorem 3.

The failure of the convex-programming approach seems to suggest that the problem might be hard. In fact, when $\rho \rightarrow -\infty$, CES utilities converge to Leontief utilities for which finding an approximate equilibrium is PPAD-complete [Codenotti *et al.*, 2006] and computing an actual equilibrium (to a desired precision) is FIXP-complete [Garg *et al.*, 2014]. This argument, however, is less compelling due to the fact that a market with CES

utilities converging to a Leontief market, as $\rho \rightarrow -\infty$, does not mean that the equilibria of the CES market converge to an equilibrium of the Leontief market at the limit. Actually, in [Chen *et al.*, 2013], we construct an example demonstrating that it is possible that CES markets have equilibria that converge but the Leontief market at the limit does not even have any (approximate) equilibrium.

Moreover, with respect to the problem of determining whether a market equilibrium exists, CES utilities do not behave like the Leontief limit but rather like tractable utilities. Typically, the tractability of the equilibrium existence problem conforms with that of the equilibrium computation problem (under standard sufficient conditions for existence). For example, the existence problem for linear utilities can be solved in polynomial time [Gale, 1976] (as can the computation problem [Jain, 2007]), and the same holds for Cobb-Douglas utilities [Eaves, 1985], whereas the existence problem is NP-hard for Leontief utilities [Codonotti *et al.*, 2006] and for separable piecewise-linear utilities [Vazirani and Yannakakis, 2011] (and their equilibrium computation problem under standard sufficient conditions for existence is PPAD-hard or FIXP-hard [Codonotti *et al.*, 2006; Chen *et al.*, 2009a; Vazirani and Yannakakis, 2011; Garg *et al.*, 2014]). However, the problem of whether there exists an equilibrium in a CES market can be solved in polynomial time for all (finite) values of ρ : a simple necessary and sufficient condition for the existence of an equilibrium in a CES market is based [Codonotti *et al.*, 2005a] on the decomposition of the economy graph into strongly connected components. In the same paper it was also proved that the computation of an equilibrium for the whole market (if the condition is satisfied) amounts to the computation of equilibria for the submarkets induced by the strongly connected components. Hence in this chapter we will focus on CES markets with a strongly connected economy graph.

3.1.3 Main Challenges and Statement of Results

The difficulty in resolving the complexity of the equilibrium computation problem for CES markets with a constant $\rho < -1$ is mainly due to the continuous nature of the problem. Most, if not all, of the problems shown to be PPAD-hard have a rich underlying combinatorial structure, whether it is to find an approximate Nash equilibrium in a normal-form

game [Daskalakis *et al.*, 2009; Chen *et al.*, 2009b] or to compute an approximate equilibrium in a market with Leontief utilities [Codenotti *et al.*, 2006] or with additively separable and concave piecewise-linear utilities [Chen *et al.*, 2009a; Vazirani and Yannakakis, 2011]. In contrast, given a price vector \mathbf{p} , the optimal bundle \mathbf{x} of a CES trader is a continuous function over \mathbf{p} , with an explicit algebraic form. This closed form for the demand implies that the problem of finding a market equilibrium now boils down to solving a system of polynomial equations over variables \mathbf{p} , and it is not clear how to extract a useful combinatorial structure from it.

Our main result in this chapter shows how to address this challenge and settles the complexity of finding approximate equilibria in CES markets for all values of $\rho < -1$:

Theorem 3 (Main Result). *For any rational number $\rho < -1$, ρ -CES-APPROX is PPAD-hard.*

Combining Theorem 3 with the PPAD membership that we also obtain in [Chen *et al.*, 2013], we have

Corollary 1. *For any rational number $\rho < -1$, ρ -CES-APPROX is PPAD-complete.*

3.1.4 Main Ideas behind our Reduction

Example 3.1.1 has a crucial, counter-intuitive, property:

Property 1. *At and around prices $(1, 1)$, although the supply of at least one good is exhausted, increasing the price of this good by a very small value while decreasing the other good's price by the same amount will increase the demand for the most expensive good.*

We will use this property to embed combinatorial structure in a market and simulate a polymatrix game. In more details, given a $2n \times 2n$ polymatrix game \mathbf{P} we construct a market $M_{\mathbf{P}}$ in which each trader owns at most two goods, is interested in one or two goods, and her utility is one of the following functions:

$$u(x) = x, \quad u(x_1, x_2) = (x_1^\rho + x_2^\rho)^{1/\rho}, \quad \text{or} \quad u(x_1, x_2) = (\alpha \cdot x_1^\rho + x_2^\rho)^{1/\rho} \quad (3.2)$$

where α is a positive rational constant that depends on ρ only. We then show that from any ϵ -approximate equilibrium \mathbf{p} of $M_{\mathbf{P}}$, for some polynomially small ϵ , we can recover a $(1/n)$ -well-supported Nash equilibrium in polynomial time.

3.1.4.1 Ingredients of the market

The main building blocks in the construction of $M_{\mathbf{P}}$ are the following:

Main Block. Given a positive rational number μ , we use $\mathbf{CES}(\mu, G_i, G_j)$ to denote the addition to $M_{\mathbf{P}}$ of a sub-market that consists of two traders T_i and T_j . T_i and T_j are only interested in G_i and G_j and have the same utility functions as those of the two traders in example 3.1.1. T_i has μ units of G_i and T_j has μ units of G_j .

Regulating Trader. We use $\mathbf{REG}(G_i, G_j)$ to denote a trader that owns n^4 units of each of goods G_i and G_j and her utility is $u(x_i, x_j) = (x_i^\rho + x_j^\rho)^{1/\rho}$.

Single-minded traders. We say a trader is a $(r, G_i : G_j)$ -trader, if her endowment consists of r units of G_i and she is only interested in G_j .

3.1.4.2 Setting up the market

We will now describe how $M_{\mathbf{P}}$ is built. Let $m = n^7$ be a parameter (we will justify its use when we construct the “gap amplification” part of the market).

Types of goods. The market $M_{\mathbf{P}}$ consists of the following $O(nm) = O(n^8)$ goods:

$$G_{2i-1,j} \text{ and } G_{2i,j}, \quad \text{for } i \in [n] \text{ and } j \in [0 : m].$$

To make the presentation easier to follow, we use G_i and H_i to denote $G_{i,0}$ and $G_{i,m}$ respectively and divide the goods into $n(m+1)$ groups: $\mathcal{R}_{i,j} = \{G_{2i-1,j}, G_{2i,j}\}$, for each $i \in [n]$ and $j \in [0 : m]$. Finally, we use $\pi(G)$ to denote the price of good G . We will omit for now some auxiliary goods that are not crucial for the exposition of our ideas.

Encoding the mixed strategies. We start with the part of the market that encodes 1) a vector \mathbf{x} of $2n$ variables and 2) the $2n$ linear forms $\mathbf{x}^T \cdot \mathbf{P}_{*,j}$, $j \in [2n]$. Those correspond to 1) the mixed strategies and 2) the expected pay-offs of the players in polymatrix. Here is the construction:

1. We use $\pi(H_1), \dots, \pi(H_{2n})$ to encode (a non-normalized version of) \mathbf{x} .

2. To encode the linear forms, we set up the market in such a way that the amount of money spent on good G_j corresponds to $\sum_{i=1}^{2n} \pi(H_i) \mathbf{P}_{i,j}$ which is a non-normalized version of $\mathbf{x}^T \cdot \mathbf{P}_{*,j}$. In particular, for each pair $i, j \in [n]$, we add to $M_{\mathbf{P}}$ the following four traders who trade from group $\mathcal{R}_{i,m}$ to group $\mathcal{R}_{j,0}$: one $(P_{2i-1,2j-1}, H_{2i-1} : G_{2j-1})$ -trader, one $(P_{2i-1,2j}, H_{2i-1} : G_{2j})$ -trader, one $(P_{2i,2j-1}, H_{2i} : G_{2j-1})$ trader, and one $(P_{2i,2j}, H_{2i} : G_{2j})$ -trader.

Consider the amount of money those traders spend on G_j . Since they are single-minded, they spend their whole endowment on G_j , that is, $\sum_{i=1}^{2n} \pi(H_i) \mathbf{P}_{i,j}$; which is what we wanted to achieve.

Converting prices to mixed strategies. We will set up the rest of the market in such a way that, after appropriately scaling the prices, it holds

$$\pi(H_{2i-1}) + \pi(H_{2i}) = 2 + \gamma_i$$

for some very small $\gamma_i > 0$ and all $i \in [n]$. In that case, one can write $\pi(H_{2i-1})$ and $\pi(H_{2i})$ in the form $\pi(H_{2i-1}) = 1 + y_i$ and $\pi(H_{2i}) = 1 - y_i + \gamma_i$, for some (possibly negative) y_i . Then, for a carefully chosen threshold $\theta > 0$, we can extract a vector \mathbf{x} from \mathbf{p} as follows:

$$x_{2i-1} = \begin{cases} 1 & , \text{ if } y_i \geq \theta \\ 0 & , \text{ if } y_i \leq -\theta \\ \frac{\theta + y_i}{2\theta} & , \text{ otherwise} \end{cases}$$

and $x_{2i} = 1 - x_{2i-1}$ for all $i \in [2n]$. From this, it is clear that \mathbf{x} is non-negative and $x_{2i-1} + x_{2i} = 1$ for all $i \in [n]$, as required by polymatrix.

To achieve $\pi(H_{2i-1}) + \pi(H_{2i}) = 2 + \gamma_i$, we add a market $\mathbf{CES}(\mu, H_{2i-1}, H_{2i})$ for an appropriate μ that only depends on n . Even though there are more traders in the market, as long as their total endowment of H_{2i-1} and H_{2i} is negligible compared to μ , block $\mathbf{CES}(\mu, H_{2i-1}, H_{2i})$ dominates the sub-market that contains goods H_{2i-1} and H_{2i} and any equilibrium of the whole market must be very close to one of the equilibria of $\mathbf{CES}(\mu, H_{2i-1}, H_{2i})$. Since the latter are of the form $(1 + \theta, 1 - \theta)$ for some small $|\theta| \geq 0$, any ϵ -approximate market equilibrium of $M_{\mathbf{P}}$ must satisfy $\pi(H_{2i-1}) + \pi(H_{2i}) =$

$2 + \gamma_i$ for a polynomially small γ_i (after appropriately scaling all the prices to achieve $\min_i \{\pi(H_{2i-1}) + \pi(H_{2i})\} = 2$).

Enforcing equilibrium conditions. To complete the market, we must make sure that if for the recovered \mathbf{x} it holds that $\mathbf{x}^T \cdot \mathbf{P}_{*,2j-1} > \mathbf{x}^T \cdot \mathbf{P}_{*,2j} + 1/n$ then $x_{2j} = 0$, or equivalently, $\pi(H_{2i-1}) \geq 1 + \theta$. To achieve this, we are going to enforce the following two conditions for some $0 < \alpha < \theta$:

Condition 1: $\mathbf{x}^T \cdot \mathbf{P}_{*,2j-1} > \mathbf{x}^T \cdot \mathbf{P}_{*,2j} + 1/n$ implies $\pi(G_{2j-1}) \approx 1 + \alpha$.

Condition 2: $\pi(G_{2j-1}) \approx 1 + \alpha$ implies $\pi(H_{2i-1}) \geq 1 + \theta$.

Enforcing Condition 1. Since \mathbf{x} is a normalization of the prices of goods H_k , if we carefully choose θ when recovering \mathbf{x} then whenever $\mathbf{x}^T \cdot \mathbf{P}_{*,2j-1} > \mathbf{x}^T \cdot \mathbf{P}_{*,2j} + 1/n$ it must also hold that $\sum_{i \in [2n]} \pi(H_i) \mathbf{P}_{i,2j-1} > \sum_{i \in [2n]} \pi(H_i) \mathbf{P}_{i,2j} + \frac{1}{\Theta(n)}$. Therefore, the single-minded type- $(P_{i,j}, H_i : G_j)$ -traders spend at least $\frac{1}{\Theta(n)}$ more money on G_{2j-1} than on G_{2j} . Since the supplies of G_{2j-1} and G_{2j} are the same, it follows that in any approximate equilibrium we must have $\pi(G_{2j-1}) \geq \pi(G_{2j}) + \frac{1}{\Theta(n)}$. Hence, all we need is to enforce that in this case $\pi(G_{2j-1}) \approx 1 + \alpha \approx \pi(G_{2j}) + \frac{1}{\Theta(n)}$. For this purpose, we add a **REG** (G_{2i-1}, G_{2i}) trader for all $i \in [n]$. According to equation 3.1, such a trader will buy goods G_{2i-1} and G_{2i} in a ratio of $\left(\frac{\pi(G_{2i-1})}{\pi(G_{2i})}\right)^{\frac{1}{1-\rho}}$ to 1. By definition of ϵ -approximate equilibrium, and because her initial supply dominates that of the rest of the traders, it must therefore hold that

$$1 - \alpha \leq \pi(G_{2j-1}), \pi(G_{2j}) \leq 1 + \alpha \quad (3.3)$$

for a polynomially small $0 < \alpha < \theta$. Combining this with $\pi(G_{2j-1}) \geq \pi(G_{2j}) + \frac{1}{\Theta(n)}$ we get that $\pi(G_{2j-1}) \approx 1 + \alpha$ which is what we wanted.

Enforcing Condition 2. It now remains to enforce that $\pi(G_{2j-1}) \approx 1 + \alpha$ implies $\pi(H_{2i-1}) \geq 1 + \theta$. We achieve that through a sequence of connected sub-markets over groups $\mathcal{R}_{i,j}$. In particular, to finish the construction of the market, we set $m = 4t = n^7$ and for each $i \in [n]$ and $j \in [m]$:

1. We add two traders who trade from group $\mathcal{R}_{i,j-1}$ to group $\mathcal{R}_{i,j}$:
one $(n, G_{2i-1,j-1} : G_{2i-1,j})$ -trader and one $(n, G_{2i,j-1} : G_{2i,j})$ -trader




Figure 3.1: A chain of markets over groups $\mathcal{R}_{i,j}$ of goods. Arrows correspond to $(n, G_i : G_j)$ -traders.

2. We add a market $\mathbf{CES}(\mu, G_{2i-1,j}, G_{2i,j})$

This part of the construction is shown in figure 3.1.

Assume now that for group $\mathcal{R}_{i,j-1}$ it is the case that $\pi(G_{2i-1,j-1}) \approx 1 + \alpha_{j-1}$ and $\pi(G_{2i,j-1}) \approx 1 - \alpha_{j-1}$ for some α_{j-1} . Notice that in any equilibrium $\pi(G_{2i-1,j})$ and $\pi(G_{2i,j})$ must be close to $(1, 1)$, otherwise one of the single-minded traders from $\mathcal{R}_{i,j-1}$ to $\mathcal{R}_{i,j}$ will ask for a large amount of the cheapest of $G_{2i-1,j}$ and $G_{2i,j}$. This is where property 1 comes into play. Since $\pi(G_{2i-1,j})$ and $\pi(G_{2i,j})$ are close to $(1, 1)$, sub-market $\mathbf{CES}(\mu, G_{2i-1,j}, G_{2i,j})$ has two options: Either push towards prices $(1 - \alpha_j, 1 + \alpha_j)$ or towards $(1 + \alpha_j, 1 - \alpha_j)$ for some small $\alpha_j > 0$. If it pushes towards prices $(1 - \alpha_j, 1 + \alpha_j)$ then from property 1 the demand for good $G_{2i-1,j}$ from $\mathbf{CES}(\mu, G_{2i-1,j}, G_{2i,j})$ will go down; but because it will now be cheaper, the demand from the single-minded traders from $\mathcal{R}_{i,j-1}$ to $\mathcal{R}_{i,j}$ will increase far more and equilibrium constraints will be violated. On the other hand, if prices are pushed towards $(1 + \alpha_j, 1 - \alpha_j)$ then, because of property 1 again, the demand from $\mathbf{CES}(\mu, G_{2i-1,j}, G_{2i,j})$ will increase a lot and go above the amount of $G_{2i-1,j}$ it brings in initially. To maintain equilibrium constraints, the demand for $G_{2i-1,j}$ from the single-minded traders from $\mathcal{R}_{i,j-1}$ to $\mathcal{R}_{i,j}$ must decrease significantly below the supply of $G_{2i-1,j}$ that comes from traders from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i,j+1}$. The only way for this to happen is to set $\alpha_{j-1} \ll \alpha_j$ so that $\pi(G_{2i-1,j}) \approx 1 + \alpha_j \gg 1 + \alpha_{j-1} \approx \pi(G_{2i-1,j-1})$. This argument implies the following lemma that is essentially condition 2 and concludes the sketch of our reduction.

Figure 3.2 shows the completed $M\mathbf{P}$.




Figure 3.2: Market $M_{\mathbf{P}}$: Black (solid) arrows correspond to single minded-traders from H_i 's to G_j 's. Blue (dashed) arrows correspond to the chains of $\mathcal{R}_{i,j}$ markets of Figure 3.1.

Lemma 1. *For all $j \in [0 : m]$ there are values $\alpha = \alpha_0 \ll \alpha_1 \ll \dots \ll \alpha_m$, with $\theta \ll \alpha_m$ such that whenever $\pi(G_{2i-1,j}) \approx 1 + \alpha_j$ and $\pi(G_{2i,j}) \approx 1 - \alpha_j$ then it must hold that $\pi(G_{2i-1,j+1}) \approx 1 + \alpha_{j+1}$ and $\pi(G_{2i,j+1}) \approx 1 - \alpha_{j+1}$ in any ϵ -approximate equilibrium.*

3.2 Analysing the Market of Example 3.1.1

In this section we analyse example 3.1.1 and prove many useful properties that play an important role in our reduction. This section is very technical and a reader who only aims to a high-level understanding of our reduction should feel free to proceed to Section 3.3 and use this section as reference for the statement of the lemmas, whenever they are used.

3.2.1 The Excess Spending Function

We need the following notion of *excess spending*. Let S denote a subset of traders. Given \mathbf{p} and a good G , the excess spending on G from traders in S is the product of $\pi(G)$ and the excess demand of G from S :

$$(\text{total demand of } G \text{ from } S - \text{total supply of } G \text{ from } S) \times \pi(G)$$

For convenience we always use $r > 1$ to denote $-\rho$.

Let M denote the following market described in Example 3.1.1 with two goods G_1, G_2 and two traders T_1, T_2 : T_1 has 1 unit of G_1 , T_2 has 1 unit of G_2 , and their utilities are

$$u_1(x_1, x_2) = (\alpha \cdot x_1^\rho + x_2^\rho)^{1/\rho} \quad \text{and} \quad u_2(x_1, x_2) = (x_1^\rho + \alpha \cdot x_2^\rho)^{1/\rho}$$

for some rational number $\alpha > 0$. By (3.1), one can show that given any positive prices π_1 and π_2 , the optimal bundles $(x_{1,1}, x_{1,2})$ and $(x_{2,1}, x_{2,2})$ are unique and must satisfy

$$\frac{x_{1,1}}{x_{1,2}} = \left(\alpha \cdot \frac{\pi_2}{\pi_1} \right)^{1/(1+r)} \quad \text{and} \quad \frac{x_{2,1}}{x_{2,2}} = \left(\frac{1}{\alpha} \cdot \frac{\pi_2}{\pi_1} \right)^{1/(1+r)} \quad (3.4)$$

It is clear that $(1, 1)$ is a market equilibrium of M .

From now on we assume that α is a positive rational number such that $a = \alpha^{1/(r+1)}$ is rational as well. We are interested in the excess spending $f(x)$ on G_1 from T_1 and T_2 when the prices $\pi_1 = 1 + x$ and $\pi_2 = 1 - x$ with $x \in (-1, 1)$. Let $m_{i,j}$ denote the amount of money T_i spends on G_j , then

$$\frac{m_{1,1}}{m_{1,2}} = a \left(\frac{\pi_1}{\pi_2} \right)^{r/(1+r)} \quad \text{and} \quad \frac{m_{2,1}}{m_{2,2}} = \frac{1}{a} \left(\frac{\pi_1}{\pi_2} \right)^{r/(1+r)}$$

We also have $m_{1,1} + m_{1,2} = \pi_1$. This gives us an explicit form of $m_{1,1}$ as a function of x :

$$m_{1,1}(x) = \frac{\pi_1}{1 + \frac{1}{a} \left(\frac{\pi_2}{\pi_1} \right)^{\frac{r}{1+r}}} = \frac{1 + x}{1 + \frac{1}{a} \left(\frac{1-x}{1+x} \right)^{\frac{r}{1+r}}}$$

Similarly, we have the following explicit form of $m_{2,1}$, as a function of x :


$$m_{2,1}(x) = \frac{\pi_2}{1 + a \left(\frac{\pi_2}{\pi_1} \right)^{\frac{r}{1+r}}} = \frac{1 - x}{1 + a \left(\frac{1-x}{1+x} \right)^{\frac{r}{1+r}}}$$

The excess spending function $f(x)$ on G_1 from T_1 and T_2 is then

$$f(x) = m_{1,1}(x) + m_{2,1}(x) - (1 + x), \quad \text{for } x \in (-1, 1).$$

It is easy to show that $f(0) = 0$ and $f(x) = -f(-x)$ for any $x \in (-1, 1)$. By symmetry,

$$f(x) = -f(-x) \Rightarrow f'(x) = f'(-x) \Rightarrow f''(x) = -f''(-x) \Rightarrow f''(0) = 0$$

Figure 3.3: The excess spending function f .

3.2.2 Properties of the Excess Spending function

Our first goal is to prove the following properties about f :

Lemma 3.2.1. *When $a > (r + 1)/(r - 1)$ is rational, $f'(0) > 0$ is also rational, and f has three roots in $(-1, 1)$. Let $\{-\theta, 0, \theta\}$ denote these roots, with $\theta > 0$. Then $f'(\theta) < 0$.*

Proof. First we replace x by the following variable y . Let

$$y^{1+r} = \frac{1-x}{1+x} \quad \text{and} \quad x = \frac{1-y^{1+r}}{1+y^{1+r}} \quad (3.5)$$

It suffices to show that, when $a > (r + 1)/(r - 1)$, the following function $p(y)$ has three roots over $(0, +\infty)$:

$$p(y) = \frac{2}{(1+y^{1+r})(1+y^r/a)} + \frac{2y^{1+r}}{(1+y^{1+r})(1+ay^r)} - \frac{2}{1+y^{1+r}}$$

Let $q(y) = (1+y^{1+r})(1+y^r/a)(1+ay^r) \cdot p(y)$. Then it suffices to show that

$$q(y) = 2(1+ay^r) + 2y^{1+r}(1+y^r/a) - 2(1+y^r/a)(1+ay^r) = \frac{2}{a}y^r(y^{1+r} - ay^r + ay - 1)$$

has three roots. Taking the derivative of $h(y) = y^{1+r} - ay^r + ay - 1$, we get

$$h'(y) = (r+1)y^r - ary^{r-1} + a$$

It is easy to see that $h(0) = -1 < 0$, $h(1) = 0$, and $h(y) \rightarrow +\infty$ when $y \rightarrow +\infty$. Moreover,

$$h'(1) = (r+1) - ar + a = (r+1) - a(r-1) < 0$$

when $a > (r+1)/(r-1)$. This implies that h has at least three roots in $(0, +\infty)$ and thus, f has at least three roots in $(-1, 1)$. Next we show that h has at most three roots. Indeed:

$$h''(y) = r(r+1)y^{r-1} - ar(r-1)y^{r-2} = ry^{r-2}((r+1)y - a(r-1))$$

Therefore, there is a threshold $b = a(r-1)/(r+1) > 0$ such that $h''(y) > 0$ when $y > b$; and $h''(y) < 0$ when $y < b$. This implies that $h'(b)$ is the minimum of h' over $[0, +\infty)$. It follows from $h'(b) \leq h'(1) < 0$ that h' has exactly one root in $(0, b)$ and exactly one root in $(b, +\infty)$. This implies that h has at most three roots in $(0, +\infty)$ and thus, f has at most three roots in $(-1, 1)$. As a result, f has exactly three roots.

Let $\{-\theta, 0, \theta\}$ denote the three roots of f with $\theta > 0$. Then $\{y(-\theta), 1, y(\theta)\}$ are exactly the three roots of h . From the proof we also have $f'(0) > 0$ and $f'(\theta) < 0$. To see this,

$$f(x) = p(y(x)) \Rightarrow f'(x) = p'(y) \cdot \left(\frac{1}{1+r}\right) \cdot \left(\frac{1-x}{1+x}\right)^{\frac{-r}{1+r}} \cdot \frac{-2}{(1+x)^2}$$

This implies that $f'(0) = -2p'(1)/(1+r)$. Taking the derivative of

$$(1+y^{1+r})(1+y^r/a)(1+ay^r) \cdot p(y) = 2y^r h(y)/a$$

and plugging in $h(1) = p(1) = 0$, we get

$$p'(1) = h'(1)/(1+a)^2 < 0$$

and thus, $f'(0) > 0$ is rational. By using $h(y(\theta)) = 0$ and $h'(y(\theta)) > 0$, we can similarly show that $f'(\theta) < 0$. The lemma follows. \square

From now on, we assume that $a > (r+1)/(r-1)$, and let $\{-\theta, 0, \theta\}$ denote the three roots of f over $(-1, 1)$, with $\theta > 0$. Let $\lambda = f'(0)$, which is rational and positive. Let

$$g(x) = f(x) - \lambda x, \quad \text{for } x \in (-1, 1).$$

From the definition of $g(x)$, we have $g(0) = 0$, $g'(0) = 0$, and $g''(0) = 0$.

Next we show that when a is chosen carefully, g satisfies the following property:

Lemma 3.2.2. *Given any rational number $r > 1$, there is a rational number a such that $a > (r+1)/(r-1)$, $\alpha = a^{1+r}$ is rational, and $g(x) < 0$ for all $x \in (0, 1)$. From the symmetry of g , $g(x) > 0$ for all $x \in (-1, 0)$.*




Figure 3.4: The function g and the line $-\lambda x$, where $\lambda = f'(0)$.

Proof. Assume for contradiction that there is an $x^* \in (0, 1)$ such that $g(x^*) \geq 0$, and $g(-x^*) \leq 0$. Similar to the proof of Lemma 3.2.1, we use y in (3.5) to replace x . We are interested in $p(y)$ over $y \in (0, +\infty)$:

$$p(y) = \frac{2}{(1+y^{1+r})(1+y^r/a)} + \frac{2y^{1+r}}{(1+y^{1+r})(1+ay^r)} - \frac{2}{1+y^{1+r}} - \lambda \cdot \frac{1-y^{1+r}}{1+y^{1+r}}$$

By the definition of $p(y)$, we have

$$g(x) = p(y(x)) \Rightarrow p(1) = 0, p'(1) = 0 \text{ and } p''(1) = 0 \quad (3.6)$$

using the chain rule as well as the fact that $y'(x)$ is nonzero at $x = 0$. Let $y_1 = y(x^*)$ and $y_2 = y(-x^*)$. Then we have $0 < y_1 < 1 < y_2$, $p(y_1) \geq 0$ and $p(y_2) \leq 0$. Next we use $q(y)$ to denote the following function:

$$q(y) = (1+y^{1+r})(1+y^r/a)(1+ay^r) \cdot p(y)/2.$$

Then we have

$$q(y) = (1+ay^r) + y^{1+r}(1+y^r/a) - (1+y^r/a)(1+ay^r) - (\lambda/2)(1-y^{1+r})(1+y^r/a)(1+ay^r)$$

By the definition of $q(y)$, we have $q(y_1) \geq 0$ and $q(y_2) \leq 0$. We use $u, v, w > 0$ to denote

$$u = \frac{\lambda}{2}, \quad v = \frac{a\lambda}{2} + \frac{\lambda}{2a} + \frac{1}{a} \quad \text{and} \quad w = 1 + \frac{\lambda}{2}$$

then we can rewrite $q(y)$ as follows:

$$q(y) = u \cdot y^{1+3r} + v \cdot y^{1+2r} - w \cdot y^{2r} + w \cdot y^{1+r} - v \cdot y^r - u$$

Taking its derivative, we get

$$q'(y) = u(1 + 3r) \cdot y^{3r} + v(1 + 2r) \cdot y^{2r} - 2wr \cdot y^{2r-1} + w(1 + r) \cdot y^r - vr \cdot y^{r-1}$$

Let $q'(y) = y^{r-1} \cdot s(y)$, then we have

$$s(y) = u(1 + 3r) \cdot y^{1+2r} + v(1 + 2r) \cdot y^{1+r} - 2wr \cdot y^r + w(1 + r) \cdot y - vr$$

Taking its derivative, we get

$$s'(y) = u(1 + 3r)(1 + 2r) \cdot y^{2r} + v(1 + 2r)(1 + r) \cdot y^r - 2wr^2 \cdot y^{r-1} + w(1 + r) \quad (3.7)$$

and its second-order derivative

$$s''(y) = 2ur(1 + 3r)(1 + 2r) \cdot y^{2r-1} + vr(1 + 2r)(1 + r) \cdot y^{r-1} - 2wr^2(r - 1) \cdot y^{r-2}$$

Let $s''(y) = y^{r-2} \cdot t(y)$, then we have

$$t(y) = 2ur(1 + 3r)(1 + 2r) \cdot y^{r+1} + vr(1 + 2r)(1 + r) \cdot y - 2wr^2(r - 1) \quad (3.8)$$

We prove some useful properties about these functions. First, we show that $s''(1)$ is indeed positive when a is close enough to $(r + 1)/(r - 1)$. By (3.7), we have

$$s''(1) = 2ur(1 + 3r)(1 + 2r) + vr(1 + 2r)(1 + r) - 2wr^2(r - 1)$$

Let $c = a + 1/a$. Plugging in $v = cu + 1/a$ and $w = 1 + u$, we have

$$s''(1) = 2ur(1 + 3r)(1 + 2r) + (cu + 1/a)(1 + 2r)(1 + r)r - 2(1 + u)r^2(r - 1)$$

The trouble here is that λ (and u) depends on the choice of a . But note that the coefficient of u in $s''(1)$ is

$$2r(1 + 3r)(1 + 2r) + cr(1 + 2r)(1 + r) - 2r^2(r - 1) > 0$$

and u is positive when $a > (r + 1)/(r - 1)$. The rest of $s''(1)$ is the following:

$$(1/a)(1 + 2r)(1 + r)r - 2r^2(r - 1)$$

Let $a = (1 + \epsilon)(r + 1)/(r - 1)$. When ϵ goes to 0, the expression above converges to

$$r(r - 1)(1 + 2r) - 2r^2(r - 1) = r(r - 1)(1 + 2r - 2r) = r(r - 1) > 0$$

Therefore there exists a positive rational number $a > (r + 1)/(r - 1)$ such that $s''(1) > 0$ and $\alpha = a^{1+r}$ is rational (note that we do not care about the number of bits needed to encode it). We use such an a from now on. From the definition of q and s from p as well as the chain rule, one can show that $p(1) = p'(1) = p''(1) = 0$ (equation 3.6) implies that $q(1) = q'(1) = q''(1) = 0$ and $s(1) = s'(1) = 0$; Furthermore, since $s''(1) > 0$ we have $p'''(1) > 0$. Together with (3.6), we know there is a small enough $\epsilon > 0$ that satisfies:

$$p(1 + \epsilon) > 0, \quad p(1 - \epsilon) < 0 \quad \text{and} \quad y_1 < 1 - \epsilon < 1 + \epsilon < y_2$$

Recall that $p(y_1) \geq 0$ and $p(y_2) \leq 0$. By the definition of $q(y)$ from $p(y)$, we have

$$q(y_1) \geq 0, \quad q(1 - \epsilon) < 0, \quad q(1 + \epsilon) > 0 \quad \text{and} \quad q(y_2) \leq 0 \quad (3.9)$$

In the rest of the proof we show that this cannot happen.

First it is easy to check that $t(0) < 0$; $t(y) > 0$ when $y \rightarrow +\infty$; and $t'(y) > 0$ for any $y > 0$. This shows that there is a unique $b \in (0, \infty)$ such that $t(y) < 0$ for any $y < b$, $t(b) = 0$, and $t(y) > 0$ for any $y > b$. Using $s''(y) = y^{r-2} \cdot t(y)$, the same holds for $s''(y)$.

Now we examine $s'(y)$. Note that $s'(0) > 0$ and $s'(y) > 0$ when $y \rightarrow \infty$. It follows from the property of $s''(y)$ that going from $y = 0$ to $+\infty$, the sign of $s'(y)$ can change at most twice from positive to negative and then back to positive.

Finally regarding $s(y)$, we have $s(0) < 0$ and $s(y) > 0$ when $y \rightarrow +\infty$. By the property of $s'(y)$ we know $s(y)$ can have at most three roots in $(0, +\infty)$. From $q'(y) = y^{r-1} \cdot s(y)$, the same statement also holds for $q'(y)$. However, this contradicts with (3.9) because

1. From $q(0) < 0$ and $q(y_1) \geq 0$, there exists a $y \in (0, y_1)$ such that $q'(y) > 0$;
2. From $q(y_1) \geq 0$ and $q(1 - \epsilon) < 0$, there exists a $y \in (y_1, 1 - \epsilon)$ such that $q'(y) < 0$;
3. From $q(1 - \epsilon) < 0$ and $q(1 + \epsilon) > 0$, there exists a $y \in (1 - \epsilon, 1 + \epsilon)$ such that $q'(y) > 0$;
4. From $q(1 + \epsilon) > 0$ and $q(y_2) \leq 0$, there exists a $y \in (1 + \epsilon, y_2)$ such that $q'(y) < 0$;
5. From $q(y_2) \leq 0$ and $q(y) > 0$ when $y \rightarrow +\infty$, there exists a $y \in (y_2, +\infty)$ such that $q'(y) > 0$.

It follows that $q'(y)$ has at least four roots in $(0, +\infty)$, a contradiction. \square

From now on we always assume that a is positive and rational such that $\alpha = a^{1+r}$ is rational, f satisfies conditions of Lemma 3.2.1, and g satisfies conditions of Lemma 3.2.2. We remind the reader that λ denotes $f'(0)$, a positive rational number. We also use θ to denote the positive root of f . While θ is not rational in general, we can use f (and h in the proof of Lemma 3.2.1) to compute a γ -rational approximation θ^* of θ , i.e. $|\theta^* - \theta| \leq \gamma$, in time polynomial in $1/\gamma$. Let σ be $f'(\theta) < 0$. The following corollaries follow from Lemma 3.2.1 and 3.2.2.

Corollary 3.2.1. *We have $g(x) < -\lambda x < -\lambda\theta$ for any $x \in (\theta, 1)$; $g(x) > -\lambda x > -\lambda\theta$ for any $x \in (0, \theta)$.*

Proof. By Lemma 3.2.1 we have $f(x) < 0$ for any $x \in (\theta, 1)$ thus, $g(x) = f(x) - \lambda x < -\lambda x$. By Lemma 3.2.1 we have $f(x) > 0$ for any $x \in (0, \theta)$ thus, $g(x) = f(x) - \lambda x > -\lambda x$. \square

Corollary 3.2.2. *$g(\theta) = -\lambda\theta$ and $g'(\theta) = \sigma - \lambda < -\lambda$, where $\sigma = f'(\theta)$.*

Corollary 3.2.3. *There exists a positive constant c such that for any $x \in [-c, c]$:*

$$|f(x) - \lambda x| \leq |\lambda x/2| \quad \text{and} \quad |f(\theta + x) - \sigma x| \leq |\sigma x/2|$$

Given a sufficiently large positive integer N , we are interested in f and g over:

$$A_N = [-\delta, \delta], \quad B_N = [\delta, \theta - \delta], \quad C_N = [\theta - \delta, \theta + \delta] \quad (3.10)$$

$$B'_N = [-\theta + \delta, -\delta], \quad C'_N = [-\theta - \delta, -\theta + \delta] \quad \text{and} \quad S_N = [-\theta - \delta, \theta + \delta]$$

where $\delta = 1/N$. We use Lemma 3.2.1 and Lemma 3.2.2 to prove the following lemmas:

Lemma 3.2.3. *When N is sufficiently large, we have $|g(x)| \leq |\lambda x/2|$ for any $x \in A_N$.*

Proof. The lemma follows directly from the first part of Corollary 3.2.3. \square

Lemma 3.2.4. *When N is sufficiently large, $f(x) \geq \min(\lambda, |\sigma|)\delta/2$ for all $x \in B_N$.*

Proof. Assume for contradiction that this is not the case, meaning that there is an infinite sequence of N and x_N such that $x_N \in B_N$ but $f(x_N) < \min(\lambda, |\sigma|)\delta/2$. As $x_N \in [0, \theta]$ is compact, there is a subsequence of x_N that converges to a root x^* of f in $[0, \theta]$. As 0 and θ are the only nonnegative roots of f , $x^* = 0$ or θ . But no matter which case it is, the derivative of f at x^* is smaller than we expect and we get a contradiction. \square

Using Lemma 3.2.4, we prove the following lemma:

Lemma 3.2.5. *Assume that N is sufficiently large. If*

$$g(x) = -\lambda\theta \pm \Delta$$

where $\Delta = \delta(\lambda - \sigma/2)$, then we must have that $x \in C_N$.

Proof. First $g(x) < 0$ when N is sufficiently large. From Lemma 3.2.2 we have $x > 0$. Replacing x by $\theta + y$, we have

$$f(\theta + y) - \lambda(\theta + y) = -\lambda\theta \pm \Delta \Rightarrow f(\theta + y) = \lambda y \pm \Delta$$

As $f(\theta + y) < 0$ when $y > 0$, and $f(\theta + y) > 0$ when $y < 0$ (and $x = \theta + y > 0$), we have $|y| < \Delta/\lambda$ and thus Corollary 3.2.3 applies when N is sufficiently large: If $y > 0$ we have

$$3\sigma y/2 \leq \lambda y \pm \Delta = f(\theta + y) \leq \sigma y/2,$$

which implies $0 < y \leq \Delta/(\lambda - \sigma/2) = \delta$. The case when $y < 0$ is similar. \square

Note that by the symmetry of f and g , similar lemmas can be proved for B'_N, C'_N .

3.3 Markets with CES Utilities are PPAD-hard

In this section we prove Theorem 3. Let $\rho < -1$ be a fixed rational number, and let $r = |\rho|$. Given any normalized $2n \times 2n$ polymatrix game \mathbf{P} , we construct a market $M_{\mathbf{P}}$ in which each trader has a CES utility function of parameter ρ .

3.3.1 Our Construction

The main building block in the construction is the following:

CES Market Block: We use M to denote the CES market discussed in Example 3.1.1 and Section 3.2, with rational constants α and a satisfying all conditions of Lemma 3.2.1 and Lemma 3.2.2. We use the following notation. Given a positive rational number μ , we use $\text{CES}(\mu, G_1, G_2)$ to denote the addition to $M_{\mathbf{P}}$ of a sub-market that consists of two traders T_i and T_j . T_1 and T_2 are only interested in G_1 and G_2 and have the same utility functions

as those of the two traders in M . T_1 has μ units of G_1 and T_2 has μ units of G_2 . We let $f_\mu(x)$ denote the excess spending function (see section 3.2) on G_1 from these two traders when the prices of G_1 and G_2 are $1+x$ and $1-x$. Then $f_\mu(x) = \mu \cdot f(x)$.

Recall $\lambda = f'(0)$ is positive and rational, θ is the positive root of f , and $\sigma = f'(\theta) < 0$. Let $m = n^7$.

Construction of $M_{\mathbf{P}}$. The market $M_{\mathbf{P}}$ consists of the following $O(nm) = O(n^8)$ goods:

$$AUX_i, G_{2i-1,j} \text{ and } G_{2i,j}, \quad \text{for } i \in [n] \text{ and } j \in [0 : m].$$

We divide the goods into $n(m+1)$ groups: $\mathcal{R}_{i,j} = \{G_{2i-1,j}, G_{2i,j}\}$, $i \in [n]$ and $j \in [0 : m]$.

First for each $i \in [n]$, we add a trader with $\tau = n^4$ units of $G_{2i-1,0}$ and $G_{2i,0}$ each, and set her utility to be

$$u(x_1, x_2) = (x_1^\rho + x_2^\rho)^{1/\rho}$$

where x_1 (or x_2) denotes the amount of $G_{2i-1,0}$ (or $G_{2i,0}$, respectively) she obtains.

Next for each $\mathcal{R}_{i,j}$, $i \in [n]$ and $j \in [m]$, we create a market $\mathbf{CES}(\mu, G_{2i-1,j}, G_{2i,j})$ with $\mu = n/\lambda$.

Now we add a number of single-minded traders who trade between different groups. We say a trader is a $(r, G_1 : G_2)$ -trader, if her endowment consists of r units of G_1 and she is only interested in G_2 . We say a trader is a $(r, G_1, G_2 : G_3)$ -trader, if her endowment consists of r units of G_1 and G_2 each, and she is only interested in G_3 .

At the same time we construct a weighted directed graph $\mathcal{G} = (V, E)$ which will be used in the proof of correctness only. Here each group of goods $\mathcal{R}_{i,j}$ corresponds to a vertex in the graph \mathcal{G} so $|V| = n(m+1)$. Given two groups $\mathcal{R}_{i,j}$ and $\mathcal{R}_{i',j'}$, we add an edge from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i',j'}$ in \mathcal{G} whenever we create a set of traders who trade from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i',j'}$.

Our construction below always makes sure that, whenever we create a set of traders who trade from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i',j'}$, the total initial endowment of these traders consists of the same amount, say $w > 0$, of $G_{2i-1,j}$ and $G_{2i,j}$. We then set w as the weight of this edge. We will prove by the end of the construction that \mathcal{G} is a strongly connected graph and for each group $\mathcal{R}_{i,j}$ the total in-weight is the same as its total out-weight.

Here is the construction:

1. For each $i \in [2n]$, we use G_i to denote $G_{i,0}$ and H_i to denote $G_{i,m}$ for convenience. For each pair $i, j \in [n]$, we add to $M_{\mathbf{P}}$ the following four traders who trade from group $\mathcal{R}_{i,m}$ to group $\mathcal{R}_{j,0}$: one $(P_{2i-1,2j-1}, H_{2i-1} : G_{2j-1})$ -trader, one $(P_{2i-1,2j}, H_{2i-1} : G_{2j})$ -trader, one $(P_{2i,2j-1}, H_{2i} : G_{2j-1})$ -trader, and one $(P_{2i,2j}, H_{2i} : G_{2j})$ -trader. Since \mathbf{P} is normalized, we have

$$P_{2i-1,2j-1} + P_{2i-1,2j} = P_{2i,2j-1} + P_{2i,2j} = 1$$

Thus, the total endowment of these four traders consists of one unit of H_{2i-1} and H_{2i} each, so we add an edge in \mathcal{G} from $\mathcal{R}_{i,m}$ to $\mathcal{R}_{j,0}$ with weight 1. At this moment, the total out-weight of each $\mathcal{R}_{i,m}$ in \mathcal{G} (a complete bipartite graph) is n , and the total in-weight of each $\mathcal{R}_{j,0}$ in \mathcal{G} is n .

2. Next for each $i \in [n]$ and $j \in [m]$, we add two traders who trade from group $\mathcal{R}_{i,j-1}$ to group $\mathcal{R}_{i,j}$: one $(n, G_{2i-1,j-1} : G_{2i-1,j})$ -trader and one $(n, G_{2i,j-1} : G_{2i,j})$ -trader. As their total endowment consists of n units of $G_{2i-1,j-1}$ and $G_{2i,j-1}$ each, we add an edge from $\mathcal{R}_{i,j-1}$ to $\mathcal{R}_{i,j}$ of weight n .

This finishes the construction of \mathcal{G} . It is also easy to verify that \mathcal{G} is strongly connected and each vertex has both its total in-weight and out-weight equal to n .

Finally, we add traders between AUX_j and $\mathcal{R}_{j,0}$ for each $j \in [n]$. Let

$$r_{2j-1} = 2n - \sum_{i \in [2n]} P_{i,2j-1} > 0 \quad \text{and} \quad r_{2j} = 2n - \sum_{i \in [2n]} P_{i,2j} > 0 \quad (3.11)$$

Because the polymatrix game \mathbf{P} is normalized, note that

$$r_{2j-1} + r_{2j} = 2n, \quad \text{for any } j \in [n].$$

Let θ^* denote a γ -rational approximation of θ (see definition 7), the positive root of f , where $\gamma = 1/n^7$. Then we add the following three traders: one $((1 - \theta^*)r_{2j-1}, AUX_j : G_{2j-1})$ -trader, one $((1 - \theta^*)r_{2j}, AUX_j : G_{2j})$ trader, and one $((1 - \theta^*)n, G_{2j-1}, G_{2j} : AUX_j)$ -trader. Note that $r_{2j-1} + r_{2j} = 2n$ as \mathbf{P} is normalized.

This finishes the construction of $M_{\mathbf{P}}$. It follows immediately from the strong connectivity of \mathcal{G} that the economy graph of $M_{\mathbf{P}}$ is strongly connected as well. Thus, $M_{\mathbf{P}}$ is a valid input

of problem ρ -CES-APPROX and can be constructed from \mathbf{P} in polynomial time. We also record the following properties of $M_{\mathbf{P}}$:

Lemma 3.3.1. *For each $i \in [n]$, the total supply of good AUX_i is $2n(1 - \theta^*)$;*

For each $i \in [2n]$, the total supply of good $G_{i,0}$ is $\tau + (2 - \theta^)n$; and*

For each $i \in [2n]$ and $j \in [m]$, the total supply of good $G_{i,j}$ is $\mu + n = \Theta(n)$.

Finally, we give some intuition about the choice of $m = n^7$ here. The key challenge for the reduction is to make sure that in any approximate equilibrium, a gap between the prices of $G_{2i-1,0}$ and $G_{2i,0}$ gets amplified in prices of $G_{2i-1,m}$ and $G_{2i,m}$. More precisely, whenever the ratio of the price of $G_{2i-1,0}$ to that of $G_{2i,0}$ is large (or small), the ratio of the price of $G_{2i-1,m}$ to that of $G_{2i,m}$ must be even larger (or smaller). This is achieved in our construction by $m = n^7$ rounds of minor amplifications, from $G_{2i-1,j}, G_{2i,j}$ to $G_{2i-1,j+1}, G_{2i,j+1}$, for each $j \in [0 : m]$.

3.3.2 Proof of Correctness

We introduce additively approximate market equilibria to simplify the presentation:

Definition 10. *We say \mathbf{p} is an ϵ -additively approximate market equilibrium of a market M , for some $\epsilon \geq 0$, if there exists a vector $\mathbf{z} \in Z(\mathbf{p})$ such that $z_j \leq \epsilon$ for all j .*

From the definitions, if the total supply of each good in M is bounded from above by L , then any ϵ -approximate equilibrium of M must be an (ϵL) -additively approximate equilibrium as well.

Now let \mathbf{p} denote an ϵ -additively approximate market equilibrium of $M_{\mathbf{P}}$ where $\epsilon = 1/n^{14}$. We show in the rest of this section that given \mathbf{p} , one can compute a $(1/n)$ -well-supported Nash equilibrium of \mathbf{P} efficiently in polynomial time. Theorem 3 then follows. In the proof below we use $\pi(G)$ to denote the price of a good G in the price vector \mathbf{p} . For each $\mathcal{R}_{i,j}$, we let $\pi_{i,j} = \pi(G_{2i-1,j}) + \pi(G_{2i,j})$. Finally, we use $a = b \pm c$, where $c > 0$, to denote the inequality $b - c \leq a \leq b + c$.

Moving on with the proof, note that only one trader is interested in AUX_j and thus we obtain the following lower bound.

Lemma 3.3.2. *Let \mathbf{p} denote an ϵ -additively approximate equilibrium of $M_{\mathbf{P}}$, where $\epsilon = 1/n^{14}$. If we scale \mathbf{p} so that $\pi_{j,0} = \pi(G_{2j-1}) + \pi(G_{2j}) = 2$ for some $j \in [n]$, then we have $\pi(AUX_j) \geq 1 - O(\epsilon/n)$.*

Second, by using the strong connectivity of \mathcal{G} and the property that every vertex in \mathcal{G} has the same total in-weight and out-weight, we get the following lemma:

Lemma 3.3.3. *Let \mathbf{p} denote an ϵ -additively approximate equilibrium of $M_{\mathbf{P}}$. Let*

$$\pi_{\max} = \max_{i,j} \pi_{i,j} \quad \text{and} \quad \pi_{\min} = \min_{i,j} \pi_{i,j}$$

both over $i \in [n]$ and $j \in [0 : m]$. If we scale \mathbf{p} so that $\pi_{\min} = 2$, then $\pi_{\max} = 2 + O(m\epsilon)$.

Proof. For convenience, we use u and v to denote vertices (groups) in \mathcal{G} . For each u in \mathcal{G} , we use π_u to denote $\pi_{i,j}$ if u corresponds to $\mathcal{R}_{i,j}$. An edge from u to v of weight w means traders from u to v spend $w\pi_u$ on v .

Now fix a vertex v and let \mathcal{R} denote its corresponding group of goods. As \mathbf{p} is an ϵ -approximate market equilibrium, we must have

$$\text{total money spent on goods in } \mathcal{R} - \text{total worth of goods in } \mathcal{R} \leq O(\epsilon\pi_v) \quad (3.12)$$

For those traders in the closed economy over \mathcal{R} , the money they spend on \mathcal{R} is equal to the total worth of their initial endowments of \mathcal{R} . So they cancel each other in (3.12). Below we enumerate all other traders in $M_{\mathbf{P}}$ who either own goods in \mathcal{R} at the beginning or are interested in goods in \mathcal{R} :

1. Let $N^-(v)$ denote the set of predecessors of v . Then for each $u \in N^-(v)$, the amount of money that traders from u to v spend on \mathcal{R} is $w_{u,v} \cdot \pi_u$, where $w_{u,v}$ denotes the weight of edge (u, v) .
2. Let $N^+(v)$ denote the set of successors of v . Then for each $u \in N^+(v)$, the total worth of goods in \mathcal{R} owned by traders from v to u at the beginning is $w_{v,u} \cdot \pi_v$.
3. For the special case when $\mathcal{R} = \mathcal{R}_{j,0}$ for some $j \in [n]$, we have three more traders: one $((1 - \theta^*)r_{2j-1}, AUX_j : G_{2j-1})$ -trader, one $((1 - \theta^*)r_{2j}, AUX_j : G_{2j})$ -trader, and one $((1 - \theta^*)n, G_{2j-1}, G_{2j} : AUX_j)$ -trader.

Since these are all the traders in $M_{\mathbf{P}}$ relevant to goods in \mathcal{R} , from Lemma 3.3.2 and (3.12),

$$\sum_{u \in N^-(v)} w_{u,v} \cdot \pi_u - \sum_{u \in N^+(v)} w_{v,u} \cdot \pi_v \leq O(\epsilon \pi_v), \quad \text{for each } v \in V. \quad (3.13)$$

Now we use (3.13) to prove the lemma:

1. First, each group $\mathcal{R}_{i,j}$, where $i \in [n]$ and $j \in [m-1]$, has exactly one predecessor $\mathcal{R}_{i,j-1}$ and one successor $\mathcal{R}_{i,j+1}$, both with weight n . From (3.13), we have

$$\pi_{i,j-1} - \pi_{i,j} \leq O(\epsilon \pi_{i,j}/n), \quad \text{for all } i \in [n] \text{ and } j \in [m-1]. \quad (3.14)$$

2. Next, each group $\mathcal{R}_{i,m}$, where $i \in [n]$, has only one predecessor $\mathcal{R}_{i,m-1}$ with weight n , and n successors each with weight 1. From (3.13), we have

$$\pi_{i,m-1} - \pi_{i,m} \leq O(\epsilon \pi_{i,m}/n), \quad \text{for all } i \in [n]. \quad (3.15)$$

3. Finally, each group $\mathcal{R}_{i,0}$, where $i \in [n]$, has n predecessors $\{\mathcal{R}_{\ell,m}\}_{\ell \in [n]}$, all of weight 1, and has one successor $\mathcal{R}_{i,1}$ with weight n . From (3.13), we have

$$\sum_{\ell \in [n]} \pi_{\ell,m} - n\pi_{i,0} \leq O(\epsilon \pi_{i,0}), \quad \text{for all } i \in [n]. \quad (3.16)$$

Let $\pi_{i,j} = \pi_{\min} = 2$ after scaling and $\pi_{x,y} = \pi_{\max}$. Using (3.14) and (3.15), we have

$$\pi_{i,0} \leq (1 + O(\epsilon/n))^m \cdot \pi_{i,j} = 2(1 + O(\epsilon m/n)) = 2 + O(\epsilon m/n),$$

where we used the fact that $\epsilon m/n = 1/n^8 \ll 1$. Similarly, we also have

$$\pi_{x,m} \geq (1 + O(\epsilon/n))^{-m} \cdot \pi_{\max} \geq (1 - O(\epsilon m/n)) \pi_{\max}.$$

Combining these two bounds with (3.16), we get

$$(n + O(\epsilon))(2 + O(\epsilon m/n)) \geq (n + O(\epsilon))\pi_{i,0} \geq \sum_{\ell \in [n]} \pi_{\ell,m} \geq 2(n-1) + (1 - O(\epsilon m/n))\pi_{\max}$$

Solving it for π_{\max} gives us

$$\pi_{\max} \leq \frac{2n + O(\epsilon) + O(\epsilon m) + O(\epsilon^2 m/n) - 2(n-1)}{1 - O(\epsilon m/n)} = \frac{2 + O(\epsilon m)}{1 - O(\epsilon m/n)} = 2 + O(\epsilon m).$$

This finishes the proof of the lemma. \square

We can now prove the following bound on $\pi(AUX_j)$:

Lemma 3.3.4. *Let \mathbf{p} denote an ϵ -additively approximate equilibrium of $M_{\mathbf{P}}$ with $\epsilon = 1/n^{14}$. If we scale \mathbf{p} so that $\pi_{j,0} = 2$ for some $j \in [n]$, then we have $\pi(AUX_j) \leq 1 + O(m\epsilon)$.*

Proof. We revisit (3.12). Let v denote the vertex that corresponds to $\mathcal{R}_{j,0}$.

Plugging in (3.12) the list of traders enumerated in the proof of Lemma 3.3.3, we have

$$\sum_{\ell \in [n]} \pi_{\ell,m} + 2n(1 - \beta) \cdot \pi(AUX_j) - n\pi_{j,0} - (1 - \beta)n\pi_{j,0} \leq O(\epsilon\pi_{j,0})$$

The lemma then follows directly from Lemma 3.3.3. \square

From now on, we use $x_{i,j}$ to denote the unique number that satisfies

$$\frac{1 + x_{i,j}}{1 - x_{i,j}} = \frac{\pi(G_{2i-1,j})}{\pi(G_{2i,j})}, \quad \text{for each } i \in [n] \text{ and } j \in [0 : m].$$

Note that the $x_{i,j}$'s are invariant under scaling of \mathbf{p} . If we scale \mathbf{p} so that $\pi_{i,j} = 2$, for some i and j , then we must have $\pi(G_{2i-1,j}) = 1 + x_{i,j}$ and $\pi(G_{2i,j}) = 1 - x_{i,j}$. Moreover, even if we scale \mathbf{p} so that the sum of prices of another group becomes 2, we still have the following estimations by Lemma 3.3.2, 3.3.3 and 3.3.4:

$$\pi(G_{2i-1,j}) = 1 + x_{i,j} \pm O(m\epsilon), \quad \pi(G_{2i,j}) = 1 - x_{i,j} \pm O(m\epsilon) \quad \text{and} \quad \pi(AUX_j) = 1 \pm O(m\epsilon) \quad (3.17)$$

Next we show that $x_{i,0}$ must be very close to 0 for all $i \in [n]$.

Lemma 3.3.5. *If \mathbf{p} is an ϵ -additively approximate equilibrium, then $|x_{i,0}| = O(1/n^3)$ for all $i \in [n]$.*

Proof. Fix an $i \in [n]$. We first scale \mathbf{p} so that $\pi_{i,0} = 2$, and use x to denote $x_{i,0}$. Let T denote the trader with τ units of G_{2i-1} and G_{2i} each. We let y_1 denote the demand of G_{2i-1} , and y_2 to denote the demand of G_{2i} from T . Then

$$y_1(1 + x) + y_2(1 - x) = 2\tau$$

and by (3.4) we have

$$\frac{y_1}{y_2} = \left(\frac{1 - x}{1 + x} \right)^{1/(1+r)}$$

Assume without loss of generality $x > 0$, we will show that $x = O(1/n^3)$. To this end,

$$y_2 = \frac{2\tau}{(1-x) + (1-x)^{1/(1+r)}(1+x)^{r/(1+r)}} \leq \tau + O(n),$$

which follows from \mathbf{p} being an additively approximate equilibrium. It implies that

$$(1-x)^{1/(1+r)} \geq (1+x)^{1/(1+r)} - O(1/n^3) > 1 - O(1/n^3).$$

Since r is a positive constant, we have $x = O(1/n^3)$ and the lemma follows. \square

From now on we set $N = n^6$. Recall the definition of $A_N, B_N, C_N, B'_N, C'_N, S_N$ in (3.10). Using Lemma 3.3.5, we have $x_{i,0} \in S_N$. Next we show that $x_{i,j} \in S_N$ for all i and j .

Lemma 3.3.6. *If \mathbf{p} is an ϵ -additively approximate equilibrium, then $x_{i,j} \in S_N$ for all $i \in [n]$ and $j \in [m]$.*

Lemma 3.3.6 follows directly from the following three lemmas by induction:

Lemma 3.3.7. *For any $i \in [n]$ and $j \in [m]$, if $x_{i,j-1} \in A_N$, then $x_{i,j} \in A_N \cup B_N \cup B'_N$.*

Lemma 3.3.8. *For any $i \in [n]$ and $j \in [m]$, if $x_{i,j-1} \in B_N$, then $x_{i,j} \in B_N \cup C_N$; and if $x_{i,j-1} \in B'_N$ then $x_{i,j} \in B'_N \cup C'_N$.*

Lemma 3.3.9. *For any $i \in [n]$ and $j \in [m]$, if $x_{i,j-1} \in C_N$, then $x_{i,j} \in C_N$; and if $x_{i,j-1} \in C'_N$, then $x_{i,j} \in C'_N$.*

Proof of Lemma 3.3.7. First we scale \mathbf{p} so that $\pi_{i,j} = 2$. We use x to denote $x_{i,j}$ so that the prices of $\pi(G_{2i-1,j})$ and $\pi(G_{2i,j})$ are $1+x$ and $1-x$, respectively. We also let

$$y = x_{i,j-1}, \quad \pi(G_{2i-1,j-1}) = 1+y_1 \quad \text{and} \quad \pi(G_{2i,j-1}) = 1-y_2.$$

From Lemma 3.3.3, y_1 and y_2 are both $y \pm O(m\epsilon)$. The excess spending of $G_{2i-1,j}$ of the whole market is

$$\mu \cdot f(x) + n(1+y_1) - n(1+x) = n(1/\lambda)(f(x) - \lambda x + \lambda y_1) = n(1/\lambda)(g(x) + \lambda y_1) \quad (3.18)$$

while the excess spending of $G_{2i,j}$ of the whole market is

$$-\mu \cdot f(x) + n(1-y_2) - n(1-x) = -n(1/\lambda)(g(x) + \lambda y_2) \quad (3.19)$$

As $\pi_{i,j} = 2$ and \mathbf{p} is an ϵ -additively approximate equilibrium, both (3.18) and (3.19) are at most $O(\epsilon)$. As a result, we have

$$|n(1/\lambda)(g(x) + \lambda y)| = O(nm\epsilon) \Rightarrow |g(x) + \lambda y| = O(m\epsilon) \quad (3.20)$$

since λ is a positive constant. As $|y| = |x_{i,j-1}| \leq 1/N = 1/n^6$ and $m\epsilon = 1/n^7$, we have $|g(x)| = O(1/N)$. The lemma now follows from Corollary 3.2.1 and 3.2.3. \square

Proof of Lemma 3.3.8. Assume $x_{i,j-1} \in B_N$; the proof of the other case is similar. Using the same notation and argument of Lemma 3.3.7, we start with (3.20) and get

$$g(x) \geq -\lambda y - O(m\epsilon) \geq -\lambda(\theta - 1/N + O(m\epsilon)) > -\lambda\theta,$$

where the second inequality used $y = x_{i,j-1} \in B_N$ and thus, $y \leq \theta - 1/N$. We also have

$$g(x) \leq -\lambda y + O(m\epsilon) \leq -\lambda/N + O(m\epsilon) < 0,$$

where the second inequality used $y = x_{i,j-1} \geq 1/N$. By Corollary 3.2.1 $x \in A_N \cup B_N \cup C_N$.

Assume for contradiction that $x \in A_N$. Then by Corollary 3.2.3 we have

$$-\lambda/N + O(m\epsilon) \geq g(x) \geq -\lambda x/2$$

Thus, $x \geq 2/N - O(m\epsilon) \notin A_N$ and we get a contradiction. \square

Proof of Lemma 3.3.9. Assume $x_{i,j-1} \in C_N$; the proof of the other case is similar.

Using the same notation and argument of Lemma 3.3.7, we start with (3.20) and get $O(m\epsilon) = |g(x) + \lambda y| = |g(x) + \lambda\theta \pm \lambda/N|$, which implies

$$|g(x) + \lambda\theta| \leq \lambda/N + O(m\epsilon)$$

The right side is smaller than $(\lambda - \sigma/2)/N$ as $N = n^6$, $\epsilon = 1/n^{14}$ and $m = n^7$. It follows from Lemma 3.2.5 that $x \in C_N$. The lemma follows directly. \square

We construct a $2n$ -dimensional vector \mathbf{y} from \mathbf{p} as follows. Recall θ^* is a γ -rational approximation of θ with $\gamma = 1/n^7$. Let $\delta = 1/N$. For each $i \in [n]$, if $x_{i,m} \geq \theta^* - 2\delta$, then we set $y_{2i-1} = 1$ and $y_{2i} = 0$; if $x_{i,m} \leq -(\theta^* - 2\delta)$, then we set $y_{2i-1} = 0$ and $y_{2i} = 1$; otherwise, we set y_{2i-1} and y_{2i} to be

$$y_{2i-1} = \frac{\theta^* + x_{i,m}}{2\theta^*} \quad \text{and} \quad y_{2i} = \frac{\theta^* - x_{i,m}}{2\theta^*}.$$

By definition, \mathbf{y} is a nonnegative vector and $y_{2i-1} + y_{2i} = 1$ for all $i \in [n]$. Note that when $x_{i,m} \in C_N$, we have $x_{i,m} \geq \theta^* - 2\delta$ since $\gamma < \delta$, and hence $y_{2i-1} = 1$ and $y_{2i} = 0$. Similarly, if $x_{i,m} \in C'_N$ then $y_{2i-1} = 0$ and $y_{2i} = 1$. By Lemma 3.3.3, for every $i \in [2n]$,

$$y_i = \frac{\theta^* + \pi(G_i) - 1}{2\theta^*} \pm (O(\gamma + m\epsilon + 1/N)) \Rightarrow \pi(G_i) = 2\theta^* y_i + (1 - \theta^*) \pm O(1/N).$$

To finish the proof of Theorem 3, we prove the following theorem:

Theorem 4. *When n is sufficiently large, \mathbf{y} built above is a $(1/n)$ -well-supported Nash equilibrium of \mathbf{P} .*

To prove the theorem we will need the following two lemmas.

Lemma 3.3.10. *For any $j \in [m]$, if $x_{i,j-1}, x_{i,j} \in B_N$, then $x_{i,j} = x_{i,j-1} + \Omega(1/N)$.*

Proof. Using the same notation and argument of Lemma 3.3.7, we start with (3.20) and get $g(x_{i,j}) = -\lambda x_{i,j-1} \pm O(m\epsilon)$. From Lemma 3.2.4, $g(x_{i,j}) + \lambda x_{i,j} = f(x_{i,j}) = \Omega(1/N)$ since $x_{i,j} \in B_N$. As a result, we have

$$-\lambda x_{i,j} + \Omega(1/N) = g(x_{i,j}) = -\lambda x_{i,j-1} \pm O(m\epsilon)$$

and thus, $x_{i,j} = x_{i,j-1} + \Omega(1/N)$ using $m = n^7$ and $\epsilon = 1/n^{14}$. The lemma follows. \square

We are now ready to prove the following key lemma.

Lemma 3.3.11. *For every $i \in [n]$, if $x_{i,0} \in B_N \cup C_N$, then we have $x_{i,m} \in C_N$ and $y_{2i-1} = 1, y_{2i} = 0$. Similarly if $x_{i,0} \in B'_N \cup C'_N$, then $x_{i,m} \in C'_N$ and $y_{2i-1} = 0, y_{2i} = 1$.*

Proof. By Lemma 3.3.9, we assume that $x_{i,0} \in B_N$ without loss of generality.

Now assume for contradiction that $x_{i,m} \notin C_N$. By Lemma 3.3.9 again, we have $x_{i,j} \in B_N$ for all $j \in [m]$. This contradicts Lemma 3.3.10 as $m = n^7, N = n^6$. Lemma 3.3.11 follows. \square

Finally we prove Theorem 4:

Proof of Theorem 4. We assume for contradiction that \mathbf{y} is not a $(1/n)$ -well-supported Nash equilibrium of \mathbf{P} . Without loss of generality, we assume that

$$\mathbf{y}^T \cdot \mathbf{P}_1 > \mathbf{y}^T \cdot \mathbf{P}_2 + 1/n \tag{3.21}$$

where \mathbf{P}_1 and \mathbf{P}_2 denote the first and second columns of \mathbf{P} , but $y_2 > 0$. For a contradiction, by Lemma 3.3.11, it suffices to show that (3.21) implies that $x_{1,0} \in B_N \cup C_N$.

To this end, we first scale \mathbf{p} so that $\pi(G_1) + \pi(G_2) = 2$, and use x to denote $x_{1,0}$. By Lemma 3.3.5, we have $\pi(G_1), \pi(G_2) = 1 \pm O(1/n^3)$ are very close to 1. By applying Walras' law over the whole market $M_{\mathbf{P}}$ and using the assumption that \mathbf{p} is an ϵ -additively approximate equilibrium, we have

$$\epsilon \geq \text{the excess demand of } G_1 \text{ (or } G_2) \geq -O(mn\epsilon). \quad (3.22)$$

Now we compare the total money spent on G_1 and G_2 , by all traders in $M_{\mathbf{P}}$ except the one, denoted by T , who owns τ units of G_1 and G_2 each. We list all such traders:

1. For each $i \in [2n]$, there is a $(P_{i,1}, H_i : G_1)$ -trader. The total money these traders spend on G_1 is

$$\sum_{i \in [2n]} P_{i,1} \cdot \pi(H_i) = \sum_{i \in [2n]} P_{i,1} \cdot \left(2\theta^* y_i + (1 - \theta^*) \pm O(1/N) \right)$$

2. For each $i \in [2n]$, there is a $(P_{i,2}, H_i : G_2)$ -trader. The total money these traders spend on G_2 is

$$\sum_{i \in [2n]} P_{i,2} \cdot \pi(H_i) = \sum_{i \in [2n]} P_{i,2} \cdot \left(2\theta^* y_i + (1 - \theta^*) \pm O(1/N) \right)$$

3. There is one $((1 - \theta^*)r_1, AUX_1 : G_1)$ -trader and one $((1 - \theta^*)r_2, AUX_1 : G_2)$ trader.

Recall r_1 and r_2 in (3.11). The total money these traders spend on G_1 is

$$M_1 = 2\theta^* \cdot \mathbf{y}^T \cdot \mathbf{P}_1 + 2n(1 - \theta^*) \pm O(n/N)$$

using $N = n^6$ and $m\epsilon = 1/n^7$, and the total money these traders spend on G_2 is

$$M_2 = 2\theta^* \cdot \mathbf{y}^T \cdot \mathbf{P}_2 + 2n(1 - \theta^*) \pm O(n/N)$$

Thus $M_1 - M_2 = \Omega(1/n)$ and the demand for G_1 is larger than the demand for G_2 , from these traders, by

$$\frac{M_1}{\pi(G_1)} - \frac{M_2}{\pi(G_2)} \geq \frac{M_2 + \Omega(1/n)}{\pi(G_1)} - \frac{M_2}{(1 - O(1/n^3)) \cdot \pi(G_1)} = \Omega(1/n)$$

where both inequalities used $\pi(G_1), \pi(G_2) = 1 \pm O(1/n^3)$ and $M_1, M_2 = O(n)$.

Let d_1 (or d_2) denote the demand of G_1 (or G_2 , respectively) from T . Using (3.22) and $mn\epsilon \ll 1/n$ we must have $d_2 - d_1 = \Omega(1/n)$. On the other hand, we have from (3.4):

$$\frac{d_1}{d_2} = \left(\frac{1-x}{1+x} \right)^{1/(1+r)}$$

As $\pi(G_1), \pi(G_2)$ are close to 1, d_1 and d_2 are $O(\tau)$ even if T spends all the budget on one of them. As a result, we have

$$\left(\frac{1+x}{1-x} \right)^{1/(1+r)} = \frac{d_2}{d_1} = 1 + \frac{d_2 - d_1}{d_1} = 1 + \Omega(1/n^5)$$

and thus, $x = \Omega(1/n^5)$. It follows from Lemma 3.3.5 and $N = n^6$ that $x \in B_N$. This finishes the proof of the theorem. \square

Chapter 4

The Complexity of Non-Monotone Markets

4.1 Introduction

As mentioned earlier, several market equilibrium problems have been shown to be PPAD-hard, however, the reduction techniques developed in these proofs make heavy use of the properties of the particular utilities under consideration. On the other hand, our PPAD-hardness proof for CES markets is not really fine-tuned to CES utilities. Although it uses properties of the CES utilities to ensure that no spurious equilibria appear, the main building block of the reduction relies on an interesting behavior of example 3.1.1.

Crucial Behaviour. There is a special good (in example 3.1.1 either G_1 or G_2) and a price vector $\mathbf{p} > 0$ (in example 3.1.1 $\mathbf{p} = (1, 1)$) such that at \mathbf{p} , the excess demand of G is nonnegative and raising the price of G , while keeping all other prices the same, strictly increases the demand of G .

This observation inspired us to ask the following question:

Can we prove a complexity dichotomy for any given family of utility functions?

In more details, let \mathcal{U} denote a generic family of utilities that satisfy certain mild conditions (e.g., they are continuous, quasi-concave). The question now becomes the following:

Does there exist a mathematically well-defined property of families of functions such that: For any \mathcal{U} satisfying this property the equilibrium problem it defines is in polynomial time, while for any \mathcal{U} that violates this property the problem is hard, e.g., PPAD-hard?

In this section, we obtain a PPAD-hardness result that is widely applicable to any generic family \mathcal{U} of utility functions, as long as it satisfies the following condition:

[Informal]: *There exists a market M with utilities from \mathcal{U} , a special good G in M , and a price vector $\mathbf{p} > 0$ such that at \mathbf{p} , the excess demand of G is nonnegative and raising the price of G , while keeping all other prices the same, strictly increases the demand of G .*

If such a market exists, we call it *non-monotone* and \mathcal{U} a *non-monotone* family.

In Section 4.2 we provide examples of simple non-monotone markets, constructed from various families of utilities. All families for which we have hardness results for the (approximate) equilibrium problem are non-monotone. This includes in particular the family of separable piecewise-linear functions, the family of Leontief functions, and the family of CES functions for any value of the parameter $\rho < -1$. Of course, if a family \mathcal{U} is non-monotone, then so is any superset of \mathcal{U} . We show that the existence of a non-monotone market implies the following hardness result:

Theorem 5 (Informal). *If \mathcal{U} is non-monotone, then the following problem is PPAD-hard: Given a market in which the utility of each trader is either linear or from \mathcal{U} , find an approximate market equilibrium.*

The theorem implies in particular the known PPAD-hardness of the (approximate) equilibrium problem for Arrow-Debreu markets with separable piecewise-linear utility functions [Chen *et al.*, 2009a]. In fact the proof shows that the problem is hard even in the special case where the utility function of every trader for each good is either linear or linear with a threshold after which it gets saturated and stops increasing. The theorem by itself, however, does not imply the hardness result for CES markets (Theorem 3) or for Leontief markets [Codenotti *et al.*, 2006] (even though these families are non-monotone), because of the use of linear functions.

Comparing Theorem 5 with the major known positive case of WGS, it is easy to see that if a market satisfies WGS then it cannot be non-monotone: raising the price of a good G causes the demands for all other goods to increase or stay the same, and hence by Walras' law, the demand for G cannot also increase. There remains a gap however between WGS and the complement of non-monotonicity, mainly for two reasons: 1) in the definition of non-monotone markets, the excess demand of G is required to be nonnegative at \mathbf{p} but WGS does not make such an assumption, and 2) the definition of non-monotonicity constrains the change in the demand of the good G whose price is increased whereas WGS constrains the change in the demand of all other goods. If there are only two goods the two constraints are related both ways (by Walras' law), but if there are more than two goods the implication is only in one direction. It remains an open problem whether we can further reduce the gap, and whether we can remove the use of linear functions in the theorem.

The rest of this chapter is organized as follows. In Section 4.2, we give all the necessary definitions and state formally our results. Section 4.3 contains the PPAD-hardness proof for general non-monotone utilities.

4.2 Definitions and Statement of Results

We use \mathcal{U} to denote a generic family of *continuous, quasi-concave and locally non-satiated* functions, e.g., linear functions, piecewise-linear functions (see Example 4.2.3), CES functions for a specific parameter of ρ , e.g. $\rho = -3$, or even the finite set of three functions given in (3.2). Throughout this chapter we will make the following two assumptions for \mathcal{U} .

Assumption 1. \mathcal{U} is countable and each function $g \in \mathcal{U}$ corresponds to a unique binary string \mathbf{s} so that a trader can specify a function $g \in \mathcal{U}$ using \mathbf{s} . In a market with m goods, we say a trader “*applies*” a function $g \in \mathcal{U}$ if her utility function u is of the form

$$u(x_1, \dots, x_m) = g\left(\frac{x_{\ell_1}}{b_1}, \dots, \frac{x_{\ell_k}}{b_k}\right),$$

where $g \in \mathcal{U}$ has $k \leq m$ variables; $\ell_1, \dots, \ell_k \in [m]$ are distinct indices; and b_1, \dots, b_k are positive rational numbers used to change the units. In this way, each trader can be described by a finite binary string.

Assumption 2. There exists a univariate function $g^* \in \mathcal{U}$ that is strictly monotone.

Remark. These two assumptions seem to be natural and we only need them for technical reasons that will become clear later. When a trader applies a function from \mathcal{U} , she can always change units by scaling with the appropriate values for b_1, \dots, b_k . The second assumption allows us to add single-minded traders who spend all their budget on one specific good.

We use $\mathcal{M}_{\mathcal{U}}$ to denote the set of all markets in which every trader has a rational initial endowment and applies a utility function from \mathcal{U} . We also use $\mathcal{M}_{\mathcal{U}}^*$ to denote the set of markets in which every trader has a rational initial endowment and applies either a utility function from \mathcal{U} or a linear utility with rational coefficients.

Next, we define *non-monotone* markets as well as *non-monotone* families of utilities:

Definition 11 (Non-monotone Markets and Families of Utilities). *Let M be a market with $k \geq 2$ goods. We say M is non-monotone at a price vector \mathbf{p} if the following conditions hold:*

1. $\pi_j > 0$ for all $j \in [k]$
2. For some $c > 0$, the excess demand $Z_1(y_1, \dots, y_k)$ of G_1 is a continuous function over $\mathbf{y} \in B(\mathbf{p}, c)$ with $Z_1(\mathbf{p}) \geq 0$. The partial derivative of Z_1 with respect to y_1 exists and is continuous over $B(\mathbf{p}, c)$ and is (strictly) positive at \mathbf{p} .

We call M a non-monotone market if there exists such a price vector \mathbf{p} . We also call \mathcal{U} a non-monotone family of utilities if there exists a non-monotone market in $\mathcal{M}_{\mathcal{U}}$.

Remark. By definition, M being non-monotone at \mathbf{p} means that raising the price of G_1 , while keeping the prices of all other goods the same, would actually increase the total demand of G_1 . Also note that using the continuity of Z_1 as well as its partial derivative with respect to y_1 , we can indeed require, without loss of generality, the price vector \mathbf{p} to be rational in Definition 11: if M is a non-monotone market at \mathbf{p} but \mathbf{p} is not rational, then a rational vector \mathbf{p}^* close enough to \mathbf{p} would have the same property. Therefore, whenever \mathcal{U} is non-monotone, there is a market $M \in \mathcal{M}_{\mathcal{U}}$ that is non-monotone at a rational price vector \mathbf{p} . We would like to mention that M is not necessarily strongly connected and the

excess demand $Z_1(\mathbf{p})$ of G_1 and the partial derivative of Z_1 with respect to y_1 at \mathbf{p} do not have to be rational.

We now state the main result of this chapter, a PPAD-hardness reduction to any non-monotone family \mathcal{U} of functions. We use \mathcal{U} -MARKET to denote the following problem: the input is a pair (k, M) , where k is a positive integer in unary and M is a strongly connected market from $\mathcal{M}_{\mathcal{U}}^*$ encoded in binary. The goal is to output an ϵ -approximate equilibrium of M with $\epsilon = 1/k$. Here is the formal statement.

Theorem 6 (Main Result). *Let \mathcal{U} denote a non-monotone family of utility functions. If there exists a market $M \in \mathcal{M}_{\mathcal{U}}$ such that M is non-monotone at a rational price vector \mathbf{p} and the excess demand $Z_1(\mathbf{p})$ of G_1 at \mathbf{p} is moderately computable, then \mathcal{U} -MARKET is PPAD-hard.*

Remark. By definition, \mathcal{U} being non-monotone implies the existence of M and \mathbf{p} . The other assumption made in Theorem 6 only requires that there exists one such pair (M, \mathbf{p}) for which $Z_1(\mathbf{p})$ as a specific positive number is moderately computable. We also point out that, when the assumptions of Theorem 6 hold, such a pair M and \mathbf{p} is considered a constant, therefore, we can later use it in the proof of Theorem 6 as a gadget to give a polynomial-time reduction from Polymatrix to \mathcal{U} -MARKET. As a result, all components of M are considered constants encoded by binary strings of constant length. This includes the number of goods and traders, the endowments of traders, the binary strings that specify their utility functions from \mathcal{U} , and the positive rational vector \mathbf{p} .

4.2.1 Examples of non-monotone Markets

Now we present three examples of non-monotone markets, one with CES utilities (see 3.1.1) of parameter $\rho < -1$, one with Leontief utilities, and one with additively separable and piecewise-linear utilities.

Example 4.2.1 (A Non-Monotone Market with CES Utilities of $\rho < -1$ [Gjerstad, 1996]). *Recall the market from example 3.1.1. There are two goods G_1 and G_2 and two traders T_1 and T_2 , each with 1 unit of G_1 and G_2 respectively and their utilities are $u_1(x_1, x_2) = (\alpha \cdot x_1^\rho + x_2^\rho)^{1/\rho}$ and $u_2(x_1, x_2) = (x_1^\rho + \alpha \cdot x_2^\rho)^{1/\rho}$.*




Figure 4.1: The excess demand function $Z_1(x)$ of Example 3.1.1.

From 3.1, it follows that the excess demand function and its derivative is defined and continuous anywhere except for $\mathbf{0}$. When $\rho < -1$ and α is large enough, this market is non-monotone at $(1, 1)$. To see this, we let $Z_1(x)$ denote the excess demand function of G_1 , when the price of G_1 is $1 + x$ and the price of G_2 is $1 - x$ and plot Z_1 in Figure 4.1. From the picture, it is clear that the curve has three roots (equilibria) and is non-monotone at $(1, 1)$, or equivalently, when $x = 0$.

Example 4.2.2 (A Non-Monotone Market with Leontief Utilities). Let M denote the Leontief market consisting of two traders T_1 and T_2 . T_1 has 1 unit of G_1 , T_2 has 1 unit of G_2 , and their utility functions are

$$u_1(x_1, x_2) = \min \{x_1/2, x_2\} \quad \text{and} \quad u_2(x_1, x_2) = \min \{x_1, x_2/2\}$$

respectively. It is easy to show that M is non-monotone at $(1, 1)$. □

Example 4.2.3. (A Non-Monotone Market with Additively Separable and Piecewise-Linear Utilities) We say u is additively separable and piecewise-linear if

$$u(x_1, \dots, x_k) = f_1(x_1) + \dots + f_k(x_k) \tag{4.1}$$

where f_1, \dots, f_k are all piecewise-linear functions. Consider the following market M with two goods G_1, G_2 and two traders T_1, T_2 . T_1 has 1 unit of G_1 and T_2 has 1 unit of G_2 .

Their utility functions are

$$u_1(x_1, x_2) = x_1 + f(x_2) \quad \text{and} \quad u_2(x_1, x_2) = f(x_1) + x_2 \quad \text{with} \quad f(x) = \begin{cases} 2x & \text{if } x \leq 1/3 \\ 2/3 & \text{if } x > 1/3 \end{cases}$$

It can be shown that M has $(1, 1)$ as an equilibrium and is non-monotone at $(1, 1)$. Note that in general, the excess demand of a market with such utilities is a correspondence instead of a map, and partial derivatives may not always exist. But in the definition of non-monotone markets, we only need these properties in a local neighborhood of \mathbf{p} , like $(1, 1)$ here. \square

Since linear functions are special cases of additively separable and piecewise-linear functions, a corollary of Theorem 6 and Example 4.2.3 is that finding an approximate equilibrium in a market with additively separable and concave piecewise-linear utilities is PPAD-hard. This result was shown earlier in [Chen *et al.*, 2009a] but with the use of three-segment piecewise-linear utilities. Combining this with the membership in PPAD [Vazirani and Yannakakis, 2011], we have:

Corollary 4.2.1. *The problem of computing an approximate market equilibrium in a market with additively separable and concave piecewise-linear utilities is PPAD-complete, even when each univariate function f_j in (4.1) is either linear or a linear function with a threshold.*¹

4.3 Markets with Non-Monotone and Linear Utilities are PPAD-hard

We prove Theorem 6 in this section. Let \mathcal{U} denote a non-monotone family of utilities and $M \in \mathcal{M}_{\mathcal{U}}$ denote a market that is non-monotone at a rational price vector \mathbf{p} . We let $k \geq 2$ denote the number of goods in M . We also assume that the excess demand $Z_1(\mathbf{p})$ of G_1 at \mathbf{p} is moderately computable. As discussed earlier, $M, k, \mathbf{p}, Z_1(\mathbf{p})$ (including the total supply of each good in M) are all considered constants, independent of the size of the polymatrix game we reduce from.

¹The second part of the statement follows from the construction used in the proof of Theorem 6.

4.3.1 Normalized Non-Monotone Markets

Note that in examples 3.1.1, 4.2.2, and 4.2.3, the market we construct not only is non-monotone at $\mathbf{1} = (1, 1)$ but also has $Z_1(\mathbf{1}) = 0$ (indeed $\mathbf{1}$ is an equilibrium in all three examples). The lemma below shows that this is not really a coincidence since we can always convert a non-monotone market into one that is non-monotone at $\mathbf{1}$, as shown below. Recall that $M \in \mathcal{M}_{\mathcal{U}}$ is a market that is non-monotone at a rational vector \mathbf{p} , with $k \geq 2$ goods, such that $Z_1(\mathbf{p})$ is moderately computable. We use M and \mathbf{p} to prove the following lemma:

Lemma 4.3.1 (Normalized Non-Monotone Markets). *There exist two (not necessarily rational) positive constants c and d such that, given any $\gamma > 0$, one can build a market $M_\gamma \in \mathcal{M}_{\mathcal{U}}$ with $k \geq 2$ goods G_1, \dots, G_k , in time polynomial in $1/\gamma$ with the following properties.*

1. Let $f_\gamma(x)$ denote the excess demand function of G_1 when the price of G_1 is $1+x$ and the prices of all other $(k-1)$ goods are $1-x$. Then f_γ is well defined over $[-c, c]$ with $|f_\gamma(0)| \leq \gamma$ and its derivative $f'_\gamma(0) = d > 0$. For any $x \in [-c, c]$, $f_\gamma(x)$ satisfies

$$|f_\gamma(x) - f_\gamma(0) - dx| \leq |x/D|, \quad \text{where } D = \max\{20, 20/d\}.$$
²

2. The total supply of each of the k goods in M_γ is $O(1)$.

Proof. First we construct M' from M by scaling: For each trader with utility u and initial endowment vector $\mathbf{w} \in \mathbb{Q}_+^k$, replace them by $w'_j = w_j \cdot \pi_j$ for every $j \in [k]$ and

$$u'(x_1, \dots, x_k) = u\left(\frac{x_1}{\pi_1}, \dots, \frac{x_k}{\pi_k}\right)$$

Since \mathbf{p} is rational and positive, we have $M' \in \mathcal{M}_{\mathcal{U}}$. It is also easy to verify that M' now is non-monotone at $\mathbf{1}$. Let $g(x)$ denote the excess demand function of G_1 when the price of G_1 is $1+x$ and the prices of all other goods are $1-x$. Then, by the definition of non-monotone markets, there exist two positive constants c and d such that g is well defined over $[-c, c]$, $g(0) \geq 0$ and $g'(0) = d > 0$. The latter follows from the fact that the excess demand at

²As it will become clear in the proof of Lemma 4.3.1, one can choose D to be any positive constant (by picking a small enough constant c accordingly). Our choice of $D = \max\{20, 20/d\}$ (and the constant 20) just makes sure that D is large enough for the proof of correctness of our reduction to work later.

$(1+x, 1-x, \dots, 1-x)$ is the same as that at $((1+x)/(1-x), 1, \dots, 1)$. As d (and thus, D) is a constant, setting c to be a small enough constant, it follows from $g'(0) = d$ that:

$$|g(x) - g(0) - dx| \leq |x/D|, \quad \text{for all } x \in [-c, c].$$

Next, let $Z'_1 = g(0)$ denote the excess demand of G_1 in M' at $\mathbf{1}$. Then $Z'_1 = \pi_1 \cdot Z_1(\mathbf{p})$ and thus, Z'_1 is also moderately computable. Given any $\gamma > 0$, we compute a γ -rational approximation z of Z'_1 . We assume, without loss of generality, that z is nonnegative; otherwise, simply set $z = 0$. Finally, we construct M_γ from M' by adding a trader with z units of G_1 who is only interested in G_k . It is clear that the total supply of each good in M_γ remains $O(1)$ as both \mathbf{p} and $Z_1(\mathbf{p})$ are constants.

Let $f_\gamma(x)$ denote the excess demand function of G_1 in M_γ when the price of G_1 is $1+x$ and all other goods have price $1-x$. The construction of M_γ then implies that $f_\gamma(x) = g(x) - z$ and thus, $|f_\gamma(0)| \leq \gamma$. It follows that M_γ and f_γ satisfy all the desired properties with respect to constants c and d above. \square

4.3.2 Our Construction

Given a normalized $2n \times 2n$ polymatrix game \mathbf{P} , we construct a market $M_{\mathbf{P}} \in \mathcal{M}_{\mathcal{U}}^*$ in polynomial time. We begin by describing the two building blocks of $M_{\mathbf{P}}$ and introduce some useful notation for them.

Normalized Non-Monotone Market: Given two positive rational numbers μ and γ , we use $\text{NM}(\mu, \gamma, G_1, \dots, G_k)$ to denote the addition to $M_{\mathbf{P}}$ of the following sub-market. First, we make a copy of M_γ in which the k goods that they are interested in are G_1, \dots, G_k . Then for each trader in M_γ with utility function $u(x_1, \dots, x_k)$ and endowment $\mathbf{w} = (w_1, \dots, w_k)$, we replace \mathbf{w} by $\mu\mathbf{w}$ and u by

$$u'(x_1, \dots, x_k) = u\left(\frac{x_1}{\mu}, \dots, \frac{x_k}{\mu}\right)$$

When both parameters μ and $1/\gamma$ are bounded from above by a polynomial in n , it takes time polynomial in n to create these traders. Let $f_{\mu, \gamma}(x)$ denote the excess demand of G_1 when the price of G_1 is $1+x$ and the prices of all other goods are $1-x$, then we have

$f_{\mu,\gamma}(x) = \mu \cdot f_\gamma(x)$. From the properties of f_γ stated in Lemma 4.3.1, $f_{\mu,\gamma}$ is well defined over $[-c, c]$, satisfies $|f_{\mu,\gamma}(0)| \leq \mu\gamma$, and

$$|f_{\mu,\gamma}(x) - f_{\mu,\gamma}(0) - \mu dx| \leq |\mu x/D|, \quad \text{for } x \in [-c, c], \text{ with } D = \max\{20, 20/d\}. \quad (4.2)$$

Recall that c and d are positive constants from Lemma 4.3.1, which do not depend on γ or μ . Also note that the total supply of each good in $\mathbf{NM}(\mu, \gamma, G_1, \dots, G_k)$ is $O(\mu)$.

Price-Regulating Market: Let G_1, \dots, G_ℓ denote $\ell \geq 2$ goods in $M_{\mathbf{P}}$ (where ℓ is k or 2 below), and let λ and α denote two positive rational numbers, where $\alpha < 1$. We use $\mathbf{PR}(\lambda, \alpha, G_1, \dots, G_\ell)$ below to denote the creation of two traders T_1, T_2 , and refer to the submarket they form as a *price-regulating market* [Chen *et al.*, 2009a; Vazirani and Yannakakis, 2011].

The endowment of T_1 is $(\ell - 1)\lambda$ units of G_1 and the endowment of T_2 is λ units of each of G_2, \dots, G_ℓ . We set their utility functions u_1 and u_2 , both of which are linear, as follows:

$$u_1(x_1, \dots, x_\ell) = (1 + \alpha)x_1 + \sum_{2 \leq j \leq \ell} (1 - \alpha)x_j \quad \text{and}$$

$$u_2(x_1, \dots, x_\ell) = (1 - \alpha)x_1 + \sum_{2 \leq j \leq \ell} (1 + \alpha)x_j$$

We will see that in any approximate market equilibrium, when λ is large enough and certain conditions are satisfied, a price-regulating market basically requires the prices of G_2, \dots, G_ℓ to be the same when $\ell > 2$; and the ratio of prices of G_1 and G_2 to be between $(1 - \alpha)/(1 + \alpha)$ and $(1 + \alpha)/(1 - \alpha)$.

Except for these two building blocks, all other traders in $M_{\mathbf{P}}$ are indeed *single-minded*: Each of them is only interested in one specific good and spends all her budget on it. We use the following notation. First we say a trader is a $(\tau, G_1 : G_2)$ -trader if her endowment consists of τ units of G_1 and she is only interested in G_2 . Second we say a trader is a $(\tau, G_1, G_2 : G_3)$ -trader if her endowment consists of τ units of G_1 and G_2 each, and she is only interested in G_3 .

Now we describe the construction of $M_{\mathbf{P}}$. We start with its set of goods. Without loss of generality, we always assume that $n = 2^t$ for some integer t . Then the market $M_{\mathbf{P}}$ consists

of the following $O(ntk) = O(n \log n)$ goods:

$$AUX_i, G_{2i-1,j}, G_{2i,j}, \text{ and } S_{i,\ell,r}, \quad \text{for } i \in [n], j \in [0 : 4t], \ell \in [4t] \text{ and } r \in [3 : k].$$

Recall that k is the number of goods in the non-monotone market M . The main goods in $M_{\mathbf{P}}$ are $G_{2i-1,j}$ and $G_{2i,j}$, while AUX_i and $S_{i,\ell,r}$ are auxiliary. Informally, AUX_i 's are introduced to balance the total money spent on $G_{2i-1,0}$ and $G_{2i,0}$ (see the proof of Theorem 7). On the other hand, we need $S_{i,\ell,r}$'s only when $k \geq 3$: When we need to add an NM market with $G_{2i-1,\ell}$ and $G_{2i,\ell}$ as its goods 1 and 2 we use $S_{i,\ell,3}, \dots, S_{i,\ell,k}$ as goods 3, \dots , k . When $k = 2$, we do not need $S_{i,\ell,r}$'s in $M_{\mathbf{P}}$.

We divide all the goods, except the AUX_i 's, into the following $n(4t + 1)$ groups $\{\mathcal{R}_{i,j}\}$, where $i \in [n]$ and $j \in [0 : 4t]$. For each $i \in [n]$ and $j \in [4t]$, we use $\mathcal{R}_{i,j}$ to denote

$$\mathcal{R}_{i,j} = \{G_{2i-1,j}, G_{2i,j}, S_{i,j,3}, \dots, S_{i,j,k}\},$$

a group of k goods; for each $i \in [n]$, we use $\mathcal{R}_{i,0}$ to denote $\{G_{2i-1,0}, G_{2i,0}\}$.

Next we list all the parameters used in the construction. We use α_j to denote $2^j/n^5$, for each $j \in [0 : 4t]$ (thus $\alpha_0 = 1/n^5$ and $\alpha_{4t} = 1/n$). Recall the positive constant d from Lemma 4.3.1. We let d^* denote a positive rational number (a constant) that satisfies

$$1 - 1/D \leq d^*d \leq 1, \quad \text{where } D = \max\{20, 20/d\}.$$

The rest of the parameters are

$$\beta = \alpha_{4t} = 1/n, \quad \mu = d^*n, \quad \tau = n^2, \quad \gamma = 1/n^6, \quad \xi = \epsilon nt, \quad \delta = \epsilon t \quad \text{and} \quad \epsilon = 1/n^8.$$

We explain some of the key parameters. First, ϵ is the approximation parameter of market equilibria we are interested in. Next, as for CES utilities, we need to amplify the gap between couples of prices over a chain of goods; each α_j specifies the gap between prices of $G_{2i-1,j}$ and $G_{2i,j}$ in this amplification. In more details, if the ratio of the price of $G_{2i-1,j}$ to that of $G_{2i,j}$ is $(1 + \alpha_j)/(1 - \alpha_j)$ (or $(1 - \alpha_j)/(1 + \alpha_j)$), then the ratio of $G_{2i-1,j+1}$ to $G_{2i,j+1}$ must be $(1 + \alpha_{j+1})/(1 - \alpha_{j+1})$ (or $(1 - \alpha_{j+1})/(1 + \alpha_{j+1})$). So if there is an α_0 -gap between $G_{2i-1,0}$ and $G_{2i,0}$, it would be amplified to a β -gap between $G_{2i-1,4t}$ and $G_{2i,4t}$. Finally, μ , τ and γ are parameters used in the NM and \mathbf{PR} markets that we add to $M_{\mathbf{P}}$; ξ and δ are parameters used in the proof of correctness only.

Construction of $M_{\mathbf{P}}$. The construction follows closely the one for CES markets. The difference here is that we use the non-monotone market as a black-box and, therefore, we cannot prove properties similar to the ones proved for example 3.1.1 in Section 3.2. To overcome this, we use **PR** to ensure that all equilibria are within the range of prices that **NM** is non-monotone.

First we use **NM** and **PR** to build a *closed* economy over each group $\mathcal{R}_{i,j}$. Here by a closed economy over a group of goods, we mean a set of traders whose endowments consist of goods from this group only and they are interested in goods from this group only.

1. For each group $\mathcal{R}_{i,j}$, where $i \in [n]$ and $j \in [4t]$, we add a price-regulating market

$$\mathbf{PR}(\tau, \alpha_j, G_{2i-1,j}, G_{2i,j}, S_{i,j,3}, \dots, S_{i,j,k})$$

We also add a non-monotone market

$$\mathbf{NM}(\mu, \gamma, G_{2i-1,j}, G_{2i,j}, S_{i,j,3}, \dots, S_{i,j,k})$$

We will refer to them as *the PR* market and *the NM* market over $\mathcal{R}_{i,j}$, respectively.

2. For each group $\mathcal{R}_{i,0}$ of $\{G_{2i-1,0}, G_{2i,0}\}$, $i \in [n]$, we add a price-regulating market

$$\mathbf{PR}(\tau, \alpha_0, G_{2i-1,0}, G_{2i,0})$$

We will refer to it as *the PR* market over $\mathcal{R}_{i,0}$.

Next we add a number of single-minded traders who trade between different groups. The initial endowment of each such trader consists of $G_{2i-1,j}$ and $G_{2i,j}$ of a group $\mathcal{R}_{i,j}$ (one of them or both) and she is only interested in either $G_{2i'-1,j'}$ or $G_{2i',j'}$ of another group $\mathcal{R}_{i',j'}$, where $(i, j) \neq (i', j')$. We will refer to her as a trader who trades from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i',j'}$.

At the same time we construct a weighted directed graph $\mathcal{G} = (V, E)$ as defined in Section 3.3. Here is the construction:

1. For each $i \in [2n]$, we use G_i to denote $G_{i,0}$ and H_i to denote $G_{i,4t}$ for convenience.

For each pair $i, j \in [n]$, we add to $M_{\mathbf{P}}$ the following four traders who trade from group $\mathcal{R}_{i,4t}$ to group $\mathcal{R}_{j,0}$: one $(P_{2i-1,2j-1}, H_{2i-1} : G_{2j-1})$ -trader, one $(P_{2i-1,2j},$

$H_{2i-1} : G_{2j}$ -trader, one $(P_{2i,2j-1}, H_{2i} : G_{2j-1})$ -trader, and one $(P_{2i,2j}, H_{2i} : G_{2j})$ -trader. Since \mathbf{P} is normalized, we have

$$P_{2i-1,2j-1} + P_{2i-1,2j} = P_{2i,2j-1} + P_{2i,2j} = 1$$

Thus, the total endowment of these four traders consists of one unit of H_{2i-1} and H_{2i} each, so we add an edge in \mathcal{G} from $\mathcal{R}_{i,4t}$ to $\mathcal{R}_{j,0}$ with weight 1. At this moment, the total out-weight of each $\mathcal{R}_{i,4t}$ in \mathcal{G} is n , and the total in-weight of each $\mathcal{R}_{i,0}$ in \mathcal{G} is n .

2. For each $i \in [n]$ and $j \in [0 : 4t - 1]$, we add two traders who trade from group $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i,j+1}$: one $(n, G_{2i-1,j} : G_{2i-1,j+1})$ -trader and one $(n, G_{2i,j} : G_{2i,j+1})$ -trader. We also add an edge in graph \mathcal{G} from $\mathcal{R}_{i,j}$ to $\mathcal{R}_{i,j+1}$ with weight n .

This finishes the construction of \mathcal{G} . It is also easy to check that \mathcal{G} is strongly connected and every vertex (group) has both its total in-weight and out-weight equal to n .

Finally, we add traders between AUX_j and $\mathcal{R}_{j,0}$ for each $j \in [n]$. Let r_{2j-1} and r_{2j} be the two numbers defined in (3.11). Recall that $\beta = \alpha_{4t} = 1/n$. We add to $M_{\mathbf{P}}$ three traders: one $((1 - \beta)r_{2j-1}, AUX_j : G_{2j-1})$ -trader, one $((1 - \beta)r_{2j}, AUX_j : G_{2j})$ -trader, and one $((1 - \beta)n, G_{2j-1}, G_{2j} : AUX_j)$ -trader.

This finishes the construction of $M_{\mathbf{P}}$. It follows immediately from the strong connectivity of \mathcal{G} that the economy graph of $M_{\mathbf{P}}$ is strongly connected and thus, $M_{\mathbf{P}}$ is a valid input of problem \mathcal{U} -MARKET and can be constructed from \mathbf{P} in polynomial time. We also record the following properties of $M_{\mathbf{P}}$:

Lemma 4.3.2. *For each $i \in [n]$, the total supply of AUX_i is $2(1 - \beta)n$;*

For each $i \in [2n]$, the total supply of $G_{i,0}$ is $n^2 + O(n)$;

For each $i \in [n]$ and $j \in [4t]$, the total supply of $G_{2i-1,j}$ is $(k - 1)n^2 + O(n)$; and

For each $i \in [n]$, $j \in [4t]$ and $\ell \in [3 : k]$, the total supply of $G_{2i,j}$ and $S_{i,j,\ell}$ is $n^2 + O(n)$.

4.3.3 Proof of Correctness

In this section, we prove in that given an ϵ -additively approximate equilibrium \mathbf{p} of $M_{\mathbf{P}}$, we can compute a $(1/n)$ -well-supported Nash equilibrium of \mathbf{P} in polynomial time. Theorem 6 then follows.

First, from the **PR** markets in $M_{\mathbf{p}}$, we prove the following lemma:

Lemma 4.3.3. *Let \mathbf{p} denote an ϵ -additively approximate equilibrium of $M_{\mathbf{p}}$. Then*

$$\frac{1 - \alpha_j}{1 + \alpha_j} \leq \frac{\pi(G_{2i-1,j})}{\pi(G_{2i,j})} \leq \frac{1 + \alpha_j}{1 - \alpha_j}, \quad \text{for all } i \in [n] \text{ and } j \in [0 : 4t].$$

Furthermore, we have $\pi(G_{2i,j}) = \pi(S_{i,j,3}) = \dots = \pi(S_{i,j,k})$ for all $i \in [n]$ and $j \in [4t]$.

Proof. We consider the case when $i \in [n]$ and $j \in [4t]$ since the case $j = 0$ is simpler. We denote the two traders in the **PR** market over $\mathcal{R}_{i,j}$ by T_1 and T_2 and let

$$p_{\min} = \min \{ \pi(G_{2i,j}), \pi(S_{i,j,3}), \dots, \pi(S_{i,j,k}) \}$$

$$p_{\max} = \max \{ \pi(G_{2i,j}), \pi(S_{i,j,3}), \dots, \pi(S_{i,j,k}) \}$$

First, assume for contradiction that

$$\frac{1 + \alpha_j}{\pi(G_{2i-1,j})} < \frac{1 - \alpha_j}{p_{\min}}$$

It follows that neither T_1 nor T_2 is interested in $G_{2i-1,j}$ and they only buy goods from $\mathcal{R}_{i,j}$ that are priced at p_{\min} . Let $F_{\min} \subset \mathcal{R}_{i,j}$ denote the set of such goods, then we have $G_{2i-1,j} \notin F_{\min}$. On the other hand, by the definition of p_{\min} , the budget of both T_1 and T_2 is at least $(k-1)\tau p_{\min}$. It follows that the total demand for goods in F_{\min} is at least $2(k-1)\tau$. However, the total supply of goods in F_{\min} is at most $(k-1)\tau + O(n)$, contradicting with the assumption that \mathbf{p} is an ϵ -additively approximate equilibrium.

Next, assume for contradiction that

$$\frac{1 - \alpha_j}{\pi(G_{2i-1,j})} > \frac{1 + \alpha_j}{p_{\max}}$$

and we let $F_{\max} \subset \mathcal{R}_{i,j}$ denote the set of goods priced at p_{\max} . Then neither T_1 nor T_2 is interested in goods from F_{\max} and they only buy goods from $\mathcal{R}_{i,j} - F_{\max}$. In particular, T_2 spends the part of budget she earns from selling F_{\max} on goods in $\mathcal{R}_{i,j} - F_{\max}$ as well. As goods in F_{\max} are the most expensive among $\mathcal{R}_{i,j}$, the demand for one of the goods in $\mathcal{R}_{i,j} - F_{\max}$ must be larger than the supply by $\Omega(\tau)$, contradicting with the assumption that \mathbf{p} is an ϵ -additively approximate equilibrium.

Combining these two steps, we immediately get

$$\frac{1 - \alpha_j}{1 + \alpha_j} \leq \frac{\pi(G_{2i-1,j})}{p_{\max}} \leq \frac{\pi(G_{2i-1,j})}{\pi(G_{2i,j})} \leq \frac{\pi(G_{2i-1,j})}{p_{\min}} \leq \frac{1 + \alpha_j}{1 - \alpha_j} \quad (4.3)$$

In the rest of the proof, we show that $\pi(G_{2i,j}) = \pi(S_{i,j,3}) = \dots = \pi(S_{i,j,k})$.

Assume for contradiction that this is not the case. Then $p_{\max} > p_{\min}$ which implies that neither T_1 nor T_2 is interested in F_{\max} . This leads us to the same contradiction, following the argument of the second step. The only difference is that $\pi(G_{2i-1,j})$ now might be larger than p_{\max} but can be bounded using (4.3). \square

Combining Lemma 4.3.3 (both $\pi(G_{2i,j}) = \pi(S_{i,j,3}) = \dots = \pi(S_{i,j,k})$ and the bounds on the ratio of $\pi(G_{2i-1,j})$ to $\pi(G_{2i,j})$) and $\alpha_j = o(1) \ll c$, we can now use $f_{\mu,\lambda}$ to derive the excess demand of $G_{2i-1,j}$ from the NM market over $\mathcal{R}_{i,j}$, given $\pi(G_{2i-1,j})$ and $\pi(G_{2i,j})$. From now on, for each group $\mathcal{R}_{i,j}$, $i \in [n]$ and $j \in [0 : 4t]$, we let

$$\pi_{i,j} = \pi(G_{2i-1,j}) + \pi(G_{2i,j}).$$

Next note that only one trader is interested in AUX_j and her budget is $(1 - \beta)n\pi_{j,0}$. From this we have the following lemma.

Lemma 4.3.4. *Let \mathbf{p} be an ϵ -additively approximate market equilibrium of $M_{\mathbf{P}}$ with $\epsilon = 1/n^8$. If we scale \mathbf{p} so that $\pi_{j,0} = 2$ for some $j \in [n]$, then $\pi(AUX_j) \geq 1 - O(\epsilon/n)$.*

Proof. As the total supply of AUX_j is $2n(1 - \beta)$, we have

$$2n(1 - \beta) \leq (2n(1 - \beta) + \epsilon)\pi(AUX_j).$$

This finishes the proof of the lemma. \square

Second, by using the strong connectivity of \mathcal{G} and the property that every vertex in \mathcal{G} has the same total in-weight and out-weight, we can follow the proof of Lemma 3.3.3 (replacing m with $4t$) to prove the following.

Lemma 4.3.5. *Let \mathbf{p} denote an ϵ -additively approximate equilibrium of $M_{\mathbf{P}}$. Let*

$$\pi_{\max} = \max_{i,j} \pi_{i,j} \quad \text{and} \quad \pi_{\min} = \min_{i,j} \pi_{i,j}$$

where the max and min are both taken over all $i \in [n]$ and $j \in [0 : 4t]$. If we scale \mathbf{p} so that $\pi_{\min} = 2$, then we must have $\pi_{\max} = 2 + O(\epsilon t)$.

Then we can follow the proof of Lemma 3.3.4 to prove the following bound on $\pi(AUX_j)$:

Lemma 4.3.6. *Let \mathbf{p} denote an ϵ -additively approximate market equilibrium of $M_{\mathbf{P}}$ with $\epsilon = 1/n^8$. If we scale \mathbf{p} so that $\pi_{j,0} = 2$ for some $j \in [n]$, then $\pi(AUX_j) \leq 1 + O(\epsilon t)$.*

From now on, we always use \mathbf{p} to denote the scaled price vector with $\pi_{\min} = 2$. Using Lemma 4.3.4, 4.3.5 and 4.3.6 together, we have

$$2 \leq \pi_{i,j} = \pi(G_{2i-1,j}) + \pi(G_{2i,j}) \leq 2 + O(\epsilon t) \quad \text{and} \quad \pi(AUX_i) = 1 \pm O(\epsilon t), \quad (4.4)$$

for all $i \in [n]$ and $j \in [0 : 4t]$, where the last equation follows from

$$(\pi_{i,0}/2)(1 - O(\epsilon/n)) \leq \pi(AUX_i) \leq (\pi_{i,0}/2)(1 + O(\epsilon t)).$$

For convenience, we let $\delta = \epsilon t$.

Recall that we use H_i to denote the good $G_{i,4t}$. For each $i \in [n]$, we let

$$\theta_i = \frac{\pi(H_{2i-1}) + \pi(H_{2i})}{2}$$

From (4.4) we get the following corollary:

Corollary 4.3.1. *For every $i \in [n]$, we have $1 \leq \theta_i \leq 1 + O(\delta)$.*

Next we use Walras' law to show that the excess demand of each good is close to 0 from both sides:

Lemma 4.3.7. *If \mathbf{p} is an ϵ -additively approximate equilibrium of $M_{\mathbf{P}}$. Then there is a vector $\mathbf{z} \in Z(\mathbf{p})$ such that $|\mathbf{z}|_{\infty} \leq O(\epsilon nt)$.*

Proof. Given a vector $\mathbf{z} \in Z(\mathbf{p})$ and a good G in $M_{\mathbf{P}}$, we let $z(G)$ denote the excess demand of G in \mathbf{z} . By definition, we know there is a vector $\mathbf{z} \in Z(\mathbf{p})$ such that $z(G) \leq \epsilon$ for all G , thus $|z(G)| \leq \epsilon$ for goods G with positive excess demand. By Walras' law we also have $\mathbf{z} \cdot \mathbf{p} = 0$. By Lemma 4.3.3, 4.3.4, 4.3.5 and 4.3.6, we know that all prices are close to each other. As the total number of goods in $M_{\mathbf{P}}$ is $O(nt)$ and $z(G) \leq \epsilon$ for all G , it follows from Walras' law that $|z(G)| \leq O(\epsilon nt)$ for all G with negative excess demand. \square

From now on, we let $\xi = \epsilon nt = \log n/n^7$. Now we are ready to recover a $(1/n)$ -well-supported Nash equilibrium of the polymatrix game \mathbf{P} from the price vector \mathbf{p} . Set \mathbf{x} to be the following $2n$ -dimensional nonnegative vector:

$$x_{2i-1} = \frac{\pi(H_{2i-1}) - (1 - \beta)\theta_i}{2\beta\theta_i} \quad \text{and} \quad x_{2i} = \frac{\pi(H_{2i}) - (1 - \beta)\theta_i}{2\beta\theta_i} \quad (4.5)$$

Recall that $\beta = \alpha_{4t} = 1/n$. It is easy to verify that $x_{2i-1} + x_{2i} = 1$ for each $i \in [n]$. Here $x_i \geq 0$ follows from Lemma 4.3.3. To finish the proof, we prove the following theorem:

Theorem 7. *When n is sufficiently large, \mathbf{x} from (4.5) is a $(1/n)$ -well-supported Nash equilibrium of \mathbf{P} .*

We need the following key lemma to establish Theorem 7. Recall that G_i is $G_{i,0}$.

Lemma 4.3.8. *For every $i \in [n]$, we have*

$$\begin{aligned} \frac{1 + \alpha_0}{\pi(G_{2i-1})} = \frac{1 - \alpha_0}{\pi(G_{2i})} &\Rightarrow \frac{1 + \beta}{\pi(H_{2i-1})} = \frac{1 - \beta}{\pi(H_{2i})} \quad \text{and} \\ \frac{1 - \alpha_0}{\pi(G_{2i-1})} = \frac{1 + \alpha_0}{\pi(G_{2i})} &\Rightarrow \frac{1 - \beta}{\pi(H_{2i-1})} = \frac{1 + \beta}{\pi(H_{2i})} \end{aligned}$$

Before proving Lemma 4.3.8, we use it to prove Theorem 7:

Proof of Theorem 7. Assume for contradiction that the vector \mathbf{x} we construct from \mathbf{p} in (4.5) is not a $(1/n)$ well-supported Nash equilibrium of \mathbf{P} . Then without loss of generality, we assume that

$$\mathbf{x}^T \cdot \mathbf{P}_1 > \mathbf{x}^T \cdot \mathbf{P}_2 + 1/n \quad (4.6)$$

where \mathbf{P}_1 and \mathbf{P}_2 denote the first and second columns of \mathbf{P} , respectively, but $x_2 > 0$.

To reach a contradiction, by Lemma 4.3.8, it suffices to show that (4.6) implies that

$$\frac{1 + \alpha_0}{\pi(G_1)} = \frac{1 - \alpha_0}{\pi(G_2)} \quad (4.7)$$

because it then implies that $(1 + \beta)/\pi(H_1) = (1 - \beta)/\pi(H_2)$ and thus, $x_2 = 0$ by (4.5).

To prove (4.7), we first compare the total money spent on goods G_1 and G_2 from all traders in $M_{\mathbf{P}}$ except the two traders in the \mathbf{PR} market over G_1 and G_2 , and show that the money spent on G_1 is considerably larger. Given that the prices of G_1 and G_2 are very close, it implies that the demand of G_1 from these traders is strictly larger than that of G_2 . As \mathbf{p} is an approximate market equilibrium and G_1, G_2 have the same total supply in $M_{\mathbf{P}}$, we have that the \mathbf{PR} market over G_1 and G_2 must demand strictly more G_2 than G_1 to have things balanced, which can happen only when (4.7) holds.

We start by enumerating all traders that are interested in G_1 and G_2 except the two traders in the \mathbf{PR} market over G_1 and G_2 :

1. For each $i \in [2n]$, there is a $(P_{i,1}, H_i : G_1)$ -trader. The total money these traders spend on G_1 is given by

$$\sum_{i \in [2n]} P_{i,1} \cdot \pi(H_i) = \sum_{i \in [2n]} P_{i,1} \cdot (1 - \beta + 2\beta \cdot x_i) \cdot \theta_{\lceil i/2 \rceil}$$

2. For each $i \in [2n]$, there is a $(P_{i,2}, H_i : G_2)$ -trader. The total money these traders spend on G_2 is given by

$$\sum_{i \in [2n]} P_{i,2} \cdot \pi(H_i) = \sum_{i \in [2n]} P_{i,2} \cdot (1 - \beta + 2\beta \cdot x_i) \cdot \theta_{\lceil i/2 \rceil}$$

3. Recall r_{2j-1} and r_{2j} in (3.11). There is one $((1 - \beta)r_1, AUX_1 : G_1)$ -trader, who spends her budget $(1 - \beta)r_1 \cdot \pi(AUX_1)$ on G_1 . There is one $((1 - \beta)r_2, AUX_1 : G_2)$ -trader, who spends her budget $(1 - \beta)r_2 \cdot \pi(AUX_1)$ on G_2 .

We denote by M_1 (respectively, M_2) the total money these traders spend on G_1 (resp. G_2).

Then

$$M_1 = \sum_{i \in [2n]} P_{i,1} \cdot (1 - \beta + 2\beta \cdot x_i) \cdot \theta_{\lceil i/2 \rceil} + (1 - \beta)r_1 \cdot \pi(AUX_1)$$

Plugging in $\theta_{\lceil i/2 \rceil} \geq 1$, $\pi(AUX_1) \geq 1 - O(\delta)$ and the definition of r_1 , we get

$$M_1 \geq 2n(1 - \beta) + 2\beta \cdot \mathbf{x}^T \cdot \mathbf{P}_1 - O(n\delta)$$

Similarly, we also have the total money spent on G_2 is

$$M_2 = \sum_{i \in [2n]} P_{i,2} \cdot (1 - \beta + 2\beta \cdot x_i) \cdot \theta_{\lceil i/2 \rceil} + (1 - \beta)r_2 \cdot \pi(AUX_1)$$

Plugging in $\theta_{\lceil i/2 \rceil} \leq 1 + O(\delta)$, $\pi(AUX_2) \leq 1 + O(\delta)$ and the definition of r_2 , we get

$$M_2 \leq 2n(1 - \beta) + 2\beta \cdot \mathbf{x}^T \cdot \mathbf{P}_2 + O(n\delta)$$

Combining these two bounds with (4.6), we get

$$M_1 \geq M_2 + 2\beta \cdot (1/n) - O(n\delta) = M_2 + \Theta(\beta/n)$$

since $\beta/n = 1/n^2 \gg n\delta$. Hence the difference between the demands for G_1 and G_2 from these traders is:

$$\frac{M_1}{\pi(G_1)} - \frac{M_2}{\pi(G_2)} \geq \frac{M_2 + \Theta(\beta/n)}{\pi(G_1)} - \frac{M_2(1 + \alpha_0)}{\pi(G_1)(1 - \alpha_0)} = \frac{\Theta(\beta/n)}{\pi(G_1)} - \frac{M_2}{\pi(G_1)} \cdot \frac{2\alpha_0}{1 - \alpha_0} = \omega(\xi)$$

where the last inequality used $M_2 = O(n)$, $\alpha_0 = 1/n^5$, $\beta = 1/n$, and $\xi = \log n/n^7$.

The only other traders that are interested in G_1, G_2 are the two traders in the price-regulating market over $\mathcal{R}_{1,0}$ denoted by T_1 and T_2 . Also from the construction of $M_{\mathbf{P}}$, the total supply of G_1 is exactly the same as that of G_2 . By Lemma 4.3.7, we know that the total demand of G_1 from T_1 and T_2 must be strictly smaller than the total demand of G_2 from them, which in turn implies that the total demand of G_1 from T_1 and T_2 must be strictly smaller than the total supply of G_1 from T_1 and T_2 by Walras' law.

Assume (4.7) does not hold, then by Lemma 4.3.3 we must have

$$\frac{1 + \alpha_0}{\pi(G_1)} > \frac{1 - \alpha_0}{\pi(G_2)}$$

This implies that the (unique) optimal bundle of T_1 is to buy back her initial endowment of G_1 and thus, the total demand of T_1 and T_2 for G_1 is at least as much as the total supply of G_1 from T_1 and T_2 , contradicting with Lemma 4.3.7. The theorem then follows. \square

Finally we prove Lemma 4.3.8, which crucially relies on properties (the function $f_{\mu,\gamma}$ in particular) of the NM markets added in $M_{\mathbf{P}}$. By induction it suffices to prove

Lemma 4.3.9. *For every $i \in [n]$ and $j \in [4t]$, we have*

$$\begin{aligned} \frac{1 + \alpha_{j-1}}{\pi(G_{2i-1,j-1})} = \frac{1 - \alpha_{j-1}}{\pi(G_{2i,j-1})} &\Rightarrow \frac{1 + \alpha_j}{\pi(G_{2i-1,j})} = \frac{1 - \alpha_j}{\pi(G_{2i,j})} \quad \text{and} \\ \frac{1 - \alpha_{j-1}}{\pi(G_{2i-1,j-1})} = \frac{1 + \alpha_{j-1}}{\pi(G_{2i,j-1})} &\Rightarrow \frac{1 - \alpha_j}{\pi(G_{2i-1,j})} = \frac{1 + \alpha_j}{\pi(G_{2i,j})} \end{aligned}$$

To this end, we examine a group $\mathcal{R}_{i,j}$, $i \in [n]$ and $j \in [4t]$, more closely. For convenience, we scale the price vector \mathbf{p} again so that $\pi_{i,j} = \pi(G_{2i-1,j}) + \pi(G_{2i,j}) = 2$. Note that what we need to prove in Lemma 4.3.9 remains the same after scaling. We are interested in the total demand of $G_{2i-1,j}$ from all traders in $M_{\mathbf{P}}$ except those two traders in the price-regulating market \mathbf{PR} over $\mathcal{R}_{i,j}$.

First of all, for the NM market over $\mathcal{R}_{i,j}$, we use $f(x)$ to denote the excess demand (within the NM market only) for $G_{2i-1,j}$ when the price of $G_{2i-1,j}$ is $1 + x$ and the prices

of $G_{2i,j}, S_{i,j,3}, \dots, S_{i,j,k}$ are $1 - x$. Let $\mu = d^*n = O(n)$ and $\gamma = 1/n^6$. Then $f \equiv f_{\mu,\gamma}$ in (4.2) and hence has the following properties:

$$|f(0)| = O(\mu\gamma) \quad \text{and} \quad |f(x) - f(0) - \mu dx| \leq |\mu x/D|, \quad \text{for all } x \in [-c, c] \quad (4.8)$$

where $D = \max\{20, 20/d\}$ and $c > 0$ are both constants independent of n . So when n is sufficiently large, we have $\beta = \alpha_{4t} = 1/n \ll c$. Next we use $h(x, y)$ to denote

$$h(x, y) = \text{excess demand of } G_{2i-1,j} \text{ from all traders except those in the PR over } \mathcal{R}_{i,j}$$

when the price of $G_{2i-1,j-1}$ is $1 + y$, the price of $G_{2i-1,j}$ is $1 + x$, and the prices of $G_{2i,j}, S_{i,j,3}, \dots, S_{i,j,k}$ are $1 - x$. By Lemma 4.3.3 and 4.3.5, we are interested in x, y satisfying

$$|x| \leq \alpha_j \quad \text{and} \quad |y| \leq \alpha_{j-1} + O(\delta).$$

Using f , we obtain the following more explicit form of h since other than the NM and PR markets over $\mathcal{R}_{i,j}$, there are n units of $G_{2i-1,j}$ and only one $(n, G_{2i-1,j-1} : G_{2i-1,j})$ -trader interested in $G_{2i-1,j}$:

$$h(x, y) = f(x) + \frac{n(1+y)}{1+x} - n = f(x) - \frac{nx}{1+x} + \frac{ny}{1+x}$$

We now use (4.8) to prove the following useful lemma about $h(x, y)$:

Lemma 4.3.10. *For all x and y with $|x| \leq 3|y|$ and $|y| = \alpha_{j-1} \pm O(\delta)$, we have*

$$h(x, y) > ny/2 \quad \text{if } y > 0 \quad \text{and} \quad h(x, y) < ny/2 \quad \text{if } y < 0$$

Proof. For $x/(1+x)$, we can approximate it by x when $|x|$ is small:

$$|x/(1+x) - x| = x^2/(1+x) \leq 2x^2$$

For $f(x)$, by (4.8) we can approximate it by μdx :

$$|f(x) - \mu dx| \leq |f(0)| + |\mu x/D| = O(\mu\gamma) + |nx/20|$$

where we used $D = \max\{20, 20/d\}$ and the assumption that $1 - 1/D \leq d^*d \leq 1$.

Thus, we can approximate $f(x) - nx/(1+x)$ using $(\mu d - n)x$ where the absolute value of error is bounded by $2nx^2 + O(\mu\gamma) + |nx/20|$. On the other hand, by the choice of d^* ,

$$-nx/20 \leq -nx/D \leq (\mu d - n)x \leq 0$$

Therefore, we can bound the absolute value $|f(x) - nx/(1+x)|$ by

$$2nx^2 + O(\mu\gamma) + |nx/10|$$

From $\mu = O(n)$, $\gamma = 1/n^6$, $|x| \leq 3|y|$ and $|y| = \alpha_{j-1} \pm O(\delta)$, this can be trivially bounded from above by $|ny/3|$. The lemma then follows since $|ny/(1+x)| > |5ny/6|$. \square

We are now ready to prove Lemma 4.3.9:

Proof of Lemma 4.3.9. We scale \mathbf{p} so that $\pi(G_{2i-1,j}) + \pi(G_{2i,j}) = 2$. Assume that

$$\frac{1 + \alpha_{j-1}}{\pi(G_{2i-1,j-1})} = \frac{1 - \alpha_{j-1}}{\pi(G_{2i,j-1})} \quad \text{or} \quad \frac{1 - \alpha_{j-1}}{\pi(G_{2i-1,j-1})} = \frac{1 + \alpha_{j-1}}{\pi(G_{2i,j-1})}. \quad (4.9)$$

We refer to the case that the first equation of (4.9) holds as Case 1, and the case that the second equation holds as Case 2. In Case 1 we have $y = \alpha_{j-1} \pm O(\delta)$ and in Case 2 we have $y = -\alpha_{j-1} \pm O(\delta)$ by Lemma 4.3.3 and Lemma 4.3.5. Moreover, from Lemma 4.3.3 we have $|x| \leq \alpha_j$ and thus, it always holds that $|x| \leq 3|y|$, since $\alpha_j = 2\alpha_{j-1} = \omega(\delta)$. Therefore, we can conclude from Lemma 4.3.10 that

$$h(x, y) > ny/2 \quad (\text{in Case 1}) \quad \text{or} \quad h(x, y) < ny/2 \quad (\text{in Case 2})$$

holds respectively. Because $n\alpha_{j-1} \geq n\alpha_{4t} \gg \xi$, Lemma 4.3.7 implies that the excess demand of $G_{2i-1,j}$, within the price-regulating market \mathbf{PR} over $\mathcal{R}_{i,j}$, is either strictly negative or strictly positive, respectively.

When it is strictly negative (i.e. in Case 1), we know that the first trader T_1 of the price-regulating market does not spend all her budget on $G_{2i-1,j}$. This combined with Lemma 4.3.3 implies

$$\frac{1 + \alpha_j}{\pi(G_{2i-1,j})} = \frac{1 - \alpha_j}{\pi(G_{2i,j})}$$

Similarly when it is strictly positive (in Case 2), we conclude that the second trader T_2 must be interested in $G_{2i-1,j}$ as well. This combined with Lemma 4.3.3 implies that

$$\frac{1 - \alpha_j}{\pi(G_{2i-1,j})} = \frac{1 + \alpha_j}{\pi(G_{2i,j})}$$

This finishes the proof of the lemma. \square

This page intentionally left blank.

Part II

Revenue Maximization

This page intentionally left blank.

Chapter 5

Preliminaries

5.1 Introduction

Revenue Maximization is a major goal in mathematical economics; and a central problem of the field relevant to this goal is *Multi-dimensional Bayesian Item Pricing*.

Item Pricing. There is a buyer interested in obtaining one of n heterogeneous items offered by a seller, indexed by $[n] = \{1, \dots, n\}$. We focus on the case of a *quasi-linear* buyer; her utility if she buys item i is $v_i - p_i$, where v_i and p_i are respectively her value and the price for item $i \in S$. Given prices p_1, \dots, p_n for the items, the buyer will receive an item that maximizes her utility, or nothing if all items have negative utility. The seller has full access to a probability distribution \mathcal{D} from which the buyer's valuations $\mathbf{v} = (v_1, \dots, v_n)$ for the items are drawn, and wants to assign prices to the items in a way that maximizes her expected revenue (payment from the buyer). Given \mathcal{D} as input, our goal is to find prices for the items that maximize the seller's expected revenue.

Although the deterministic setting favours simplicity, it turns out that it is possible to improve the seller's expected revenue [Thanassoulis, 2004] using randomization. This is achieved through a generalization known as the *Multi-dimensional Mechanism Design* or *Lottery Pricing* setting which is defined as follows:

Lottery Pricing. Here, the seller, instead of prices, offers a *menu* (set) M of *lotteries*. A lottery is of the form (\mathbf{x}, p) , where $p \in \mathbb{R}^+$ is its price and $\mathbf{x} = (x_1, \dots, x_n)$ is a non-

negative *allocation vector* with each x_i being the probability that the buyer receives item i and $\sum_{i \in [n]} x_i \leq 1$. After a menu M is chosen, the buyer draws a valuation \mathbf{v} from \mathcal{D} and receives a lottery that maximizes her expected utility $\sum_i x_i \cdot v_i - p$, or nothing if every lottery in M has a negative utility. Let $\Pr[\mathbf{v}]$ be the probability of $\mathbf{v} \sim \mathcal{D}$; and $\text{REV}_M(\mathbf{v})$ be the price of the lottery the buyer receives when her valuation is \mathbf{v} . Given \mathcal{D} as input, our goal is to find a menu M of lotteries that maximizes the seller's expected revenue $\text{REV}(M) = \sum_{\mathbf{v} \in \mathcal{D}} \Pr[\mathbf{v}] \cdot \text{REV}_M(\mathbf{v})$.

Notice that Item Pricing is the special case of Lottery Pricing where $x_{ij} \in \{0, 1\}$. These two settings can be further divided in terms of the buyer's preference: A *Unit-demand* buyer is interested in at most one item; as in the settings described above. This corresponds to items that are perfect substitutes of each other, e.g. when the seller is a car dealer. On the other hand, an *Additive* buyer is interested in obtaining any number of items, e.g. when visiting a mall. This is equivalent to offering lotteries where $0 \leq x_{ij} \leq 1$ for all i and j in the randomized setting and offering a menu of subsets of items in the deterministic setting. In this thesis, we will concentrate on Unit-demand buyers.

For the rest of this part, we focus on the well-studied case [Chawla *et al.*, 2007; Chawla *et al.*, 2010a; Cai and Daskalakis, 2011] that $\mathcal{D} = \times_{i=1}^n \mathcal{D}_i$ is a *product distribution*, i.e., the buyer's valuations for the items are mutually independent random variables. We assume that the n (marginal) distributions \mathcal{D}_i are discrete and the values of their supports with the corresponding probabilities are rational numbers given explicitly in the input. This seemingly simple computational problem exhibits a very rich underlying structure that was not well understood prior to our work. For example, even for Item Pricing, the optimal prices are not necessarily in the support of \mathcal{D} (see [Cai and Daskalakis, 2011] for a simple example with two items and distributions of support 2). So, a priori, it is not even clear if optimal prices can be described with polynomially many bits in the size of the input.

5.2 Related Work

Optimal mechanism design is well-understood in single-parameter settings, where there is only one item for sale, for which the seminal work of Myerson [Myerson, 1981] gives a

closed-form characterization for the optimal prices and proves that it can achieve as much revenue as any sophisticated, randomized mechanism, therefore randomization does not improve revenue in this case. However, the general case with multiple items turns out to be more complex. The latter, has been extensively studied by economists (see, e.g., [Wilson, 1996] for a survey and [McAfee *et al.*, 1989] for a simple additive case with two items) and recently has also drawn the attention of the theoretical computer science community; which concentrated on understanding both the structure and the complexity of optimal mechanisms, as well as developing simple and computationally efficient mechanisms that are approximately optimal.

5.2.1 Work on Item Pricing

For the Unit-demand Item Pricing setting, Chawla, Hartline and Kleinberg [Chawla *et al.*, 2007] show that techniques from Myerson’s work can be used to obtain an analogous closed-form characterization for prices that extract a factor 3 approximation of the optimal expected revenue (subsequently improved to 2 [Chawla *et al.*, 2010a]). Cai and Daskalakis [Cai and Daskalakis, 2011] study the case that the distributions are restricted to be monotone hazard-rate and obtain a polynomial-time approximation scheme. On the lower bound side, Guruswami *et al.* [Guruswami *et al.*, 2005] and subsequently Briest [Briest, 2008] studied Unit-demand Item Pricing when the values for the items are *correlated*, respectively obtaining APX-hardness and $\Omega(n^\epsilon)$ inapproximability, for some constant $\epsilon > 0$. Similarly, Papadimitriou and Pierrakos [Papadimitriou and Pierrakos, 2011] show that the extension of Myerson’s single-parameter setting to bidders with correlated valuations is APX-hard, even for the case of 3 bidders. However, unlike the Myerson setting, randomization improves revenue and the optimal menu can be found in polynomial time via linear programming [Dobzinski *et al.*, 2011]. Regarding Unit-demand item pricing with product distributions, Daskalakis, Deckelbaum and Tzamos [Daskalakis *et al.*, 2014a] obtain SQRT-SUM-hardness when either the support values or the probabilities are irrational. Notice that their reduction relies on the fact that, if irrationality is allowed, it is SQRT-SUM-hard to compare the revenue of two price-vectors. This does not extend to the standard discrete model we consider, for which the exact revenue of a price-vector can be computed efficiently. Prior to

the work presented in this thesis, the existence of closed-form characterizations, or efficient algorithms, for the *optimal* pricing was a well-known open problem [Chawla *et al.*, 2007; Cai and Daskalakis, 2011].

A different line of work considers the case of an additive buyer. Hart and Nisan [Hart and Nisan, 2012] studied two simple deterministic mechanisms for product distributions: selling items separately or selling a full bundle (grand bundling). They showed that selling items separately and grand bundling are respectively $\Omega(1/\log^2 n)$ and $\Omega(1/\log n)$ approximations of the optimal revenue (later improved by Li and Yao [Li and Yao, 2013] to $\Omega(1/\log n)$ for both schemes, which is tight [Hart and Nisan, 2012]). Although neither of these two mechanisms can achieve a constant factor approximation of the optimal revenue for product distributions, Babaioff *et al.* [Babaioff *et al.*, 2014] showed that the *best* of selling separately and grand bundling, gives a $(1/6)$ -approximation. Extending this work, Daskalakis *et al.* [Daskalakis *et al.*, 2013; Daskalakis *et al.*, 2014b] studied conditions for grand-bundling mechanisms to be optimal, and showed that this happens if and only if two stochastic dominance conditions hold. Considering the case of many buyers, Yao [Yao, 2015] introduced a new approach for reducing the k -item n -bidder auction to k -item 1-bidder auctions. He gave a deterministic mechanism that yields at least a constant fraction of the optimal revenue for k -item n -bidder auctions with arbitrary independent valuation distributions; and a closed form for the optimal revenue when all the nk values are i.i.d.. Finally, Rubinstein [Rubinstein, 2016] considered the restriction of the problem to simple mechanisms and defined the notion of partition mechanisms, where the seller partitions the items into disjoint bundles and posts a price for each bundle; allowing the buyer to buy any number of bundles. He obtained a PTAS for the problem of finding a revenue maximizing partition mechanism (and also proved that the problem is strongly NP-hard).

5.2.2 Work on Lottery Pricing

A lot of work has also been done for the general setting of Lottery Pricing. For the Unit-demand case, Thanassoulis [Thanassoulis, 2004] showed that unlike the single-parameter setting, where the optimal mechanism is deterministic [Myerson, 1981], an optimal mechanism for two items drawn uniformly and independently from $[5, 6]$ must involve randomiza-

tion. However, Chawla et al. [Chawla *et al.*, 2010b] showed that for independent valuations randomization can improve the optimal revenue by at most a factor of 4. In contrast, Briest et al. [Briest *et al.*, 2010] showed that when \mathcal{D} is correlated the improvement in revenue by randomization can be unbounded even for instances with four items; they also showed that if \mathcal{D} is given explicitly, by listing the probability of every valuation vector, then one can solve a linear program and find an optimal mechanism in polynomial time. Pavlov [Pavlov, 2010] completely characterized optimal mechanisms under both the unit-demand and additive settings when there are two items and their values are drawn independently from distributions that meet certain conditions. On a similar note, Wang and Tang [Wang and Tang, 2014] studied conditions under which the optimal randomized mechanism has “simple” menus, i.e., menus that either satisfy a monotonicity property (allocation probabilities of items increase along with their prices), or consist of a small number of lotteries. For the additive setting, Manelli and Vincent [Manelli and Vincent, 2006] gave an example where randomization strictly improves revenue and also provided certain sufficient conditions under which deterministic mechanisms are optimal.

On the other hand, work on lower-bounds is limited. Hart and Nisan [Hart and Nisan, 2013] introduced the notion of *menu size*, the minimum number of lotteries needed to achieve the optimal revenue. They showed that for an additive buyer there exists a correlated continuous distribution for which no mechanism of finite menu size can achieve a positive fraction of the optimal revenue. In our work [Chen *et al.*, 2015], we obtain a similar result for a unit-demand buyer and prove that, even for discrete product distributions with support size 2 for each of the items, there are instances that require an exponential number of lotteries to achieve the optimal revenue. These results do not preclude the existence of an efficient algorithm for the mechanism design problem, as we will see in chapter 7. However, for an additive buyer, Daskalakis et al. [Daskalakis *et al.*, 2014a] showed that, unless $P^{\#P} \subseteq ZPP$, there is no efficient algorithm that implements an optimal mechanism for product distributions. Note that this result does not extend to the unit-demand case, which prior to the work presented in this thesis remained a well-known open problem.

We close this section with table 5.1, where we summarize the computational landscape of Revenue Maximization prior to our work and indicate the problems resolved by us.

Problem	Best Algorithm	Complexity	Menu-Size
<i>Additive Lottery Pricing</i>	1/6-Approximation	#P-hard	Exponential
<i>Unit-demand Lottery Pricing</i>	Open	Open (R)	Open (R)
<i>Additive Item Pricing</i>	1/6-Approximation	Open	Open
<i>Unit-demand Item Pricing</i>	1/2-Approximation	Open (R)	n

Table 5.1: Computational landscape of Revenue Maximization before this thesis. We use “R” to indicate that the problem is resolved in this thesis.

5.3 Our Results and Organization of this part

For the rest of this part, we concentrate on the *Computational Complexity* of both Item and Lottery Pricing with a Unit-demand buyer. Furthermore, we are also interested in the *Menu-size Complexity* of Unit-demand Lottery pricing.

In Chapter 6, we *resolve the computational complexity* of Unit-demand Item Pricing and prove that the problem is NP-hard even for distributions of support size at most 3. Furthermore, we obtain *a structural characterization* for the optimal prices as solutions to a Linear Program restricted to cells of a partition of the price space. This *yields NP-membership* for the problem; the previous upper-bound being NEXP. Finally, we obtain a *polynomial time algorithm* for the case that all marginal distributions have support 2.

In Chapter 7 we *resolve the computational complexity* of Unit-demand Lottery Pricing. We note that there is [Chen *et al.*, 2015] an interesting family \mathcal{D}^* of distributions with support size 2 where randomization not only improves revenue but an exponential number of lotteries is required to achieve optimality. Although \mathcal{D}^* trivially rules out any efficient algorithm that lists explicitly all lotteries in an optimal menu, there is a deterministic polynomial-time algorithm that, given any $\mathbf{v} \in \mathcal{D}^*$, outputs a lottery $\ell_{\mathbf{v}}$ such that $\{\ell_{\mathbf{v}} : \mathbf{v} \in \mathcal{D}^*\}$ is an optimal menu for \mathcal{D}^* . We prove that *no universal efficient (possibly randomized) algorithm that computes an optimal menu in this fashion exists for product distributions, unless $P^{\text{NP}} = P^{\#P}$* . We also prove that for the special case of support-size 2 and the same high value for all items *randomization does not improve revenue*, and therefore one can use our algorithm for Item Pricing and distributions of support size 2.

Chapter 6

The Complexity of Optimal Multidimensional Pricing

6.1 Introduction

In this chapter, we concentrate on Unit-demand Item Pricing with *discrete product distributions*. As we mentioned in Chapter 5, this seemingly simple computational problem has a very rich underlying structure. Prior to our work, the complexity of the problem was a well-known open problem [Chawla *et al.*, 2007; Cai and Daskalakis, 2011; Daskalakis *et al.*, 2014a]. In fact, it was not even known whether the optimal prices can be described with polynomially many bits in the size of the input; or whether the problem is in NP.

Here, we address this by proving NP membership for arbitrary distributions and NP-hardness for distributions of support-size 3. To complement this result, we present a polynomial time algorithm for arbitrary distributions of support-size 2.

6.2 Preliminaries

6.2.1 Problem Definition and Main Results

Recall that in the Unit-demand Item Pricing setting, there is one buyer and one seller with n items indexed by $[n] = \{1, 2, \dots, n\}$. The buyer is interested in buying at most one item (unit demand), and her valuations of the items are drawn from n independent discrete

distributions. In particular, we use $V_i = \{v_{i,1}, \dots, v_{i,|V_i|}\}$, $i \in [n]$, to denote the support of the value distribution of item i , where $0 \leq v_{i,1} < \dots < v_{i,|V_i|}$. We also use $q_{i,j} > 0$, $j \in [|V_i|]$, to denote the probability of item i having value $v_{i,j}$, with $\sum_j q_{i,j} = 1$. Let $V = \times_{i=1}^n V_i$. We use $\Pr[\mathbf{v}]$ to denote the probability of the valuation vector being $\mathbf{v} = (v_1, \dots, v_n) \in V$, i.e., the product of $q_{i,j}$'s over i, j such that $i \in [n]$ and $v_i = v_{i,j}$.

The n distributions, i.e., V_i and $q_{i,j}$, are given to the seller explicitly and she then assigns a price $p_i \geq 0$ to each item. Once the price vector $\mathbf{p} = (p_1, \dots, p_n) \in \mathbb{R}_+^n$ is fixed, the buyer draws her values $\mathbf{v} = (v_1, \dots, v_n)$ from the n distributions independently, i.e., $\mathbf{v} \in V$ with probability $\Pr[\mathbf{v}]$. The buyer is quasi-linear, i.e., her utility for item i equals $v_i - p_i$. Let

$$\mathcal{U}(\mathbf{v}, \mathbf{p}) = \max_{i \in [n]} (v_i - p_i).$$

If $\mathcal{U}(\mathbf{v}, \mathbf{p}) \geq 0$, the buyer selects an item $i \in [n]$ that maximizes her utility $v_i - p_i$, and the seller has revenue p_i . Otherwise, she does not select any item and the seller has revenue 0.

Knowing the value distributions and the buyer's behaviour described above, the seller's objective is to compute a price vector $\mathbf{p} \in \mathbb{R}_+^n$ that maximizes the expected revenue

$$\mathcal{R}(\mathbf{p}) = \sum_{i \in [n]} p_i \cdot \Pr [\text{buyer selects item } i].$$

We use `ITEM-PRICING` to denote the following decision problem: The input consists of n discrete distributions, with $v_{i,j}$ and $q_{i,j}$ all being rational and encoded in binary, and a rational number $t \geq 0$. The problem asks whether the supremum of the expected revenue $\mathcal{R}(\mathbf{p})$ over all price vectors $\mathbf{p} \in \mathbb{R}_+^n$ is at least t , where we use \mathbb{R}_+ to denote the set of nonnegative real numbers.

For the aforementioned decision problem to be well-defined, we need a tie-breaking rule that specifies which item the buyer selects when there are multiple items with maximum nonnegative utility. Here, we will use the *maximum price*¹ tie-breaking rule for convenience: when there are multiple items with maximum nonnegative utility, the buyer selects the item with the smallest index among items with the highest price. Notice that the critical part is to select an item with the highest price — any such selection results in the same revenue.

¹It may also be called the *maximum value* tie-breaking rule, since an item with the maximum price among a set of items with the same utility must also have the maximum value.

However we need to make such a choice so that we can talk about “the” item selected by the buyer in the proofs. We show in Section 6.2.2 that the choice of tie-breaking rule does not affect the supremum of the expected revenue (and hence, the complexity of the problem).

We are now ready to state our main results. First, we show in Section 6.3 that ITEM-PRICING is in NP.

Theorem 8. ITEM-PRICING is in NP.

Second, we present in Section 6.4 a polynomial-time algorithm for ITEM-PRICING when all the distributions have support size at most 2.

Theorem 9. ITEM-PRICING is in P when every distribution has support size at most 2.

As our main result, we resolve the computational complexity of the problem. We show that it is NP-hard even when all distributions have support size at most 3 (Section 6.5). In fact, in our paper [Chen *et al.*, 2014] we obtain the same result for identical independent distributions.

Theorem 10. ITEM-PRICING is NP-hard even when every distribution has support size 3.

6.2.2 Tie-Breaking Rules

In this section, we show that the supremum of the expected revenue over $\mathbf{p} \in \mathbb{R}_+^n$ is invariant to tie-breaking rules. Formally, a tie-breaking rule is a mapping from the set of pairs (\mathbf{v}, \mathbf{p}) with $\mathcal{U}(\mathbf{v}, \mathbf{p}) \geq 0$ to an item k such that $v_k - p_k = \mathcal{U}(\mathbf{v}, \mathbf{p})$.

We need some notation. Let B be the maximum price tie-breaking rule described earlier. We will denote by $\mathcal{R}(\mathbf{p})$ the expected revenue of \mathbf{p} under B , and by $\mathcal{R}(\mathbf{v}, \mathbf{p})$ the seller’s revenue under B when the valuation vector is $\mathbf{v} \in V$. Given a price vector \mathbf{p} and a valuation vector $\mathbf{v} \in V$, we also denote by $\mathcal{T}(\mathbf{v}, \mathbf{p})$ the set of items with maximum nonnegative utility (so $\mathcal{T}(\mathbf{v}, \mathbf{p}) = \emptyset$ iff $\mathcal{U}(\mathbf{v}, \mathbf{p}) < 0$).

We show the following:

Lemma 6.2.1. The supremum of the expected revenue over $\mathbf{p} \in \mathbb{R}_+^n$ is invariant to tie-breaking rules.

Proof. Let $v_{i,j}$ and $q_{i,j}$ denote the numbers that specify the distributions. Let B' be a tie-breaking rule. We will use $\mathcal{R}'(\mathbf{p})$ to denote the expected revenue of \mathbf{p} under B' and use $\mathcal{R}'(\mathbf{v}, \mathbf{p})$ to denote the seller's revenue under B' when the valuation vector is $\mathbf{v} \in V$.

It is clear that for any $\mathbf{p} \in \mathbb{R}_+^n$ and $\mathbf{v} \in V$, we have $\mathcal{R}(\mathbf{v}, \mathbf{p}) \geq \mathcal{R}'(\mathbf{v}, \mathbf{p})$ since B picks an item with the highest price among those that maximize the utility. Hence, it follows that $\sup_{\mathbf{p}} \mathcal{R}(\mathbf{p}) \geq \sup_{\mathbf{p}} \mathcal{R}'(\mathbf{p})$.

On the other hand, given any price vector $\mathbf{p} \in \mathbb{R}_+^n$, we consider

$$\mathbf{p}_\epsilon = (\max(0, p_1 - r_1\epsilon), \dots, \max(0, p_n - r_n\epsilon)) \in \mathbb{R}_+^n,$$

where $\epsilon > 0$ and r_i is the rank of p_i sorted in increasing order (when there are ties, the item with the smaller index is ranked higher). We claim that

$$\lim_{\epsilon \rightarrow 0^+} \mathcal{R}'(\mathbf{p}_\epsilon) = \mathcal{R}(\mathbf{p}). \quad (6.1)$$

It then follows from (6.1) that $\sup_{\mathbf{p}} \mathcal{R}'(\mathbf{p}) \geq \sup_{\mathbf{p}} \mathcal{R}(\mathbf{p})$, which gives the proof of the lemma.

To prove (6.1), we show that the following holds for any valuation vector $\mathbf{v} \in V$:

$$\lim_{\epsilon \rightarrow 0^+} \mathcal{R}'(\mathbf{v}, \mathbf{p}_\epsilon) = \mathcal{R}(\mathbf{v}, \mathbf{p}). \quad (6.2)$$

Observe that (6.1) follows from (6.2) since

$$\mathcal{R}(\mathbf{p}) = \sum_{\mathbf{v} \in V} \mathcal{R}(\mathbf{v}, \mathbf{p}) \cdot \Pr[\mathbf{v}] \quad \text{and} \quad \mathcal{R}'(\mathbf{p}_\epsilon) = \sum_{\mathbf{v} \in V} \mathcal{R}'(\mathbf{v}, \mathbf{p}_\epsilon) \cdot \Pr[\mathbf{v}].$$

To prove (6.2), we consider two cases. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) < 0$, then we have $\mathcal{U}(\mathbf{v}, \mathbf{p}_\epsilon) < 0$ when ϵ is sufficiently small, and thus, $\mathcal{R}(\mathbf{v}, \mathbf{p}) = \mathcal{R}'(\mathbf{v}, \mathbf{p}_\epsilon) = 0$. When $\mathcal{U}(\mathbf{v}, \mathbf{p}) \geq 0$, we make the following three observations. First, the utility of an item $i \in [n]$ under \mathbf{p}_ϵ is at least as high as that under \mathbf{p} . Second, if $v_i - p_i > v_j - p_j$ for some items $i, j \in [n]$, then under \mathbf{p}_ϵ the utility of item i remains strictly higher than that of item j , for ϵ sufficiently small. Third, if $v_i - p_i = v_j - p_j$ and $p_i > p_j$ (in particular, $p_i > 0$) for some $i, j \in [n]$, then under \mathbf{p}_ϵ the utility of item i is strictly higher than that of item j when $\epsilon \ll p_i$, as $r_i > r_j$. It follows from these observations that when ϵ is sufficiently small, B' must pick, given \mathbf{v} and \mathbf{p}_ϵ , an item $k \in [n]$ such that $p_k = \mathcal{R}(\mathbf{v}, \mathbf{p})$. (6.2) then follows from the definition of \mathbf{p}_ϵ . \square

We will henceforth always adopt the maximum price tie-breaking rule, and use $\mathcal{R}(\mathbf{v}, \mathbf{p})$ to denote the revenue of the seller with respect to this rule. One of the advantages of this rule is that the supremum of the expected revenue $\mathcal{R}(\mathbf{p})$ is always achievable, so it makes sense to talk about whether \mathbf{p} is optimal or not. In the following example, we point out that this does not hold for general tie-breaking rules.

Example: Suppose item 1 has value 10 with probability 1, item 2 has value 8 with probability 1/2 and value 12 with probability 1/2, and in case of tie the buyer prefers item 1. The supremum in this example is 11: set $p_1 = 10$ for item 1 and $p_2 = 12 - \epsilon$ for item 2. The buyer will buy item 1 with probability 1/2 (if her value for item 2 is 8) and item 2 with probability 1/2 (if her value for item 2 is 12). However, an expected revenue of 11 is not achievable: if we give price 12 to item 2, then the buyer will always buy item 1 and the revenue is 10. Note that the expected revenue for this tie-breaking rule is not a continuous function of the prices.

Before proving that the supremum is indeed always achievable under the maximum price rule, we start by showing that without loss of generality, we may focus the search for an optimal price vector in the set $P = \times_{i=1}^n [a_i, b_i]$, where $a_i = \min_j v_{i,j}$ and $b_i = \max_j v_{i,j}$ denote the minimum and maximum values in the support V_i , respectively.

Lemma 6.2.2. *For any price vector $\mathbf{p} \in \mathbb{R}_+^n$, there exists a $\mathbf{p}' \in P$ such that $\mathcal{R}(\mathbf{p}') \geq \mathcal{R}(\mathbf{p})$.*

Proof. First, it is straightforward that no price p_i should be above b_i ; if such a price exists, we can simply replace it by b_i and this will not decrease the expected revenue.

The non-trivial part is to argue that it is no loss of generality to assume that no price p_i is below a_i . Let $\mathbf{p} \in \times_{i=1}^n [0, b_i]$. Suppose that there exists $i \in [n]$ such that $p_i < a_i$, i.e., the set $L(\mathbf{p}) = \{i \in [n] : p_i < a_i\}$ is nonempty; otherwise, there is nothing to prove.

Fix an $i \in L(\mathbf{p})$ arbitrarily and let $S_i = \{j \in [n] : p_j < a_i\}$. We consider the price vector $\tilde{\mathbf{p}}$ defined by $\tilde{p}_j = \min\{b_j, a_i\}$ for $j \in S_i$ and $\tilde{p}_j = p_j$ otherwise. As $i \in S_i$, it follows that $S_i \neq \emptyset$ and therefore $\tilde{\mathbf{p}} \neq \mathbf{p}$ (in particular, $\tilde{p}_i = a_i$ now). It is also clear that $\tilde{\mathbf{p}} \in \times_{i=1}^n [0, b_i]$. It suffices to show that $\mathcal{R}(\tilde{\mathbf{p}}) \geq \mathcal{R}(\mathbf{p})$.

Indeed, note that $|L(\tilde{\mathbf{p}})| < |L(\mathbf{p})|$ so this process will terminate in at most n stages. After the last stage we will obtain a vector $\mathbf{p}' \in P$ whose expected revenue is lower bounded

by all the previous ones.

To prove that $\mathcal{R}(\tilde{\mathbf{p}}) \geq \mathcal{R}(\mathbf{p})$, we proceed as follows. Given any valuation vector $\mathbf{v} \in V$, we compare the revenue $\mathcal{R}(\mathbf{v}, \mathbf{p})$ to $\mathcal{R}(\mathbf{v}, \tilde{\mathbf{p}})$ and consider the following two cases:

- **Case 1:** On input (\mathbf{v}, \mathbf{p}) , the item selected by the buyer is not from S_i . We claim that the same item is selected on input $(\mathbf{v}, \tilde{\mathbf{p}})$. Indeed, we did not decrease prices of items in S_i , hence their utilities did not go up, while the utilities of the remaining items did not change. Therefore, the revenue does not change in this case, i.e., $\mathcal{R}(\mathbf{v}, \tilde{\mathbf{p}}) = \mathcal{R}(\mathbf{v}, \mathbf{p})$.
- **Case 2:** On input (\mathbf{v}, \mathbf{p}) , the item selected is from S_i . Then by the definition of S_i , the revenue $\mathcal{R}(\mathbf{v}, \mathbf{p})$ we get is certainly less than a_i . On input $(\mathbf{v}, \tilde{\mathbf{p}})$, we know that $\mathcal{U}(\mathbf{v}, \tilde{\mathbf{p}}) \geq 0$ (since item i must have nonnegative utility, i.e., $v_i - \tilde{p}_i = v_i - a_i \geq 0$) and thus, $\mathcal{T}(\mathbf{v}, \tilde{\mathbf{p}}) \neq \emptyset$. We claim that $\mathcal{R}(\mathbf{v}, \tilde{\mathbf{p}}) \geq a_i > \mathcal{R}(\mathbf{v}, \mathbf{p})$. To see this, we consider two sub-cases. If $\mathcal{U}(\mathbf{v}, \tilde{\mathbf{p}}) = 0$, then we must have $i \in \mathcal{T}(\mathbf{v}, \tilde{\mathbf{p}})$ and the claim follows from our choice of the maximum price tie-breaking rule. If $\mathcal{U}(\mathbf{v}, \tilde{\mathbf{p}}) > 0$, then every $j \in \mathcal{T}(\mathbf{v}, \tilde{\mathbf{p}})$ must satisfy $\tilde{p}_j \geq a_i$; otherwise, by definition of $\tilde{\mathbf{p}}$ we have $\tilde{p}_j = b_j$ and $v_j - \tilde{p}_j \leq 0$, a contradiction. From $\tilde{p}_j \geq a_i$ and $j \in \mathcal{T}(\mathbf{v}, \tilde{\mathbf{p}})$, we have $\mathcal{R}(\mathbf{v}, \tilde{\mathbf{p}}) \geq a_i$.

The lemma follows by combining the two cases. □

Now we show that the supremum is always achievable under the maximum price rule B .

Lemma 6.2.3. *There exists a price vector $\mathbf{p}^* \in P$ such that $\mathcal{R}(\mathbf{p}^*) = \sup_{\mathbf{p}} \mathcal{R}(\mathbf{p})$.*

Proof. By the compactness of P , it suffices to show that if a sequence of vectors $\{\mathbf{p}_i\}$ approaches \mathbf{p} , then

$$\mathcal{R}(\mathbf{p}) \geq \lim_{i \rightarrow \infty} \mathcal{R}(\mathbf{p}_i).$$

To this end, it suffices to show that, for any valuation vector $\mathbf{v} \in V$,

$$\mathcal{R}(\mathbf{v}, \mathbf{p}) \geq \lim_{i \rightarrow \infty} \mathcal{R}(\mathbf{v}, \mathbf{p}_i). \quad (6.3)$$

Given any valuation $\mathbf{v} \in V$, it is easy to check that $\mathcal{T}(\mathbf{v}, \mathbf{p}_i) \subseteq \mathcal{T}(\mathbf{v}, \mathbf{p})$ when i is sufficiently large. (Again consider two cases: $\mathcal{U}(\mathbf{v}, \mathbf{p}) < 0$ and $\mathcal{U}(\mathbf{v}, \mathbf{p}) \geq 0$.) (6.3) then follows, since $\mathcal{R}(\mathbf{v}, \mathbf{p})$ is the highest price of all items in $\mathcal{T}(\mathbf{v}, \mathbf{p})$ under the maximum price tie-breaking rule. □

6.3 Membership in NP

In this section we prove Theorem 8, i.e., ITEM-PRICING is in NP.

Proof of Theorem 8. We start with some notation. Given a price vector $\mathbf{p} \in \mathbb{R}_+^n$ and a valuation $\mathbf{v} \in V$, let $\mathcal{I}(\mathbf{v}, \mathbf{p}) \in [n] \cup \{\text{nil}\}$ denote the item picked by the buyer under the maximum price tie-breaking rule, with $\mathcal{I}(\mathbf{v}, \mathbf{p}) = \text{nil}$ iff $\mathcal{U}(\mathbf{v}, \mathbf{p}) < 0$. We will partition $P = \times_{i=1}^n [a_i, b_i]$ into equivalence classes so that two price vectors \mathbf{p}, \mathbf{p}' from the same class yield the same outcome for all valuations: $\mathcal{I}(\mathbf{v}, \mathbf{p}) = \mathcal{I}(\mathbf{v}, \mathbf{p}')$ for all \mathbf{v} .

Consider the partition of P induced by the following set of hyperplanes. For each item $i \in [n]$ and each value $s_i \in V_i$, we have a hyperplane $p_i = s_i$. For each pair of items $i, j \in [n]$ and pair of values $s_i \in V_i$ and $t_j \in V_j$, we have a hyperplane $s_i - p_i = t_j - p_j$, i.e., $p_i - p_j = s_i - t_j$. These hyperplanes partition our search space P into polyhedral cells, where the points in each cell lie on the same side of each hyperplane (either on the hyperplane or in one of the two open-halfspaces).

We claim that, for every valuation $\mathbf{v} \in V$, all the vectors in each cell yield the same outcome. Consider any cell C . It is defined by a set of equations and inequalities. Given any price vector $\mathbf{p} \in C$ and any value $s_i \in V_i$, let $V(\mathbf{p}, s_i)$ be the set of valuation vectors $\mathbf{v} \in V$ such that $v_i = s_i$ and the buyer ends up buying item i on (\mathbf{v}, \mathbf{p}) . We claim that $V(\mathbf{p}, s_i)$ does not depend on \mathbf{p} , i.e., it is the same set $V(s_i) = V(\mathbf{p}, s_i)$ over all $\mathbf{p} \in C$. To this end, first, if the points of C satisfy $p_i > s_i$ then $V(\mathbf{p}, s_i) = \emptyset$. So suppose that C satisfies $\mathbf{p} \leq s_i$. Consider any valuation vector $\mathbf{v} \in V$ with $v_i = s_i$. The valuation \mathbf{v} is in $V(\mathbf{p}, s_i)$ iff for all $j \neq i$, we have $s_i - p_i \geq v_j - p_j$, and in case of equality we have $s_i \geq v_j$ (iff $p_i \geq p_j$ due to the equality), and in case of further equality $s_i = v_j$ we have $i < j$. Because all points of the cell C lie on the same side of each hyperplane $s_i - p_i = v_j - p_j$, it follows that $V(\mathbf{p}, s_i)$ does not depend on \mathbf{p} . As a result, for any cell C and any $\mathbf{v} \in V$, all the points $\mathbf{p} \in C$ yield the same outcome $\mathcal{I}(\mathbf{v}, \mathbf{p})$.

Next, we show that it is easy to compute the supremum of the expected revenue $\mathcal{R}(\mathbf{p})$ over $\mathbf{p} \in C$, for each cell C . To this end, let $W_i = \cup_{s_i \in V_i} V(s_i) \subseteq V$ denote the set of valuations for which the buyer picks item i if the prices lie in the cell C , and let γ_i be the probability of W_i : $\gamma_i = \sum_{\mathbf{v} \in W_i} \Pr[\mathbf{v}]$. It turns out that γ_i can be computed efficiently, since

the probability of $V(s_i)$ can be computed efficiently as shown below (and W_i is the disjoint union of $V(s_i)$, $s_i \in V_i$).

Given $s_i \in V_i$, to compute the probability of $V(s_i)$, we note that $V(s_i)$ is actually the Cartesian product of subsets of V_j , $j \in [n]$. For each $j \neq i$, we can determine efficiently the subset of values $L_j \subseteq V_j$ such that the buyer prefers item i to j if i has value s_i and j has value from L_j . As a result, we have

$$V(s_i) = L_1 \times \cdots \times L_{i-1} \times \{s_i\} \times L_{i+1} \times \cdots \times L_n,$$

and thus, we multiply the probabilities of these subsets L_j , for all j , and the probability of s_i . Summing up the probabilities of $V(s_i)$ over $s_i \in V_i$ gives us γ_i , the probability of W_i .

Finally, the supremum of the expected revenue $\mathcal{R}(\mathbf{p})$ over all $\mathbf{p} \in C$ is the maximum of $\sum_{i \in [n]} \gamma_i \cdot p_i$ over all \mathbf{p} in the closure of C . Let C' denote the closure of C ; this is the polyhedron obtained by changing all the strict inequalities of C into weak inequalities. The supremum of $\sum_i \gamma_i \cdot p_i$ over all points $\mathbf{p} \in C'$ can be computed in polynomial time by solving the linear program that maximizes $\sum \gamma_i \cdot p_i$ subject to $\mathbf{p} \in C'$. In fact, as we will show below after the proof of Theorem 8, this LP has a special form: The question of whether a set of equations and inequalities with respect to a set of hyperplanes of the form $p_i = s_i$ and $p_i - p_j = s_i - t_j$ is consistent, i.e., defines a nonempty cell, can be formulated as a negative weight cycle problem, and the optimal solution for a nonempty cell can be computed by solving a single-source shortest path problem. It follows that the specification of a cell C in the partition is an appropriate *yes* certificate for the decision problem ITEM-PRICING, and the theorem is proved. \square

Next we describe in more detail how to determine whether a set of equations and inequalities defines a nonempty cell, and how to compute the optimal solution over a nonempty cell. The description of a (candidate) cell C consists of equations and inequalities specifying (1) for each item i , the relation of p_i to every value $s_i \in V_i$, and (2) for each pair of items i, j and each pair of values $s_i \in V_i$ and $t_j \in V_j$, the relation of $p_i - p_j$ to $s_i - t_j$. Construct a weighted directed graph $G = (N, E)$ over $n+1$ nodes $N = \{0, 1, \dots, n\}$ where nodes $1, \dots, n$ correspond to the n items. For each inequality of the form $p_i < s_i$ or $p_i \leq s_i$, include an edge $(0, i)$ with weight s_i , and call the edge strict or weak accordingly as the inequality is

strict or weak. In fact, there is a tightest such inequality (i.e., with the smallest value s_i) since the cell is in P , and it suffices to include the edge for this inequality only. Similarly, for each inequality of the form $p_i > s_i$ or $p_i \geq s_i$ (or only for the tightest such inequality, i.e. the one with the largest value s_i) include an edge $(i, 0)$ with weight $-s_i$. For each inequality of the form $p_i - p_j < s_i - t_j$ or $p_i - p_j \leq s_i - t_j$ (or only for the tightest such inequality) include a (strict or weak) edge (j, i) with weight $s_i - t_j$. Similarly, for every inequality of the form $p_i - p_j > s_i - t_j$ or $p_i - p_j \geq s_i - t_j$ (or only for the tightest such inequality) include a (strict or weak) edge (i, j) with weight $t_j - s_i$.

We prove the following connections between $G = (N, E)$ and the cell C :

Lemma 6.3.1. 1. *A set of equations and inequalities defines a nonempty cell if and only if the corresponding graph G does not contain a negative weight cycle or a zero weight cycle with a strict edge.*

2. *The supremum of the expected revenue for a nonempty cell is achieved by the price vector \mathbf{p} that consists of the distances from node 0 to the other nodes of the graph G .*

Proof. 1. Considering node 0 as having an associated variable p_0 with fixed value 0, the given set of equations (i.e., pairs of weak inequalities) and (strict) inequalities can be viewed as a set of difference constraints on the variables (p_0, p_1, \dots, p_n) , and it is well known that the feasibility of such a set of constraints can be formulated as a negative weight cycle problem. If there is a cycle with negative weight w , then adding all the inequalities corresponding to the edges of the cycle yields the constraint $0 \leq w$ (which is false); if there is a cycle with zero weight but also a strict edge, then summing the inequalities yields $0 < 0$.

Conversely, suppose that G does not contain a negative weight cycle or a zero weight cycle with a strict edge. For each strict edge e , replace its weight $w(e)$ by $w'(e) = w(e) - \epsilon$ for a sufficiently small $\epsilon > 0$ (we can treat ϵ symbolically), and let $G(\epsilon)$ be the resulting weighted graph. Note that $G(\epsilon)$ does not contain any negative weight cycle, hence all shortest paths are well-defined in $G(\epsilon)$. Compute the shortest (minimum weight) paths from node 0 to all the other nodes in $G(\epsilon)$, and let $\mathbf{p}(\epsilon)$ be the vector of distances from 0. For each edge (i, j) the distances $p_i(\epsilon)$ and $p_j(\epsilon)$ (where $p_0(\epsilon) = 0$) must satisfy $p_j(\epsilon) \leq p_i(\epsilon) + w'(i, j)$, hence all the (weak and strict) inequalities are satisfied.

To determine if a set of equations and inequalities defines a nonempty cell, we can form

the graph $G(\epsilon)$ and test for the existence of a negative weight cycle using for example the Bellman-Ford algorithm.

2. Suppose that cell C specified by the constraints is nonempty. Then we claim that the vector $\mathbf{p} = \mathbf{p}(0)$ of distances from node 0 to the other nodes in the graph G is greater than or equal to any vector $\mathbf{p}' \in C$ in all coordinates. We can show this by induction on the depth of a node in the shortest path tree T of G rooted at node 0. Letting $p'_0 = p_0 = 0$, the basis is trivial. For the induction step, consider a node j with parent i in T . By the inductive hypothesis $p'_i \leq p_i$. The edge (i, j) implies that $p'_j - p'_i \leq w(i, j)$ or $< w(i, j)$, and the presence of the edge (i, j) in the shortest path tree implies that $p_j = p_i + w(i, j)$. Therefore, $p'_j \leq p_j$.

The supremum of the expected revenue $\mathcal{R}(\mathbf{p}')$ over the cell C is given by the optimal value of the linear program that maximizes $\sum_{i \in [n]} \gamma_i \cdot p'_i$ subject to $\mathbf{p}' \in C'$, where C' is the closure of the cell C . Observe that all the coefficients γ_i of the objective function are nonnegative, and clearly \mathbf{p} is in the closure C' . Therefore \mathbf{p} achieves the supremum of the expected revenue over C . \square

The NP characterization of ITEM-PRICING and the corresponding structural characterization of the optimal price vector $\mathbf{p} = \mathbf{p}(0)$ of each cell have several easy and useful consequences.

First, we get an alternative proof of Lemma 6.2.3 for the maximum tie-breaking rule:

Second Proof of Lemma 6.2.3. Suppose that the supremum of the expected revenue is achieved in cell C . Let G be the corresponding graph, and let \mathbf{p} be the price vector of the distances from node 0 to the other nodes. If $\mathbf{p} \in C$ then the conclusion is immediate, so assume $\mathbf{p} \notin C$. From the proof of the above lemma we have that $\mathbf{p} \geq \mathbf{p}'$ coordinate-wise for all $\mathbf{p}' \in C$.

We claim that for any valuation $\mathbf{v} \in V$, the revenue $\mathcal{R}(\mathbf{v}, \mathbf{p})$ is at least as large as the revenue $\mathcal{R}(\mathbf{v}, \mathbf{p}')$ under any $\mathbf{p}' \in C$. Suppose that the buyer selects item i under \mathbf{v} for prices \mathbf{p}' . Then $p'_i \leq v_i$ and thus also $p_i \leq v_i$ (since \mathbf{p} is in the closure of C) and thus i is also eligible for selection under \mathbf{p} . If the buyer selects i under \mathbf{p} then we know that $p_i \geq p'_i$ and the conclusion follows. Suppose that the buyer selects another item j under \mathbf{p} and that

$p'_i > p_j$ and hence $p_i > p_j$. Then we must have $v_j - p_j > v_i - p_i$ due to the tie-breaking rule. The facts that \mathbf{p} is in the closure of C and $v_j - p_j > v_i - p_i$ imply that $v_j - p'_j > v_i - p'_i$ for all $\mathbf{p}' \in C$, and therefore the buyer should have picked j instead of i under prices \mathbf{p}' , a contradiction.

We conclude that for any $\mathbf{v} \in V$, $\mathcal{R}(\mathbf{v}, \mathbf{p}) \geq \mathcal{R}(\mathbf{v}, \mathbf{p}')$ for any $\mathbf{p}' \in C$, and the lemma follows. \square

Another consequence suggested by the structural characterization of Lemma 6.3.1 is that the maximum of expected revenue can always be achieved by a price vector \mathbf{p} in which all prices p_i are sums of a value and differences between pairs of values of items. This implies for example the following useful corollary.

Corollary 6.3.1. *If all the values in V_i , $i \in [n]$, are integers, then there must exist an optimal price vector $\mathbf{p} \in P$ with integer coordinates.*

6.4 A polynomial-time Algorithm for support size 2

In this section, we present a polynomial-time algorithm for the case that each distribution has support size at most 2. In Section 6.4.1, we give a polynomial-time algorithm under a certain “non-degeneracy” assumption on the values. In Section 6.4.2 we generalize this algorithm to handle the general case.

6.4.1 An Interesting Special case.

In this subsection, we assume that every item has support size 2, where $V_i = \{a_i, b_i\}$ satisfies $b_i > a_i > 0$, for all $i \in [n]$. Let $q_i : 0 < q_i < 1$ denote the probability of the value of item i being b_i . For convenience, we also let $t_i = b_i - a_i > 0$. In addition, we assume in this subsection that the value-vectors $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$ satisfy the following “non-degeneracy” assumption:

Non-degeneracy assumption: $b_1 < b_2 < \dots < b_n$, $a_i \neq a_j$ and $t_i \neq t_j$ for all $i, j \in [n]$.

As we show next in Section 6.4.2, this special case encapsulates the essential difficulty of the problem.

Let OPT denote the set of optimal price vectors in $P = \times_{i=1}^n [a_i, b_i]$ that maximize the expected revenue $\mathcal{R}(\mathbf{p})$. Next we prove a sequence of lemmas to show that, given \mathbf{a} and \mathbf{b} that satisfy all the conditions above one can compute efficiently a set $A \subseteq P$ of price vectors such that $|A| = O(n^2)$ and $\text{OPT} \subseteq A$. Hence, by computing $\mathcal{R}(\mathbf{p})$ for all $\mathbf{p} \in A$, we get both the maximum of expected revenue and an optimal price vector.

We start with the following lemma:

Lemma 6.4.1. *If $\mathbf{p} \in P$ satisfies $p_i > a_i$ for all $i \in [n]$, then either $\mathbf{p} = \mathbf{b}$ or we have $\mathbf{p} \notin \text{OPT}$.*

Proof. Assume for contradiction that $\mathbf{p} \in P$ satisfies $p_i > a_i$, for all $i \in [n]$ but $\mathbf{p} \neq \mathbf{b}$. It then follows from the maximum price tie-breaking rule that $\mathcal{R}(\mathbf{v}, \mathbf{b}) \geq \mathcal{R}(\mathbf{v}, \mathbf{p})$ for all $\mathbf{v} \in V$. Moreover, there is at least one $\mathbf{v}^* \in V$ such that $\mathcal{R}(\mathbf{v}^*, \mathbf{b}) > \mathcal{R}(\mathbf{v}^*, \mathbf{p})$: If $p_i < b_i$, then consider \mathbf{v}^* with $v_i^* = b_i$ and $v_j^* = a_j$ for all other j . It follows that $\mathcal{R}(\mathbf{b}) > \mathcal{R}(\mathbf{p})$ as we assumed that $0 < q_i < 1$ for all $i \in [n]$ and thus, $\mathbf{p} \notin \text{OPT}$. \square

Next we show that there can be at most one i such that $p_i = a_i$; otherwise $\mathbf{p} \notin \text{OPT}$. We emphasize that all the conditions on V_i are assumed in the lemmas below, the non-degeneracy assumption in particular.

Lemma 6.4.2. *If $\mathbf{p} \in P$ has more than one $i \in [n]$ such that $p_i = a_i$, then we have $\mathbf{p} \notin \text{OPT}$.*

Proof. Assume for contradiction that $\mathbf{p} \in P$ has more than one i such that $p_i = a_i$. We prove the lemma by explicitly constructing a new price vector $\mathbf{p}' \in P$ from \mathbf{p} such that $\mathcal{R}(\mathbf{v}, \mathbf{p}') \geq \mathcal{R}(\mathbf{v}, \mathbf{p})$ for all $\mathbf{v} \in V$ and $\mathcal{R}(\mathbf{v}^*, \mathbf{p}') > \mathcal{R}(\mathbf{v}^*, \mathbf{p})$ for at least one $\mathbf{v}^* \in V$. This implies that $\mathcal{R}(\mathbf{p}') > \mathcal{R}(\mathbf{p})$ and thus, \mathbf{p} is not optimal. We will be using this simple strategy in most of the proofs of this section.

Let $k \in [n]$ denote *the* item with the smallest a_k among all $i \in [n]$ with $p_i = a_i$. By the non-degeneracy assumption, k is unique. Recall that $t_k = b_k - a_k = b_k - p_k$. We let S denote the set of $i \in [n]$ such that $b_i - p_i = t_k$, so $k \in S$. By the non-degeneracy assumption again, we have $p_i > a_i$ for all $i \in S - \{k\}$. We now construct $\mathbf{p}' \in P$ as follows: For each $i \in [n]$, set $p'_i = p_i$ if $i \notin S$; otherwise set $p'_i = p_i + \epsilon$ for some sufficiently small $\epsilon > 0$. Next

we show that $\mathcal{R}(\mathbf{v}, \mathbf{p}') \geq \mathcal{R}(\mathbf{v}, \mathbf{p})$ for all $\mathbf{v} \in V$. Fix a $\mathbf{v} \in V$. We consider the following three cases:

1. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) = t_k$, then $\mathcal{T}(\mathbf{v}, \mathbf{p}) \subseteq S$ by the definition of S . When ϵ is sufficiently small, we have

$$\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p}) \quad \text{and} \quad \mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}) + \epsilon > \mathcal{R}(\mathbf{v}, \mathbf{p}).$$

2. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) = 0$ and $k \in \mathcal{T}(\mathbf{v}, \mathbf{p})$, then we have $\mathcal{T}(\mathbf{v}, \mathbf{p}) \cap S = \{k\}$ since $b_i > p_i > a_i$ for all other $i \in S$. We claim that $\mathcal{R}(\mathbf{v}, \mathbf{p}) > p_k$ in this case. To see this, note that there exists an item $\ell \in [n]$ such that $p_\ell = a_\ell$ and $p_\ell > p_k$ by our choice of k . As $\mathcal{U}(\mathbf{v}, \mathbf{p}) = 0$, we must have $v_\ell = a_\ell$ and thus, $\ell \in \mathcal{T}(\mathbf{v}, \mathbf{p})$ and $\mathcal{R}(\mathbf{v}, \mathbf{p}) \geq p_\ell$ is not obtained from selling item k . Therefore, we have

$$\mathcal{U}(\mathbf{v}, \mathbf{p}') = 0, \quad \mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p}) - \{k\} \quad \text{and} \quad \mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}).$$

3. Finally, if neither of the cases above happens, then we have $\mathcal{T}(\mathbf{v}, \mathbf{p}) \cap S = \emptyset$ (note that this includes the case when $\mathcal{T}(\mathbf{v}, \mathbf{p}) = \emptyset$). For this case we have $\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p})$ and $\mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p})$.

The lemma then follows because in the second case above, we indeed showed that the following valuation vector \mathbf{v}^* in V satisfies $\mathcal{R}(\mathbf{v}^*, \mathbf{p}') > \mathcal{R}(\mathbf{v}^*, \mathbf{p})$: $v_k = b_k$ and $v_i = a_i$ for all $i \neq k$. \square

Lemma 6.4.2 reduces our search space to \mathbf{p} such that either $\mathbf{p} = \mathbf{b}$ or $\mathbf{p} \in P_k$ for some $k \in [n]$, where we use P_k to denote the set of price vectors $\mathbf{p} \in P$ such that $p_k = a_k$ and $p_i > a_i$ for all other $i \in [n]$.

The next lemma further restricts our attention to $\mathbf{p} \in P_k$ such that $p_i \in \{b_i, b_i - t_k\}$ for all $i \neq k$.

Lemma 6.4.3. *If $\mathbf{p} \in P_k$ but $p_i \notin \{b_i, b_i - t_k\}$ for some $i \neq k$, then we have $\mathbf{p} \notin \text{OPT}$.*

Proof. Assume for contradiction that $p_\ell \notin \{b_\ell, b_\ell - t_k\}$. As $\mathbf{p} \in P_k$, we also have $p_\ell > a_\ell$. Now we use S to denote the set of all $i \in [n]$ such that $b_i - p_i = b_\ell - p_\ell$. It is clear that $k \notin S$. We use \mathbf{p}' to denote the following new price vector: $p'_i = p_i$ for all $i \notin S$, and

$p'_i = p_i + \epsilon$ for all $i \in S$, where $\epsilon > 0$ is sufficiently small. We use the same proof strategy to show that $\mathcal{R}(\mathbf{p}') > \mathcal{R}(\mathbf{p})$. Fix any $\mathbf{v} \in V$. We have

1. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) < 0$, then clearly $\mathcal{U}(\mathbf{v}, \mathbf{p}') < 0$ as well and thus, $\mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}) = 0$.
2. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) = b_\ell - p_\ell$, then $\mathcal{T}(\mathbf{v}, \mathbf{p}) \subseteq S$ by the definition of S . When ϵ is sufficiently small,

$$\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p}) \quad \text{and} \quad \mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}) + \epsilon > \mathcal{R}(\mathbf{v}, \mathbf{p}).$$

3. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) \geq 0$ but $\mathcal{U}(\mathbf{v}, \mathbf{p}) \neq b_\ell - p_\ell$, then it is easy to see that $\mathcal{T}(\mathbf{v}, \mathbf{p}) \cap S = \emptyset$, because $p_i > a_i$ and $b_i - p_i = b_\ell - p_\ell$ for all $i \in S$. It follows that $\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p})$ and $\mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p})$.

The lemma follows by combining all three cases. □

As suggested by Lemma 6.4.3, for each $k \in [n]$, we use P'_k to denote the set of $\mathbf{p} \in P_k$ such that $p_k = a_k$ and $p_i \in \{b_i, b_i - t_k\}$ for all other i . In particular, p_i must be b_i if $t_i < t_k$ ($t_i \neq t_k$, by the non-degeneracy assumption). The next lemma shows that we only need to consider $\mathbf{p} \in P'_k$ such that $p_i = b_i$ for all $i < k$.

Lemma 6.4.4. *If $\mathbf{p} \in P'_k$ satisfies $p_\ell = b_\ell - t_k > a_\ell$ for some $\ell < k$, then we have $\mathbf{p} \notin \text{OPT}$.*

Proof. We construct \mathbf{p}' from \mathbf{p} as follows. Let S denote the set of all $i < k$ such that $p_i = b_i - t_k > a_i$. By our assumption, S is nonempty. Then set $p'_i = p_i$ for all $i \notin S$ and $p'_i = p_i + \epsilon$ for all $i \in S$, where $\epsilon > 0$ is sufficiently small. Similarly we show that $\mathcal{R}(\mathbf{p}') > \mathcal{R}(\mathbf{p})$ by considering the following cases:

1. If $\mathcal{U}(\mathbf{v}, \mathbf{p}) = t_k$ and $\mathcal{T}(\mathbf{v}, \mathbf{p}) \cap S \neq \emptyset$, we consider the following cases. If $\mathcal{T}(\mathbf{v}, \mathbf{p}) \subseteq S$, then

$$\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p}) \quad \text{and} \quad \mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}) + \epsilon > \mathcal{R}(\mathbf{v}, \mathbf{p}).$$

Otherwise, there exists a $j \geq k$ such that $j \in \mathcal{T}(\mathbf{v}, \mathbf{p})$. This implies that $\mathcal{R}(\mathbf{v}, \mathbf{p}) \geq p_j = b_j - t_k$ is not obtained from any item in S . As a result, $\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p}) - S$ and $\mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p})$.

2. If the case above does not happen, then we must have $\mathcal{T}(\mathbf{v}, \mathbf{p}) \cap S = \emptyset$ (this includes the case when $\mathcal{T}(\mathbf{v}, \mathbf{p}) = \emptyset$). As a result, we have $\mathcal{T}(\mathbf{v}, \mathbf{p}') = \mathcal{T}(\mathbf{v}, \mathbf{p})$ and $\mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p})$.

The lemma follows by combining the two cases. \square

Finally, we use P_k^* for each $k \in [n]$ to denote the set of $\mathbf{p} \in P$ such that $p_k = a_k$; $p_i = b_i$ for all $i < k$; $p_i = b_i$, for all $i > k$ such that $t_i < t_k$; and $p_i \in \{b_i, b_i - t_k\}$, for all other $i > k$. However, P_k^* may still be exponentially large in general. Let T_k denote the set of $i > k$ such that $t_i > t_k$. Given $\mathbf{p} \in P_k^*$, our last lemma below implies that, if i is the smallest index in T_k such that $p_i = b_i - t_k$, then $p_j = b_j - t_k$ for all $j \in T_k$ larger than i ; otherwise \mathbf{p} is not optimal. In other words, \mathbf{p} has to be monotone in setting p_j , $j \in T_k$, to be $b_j - t_k$; otherwise \mathbf{p} is not optimal. As a result, there are only $O(n^2)$ many price vectors that we need to check, and the best one among them is optimal. We use $A \subseteq \cup_k P_k^*$ to denote this set of price vectors. In other words, A contains all the price-vectors of the form $p = (b_1, \dots, b_{k-1}, a_k, b_{k+1}, \dots, b_{i-1}, p_i, \dots, \dots, p_n)$ for some k and $i > k$, where $p_i = b_i$ if $t_i < t_k$ and $p_i = b_i - t_k$ otherwise.

Lemma 6.4.5. *Given $k \in [n]$ and $\mathbf{p} \in P_k^*$, if there exist two indices $c, d \in T_k$ such that $c < d$, $p_c = b_c - t_k$ but $p_d = b_d$, then we must have $\mathbf{p} \notin \text{OPT}$.*

Proof. We use t to denote t_k for convenience. Also we may assume, without loss of generality, that there is no index between c and d in T_k ; otherwise we can use it to replace either c or d , depending on its price.

We define two vectors from \mathbf{p} . First, let \mathbf{p}' denote the vector obtained from \mathbf{p} by replacing $p_d = b_d$ by $p'_d = b_d - t$. Let \mathbf{p}^* denote the vector obtained from \mathbf{p} by replacing $p_c = b_c - t$ by $p_c^* = b_c$. In other words, the c th and d th entries of $\mathbf{p}, \mathbf{p}', \mathbf{p}^*$ are $(b_c - t, b_d), (b_c - t, b_d - t), (b_c, b_d)$, respectively, while all other $n - 2$ entries are the same. Our plan is to show that if $\mathcal{R}(\mathbf{p}) \geq \mathcal{R}(\mathbf{p}')$, then $\mathcal{R}(\mathbf{p}^*) > \mathcal{R}(\mathbf{p})$. This implies that \mathbf{p} cannot be optimal and the lemma follows.

We need some notation. Let V' denote the projection of V onto all but the c th and d th coordinates:

$$V' = \times_{i \in [n] - \{c, d\}} V_i.$$

We use $[n] - \{c, d\}$ to index entries of vectors \mathbf{u} in V' . Let $U \subseteq V'$ denote the set of vectors $\mathbf{u} \in V'$ such that $u_i - p_i < t$ for all $i > d$. (This just means that for each $i \in T_k$, if $i > d$ and $p_i = b_i - t$, then $u_i = a_i$.) Given $\mathbf{u} \in V'$, $v_c \in \{a_c, b_c\}$ and $v_d \in \{a_d, b_d\}$, we use (\mathbf{u}, v_c, v_d) to denote a n -dimensional price vector in V . Now we compare the expected revenue $\mathcal{R}(\mathbf{p})$, $\mathcal{R}(\mathbf{p}')$ and $\mathcal{R}(\mathbf{p}^*)$.

First, we claim that, if $\mathbf{v} = (\mathbf{u}, v_c, v_d) \in V$ but $\mathbf{u} \notin U$, then we have $\mathcal{R}(\mathbf{v}, \mathbf{p}) = \mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}^*)$. This is simply because there exists an item $i > d$ such that $v_i - p_i = t$, so it always dominates both items c and d . As a result, the difference among \mathbf{p} , \mathbf{p}' and \mathbf{p}^* no longer matters. Second, it is easy to show that for any $\mathbf{v} = (\mathbf{u}, a_c, a_d) \in V$, then $\mathcal{R}(\mathbf{v}, \mathbf{p}) = \mathcal{R}(\mathbf{v}, \mathbf{p}') = \mathcal{R}(\mathbf{v}, \mathbf{p}^*)$ as the utility from c and d are negative.

Now we consider a vector $\mathbf{v} = (\mathbf{u}, v_c, v_d) \in V$ such that $\mathbf{u} \in U$ and (v_c, v_d) is either (a_c, b_d) , (b_c, a_d) , or (b_c, b_d) . For convenience, for each $\mathbf{u} \in U$ we use \mathbf{u}_1^+ to denote (\mathbf{u}, a_c, b_d) ; \mathbf{u}_2^+ to denote (\mathbf{u}, b_c, a_d) ; and \mathbf{u}_3^+ to denote (\mathbf{u}, b_c, b_d) . By the definition of U , we have the following simple cases:

1. For \mathbf{p} , we have $\mathcal{R}(\mathbf{u}_2^+, \mathbf{p}) = b_c - t$ and $\mathcal{R}(\mathbf{u}_3^+, \mathbf{p}) = b_c - t$;
2. For \mathbf{p}' , we have $\mathcal{R}(\mathbf{u}_1^+, \mathbf{p}') = b_d - t$, $\mathcal{R}(\mathbf{u}_2^+, \mathbf{p}') = b_c - t$ and $\mathcal{R}(\mathbf{u}_3^+, \mathbf{p}') = b_d - t$.

We need the following equation:

$$\mathcal{R}(\mathbf{u}_1^+, \mathbf{p}) = \mathcal{R}(\mathbf{u}_1^+, \mathbf{p}^*) = \mathcal{R}(\mathbf{u}_3^+, \mathbf{p}^*) \quad (6.4)$$

as well as the following two inequalities:

$$\mathcal{R}(\mathbf{u}_1^+, \mathbf{p}^*) - (b_d - b_c) \leq \mathcal{R}(\mathbf{u}_2^+, \mathbf{p}^*) \leq \mathcal{R}(\mathbf{u}_1^+, \mathbf{p}^*) \quad (6.5)$$

Given a $\mathbf{v} \in V$, recall that $\Pr[\mathbf{v}]$ denotes the probability of the valuation vector being \mathbf{v} . Given a $\mathbf{u} \in U$, we also use $\Pr[\mathbf{u}]$ to denote the probability of the $n - 2$ items, except items c and d , taking values \mathbf{u} . Let

$$h_1 = (1 - q_c)q_d, \quad h_2 = q_c(1 - q_d) \quad \text{and} \quad h_3 = q_cq_d.$$

Clearly we have $h_1, h_2, h_3 > 0$ and $\Pr[\mathbf{u}_i^+] = \Pr[\mathbf{u}] \cdot h_i$, for all $\mathbf{u} \in U$ and $i \in [3]$.

In order to compare $\mathcal{R}(\mathbf{p})$, $\mathcal{R}(\mathbf{p}')$ and $\mathcal{R}(\mathbf{p}^*)$, we only need to compare the following three sums:

$$\sum_{i \in [3]} \sum_{\mathbf{u} \in U} \Pr[\mathbf{u}_i^+] \cdot \mathcal{R}(\mathbf{u}_i^+, \mathbf{p}), \quad \sum_{i \in [3]} \sum_{\mathbf{u} \in U} \Pr[\mathbf{u}_i^+] \cdot \mathcal{R}(\mathbf{u}_i^+, \mathbf{p}') \quad \text{and} \quad \sum_{i \in [3]} \sum_{\mathbf{u} \in U} \Pr[\mathbf{u}_i^+] \cdot \mathcal{R}(\mathbf{u}_i^+, \mathbf{p}^*).$$

For the first sum, we can rewrite it as (here all sums are over $\mathbf{u} \in U$):

$$h_1 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \mathcal{R}(\mathbf{u}_1^+, \mathbf{p}) + h_2 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_c - t) + h_3 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_c - t), \quad (6.6)$$

while the sum for $\mathcal{R}(\mathbf{p}')$ is the following:

$$h_1 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_d - t) + h_2 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_c - t) + h_3 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_d - t). \quad (6.7)$$

Since $c < d$ and $b_c < b_d$, $\mathcal{R}(\mathbf{p}) \geq \mathcal{R}(\mathbf{p}')$ would imply that

$$\sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \mathcal{R}(\mathbf{u}_1^+, \mathbf{p}) > \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_d - t). \quad (6.8)$$

On the other hand, we can also rewrite the sum for $\mathcal{R}(\mathbf{p}^*)$ as

$$h_1 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \mathcal{R}(\mathbf{u}_1^+, \mathbf{p}^*) + h_2 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \mathcal{R}(\mathbf{u}_2^+, \mathbf{p}^*) + h_3 \cdot \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \mathcal{R}(\mathbf{u}_3^+, \mathbf{p}^*). \quad (6.9)$$

The first sum in (6.9) is the same as that of (6.6). For the second sum, from (6.5), (6.4) and (6.8) we have

$$\begin{aligned} \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \mathcal{R}(\mathbf{u}_2^+, \mathbf{p}^*) &\geq \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \left(\mathcal{R}(\mathbf{u}_1^+, \mathbf{p}) - (b_d - b_c) \right) \\ &> \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot \left(b_d - t - (b_d - b_c) \right) = \sum_{\mathbf{u}} \Pr[\mathbf{u}] \cdot (b_c - t). \end{aligned}$$

The third sum in (6.9) is also strictly larger than that of (6.6) as $\mathcal{R}(\mathbf{u}_3^+, \mathbf{p}^*) = \mathcal{R}(\mathbf{u}_1^+, \mathbf{p}^*) \geq \mathcal{R}(\mathbf{u}_2^+, \mathbf{p}^*)$ while the second and third sums in (6.6) are the same, ignoring h_2 and h_3 . Thus, $\mathcal{R}(\mathbf{p}^*) > \mathcal{R}(\mathbf{p})$. \square

6.4.2 General Case

Now we deal with the general case. Let I denote an input instance with n items, in which $|V_i| \leq 2$ for all i . For each $i \in [n]$, either $V_i = \{a_i, b_i\}$ where $b_i > a_i \geq 0$, or $V_i = \{b_i\}$, where

$b_i \geq 0$. We let $D \subseteq [n]$ denote the set of $i \in [n]$ such that $|V_i| = 2$. For each item $i \in D$, we use $q_i : 0 < q_i < 1$ to denote the probability of its value being b_i . Each item $i \notin D$ has value b_i with probability 1. As permuting the items does not affect the maximum expected revenue, we may assume without loss of generality that $b_1 \leq b_2 \leq \dots \leq b_n$.

The idea is to perturb I (symbolically), so that the new instances satisfy all conditions described at the beginning of the section, which we know how to solve efficiently. For this purpose, we define a new n -item instance I_ϵ from I for any $\epsilon > 0$: For each $i \in D$, the support of item i is $V_{i,\epsilon} = \{a_i + i\epsilon, b_i + 2i\epsilon\}$, and for each $i \notin D$, the support of item i is $V_{i,\epsilon} = \{b_i + i\epsilon, b_i + 2i\epsilon\}$. For each $i \in D$, the probability of the value being $b_i + 2i\epsilon$ is still set to be q_i , while for each $i \notin D$, the probability of the value being $b_i + 2i\epsilon$ is set to be $1/2$. In the rest of the section, we use $\mathcal{R}(\mathbf{p})$ and $\mathcal{R}(\mathbf{v}, \mathbf{p})$ to denote the revenue with respect to I , and use $\mathcal{R}_\epsilon(\mathbf{p})$ and $\mathcal{R}_\epsilon(\mathbf{v}, \mathbf{p})$ to denote the revenue with respect to I_ϵ . Let $V_\epsilon = \times_{i=1}^n V_{i,\epsilon}$. Let ρ denote the following map from V_ϵ to V : ρ maps $\mathbf{u} \in V_\epsilon$ to $\mathbf{v} \in V$, where 1) $v_i = b_i$ when $i \notin D$; 2) $v_i = a_i$ if $u_i = a_i + i\epsilon$ and $v_i = b_i$ if $u_i = b_i + 2i\epsilon$ when $i \in D$.

It is easy to verify that, when $\epsilon > 0$ is sufficiently small, the new instance I_ϵ satisfies all conditions given at the beginning of the section, including the non-degeneracy assumption. Moreover, we show that

Lemma 6.4.6. *The limit of $\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p})$ exists as $\epsilon \rightarrow 0$, and can be computed in polynomial time.*

Proof. Since I_ϵ satisfies all the conditions, we know there is a set of $O(n^2)$ price vectors, denote by A_ϵ for I_ϵ , such that the best vector in A_ϵ is optimal for I_ϵ and achieves $\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p})$.

Furthermore, from the construction of A_ϵ , we know that every vector \mathbf{p}_ϵ in A_ϵ has an explicit expression in ϵ : each entry of \mathbf{p}_ϵ is indeed an affine linear function of ϵ . As a result, the limit of $\mathcal{R}_\epsilon(\mathbf{p}_\epsilon)$ as ϵ approaches 0 exists and can be computed efficiently. Since $\lim_{\epsilon \rightarrow 0} (\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p}))$ is just the maximum of these $O(n^2)$ limits, it also exists and can be computed in polynomial time in the input size of I . \square

Finally, the next two lemmas show that this limit is exactly the maximum expected revenue of I .

Lemma 6.4.7. $\max_{\mathbf{p}} \mathcal{R}(\mathbf{p}) \leq \lim_{\epsilon \rightarrow 0} (\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p}))$.

Proof. Let \mathbf{p}^* denote an optimal price vector of I . It suffices to show that, when ϵ is sufficiently small,

$$\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p}) \geq \mathcal{R}(\mathbf{p}^*) - 4n^2\epsilon. \quad (6.10)$$

The proof is similar to that of Lemma 6.2.1. Let \mathbf{p}' denote the vector in which $p'_i = \max(0, p_i^* - 4r_i n\epsilon)$, where r_i is the rank of p_i^* among $\{p_1^*, \dots, p_n^*\}$ sorted in the increasing order (when there are ties, items with lower index are ranked higher). We claim that, when $\epsilon > 0$ is sufficiently small,

$$\mathcal{R}_\epsilon(\mathbf{u}, \mathbf{p}') \geq \mathcal{R}(\rho(\mathbf{u}), \mathbf{p}^*) - 4n^2\epsilon, \quad \text{for any } \mathbf{u} \in V_\epsilon, \quad (6.11)$$

from which we get $\mathcal{R}_\epsilon(\mathbf{p}') \geq \mathcal{R}(\mathbf{p}^*) - 4n^2\epsilon$ and (6.10) follows.

To prove (6.11) we fix a $\mathbf{u} \in V_\epsilon$ and let $\mathbf{v} = \rho(\mathbf{u}) \in V$. (6.11) holds trivially if $\mathcal{R}(\mathbf{v}, \mathbf{p}^*) = 0$. Assume that $\mathcal{R}(\mathbf{v}, \mathbf{p}^*) > 0$, and let k denote the item selected in I on $(\mathbf{v}, \mathbf{p}^*)$. (6.11) also holds trivially if $p_k^* < 4n^2\epsilon$, so without loss of generality, we assume that $p_k \geq 4n^2\epsilon$. For any other item $j \in [n]$, we compare the utilities of items k and j in I_ϵ on $(\mathbf{u}, \mathbf{p}')$. We claim that

$$u_k - p'_k > u_j - p'_j \quad (6.12)$$

because 1) if $v_k - p_k^* > v_j - p_j^*$, then (6.12) holds when ϵ is sufficiently small; 2) if $v_k - p_k^* = v_j - p_j^*$ and $p_k^* > p_j^*$, then (6.12) holds because $p_k^* - p'_k - (p_j^* - p'_j) \geq 4n\epsilon > (v_k - u_k) + (u_j - v_j)$; 3) finally, the case when $v_k - p_k^* = v_j - p_j^*$, $p_k = p_j$ and $k < j$ follows similarly from $r_k > r_j$. Therefore, k remains to be the item being selected in I_ϵ on $(\mathbf{u}, \mathbf{p}')$. (6.11) then follows from the fact that $p'_k \geq p_k^* - 4n^2\epsilon$ by definition. \square

Lemma 6.4.8. $\max_{\mathbf{p}} \mathcal{R}(\mathbf{p}) \geq \lim_{\epsilon \rightarrow 0} (\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p}))$.

Proof. From the proof of Lemma 6.4.6, there is a price vector $\mathbf{p}_\epsilon \in A_\epsilon$ in which every entry is an affine linear function of ϵ , such that (as the cardinality of $|A_\epsilon|$ is bounded from above by $O(n^2)$)

$$\lim_{\epsilon \rightarrow 0} \left(\max_{\mathbf{p}} \mathcal{R}_\epsilon(\mathbf{p}) \right) = \lim_{\epsilon \rightarrow 0} \mathcal{R}_\epsilon(\mathbf{p}_\epsilon).$$

Let $\tilde{\mathbf{p}} \in \mathbb{R}_+^n$ denote the limit of \mathbf{p}_ϵ , by simply removing all the ϵ 's in the affine linear functions. Moreover, we note that $|\tilde{p}_i - p_{\epsilon,i}| = O(n\epsilon)$ by the construction of A_ϵ , where we use $p_{\epsilon,i}$ to denote the i th entry of \mathbf{p}_ϵ .

Next, let \mathbf{q}_ϵ denote the vector in which the i th entry $q_{\epsilon,i} = \max(0, \tilde{p}_i - r_i n^2 \epsilon)$ for all $i \in [n]$, where r_i is the rank of \tilde{p}_i among entries of $\tilde{\mathbf{p}}$ sorted in increasing order (again, when there are ties, items with lower index are ranked higher). To prove the lemma, it suffices to show that, when ϵ is sufficiently small,

$$\mathcal{R}(\mathbf{q}_\epsilon) \geq \mathcal{R}_\epsilon(\mathbf{p}_\epsilon) - O(n^3 \epsilon).$$

To this end, we show that for any vector $\mathbf{u} \in V_\epsilon$ with $\mathbf{v} = \rho(\mathbf{u})$,

$$\mathcal{R}(\mathbf{v}, \mathbf{q}_\epsilon) \geq \mathcal{R}_\epsilon(\mathbf{u}, \mathbf{p}_\epsilon) - O(n^3 \epsilon). \quad (6.13)$$

Finally we prove (6.13). First, we note that if $\mathcal{U}(\mathbf{v}, \tilde{\mathbf{p}}) < 0$, then $\mathcal{R}(\mathbf{v}, \mathbf{q}_\epsilon) = \mathcal{R}_\epsilon(\mathbf{u}, \mathbf{p}_\epsilon) = 0$ when $\epsilon > 0$ is sufficiently small (as \mathbf{u} approaches \mathbf{v} and $\mathbf{p}_\epsilon, \mathbf{q}_\epsilon$ approach $\tilde{\mathbf{p}}$). Otherwise, we have $\mathcal{U}(\mathbf{v}, \mathbf{q}_\epsilon) > \mathcal{U}(\mathbf{v}, \tilde{\mathbf{p}}) \geq 0$ and we use k to denote the item selected in I on $(\mathbf{v}, \mathbf{q}_\epsilon)$. To violate (6.13), the item selected in I_ϵ on $(\mathbf{u}, \mathbf{p}_\epsilon)$ must be an item ℓ different from k satisfying $\tilde{p}_\ell > \tilde{p}_k$. Below we show that this cannot happen. Consider all the cases: 1) if $v_k - \tilde{p}_k < v_\ell - \tilde{p}_\ell$, we get a contradiction since item k is dominated by ℓ in I on $(\mathbf{v}, \mathbf{q}_\epsilon)$ when ϵ is sufficiently small; 2) if $v_k - \tilde{p}_k > v_\ell - \tilde{p}_\ell$, we get a contradiction with ℓ being selected in I_ϵ on $(\mathbf{u}, \mathbf{p}_\epsilon)$ when ϵ is sufficiently small; 3) if $v_k - \tilde{p}_k = v_\ell - \tilde{p}_\ell$ and $\tilde{p}_\ell > \tilde{p}_k$, we conclude that $v_k - q_{\epsilon,k} < v_\ell - q_{\epsilon,\ell}$, contradicting again with k being selected in I on $(\mathbf{v}, \mathbf{q}_\epsilon)$. (6.13) follows by combining all these cases. \square

6.5 NP-hardness for support size 3

In this section, we give a polynomial-time reduction from PARTITION to ITEM-PRICING for distributions with support (at most) 3. Recall that in the PARTITION problem [Garey and Johnson, 1979] we are given a set $C = \{c_1, \dots, c_n\}$ of n positive integers and wish to determine whether it is possible to partition C into two subsets with equal sum. We may assume without loss of generality that $c_1 = \max(c_1, \dots, c_n)$.

Given an instance of PARTITION, we construct an instance of ITEM-PRICING as follows. We have n items. Each item $i \in [n]$ can take 3 possible integer values $0, a, b$, where $b > a > 0$, i.e., $V_i = \{0, a, b\}$ for all $i \in [n]$. Let $q_i = \Pr[v_i = b]$ and $r_i = \Pr[v_i = a]$. We set

$q_i = c_i/M$ where $M = 2^n c_1^3$ and

$$r_i = \frac{b-a}{a(1-t_i)} \cdot q_i, \quad \text{where } t_i = \frac{b}{2a} \cdot \sum_{j \neq i, j \in [n]} q_j.$$

The two parameters a and b should be thought of as universal constants (independent of the given instance of PARTITION) throughout the proof. We will eventually set these constants to be $a = 1$, $b = 3$ (this choice is not necessary, there is flexibility in our proof and indeed any values with $b > 2a$ will work). However, for the sake of the presentation, we will keep a, b as generic parameters for most of the calculations till the end.

Note that the definition of r_i implies that

$$bq_i = a(q_i + r_i) - ar_i t_i. \quad (6.14)$$

Let $N = 2^n c_1^2$. Then we have $q_i, r_i = O(1/N)$ and $t_i = O(n/N)$ for all i . Thus, each distribution assigns most of its probability mass to the point 0. This is a crucial property which allows us to get a handle on the optimal revenue. For an arbitrary general instance of the pricing problem, the expected revenue is a highly complex nonlinear function. The fact that most of the probability mass in our construction is concentrated at 0 implies that valuation vectors with many nonzero entries contribute very little to the expected revenue. As we will argue, the revenue is approximated well by its 1st and 2nd order terms with respect to $\text{poly}(n)/N$, which essentially corresponds to the contribution of all valuations in which at most two items have nonzero value. The probabilities q_i, r_i are chosen carefully so that the optimization of the expected revenue amounts to a quadratic optimization problem, which achieves its maximum possible value when the given set C of integers has a partition into two parts with equal sums.

Our main claim is that, for an appropriate value t^* , there exists a price vector with expected revenue at least t^* if and only if there exists a solution to the original instance of the Partition problem.

Before we proceed with the proof, we will need some notation. For $T_1, T_2, \epsilon \in \mathbb{R}_+$ we write $T_1 = T_2 \pm \epsilon$ to denote that $|T_1 - T_2| \leq \epsilon$.

Note that, as both the q_i 's and the t_i 's are very small positive quantities, we have that

$r_i \approx (b - a)q_i/a$. Formally, with the above notation we can write

$$r_i = \frac{b - a}{a(1 - t_i)} \cdot q_i = \frac{b - a}{a} \cdot q_i \pm 2 \frac{b - a}{a} \cdot q_i t_i = \frac{b - a}{a} \cdot q_i \pm O(n/N^2). \quad (6.15)$$

Lemma 6.2.2 and Corollary 6.3.1 imply that a revenue maximizing price vector can be assumed to have non-negative integer coefficients of magnitude at most b . The following lemma establishes the stronger statement that, for our particular instance, an optimal price vector \mathbf{p} can be assumed to have each p_i in the set $\{a, b\}$.

Lemma 6.5.1. *There is an optimal price vector $\mathbf{p} \in \{a, b\}^n$.*

Proof. By Lemma 6.2.2 and Corollary 6.3.1, there is an optimal price vector with integer coordinates in $[0 : b]$. Let \mathbf{p} be any (integer) vector in $[0 : b]^n$ that has at least one coordinate $p_j \notin \{a, b\}$. We will show below that $\mathcal{R}(\mathbf{p}) < \mathcal{R}(\mathbf{b})$, where \mathbf{b} denotes the all- b vector, and hence \mathbf{p} is not optimal.

Consider an index $i \in [n]$ with $p_i > 0$. The probability the buyer selects item i is bounded from above by $\Pr[v_i \geq p_i]$, the probability that item i has value at least p_i , and is bounded from below by

$$\Pr[v_i \geq p_i] \cdot \prod_{j \neq i, j \in [n]} (1 - q_j - r_j) \geq \Pr[v_i \geq p_i] \cdot (1 - O(n/N)).$$

Note that the second term in the LHS above is the probability that all items other than i have value 0 and the inequality uses the fact that $q_j, r_j = O(1/N)$. Applying these two bounds on \mathbf{p} and \mathbf{b} we obtain

$$\mathcal{R}(\mathbf{b}) \geq \sum_{i \in [n]} q_i (1 - O(n/N)) \cdot b \quad \text{and} \quad \mathcal{R}(\mathbf{p}) \leq \sum_{i: p_i > 0} \Pr[v_i \geq p_i] \cdot p_i.$$

So $\mathcal{R}(\mathbf{b}) \geq (\sum_{i \in [n]} q_i b) - O(n^2/N^2)$. Regarding $\mathcal{R}(\mathbf{p})$, we consider the following three cases. For $i \in [n]$ with $p_i = b$, the probability that $v_i \geq p_i$ is q_i and the contribution of such an item to the second sum is $q_i b$. Similarly, for $i \in [n]$ with $p_i = a$, the probability that $v_i \geq p_i$ is $q_i + r_i$ and the contribution to the sum is

$$(q_i + r_i)a \leq q_i b + O(n/N^2),$$

where the inequality follows from (6.15). Finally, we consider an item $i \in [n]$ with $p_i \notin \{a, b\}$. If $a < p_i < b$ then the contribution is $q_i p_i$, which is at most $q_i(b - 1) = q_i b - q_i$, since p_i is

integer. If $p_i < a$, then the contribution is $(q_i + r_i)p_i$, which is at most $(q_i + r_i)(a - 1) = q_i b + ar_i t_i - q_i - r_i = q_i b - q_i - r_i(1 - at_i)$. In both cases, the contribution to the sum is at most

$$q_i b - q_i \leq q_i b - (1/M).$$

Note that the definition of M and N implies that $1/M \gg n^2/N^2$. Because there exists at least one j with $p_j \notin \{a, b\}$, it follows that $\mathcal{R}(\mathbf{p}) < \mathcal{R}(\mathbf{b})$ which completes the proof of the lemma. \square

As a result, to maximize the expected revenue it suffices to consider price vectors in $\{a, b\}^n$. Given any price-vector $\mathbf{p} \in \{a, b\}^n$, we let $S = S(\mathbf{p}) = \{i \in [n] : p_i = a\}$ and $T = T(\mathbf{p}) = \{i \in [n] : p_i = b\}$. The main idea of the proof is to establish an appropriate *quadratic form approximation* to the expected revenue $\mathcal{R}(\mathbf{p})$ that is sufficiently accurate for the purposes of our reduction.

Approximating the Revenue. We appropriately partition the valuation space V into three events that yield positive revenue. We then approximate the probability of each and its contribution to the expected revenue up to, and including, 2nd order terms, i.e., terms of order $O(\text{poly}(n)/N^2)$, and we ignore 3rd order terms, i.e., terms of order $O(\epsilon)$ where $\epsilon = n^3/N^3$.

In particular, we consider the following disjoint events:

- *First Event:* $E_1 = \{\mathbf{v} \in V \mid \exists i \in S : v_i = b\}$.

Note that for any $\mathbf{v} \in E_1$ we have $\mathcal{R}(\mathbf{v}, \mathbf{p}) = a$. The probability of this event is

$$\Pr[E_1] = 1 - \prod_{i \in S} (1 - q_i) = \sum_{i \in S} q_i - \sum_{i \neq j \in S} q_i q_j \pm O(\epsilon).$$

- *Second Event:* $E_2 = \overline{E_1} \cap \{\mathbf{v} \in V \mid \exists i \in S : v_i = a \text{ and } \forall i \in T : v_i \in \{0, a\}\}$.

Note that for any $\mathbf{v} \in E_2$ we have $\mathcal{R}(\mathbf{v}, \mathbf{p}) = a$. The probability of this event is

$$\Pr[E_2] = \prod_{j \in T} (1 - q_j) \left[\prod_{i \in S} (1 - q_i) - \prod_{i \in S} (1 - q_i - r_i) \right]$$

Using the elementary identities

$$\begin{aligned}\prod_{j \in T} (1 - q_j) &= 1 - \sum_{j \in T} q_j + \sum_{i \neq j \in T} q_i q_j \pm O(\epsilon) \\ \prod_{i \in S} (1 - q_i) &= 1 - \sum_{i \in S} q_i + \sum_{i \neq j \in S} q_i q_j \pm O(\epsilon) \\ \prod_{i \in S} (1 - q_i - r_i) &= 1 - \sum_{i \in S} (q_i + r_i) + \sum_{i \neq j \in S} (q_i + r_i)(q_j + r_j) \pm O(\epsilon),\end{aligned}$$

we can write

$$\begin{aligned}\Pr[E_2] &= \left[1 - \sum_{j \in T} q_j + \sum_{i \neq j \in T} q_i q_j \pm O(\epsilon) \right] \cdot \left[\sum_{i \in S} r_i + \sum_{i \neq j \in S} q_i q_j - \sum_{i \neq j \in S} (q_i + r_i)(q_j + r_j) \pm O(\epsilon) \right] \\ &= \sum_{i \in S} r_i - \sum_{i \in S} r_i \sum_{j \in T} q_j + \sum_{i \neq j \in S} q_i q_j - \sum_{i \neq j \in S} (q_i + r_i)(q_j + r_j) \pm O(\epsilon).\end{aligned}$$

- *Third Event:* $E_3 = \overline{E_1} \cap \{\mathbf{v} \in V \mid \exists i \in T : v_i = b\}$.

Note that for any $\mathbf{v} \in E_3$ we have $\mathcal{R}(\mathbf{v}, \mathbf{p}) = b$. The probability of this event is

$$\begin{aligned}\Pr[E_3] &= \prod_{i \in S} (1 - q_i) \left[1 - \prod_{j \in T} (1 - q_j) \right] \\ &= \left(1 - \sum_{i \in S} q_i + \sum_{i \neq j \in S} q_i q_j \pm O(\epsilon) \right) \left(\sum_{j \in T} q_j - \sum_{i \neq j \in T} q_i q_j \pm O(\epsilon) \right) \\ &= \sum_{j \in T} q_j - \sum_{i \neq j \in T} q_i q_j - \sum_{i \in S} q_i \sum_{j \in T} q_j \pm O(\epsilon).\end{aligned}$$

Therefore, for the expected revenue $\mathcal{R}(\mathbf{p})$ we have:

$$\begin{aligned}\mathcal{R}(\mathbf{p}) &= (\Pr[E_1] + \Pr[E_2]) \cdot a + \Pr[E_3] \cdot b \\ &= a \cdot \left(\sum_{i \in S} (q_i + r_i) - \sum_{i \neq j \in S} (q_i + r_i)(q_j + r_j) - \sum_{i \in S} r_i \sum_{j \in T} q_j \right) \\ &\quad + b \cdot \left(\sum_{j \in T} q_j - \sum_{i \neq j \in T} q_i q_j - \sum_{i \in S} q_i \sum_{j \in T} q_j \right) \pm O(\epsilon).\end{aligned}$$

Using (6.14) it follows that the first order term of the revenue is

$$b \sum_{j \in T} q_j + a \sum_{i \in S} (q_i + r_i) = b \sum_{j \in [n]} q_j + \sum_{i \in S} (a(q_i + r_i) - bq_i) = b \sum_{j \in [n]} q_j + \sum_{i \in S} (ar_i t_i).$$

Observe that the first term $b \sum_{j \in [n]} q_j$ in the above expression is a constant L_1 , independent of the pricing (i.e., the partition of the items into S and T).

In the second order term, we can rewrite the expression $a \sum_{i \neq j \in S} (q_i + r_i)(q_j + r_j)$ as

$$\begin{aligned}
& \frac{1}{2} \cdot \sum_{i \in S} (q_i + r_i) \sum_{j \in S, j \neq i} a(q_j + r_j) \\
&= \frac{1}{2} \cdot \sum_{i \in S} (q_i + r_i) \sum_{j \in S, j \neq i} (bq_j + ar_j t_j) \\
&= \frac{b}{2} \cdot \sum_{i \in S} q_i \sum_{j \in S, j \neq i} q_j + \frac{b}{2} \cdot \sum_{i \in S} r_i \sum_{j \in S, j \neq i} q_j + \frac{1}{2} \cdot \sum_{i \in S} (q_i + r_i) \sum_{j \in S, j \neq i} ar_j t_j \\
&= b \sum_{i \neq j \in S} q_i q_j + \frac{b}{2} \sum_{i \in S} r_i \sum_{j \in S, j \neq i} q_j \pm O(\epsilon)
\end{aligned}$$

where in the first expression above, the double summation is multiplied by $1/2$ because each unordered pair $i \neq j \in S$ is included twice. Thus, the second order term of the expected revenue $\mathcal{R}(\mathbf{p})$ is

$$\begin{aligned}
& -a \sum_{i \neq j \in S} (q_i + r_i)(q_j + r_j) - a \sum_{i \in S} r_i \sum_{j \in T} q_j - b \sum_{i \neq j \in T} q_i q_j - b \sum_{i \in S} q_i \sum_{j \in T} q_j \\
&= -b \sum_{i \neq j \in S} q_i q_j - \frac{b}{2} \sum_{i \in S} r_i \sum_{j \in S, j \neq i} q_j - a \sum_{i \in S} r_i \sum_{j \in T} q_j - b \sum_{i \neq j \in T} q_i q_j - b \sum_{i \in S} q_i \sum_{j \in T} q_j \pm O(\epsilon) \\
&= -b \sum_{i \neq j \in [n]} q_i q_j - \frac{b}{2} \sum_{i \in S} r_i \sum_{j \in S, j \neq i} q_j - a \sum_{i \in S} r_i \sum_{j \in T} q_j \pm O(\epsilon)
\end{aligned}$$

The first term in the last expression is a constant L_2 independent of the pricing. As a result, we can rewrite the second order term as follows:

$$L_2 - \frac{b}{2} \sum_{i \in S} r_i \sum_{j \in S, j \neq i} q_j - a \sum_{i \in S} r_i \sum_{j \in T} q_j \pm O(\epsilon) = L_2 - \sum_{i \in S} r_i \left(\frac{b}{2} \sum_{j \in S, j \neq i} q_j + a \sum_{j \in T} q_j \right) \pm O(\epsilon).$$

Summing with the first order term and letting $L = L_1 + L_2$, we have:

$$\begin{aligned}
 \mathcal{R}(\mathbf{p}) &= L + \sum_{i \in S} r_i \left(at_i - \frac{b}{2} \sum_{j \in S, j \neq i} q_j - a \sum_{j \in T} q_j \right) \pm O(\epsilon) \\
 &= L + \sum_{i \in S} r_i \left(\frac{b}{2} \sum_{j \neq i} q_j - \frac{b}{2} \sum_{j \in S, j \neq i} q_j - a \sum_{j \in T} q_j \right) \pm O(\epsilon) \\
 &= L + \sum_{i \in S} r_i \cdot \left(\frac{b}{2} - a \right) \sum_{j \in T} q_j \pm O(\epsilon) \\
 &= L + \frac{b-a}{a} \cdot \left(\frac{b}{2} - a \right) \cdot \frac{1}{M^2} \cdot \sum_{i \in S} c_i \cdot \sum_{j \in T} c_j \pm O(\epsilon).
 \end{aligned}$$

Now setting $a = 1, b = 3$ in the previous expression, we have that for any $\mathbf{p} \in \{a, b\}^n$,

$$\mathcal{R}(\mathbf{p}) = L + \frac{1}{M^2} \left(\sum_{i \in S} c_i \right) \cdot \left(\sum_{j \in T} c_j \right) \pm O(\epsilon). \quad (6.16)$$

At this point, we observe that the sum of the two factors $\sum_{i \in S} c_i, \sum_{j \in T} c_j$ in (6.16) is a constant (independent of the partition). Thus, their product is maximized when they are equal. Because $\epsilon = o(1/M^2)$, it follows that the revenue is maximized when the product of the two factors is maximized. In particular, if there exists a partition of the set $C = \{c_1, \dots, c_n\}$ into two sets with equal sums $H = (\sum_{i \in [n]} c_i)/2$, then the corresponding partition of the indices into the sets S and T yields revenue $L + \frac{1}{M^2} \cdot H^2 \pm O(\epsilon)$. On the other hand, if there is no such equipartition of the set C , then for any partition of the indices, the revenue will be at most $L + \frac{1}{M^2} (H+1)(H-1) \pm O(\epsilon) = L + \frac{1}{M^2} (H^2 - 1) \pm O(\epsilon)$. Since $\epsilon = o(1/M^2)$ it follows that there exists a partition of the set $C = \{c_1, \dots, c_n\}$ into two sets with equal sums if and only if there exists a price vector $\mathbf{p} \in \{a, b\}^n$ with $\mathcal{R}(\mathbf{p}) \geq t^* = L + \frac{1}{M^2} (H^2 - \frac{1}{2})$. This completes the proof.

Remark. In the above construction, the support of the distributions includes the value 0. It is easy to modify the construction, so that the support contains only positive values: shift all the values of the distributions up by 1 (thus, the supports now become $V_i = \{1, 2, 4\}$) and add an additional $(n+1)$ -th item which has value 1 with probability 1. This transformation increases the expected revenue by 1. It is easy to see that an optimal price vector \mathbf{p}' for the new instance will give price $p'_{n+1} = 1$ to the $(n+1)$ -th item and price $p'_i = p_i + 1$ to each other item $i \in [n]$, where \mathbf{p} is an optimal vector for the original instance.

Chapter 7

On the Complexity of Optimal Lottery Pricing and Randomized Mechanisms

7.1 Introduction

In this chapter, we study the problem of Unit-demand Lottery pricing from a complexity theoretic point of view. In particular, we are interested in the following two questions:

Menu size complexity: How many lotteries are needed to achieve the optimal revenue?

Computational complexity: How difficult it is to compute an optimal menu of lotteries?

As mentioned in Chapter 5, limited work exists for the Lottery pricing problem regarding lower-bounds; and this is for additive buyers. Hart and Nisan [Hart and Nisan, 2013], who introduced the notion of menu size, showed that there exists a correlated continuous distribution for which no mechanism of finite menu size can achieve a positive fraction of the optimal revenue. This result, which by itself does not preclude the existence of an efficient mechanism, was complemented by Daskalakis et al. [Daskalakis *et al.*, 2014a] who showed that, unless $P^{\#P} \subseteq ZPP$, there is no efficient algorithm that implements an optimal

mechanism for product distributions, even when all items have support 2.

On the other hand, both the menu-size complexity and the computational complexity of the unit-demand case remained well-known open problems. For example, no instance was known previously to require exponentially many lotteries for the optimal revenue. This is addressed by our work [Chen *et al.*, 2015] with an explicit, simple product distribution \mathcal{D}^* , for which exponentially many lotteries are needed to achieve the optimal revenue. In particular, we obtain the following theorem (and refer the reader to our paper for the proof).

Theorem 11 ([Chen *et al.*, 2015]). *Let \mathcal{D}' denote the distribution supported on $\{1, 2\}$, with probabilities $(1 - p, p)$, and let \mathcal{D}'' denote the distribution supported on $\{0, n + 2\}$, with probabilities $(1 - p, p)$, where $p = 1/n^2$. When n is sufficiently large, any optimal menu for $\mathcal{D}^* = \mathcal{D}' \times \mathcal{D}' \times \cdots \times \mathcal{D}' \times \mathcal{D}''$ over n items must have $\Omega(2^n)$ many different lotteries.*

To complement this, we prove in this thesis a number of positive results regarding the menu-size complexity of special cases of Unit-demand Lottery Pricing. We start by noting that all distributions in \mathcal{D}^* are the same except for one. This is indeed necessary. We show that lotteries do not help when \mathcal{D}_i 's have support size 2 and share the same high value.

Theorem 12. *If $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ and $\text{supp}(\mathcal{D}_i) = \{a_i, b\}$ with $a_i < b$ for all $i \in [n]$, Item Pricing achieves the same expected revenue as Lottery Pricing.*

This together with Theorem 9 also implies that an optimal menu in this case can be computed in polynomial time. Furthermore, we show that for the special case of two items the condition of \mathcal{D}_1 and \mathcal{D}_2 sharing the same high value can be dropped.

Theorem 13. *If both \mathcal{D}_1 and \mathcal{D}_2 have support size 2, then an optimal item pricing for $\mathcal{D}_1 \times \mathcal{D}_2$ achieves the same expected revenue as that of an optimal menu of lotteries.*

On the other hand, in Section 7.5 we give examples of three-item support-size-2 and two-item support-size-3 instances where lotteries do achieve strictly higher revenue than item pricing.

We are now ready to present our *main result* which regards the problem of *computing* an optimal menu of lotteries. We already argued in Chapter 5 that Theorem 11 does not rule out the existence of a deterministic polynomial-time algorithm that, given any $\mathbf{v} \in \mathcal{D}$,

outputs a lottery $\ell_{\mathbf{v}}$ such that $\{\ell_{\mathbf{v}} : \mathbf{v} \in \mathcal{D}\}$ is an optimal menu for \mathcal{D}^* . The question of whether a *universal* efficient algorithm that computes an optimal menu in this fashion exists for product distributions is motivated by a folklore connection between the lottery problem and the optimal mechanism design problem. Consider the same setting, where a unit-demand buyer with values drawn from \mathcal{D} is interested in n items offered by a seller. Here a *mechanism* is a (possibly randomized) map \mathcal{B} from the set D to $([n] \cup \{\text{nil}\}) \times \mathbb{R}$, where $\mathcal{B}(\mathbf{v}) = (b, p)$ means that the buyer is assigned item b (or no item if $b = \text{nil}$) and pays p to the seller. The optimal mechanism design problem is then to find an *individually rational* and *truthful* mechanism (see definitions in Section 7.2.1) that maximizes the expected revenue of the seller.

Let $\bar{\mathcal{B}}(\mathbf{v}) = (\mathbf{x}(\mathbf{v}), \bar{p}(\mathbf{v}))$ denote the expected outcome of \mathcal{B} on \mathbf{v} , i.e., $x_i(\mathbf{v})$ is the probability of $\mathcal{B}(\mathbf{v})$ assigning item i and $\bar{p}(\mathbf{v})$ is the expected payment. It follows trivially from definitions of the two problems that, under the same \mathcal{D} , \mathcal{B} is an optimal mechanism iff $(\bar{\mathcal{B}}(\mathbf{v}) : \mathbf{v} \in D)$ is an optimal menu.

In this thesis, we show that there exists no efficient universal algorithm to implement an optimal mechanism even when \mathcal{D}_i 's have support size 3, unless $\text{P}^{\text{NP}} = \text{P}^{\#\text{P}}$:

Theorem 14. *Unless $\text{P}^{\text{NP}} = \text{P}^{\#\text{P}}$, there exists no algorithm $\mathcal{A}(\cdot, \cdot)$ with the following two properties:*

1. *\mathcal{A} is a randomized polynomial-time algorithm that always terminates in polynomial time.*
2. *Given any instance $I = (n, \mathcal{D}_1, \dots, \mathcal{D}_n)$ to the optimal mechanism design problem, where each \mathcal{D}_i has support size 3, and any $\mathbf{v} \in \text{supp}(\mathcal{D}_1) \times \dots \times \text{supp}(\mathcal{D}_n)$, $\mathcal{A}(I, \mathbf{v})$ always outputs a pair in $([n] \cup \{\text{nil}\}) \times \mathbb{R}$, such that $\mathcal{B}_I : \mathbf{v} \mapsto \mathcal{A}(I, \mathbf{v})$ is an optimal mechanism for instance I .*

We remark that the optimal solutions in the proof of Theorem 14 have the property that they allocate with probability 1 some item for all valuations; such lotteries (mechanisms) are called complete. Thus, the result holds also for the model where lotteries are required to be complete.

7.1.1 Ideas Behind the Proofs

We begin by pointing out that given $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ the optimal menu for the corresponding instance of (Unit-demand or Additive) Lottery Pricing is characterized by a linear program in which we associate with each \mathbf{v} in $D := \text{supp}(\mathcal{D})$ a set of $n + 1$ variables to capture the lottery that the buyer receives at \mathbf{v} (see Section 7.2.1). We will refer to it as the *standard linear program* for the optimal lottery problem.

The main difficulty in proving Theorem 14 is to characterize optimal solutions to the standard linear program (denoted by $\text{LP}(I)$) for certain input instances I . In particular, we need to embed an instance of a #P-hard problem in I and then show that every optimal solution to $\text{LP}(I)$ helps us solve it. However, characterizing optimal solutions to $\text{LP}(I)$ is very challenging due to its exponentially many variables and constraints, which result in a highly complex geometric object for which our current understanding is still very limited.

The high-level approach for our proof of Theorem 14 is similar to that of [Daskalakis *et al.*, 2014a]. We simplify the problem by relaxing the standard linear program $\text{LP}(I)$ to a smaller linear program $\text{LP}'(I)$ on the same set of variables $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ but only subject to a subset of carefully picked constraints of $\text{LP}(I)$. Here $\mathbf{q}(\mathbf{v})$ is a tuple of n variables with $q_i(\mathbf{v})$ being the probability of the buyer receiving item i in the lottery; $u(\mathbf{v})$ is the utility of the buyer at \mathbf{v} to replace the role of price of the lottery. Then we focus on a highly restricted family of instances I and characterize optimal solutions to $\text{LP}'(I)$, taking advantages of the relaxed, simplified $\text{LP}'(I)$ as well as special structures of I . Finally we show that every optimal solution to $\text{LP}'(I)$ is a feasible and thus, optimal solution to the standard linear program $\text{LP}(I)$ as well, and can be used to solve the #P-hard instance embedded in it.

The similarity between our proof techniques and those of [Daskalakis *et al.*, 2014a], however, stops here due to a subtle but crucial difference between the two linear programs. In our standard $\text{LP}(I)$, the allocation variables $\mathbf{q}(\mathbf{v})$ must sum to at most 1 because the buyer is unit-demand. For the additive setting, on the other hand, there is no such constraint on the sum of $q_i(\mathbf{v})$ but the only constraint is that $q_i(\mathbf{v}) \in [0, 1]$ for all i . It turns out that this difference requires a completely different set of ideas and techniques to carry out the plan described above for the unit-demand setting, which we now sketch.

Recall that the goal is to embed a subset-sum-type #P-hard problem in I . Let g_1, \dots, g_n denote the input integers of the problem (its definition does not matter for now). We consider an instance with $n+2$ items where item i is supported on $\{a_i, \ell_i, h_i\}$ for each $i \in [n]$ with $a_i \approx 1$, $\ell_i < h_i$ and $\ell_i \approx h_i \approx 2$. In particular, h_i 's are perturbed from 2 carefully to encode g_i 's. Items $n+1$ and $n+2$ are supported on $\{0, s\}$ and $\{0, t\}$, respectively, for some s, t with $t \gg s \gg 1$. The role of item $n+1$ is to make sure that $\mathbf{q}(\mathbf{a})$, where $\mathbf{a} = (a_1, \dots, a_n, 0, 0)$, is a (almost) uniform distribution, by choosing s carefully. After this step, utilities of vectors \mathbf{w} with $w_{n+1} = w_{n+2} = 0$ can be shown to encode the desired sums of subsets of $\{g_1, \dots, g_n\}$ in every optimal solution to $\text{LP}(I)$. Let $\mathbf{c} = (a_1, \dots, a_n, 0, t)$. Our characterization of optimal solutions to $\text{LP}(I)$ implies that the utility of each vector \mathbf{v} with $v_{n+1} = 0$ and $v_{n+2} = t$ satisfies $u(\mathbf{v}) = \max\{u(\mathbf{w}), u(\mathbf{c})\}$ with $\mathbf{w} = (v_1, \dots, v_n, 0, 0)$, and $u(\mathbf{c})$ is tightly controlled by our choice of t . As \mathbf{w} has $w_{n+1} = w_{n+2} = 0$, $u(\mathbf{w})$ encodes the sum of a certain subset of $\{g_1, \dots, g_n\}$. Combining all these ingredients, we show that the #P-hard problem can be solved by choosing an appropriate parameter t , and then comparing $u(\mathbf{c})$ with $u(\mathbf{v})$ at a specific \mathbf{v} with $v_{n+1} = 0$ and $v_{n+2} = t$ (the choice of \mathbf{v} depends on part of the instance of the #P-hard problem) in any optimal solution.

For Theorems 12 and 13, we identify suitable convex combinations of the revenues of item pricings which upper bound the revenues of all lotteries. Note that this proof method is not only sound, but also complete in the pricing problem in all cases where randomization does not help, by the properties of linear programming; the problem is to show the existence of suitable coefficients for the convex combinations.

7.2 Preliminaries

We now present the standard linear program and prove a few basic properties about it.

7.2.1 The Standard Linear Program

Consider an instance $I = (n, \mathcal{D}_1, \dots, \mathcal{D}_n)$ of Lottery Pricing (see Chapter 5), where a seller offers n items, indexed by $[n] = \{1, \dots, n\}$, to a unit-demand buyer, whose valuation v_1, \dots, v_n of items is drawn from n independent discrete distributions \mathcal{D}_i , $i \in [n]$. Each

distribution \mathcal{D}_i is given explicitly in I , including both its support $D_i = \text{supp}(\mathcal{D}_i)$ and the probability of each value in D_i . Let $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ and $D = D_1 \times \cdots \times D_n$.

We now give the first (not the standard one) linear program characterization of optimal solutions to the optimal lottery problem. For each $\mathbf{v} \in D$ we introduce $n + 1$ variables $\mathbf{q}(\mathbf{v}) = (q_1(\mathbf{v}), \dots, q_n(\mathbf{v}))$ and $p(\mathbf{v})$ to denote the allocation vector and price of the lottery that the buyer receives at \mathbf{v} . Then the menu is given by $M = \{(\mathbf{q}(\mathbf{v}), p(\mathbf{v})) : \mathbf{v} \in D\}$. The only conditions are to make sure the utility of the buyer is always nonnegative and that $(\mathbf{q}(\mathbf{v}), p(\mathbf{v}))$ is a lottery in M that maximizes the utility of the buyer. This gives us a linear program characterization of optimal solutions over variables $(p(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$:

$$\text{Maximize } \sum_{\mathbf{v} \in D} \Pr[\mathbf{v}] \cdot p(\mathbf{v}) \text{ subject to} \quad (7.1)$$

$$q_i(\mathbf{v}) \geq 0 \text{ and } \sum_{i \in [n]} q_i(\mathbf{v}) \leq 1, \quad \text{for all } \mathbf{v} \in D \text{ and } i \in [n].$$

$$\sum_{i \in [n]} v_i \cdot q_i(\mathbf{v}) - p(\mathbf{v}) \geq 0, \quad \text{for all } \mathbf{v} \in D. \quad (7.2)$$

$$\sum_{i \in [n]} v_i \cdot q_i(\mathbf{v}) - p(\mathbf{v}) \geq \sum_{i \in [n]} v_i \cdot q_i(\mathbf{w}) - p(\mathbf{w}), \text{ for all } \mathbf{v}, \mathbf{w} \in D. \quad (7.3)$$

To obtain the standard linear program, we use instead of the price variables $p(\mathbf{v})$, variables $u(\mathbf{v})$ for the utilities of the buyer at the valuations \mathbf{v} , replacing $p(\mathbf{v})$ by the expression $\sum_i v_i \cdot q_i(\mathbf{v}) - u(\mathbf{v})$:

$$\text{Maximize } \sum_{\mathbf{v} \in D} \Pr[\mathbf{v}] \cdot \left(\sum_{i \in [n]} v_i \cdot q_i(\mathbf{v}) - u(\mathbf{v}) \right) \text{ subject to} \quad (7.4)$$

$$u(\mathbf{v}) \geq 0, q_i(\mathbf{v}) \geq 0, \text{ and } \sum_{i \in [n]} q_i(\mathbf{v}) \leq 1, \text{ for all } \mathbf{v} \in D \text{ and } i \in [n].$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{i \in [n]} (v_i - w_i) \cdot q_i(\mathbf{v}), \quad \text{for all } \mathbf{v}, \mathbf{w} \in D. \quad (7.5)$$

We will refer to it as the standard linear program that characterizes optimal solutions to the lottery problem and denote it by $\text{LP}(I)$. Given an optimal solution $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ to $\text{LP}(I)$, we refer to the number of lotteries in the menu it induces as its *menu size*.

For the optimal mechanism design problem (with a single unit-demand buyer), the setting is exactly the same (and so are the input instances I). A randomized mechanism is a randomized algorithm \mathcal{B} that, given $\mathbf{v} \in D$, returns a pair (a, p) , where $a \in [n] \cup \{\text{nil}\}$ is the item assigned to the buyer (or no item is assigned if $a = \text{nil}$) and $p \in \mathbb{R}$ is the payment from the buyer. Given \mathcal{B} , let $\bar{\mathcal{B}}(\mathbf{v}) = (\mathbf{x}(\mathbf{v}), \bar{p}(\mathbf{v}))$ denote the expected outcome of \mathcal{B} on \mathbf{v} : $x_i(\mathbf{v})$ is the probability that $\mathcal{B}(\mathbf{v})$ assigns item i and $\bar{p}(\mathbf{v})$ is the expected payment.

We say \mathcal{B} is *individually rational* if the buyer always has a nonnegative utility if she reports truthfully:

$$\sum_{i \in [n]} v_i \cdot x_i(\mathbf{v}) - \bar{p}(\mathbf{v}) \geq 0, \quad \text{for all } \mathbf{v} \in D.$$

We say \mathcal{B} is *truthful* if the buyer has no incentive to misreport:

$$\sum_{i \in [n]} v_i \cdot x_i(\mathbf{v}) - \bar{p}(\mathbf{v}) \geq \sum_{i \in [n]} v_i \cdot x_i(\mathbf{w}) - \bar{p}(\mathbf{w}), \quad \text{for any } \mathbf{v}, \mathbf{w} \in D.$$

The goal of the optimal mechanism design problem is then to find an individually rational and truthful mechanism \mathcal{B} that maximizes the expected revenue $\sum_{\mathbf{v} \in D} \Pr[\mathbf{v}] \bar{p}(\mathbf{v})$. From the definitions \mathcal{B} is an optimal mechanism iff the set of lotteries $\{\bar{\mathcal{B}}(\mathbf{v}) : \mathbf{v} \in D\} = \{(\mathbf{x}(\mathbf{v}), \bar{p}(\mathbf{v})) : \mathbf{v} \in D\}$ is an optimal solution to the lottery problem, that is, \mathcal{B} is an optimal mechanism iff the tuple $(u(\mathbf{v}), \mathbf{x}(\mathbf{v}) : \mathbf{v} \in D)$ it induces is an optimal solution to the standard LP(I), where we similarly replace $\bar{p}(\mathbf{v})$ by the utility $u(\mathbf{v})$ of the buyer.

7.2.2 Properties of Optimal Solutions to LP(I)

Given an instance $I = (n, \mathcal{D}_1, \dots, \mathcal{D}_n)$, we let $\mathbf{a} \in D$ denote the valuation vector with a_i being the lowest value in the support of \mathcal{D}_i for each $i \in [n]$. Then we have

Lemma 7.2.1. $u(\mathbf{a}) = 0$ in any optimal solution $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ to LP(I).

Proof. Note that in any feasible solution to LP(I) we have $u(\mathbf{v}) \geq u(\mathbf{a})$ for all $\mathbf{v} \in D$ by (7.5). If $u(\mathbf{a}) > 0$, replace $u(\mathbf{v})$ by $u(\mathbf{v}) - u(\mathbf{a})$ for all $\mathbf{v} \in D$, which results in a feasible solution with a higher revenue. \square

We assume from now on that $u(\mathbf{a}) = 0$ is fixed and $u(\mathbf{a})$ is no longer a variable of LP(I).

Lemma 7.2.2. *In any feasible solution $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ to $\text{LP}(I)$, the utility function is monotonically nondecreasing, i.e. for any two valuations \mathbf{v}, \mathbf{w} , if $\mathbf{v} \leq \mathbf{w}$ then $u(\mathbf{v}) \leq u(\mathbf{w})$.*

Proof. If $\mathbf{v} \leq \mathbf{w}$ then constraint (7.5) implies that $u(\mathbf{v}) - u(\mathbf{w}) \leq 0$. □

The allocation function \mathbf{q} is not in general monotonic, but if only one coordinate of the valuation changes, then \mathbf{q} changes monotonically in that coordinate. Given $\mathbf{v} \in D$ and $b \in D_j = \text{supp}(D_j)$, we use (\mathbf{v}_{-j}, b) to denote the vector obtained from \mathbf{v} by replacing v_j with b . The following lemma shows that if $b > v_j$, then we must have $q_j(\mathbf{v}_{-j}, b) \geq q_j(\mathbf{v})$.

Lemma 7.2.3. *Let $\mathbf{v} \in D$ and $v_j < b \in D_j$. Then any feasible solution to $\text{LP}(I)$ satisfies $q_j(\mathbf{v}_{-j}, b) \geq q_j(\mathbf{v})$.*

Proof. Let $\mathbf{w} = (\mathbf{v}_{-j}, b)$. Applying (7.5) on both (\mathbf{v}, \mathbf{w}) and (\mathbf{w}, \mathbf{v}) , we get

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{i \in [n]} (v_i - w_i) \cdot q_i(\mathbf{v}) \quad \text{and} \quad u(\mathbf{w}) - u(\mathbf{v}) \leq \sum_{i \in [n]} (w_i - v_i) \cdot q_i(\mathbf{w}).$$

The lemma follows by summing them up and using $v_i = w_i$ for all $i \neq j$. □

The lotteries of an optimal menu are not necessarily complete. However, they are complete for those valuations that are in the upper boundary of the domain D , i.e., have the maximum value in some coordinate (and this value is positive). In particular, if all the item supports have size 2, then all the lotteries in the optimal menu are complete, except possibly for the allocation $\mathbf{q}(\mathbf{a})$ for the valuation \mathbf{a} where all the items have the minimum value.

Lemma 7.2.4. *Let $\mathbf{v} \in D$ be a vector in which $v_i > 0$ is the largest value in D_i for some coordinate i . Then any optimal solution $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ to $\text{LP}(I)$ satisfies $\sum_{j \in [n]} q_j(\mathbf{v}) = 1$.*

Proof. Suppose that $\sum_{j \in [n]} q_j = 1 - c$ with $c > 0$. Increase the value of $q_i(\mathbf{v})$ by c . The value of the objective function strictly increases (by $\Pr[\mathbf{v}] \cdot v_j \cdot c$). The new solution is feasible: note that in (7.5), $q_i(\mathbf{v})$ appears on the right-hand side always with a nonnegative coefficient since $v_i \geq w_i$ for all $\mathbf{w} \in D$. □

7.3 Distributions with Support $\{a_i, b\}$

In this section we prove Theorem 12. Suppose that the n items $i = 1, \dots, n$ have distributions with support $\{a_i, b\}$ of size 2, where $0 \leq a_i < b$, with the same high value b . Let q_i denote the probability that item i has value $v_i = a_i$ (and $1 - q_i$ that it has value $v_i = b$). We will show that lotteries do not offer any advantage over deterministic item pricing. A consequence of course is that in this case we can compute the optimal solution in polynomial time.

Fix an optimal set of lotteries L^* . Let N denote the set of all items $\{1, \dots, n\}$. For each subset $S \subseteq N$ of items we let $\mathbf{v}(S)$ be the valuation in which items in S have value b and the rest have value a_i . Let $\Pr(S)$ be the probability of $\mathbf{v}(S)$. Let L_S be the lottery of L^* that the buyer buys for valuation $\mathbf{v}(S)$, and let p_S be the price of L_S . Let $L_\emptyset = (x_1, \dots, x_n, p_\emptyset)$ be the lottery for the valuation $\mathbf{v}(\emptyset)$. Notice that $\sum_{i \in N} x_i \leq 1$, and $p_\emptyset \leq \sum_{i \in N} a_i x_i$ as the utility is nonnegative. For each subset $S \subseteq N$ of items let $x(S) = \sum_{i \in S} x_i$.

Let R^* be the expected revenue of the optimal set of lotteries L^* . We will show that R^* is upperbounded by a convex combination of the revenues of a set of $n + 1$ item pricings. This implies that R^* is no greater than the revenue of the optimal item pricing.

Consider a valuation $\mathbf{v}(S)$ for a subset $S \neq \emptyset$. The utility of lottery L_\emptyset for valuation $\mathbf{v}(S)$ is

$$\sum_{i \notin S} a_i x_i + b \sum_{i \in S} x_i - p_\emptyset \geq \sum_{i \notin S} a_i x_i + b \sum_{i \in S} x_i - \sum_{i \in N} a_i x_i = \sum_{i \in S} (b - a_i) x_i.$$

The utility of the lottery L_S that is bought under $\mathbf{v}(S)$ must be at least as large as that of L_\emptyset . The value of the lottery L_S is at most b , thus $b - p_S \geq \sum_{i \in S} (b - a_i) x_i$, hence $p_S \leq b - \sum_{i \in S} (b - a_i) x_i$. Therefore, the total optimal expected revenue R^* is:

$$\begin{aligned} R^* &= \sum_{\emptyset \neq S \subseteq N} p_S \Pr(S) + p_\emptyset \Pr(\emptyset) \leq \sum_{\emptyset \neq S \subseteq N} [b - \sum_{i \in S} (b - a_i) x_i] \Pr(S) + \sum_{i \in N} a_i x_i \Pr(\emptyset) \\ &= b(1 - \Pr(\emptyset)) - \sum_{i \in N} (b - a_i) x_i (1 - q_i) + \sum_{i \in N} a_i x_i \Pr(\emptyset). \end{aligned}$$

Consider now the following set of $n + 1$ item pricings: pricing π_0 assigns price b to all items; for each $i \in N$, pricing π_i assigns price a_i to item i and b to all the other items. The expected revenue R_0 of π_0 is $b(1 - \Pr(\emptyset))$. Under the pricing π_i , the revenue is b if $v_i = a_i$

	(a_1, a_2)	(b_1, a_2)	(a_1, b_2)	(b_1, b_2)
(a_1, b_2)	a_1	a_1	b_2	a_1
(b_1, a_2)	a_2	$\max\{b_1, a_2\}$	a_2	a_2
(b_1, b_2)	0	b_1	b_2	b_2
$(a_1, b_2 - t)$	δ	a_1	$b_2 - t$	$b_2 - t$

Table 7.1: Revenue for each potentially optimal pricing (rows) and each possible valuation vector (columns).

and $v_j = b$ for some $j \neq i$, and is a_i in all other cases (i.e., if $v_i = b$ or if all $v_j = a_j$). So the expected revenue R_i of π_i is $b(q_i - \Pr(\emptyset)) + a_i(1 - q_i + \Pr(\emptyset))$.

Let $x_0 = 1 - x(N)$, and consider the convex combination $\sum_{i=0}^n x_i R_i$ of the expected revenues of the $n + 1$ pricings $\pi_i, i = 0, \dots, n$. We have:

$$\begin{aligned}
\sum_{i=0}^n x_i R_i &= x_0 b(1 - \Pr(\emptyset)) + b \sum_{i \in N} x_i (q_i - \Pr(\emptyset)) + \sum_{i \in N} a_i x_i (1 - q_i + \Pr(\emptyset)) \\
&= b \sum_{i=0}^n x_i (1 - \Pr(\emptyset)) - b \sum_{i \in N} x_i (1 - q_i) + \sum_{i \in N} a_i x_i (1 - q_i) + \sum_{i \in N} a_i x_i \Pr(\emptyset) \\
&= b(1 - \Pr(\emptyset)) - \sum_{i \in N} (b - a_i) x_i (1 - q_i) + \sum_{i \in N} a_i x_i \Pr(\emptyset).
\end{aligned}$$

Thus, $R^* \leq \sum_{i=0}^n x_i R_i$, and hence for at least one of the pricings π_i , we must have $R^* \leq R_i$.

This finishes the proof of Theorem 12.

7.4 Two Items with Support Size 2

In this section we show Theorem 13, i.e., that offering lotteries does not improve the expected revenue when there are two items and both distributions \mathcal{D}_1 and \mathcal{D}_2 are of support size 2.

Let $\{a_i, b_i\}$ be the support of \mathcal{D}_i for $i \in \{1, 2\}$, where $0 \leq a_i < b_i$. Let q_i be the probability that item i has value a_i (and $1 - q_i$ that it has value b_i). Without loss of generality, we assume that $b_2 \geq b_1$ and write $t = b_1 - a_1$. We consider the following four item pricings: $(a_1, b_2), (b_1, a_2), (b_1, b_2), (a_1, b_2 - t)$ (according to the algorithm for the optimal item pricing in the support-2 case [Chen *et al.*, 2014], one of them is optimal).

In Table 7.1, we list the revenue for each of the four item pricings (the rows of the table)

		Allocation		
		Item 1	Item 2	Price
Valuation	(a_1, a_2)	w_1	w_2	p_1
	(b_1, a_2)	$1 - x$	x	p_2
	(a_1, b_2)	y	$1 - y$	p_3
	(b_1, b_2)	z	$1 - z$	p_4

Table 7.2: An optimal menu.

at each of the four possible valuations (the columns). The bottom left entry δ of the table is equal to a_1 if $a_2 < b_2 - t$ (i.e., if $t < b_2 - a_2$), and is equal to $b_2 - t$ if $a_2 \geq b_2 - t$.

Consider now an optimal menu L^* of lotteries. By Lemma 7.2.4, all the lotteries, except for the one bought for valuation (a_1, a_2) , are complete. In table 7.2 we list the allocation and price of each lottery bought.

Our plan is again to show that the revenue of L^* is upperbounded by a convex combination of revenues from the four item pricings. We use the following strategy.

Let $\alpha = (1 - q_1)/q_1$. Note that this is the ratio between probabilities of valuations (b_1, a_2) and (a_1, a_2) , and also those of (b_1, b_2) and (a_1, b_2) . The expected revenue of L^* then can be written as

$$q_1 q_2 \cdot (p_1 + \alpha p_2) + q_1 (1 - q_2) \cdot (p_3 + \alpha p_4).$$

Denote by \mathbf{C}_i the i th column vector of Table 7.1. Our goal is to find a non-negative vector $\mathbf{s} = (s_1, s_2, s_3, s_4)$ of weights (view s_i as the weight of the item pricing on the i th row of Table 7.1) with $\sum_{i=1}^4 s_i = 1$ such that

$$\mathbf{s} \cdot (\mathbf{C}_1 + \alpha \mathbf{C}_2) \geq p_1 + \alpha p_2 \quad \text{and} \quad \mathbf{s} \cdot (\mathbf{C}_3 + \alpha \mathbf{C}_4) \geq p_3 + \alpha p_4. \quad (7.6)$$

Let R^* be the revenue of L^* and R_i be the revenue of the item pricing on the i th column of Table 7.1. Such a weight vector \mathbf{s} then implies that $R^* \leq \sum_{i=1}^4 s_i \cdot R_i$, and Theorem 12 follows.

Here is the plan of the rest of the section. We in Section 7.4.1 bound the prices p_i of L^* , and then bound $p_1 + \alpha p_2$ and $p_3 + \alpha p_4$. We then choose an appropriate \mathbf{s} and use these bounds to prove (7.6) in Section 7.4.2.

7.4.1 Upper Bounds for the Prices in L^*

We start with upper bounds for p_i , $i \in \{1, 2, 3, 4\}$.

Bounding p_1 : For valuation (a_1, a_2) , the buyer buys (w_1, w_2, p_1) . Since it has non-negative utility,

$$p_1 \leq a_1 w_1 + a_2 w_2. \quad (7.7)$$

Bounding p_2 : For valuation (b_1, a_2) , the buyer prefers lottery $(1 - x, x, p_2)$ over (w_1, w_2, p_1) .

Thus,

$$b_1(1 - x) + a_2 x - p_2 \geq b_1 w_1 + a_2 w_2 - p_1 \xrightarrow{(7.7)} p_2 \leq b_1 - x(b_1 - a_2) - w_1(b_1 - a_1).$$

Bounding p_4 : For valuation (b_1, b_2) , the buyer prefers lottery $(z, 1 - z, p_4)$ over (w_1, w_2, p_1) , so

$$\begin{aligned} b_1 z + b_2(1 - z) - p_4 &\geq b_1 w_1 + b_2 w_2 - p_1 \xrightarrow{(7.7)} \\ p_4 &\leq b_1 z + b_2(1 - z) - w_1(b_1 - a_1) - w_2(b_2 - a_2). \end{aligned} \quad (7.8)$$

For valuation (b_1, b_2) , lottery $(z, 1 - z, p_2)$ is also preferred over $(1 - x, x, p_2)$, so we have

$$\begin{aligned} b_1 z + b_2(1 - z) - p_4 &\geq b_1(1 - x) + b_2 x - p_2 \xrightarrow{(7.8)} \\ p_4 &\leq b_1 z + b_2(1 - z) - w_1(b_1 - a_1) - x(b_2 - a_2). \end{aligned} \quad (7.9)$$

Hence, from (7.8) and (7.9) it follows that

$$p_4 \leq b_2 - z(b_2 - b_1) - w_1(b_1 - a_1) - \max\{w_2, x\}(b_2 - a_2). \quad (7.10)$$

Bounding p_3 : For valuation (a_1, b_2) , lottery $(y, 1 - y, p_3)$ is preferred over $(z, 1 - z, p_4)$, so we have

$$\begin{aligned} a_1 y + b_2(1 - y) - p_3 &\geq a_1 z + b_2(1 - z) - p_4 \xrightarrow{(7.10)} \\ p_3 &\leq b_2 - (b_2 - a_1)y + z(b_1 - a_1) - w_1(b_1 - a_1) - \max\{w_2, x\}(b_2 - a_2). \end{aligned} \quad (7.11)$$

Similarly, for valuation (a_1, b_2) , lottery $(y, 1 - y, p_3)$ is preferred over (w_1, w_2, p_1) , so we have

$$\begin{aligned} a_1 y + b_2(1 - y) - p_3 &\geq a_1 w_1 + b_2 w_2 - p_1 \xrightarrow{(7.7)} \\ p_3 &\leq a_1 y + b_2(1 - y) - w_2(b_2 - a_2). \end{aligned} \quad (7.12)$$

Plugging in $b_2 \geq a_1$ and $y \geq 0$, we have from (7.11) and (7.12) that

$$\begin{aligned} p_3 &\leq b_2 + z(b_1 - a_1) - w_1(b_1 - a_1) - \max\{w_2, x\}(b_2 - a_2) \\ &\quad \text{and} \\ p_3 &\leq b_2 - w_2(b_2 - a_2). \end{aligned} \tag{7.13}$$

Bounding $p_1 + \alpha p_2$: From (7.7) and (7.8) we get

$$p_1 + \alpha p_2 \leq \alpha b_1 - w_1(\alpha b_1 - (1 + \alpha)a_1) + w_2 a_2 - x\alpha(b_1 - a_2). \tag{7.14}$$

Bounding $p_3 + \alpha p_4$: Combining the first part of (7.14) and (7.10) we get

$$\begin{aligned} p_3 + \alpha p_4 &\leq (1 + \alpha)(b_2 - w_1(b_1 - a_1) - \max\{w_2, x\}(b_2 - a_2)) \\ &\quad - z(\alpha(b_2 - b_1) - (b_1 - a_1)). \end{aligned} \tag{7.15}$$

Similarly, from the second part of (7.14) and (7.10) we get:

$$\begin{aligned} p_3 + \alpha p_4 &\leq b_2(1 + \alpha) - z\alpha(b_2 - b_1) - w_1\alpha(b_1 - a_1) - w_2(b_2 - a_2) \\ &\quad - \max\{w_2, x\}\alpha(b_2 - a_2). \end{aligned} \tag{7.16}$$

Next we will prove that there are non-negative weights s_1, s_2, s_3 and s_4 that sum to 1 and satisfy (7.6).

7.4.2 Upper Bounds for the Expected Revenue

First we note the following useful inequality:

$$w_2 a_2 - x\alpha(b_1 - a_2) \leq \max\{w_2, x\} \cdot (a_2 + \alpha(\max\{b_1, a_2\} - b_1)),$$

which can be verified by checking both cases of $b_1 \geq a_2$ and $b_1 < a_2$.

We start with a sufficient condition on $\mathbf{s} = (s_1, \dots, s_4)$ to satisfy the first part of (7.6).

Lemma 7.4.1. *Suppose that $s_1, s_4 \geq 0$ satisfy $s_1 + s_4 = w_1$, s_2 satisfies $0 \leq s_2 \leq \max\{w_2, x\}$ and*

$$w_2 a_2 - x\alpha(b_1 - a_2) \leq s_2 \cdot (a_2 + \alpha(\max\{b_1, a_2\} - b_1)), \tag{7.17}$$

and $s_3 = 1 - w_1 - s_2$. Then $s_i \geq 0$ for all i , $\sum_{i=1}^4 s_i = 1$, and \mathbf{s} satisfies $\mathbf{s} \cdot (\mathbf{C}_1 + \alpha \mathbf{C}_2) \geq p_1 + \alpha p_2$.

Proof. We have $s_1, s_2, s_4 \geq 0$ by the assumption of the lemma. To see that $s_3 \geq 0$ note that by Lemma 7.2.3 $w_1 \leq 1 - x$. As $w_1 + w_2 \leq 1$, we have $1 - w_1 \geq \max\{w_2, x\}$ and thus, $s_3 \geq 0$. $\sum_{i=1}^4 s_i = 1$ is obvious.

Recall that δ in Table 7.1 is a_1 or $b_2 - t = b_2 - b_1 + a_1 \geq a_1$. Letting $A = \mathbf{s} \cdot (\mathbf{C}_1 + \alpha \mathbf{C}_2)$, we have:

$$\begin{aligned} A &= s_1(1 + \alpha)a_1 + s_2a_2 + s_2\alpha \max\{b_1, a_2\} + s_3\alpha b_1 + s_4\delta + s_4\alpha a_1 \\ &\geq s_1(1 + \alpha)a_1 + s_2a_2 + s_2\alpha \max\{b_1, a_2\} + s_3\alpha b_1 + s_4a_1 + s_4\alpha a_1 \\ &= (s_1 + s_4)(1 + \alpha)a_1 + s_2a_2 + s_2\alpha \max\{b_1, a_2\} + s_3\alpha b_1. \end{aligned}$$

From the choice of the s_i 's: $s_1 + s_4 = w_1$ and $s_4 = 1 - w_1 - s_2$, the above inequality becomes

$$\begin{aligned} A &\geq w_1(1 + \alpha)a_1 + s_2a_2 + s_2\alpha \max\{b_1, a_2\} + (1 - w_1 - s_2)\alpha b_1 \\ &= \alpha b_1 - w_1(\alpha b_1 - (1 + \alpha)a_1) + s_2a_2 + s_2\alpha(\max\{b_1, a_2\} - b_1). \end{aligned}$$

The lemma then follows directly from (7.14) and the assumption (7.17). \square

We next show that there is an \mathbf{s} that satisfies the second part of (7.6) as well as conditions of Lemma 7.4.1.

Lemma 7.4.2. *There exists an \mathbf{s} that satisfies conditions of Lemma 7.4.1 and $\mathbf{s} \cdot (\mathbf{C}_3 + \alpha \mathbf{C}_4) \geq p_3 + \alpha p_4$.*

Proof. Let $B = \mathbf{s} \cdot (\mathbf{C}_3 + \alpha \mathbf{C}_4)$. It follows from Table 7.1 that we have

$$B = s_1b_2 + s_1\alpha a_1 + s_2(1 + \alpha)a_2 + s_3(1 + \alpha)b_2 + s_4(1 + \alpha)(b_2 - t). \quad (7.18)$$

We will distinguish two cases.

Case 1: $(b_2 - b_1)\alpha \geq b_1 - a_1$. Set $s_2 = \max\{w_2, x\}$, $s_3 = 1 - w_1 - \max\{w_2, x\}$, $s_1 = 0$, and $s_4 = w_1$. Clearly, this assignment satisfies the conditions of Lemma 7.4.1. Equation (7.18) gives:

$$B = (1 + \alpha)(b_2 - w_1(b_1 - a_1) - \max\{w_2, x\}(b_2 - a_2)). \quad (7.19)$$

Furthermore, in this case $-z((b_2 - b_1)\alpha - (b_1 - a_1)) \leq 0$, therefore (7.15) and (7.19) give $p_3 + \alpha p_4 \leq B$.

Case 2: $(b_2 - b_1)\alpha < (b_1 - a_1)$. For this case we distinguish 3 subcases.

Case 2.1: $z \leq w_1$. Set $s_1 = w_1$, $s_2 = \max\{w_2, x\}$, $s_3 = 1 - w_1 - \max\{w_2, x\}$, and $s_4 = 0$.

Then

$$\begin{aligned} B &= w_1 b_2 + w_1 \alpha a_1 + \max\{w_2, x\}(1 + \alpha)a_2 + (1 - w_1 - \max\{w_2, x\})(1 + \alpha)b_2 \\ &= b_2(1 + \alpha) - w_1 \alpha(b_2 - a_1) - \max\{w_2, x\}(1 + \alpha)(b_2 - a_2). \end{aligned} \quad (7.20)$$

Using $(b_2 - b_1)\alpha < (b_1 - a_1)$ and $z \leq w_1$ in (7.15), we have

$$\begin{aligned} p_3 + \alpha p_4 &\leq (1 + \alpha)(b_2 - w_1(b_1 - a_1) - \max\{w_2, x\}(b_2 - a_2)) - w_1(\alpha(b_2 - b_1) - (b_1 - a_1)) \\ &= b_2(1 + \alpha) - w_1 \alpha(b_2 - a_1) - \max\{w_2, x\}(1 + \alpha)(b_2 - a_2) = B. \end{aligned}$$

Case 2.2: $z > w_1$ and $x \leq w_2$. Using the same assignment of \mathbf{s} as in Case 2.1, by $z > w_1$, (7.16) gives

$$p_3 + \alpha p_4 \leq b_2(1 + \alpha) - w_1 \alpha(b_2 - a_1) - w_2(b_2 - a_2) - \max\{w_2, x\}\alpha(b_2 - a_2). \quad (7.21)$$

Furthermore, $x \leq w_2$ implies that $w_2 = \max\{w_2, x\}$. It follows from (7.21) and (7.20) that $p_3 + \alpha p_4 \leq B$.

Case 2.3: $z > w_1$ and $x > w_2$. Setting $s_1 = w_1$, $s_3 = 1 - w_1 - s_2$, and $s_4 = 0$, with

$$s_2 = (w_2 + x\alpha)/(1 + \alpha).$$

Clearly $s_2 \leq \max\{w_2, x\}$. We verify that (7.17) at the end but first compare B and $p_3 + \alpha p_4$.

We have

$$\begin{aligned} B &= w_1 b_2 + w_1 \alpha a_1 + s_2(1 + \alpha)a_2 + (1 - w_1 - s_2)(1 + \alpha)b_2 \\ &= b_2(1 + \alpha) - w_1 \alpha(b_2 - a_1) - s_2(1 + \alpha)(b_2 - a_2). \end{aligned} \quad (7.22)$$

Since $x > w_2$, equation (7.21) gives

$$p_3 + \alpha p_4 \leq b_2(1 + \alpha) - w_1 \alpha(b_2 - a_1) - w_2(b_2 - a_2) - x\alpha(b_2 - a_2). \quad (7.23)$$

It follows from our choice of s_2 and $b_2 - a_2 \geq 0$ that $p_3 + \alpha p_4 \leq B$.

Finally we verify that our choice of s_2 satisfies (7.17) in this case. To see this, we have

$$\begin{aligned} &(1 + \alpha)(w_2 a_2 - x\alpha(b_1 - a_2)) - (w_2 + x\alpha)(a_2 + \alpha(\max\{b_1, a_2\} - b_1)) \\ &= w_2 \alpha a_2 - x \alpha b_1 + x \alpha^2 a_2 - w_2 \alpha \max\{b_1, a_2\} + w_2 \alpha b_1 - x \alpha^2 \max\{b_1, a_2\} \leq 0. \end{aligned}$$

The last inequality is true because $x > w_2$. The lemma follows by combining all the cases. \square

Theorem 13 follows from Lemma 7.4.1 and Lemma 7.4.2.

7.5 Small Instances where Lotteries help

In this section, we give examples where lotteries can extract a strictly higher revenue than the optimal item pricing. In the first example, there are three items and each \mathcal{D}_i has support size 2; in the second example, there are two items and each \mathcal{D}_i has support size 3.

Three items, support size 2: We consider the following instance I with three items. The three items have distributions with support $\{5, b_i\}$ for $i \in [3]$, where $b_1 = 10$ and $b_2 = b_3 = 6$. Let p_i be the probability that item i has value 5. Then set $p_1 = 0.6$, $p_2 = 0.7$ and $p_3 = 0.8$.

There are two optimal item pricings: $(10, 6, 5)$ and $(9, 6, 5)$, with expected revenue 6.744. The optimal menu for I consists of four lotteries: $\mathbf{x}_1 = (1, 0, 0)$ at price 9.5, $\mathbf{x}_2 = (0, 1, 0)$ at price 5.5, $\mathbf{x}_3 = (0, 0, 1)$ at price 5.5, and $\mathbf{x}_4 = (0, 0.5, 0.5)$ at price 5. The expected revenue of this menu is 6.806.

Two items, support size 3: Next consider the following instance J with two items and identical distributions. Each item has value 4 with probability 0.5, value 6 with probability 0.2, and value 7 with probability 0.3.

There are also two optimal item pricings: $(6, 4)$ and $(6, 6)$, with expected revenue 4.5. The optimal menu for instance J consists of three lotteries: $\mathbf{x}_1 = (1, 0)$ at price 6, $\mathbf{x}_2 = (0, 1)$ at price 5, and $\mathbf{x}_3 = (0, 0.5)$ at price 2. The expected revenue of this menu is 4.56.

7.6 Hardness of Optimal Mechanism Design

In this section we prove Theorem 14. This is done by giving a polynomial-time reduction from a #P-hard problem called COMP. We delay its definition and proof of #P-hardness to Section 7.6.2.3 and only mention for now that it is a generalization of #SUBSET-SUM.

This section is organized as follows. In Section 7.6.1, we characterize optimal solutions to a relaxation to the standard linear program $LP(I)$ when parameters of the instance I satisfy certain conditions. We call the relaxed linear program $LP'(I)$, and the characterization is summarized in Section 7.6.1.7. In Section 7.6.2 we pin down the rest of parameters of I to embed the $\#P$ -hard problem COMP. More formally, one can construct an instance I to the lottery problem from an instance of COMP in polynomial time such that a specific entry of any optimal solution to $LP'(I)$ can be used to answer COMP. Finally we show that for such instances I , any optimal solution to $LP'(I)$ must be an optimal solution to $LP(I)$. Then an efficient universal algorithm for the optimal mechanism design problem implies $P^{NP} = P^{\#P}$. This finishes the proof of Theorem 14.

7.6.1 Linear Program Relaxation

Let I denote an instance of $n + 2$ items with the following properties. Each item $i \in [n]$ is supported over $D_i = \{a_i, \ell_i, h_i\}$ with $a_i < \ell_i < h_i$. Probabilities of a_i, ℓ_i and h_i are $1 - p - r$, p and r , respectively, where

$$p = \frac{1}{2^{n^4}} \quad \text{and} \quad r = \frac{p}{2^{n^2}}. \quad (7.24)$$

So p and r satisfy $p = (r/p)^{n^2}$. Let $\beta = 1/2^n$. The support $\{a_i, \ell_i, h_i\}$ of item $i \in [n]$ satisfies

$$\ell_i = 2 + 3(n - i)\beta, \quad \ell_i + \beta \leq h_i \leq \ell_i + \left(1 + \frac{1}{2^{2n}}\right)\beta, \quad \text{and} \quad |a_i - 1| = O(np). \quad (7.25)$$

Let $d_i = \ell_i - a_i \approx 1$ and $\tau_i = h_i - \ell_i$. Our choices of ℓ_i and h_i guarantee that $\tau_i \approx \beta$ as well as $\ell_i > h_{i+1} + \beta$ (or $\ell_i \approx h_{i+1} + 2\beta$ more precisely) for all i from 1 to $n - 1$. Item $n + 1$ takes value 0 with probability $1 - \delta$, and s with probability δ ; item $n + 2$ takes value 0 with probability $1 - \delta^2$, and t with probability δ^2 . So let $D_{n+1} = \{0, s\}$, $D_{n+2} = \{0, t\}$, and $D = D_1 \times \cdots \times D_{n+2}$. We impose the following conditions on δ, s and t throughout Section 7.6.1:

$$\delta = \frac{1}{2^{n^6}}, \quad s = \Theta\left(\frac{1}{pn}\right), \quad t = O\left(\frac{\beta}{r^{m+1}m}\right), \quad t = \Omega\left(\frac{\beta}{r^{m+1}m2^n}\right), \quad \text{where } m = \lceil n/2 \rceil. \quad (7.26)$$

Note that $\delta \ll r \ll p$, and $t = 2^{\Theta(n^5)} \gg s = 2^{\Theta(n^4)} \gg 1$. Precise values of the a_i 's, the h_i 's, and s and t will be chosen later on in Section 7.6.2 after we have analyzed the

structure of the problem. In particular, the h_i 's and t will be used to reflect the instance of the #P-hard problem that we will embed in I and $\text{LP}(I)$. (7.24), (7.25) and (7.26) are sufficient for our analysis in Section 7.6.1 of the relaxed LP to be described below.

We need some notation before describing the relaxation of $\text{LP}(I)$. Given $\mathbf{v} \in D$, we use $S(\mathbf{v})$ to denote the set of $i \in [n]$ such that $v_i \in \{\ell_i, h_i\}$, $S^-(\mathbf{v})$ to denote the set of $i \in [n]$ such that $v_i = \ell_i$, and $S^+(\mathbf{v})$ to denote the set of $i \in [n]$ such that $v_i = h_i$. So we always have $S(\mathbf{v}) = S^+(\mathbf{v}) \cup S^-(\mathbf{v}) \subseteq [n]$.

Next we partition D into T_1, T_2, T_3, T_4 , where T_1 consists of vectors with $v_{n+1} = v_{n+2} = 0$, T_2 consists of vectors with $v_{n+1} = s$ and $v_{n+2} = 0$, T_3 consists of vectors with $v_{n+2} = t$ and $v_{n+1} = 0$, and T_4 consists of vectors with $v_{n+1} = s$ and $v_{n+2} = t$. We call vectors in T_i *type- i* vectors. We denote the bottom vector $(a_1, \dots, a_n, 0, 0)$ by \mathbf{a} , $(a_1, \dots, a_n, s, 0)$ by \mathbf{c}_2 , $(a_1, \dots, a_n, 0, t)$ by \mathbf{c}_3 , and (a_1, \dots, a_n, s, t) by \mathbf{c}_4 (so \mathbf{c}_i is the *bottom* of type- i vectors for i from 2 to 4). By Lemma 7.2.1, we have $u(\mathbf{a}) = 0$ in any optimal solution to $\text{LP}(I)$ so we fix it to be 0.

Given $\mathbf{v} \in D$, we write $\text{BLOCK}(\mathbf{v})$ to denote the set of $\mathbf{w} \in D$ with $S(\mathbf{w}) = S(\mathbf{v})$, $w_{n+1} = v_{n+1}$, and $w_{n+2} = v_{n+2}$; we refer to $\text{BLOCK}(\mathbf{v})$ as the *block* that contains \mathbf{v} . It would also be helpful to view each T_i as a collection of (disjoint) blocks. We say $\mathbf{v} \in D$ is *essential* if $S^+(\mathbf{v}) = \emptyset$ (here the intuition is that within each block, there is a unique essential vector with the largest mass of probability, given $r \ll p$ in (7.24)). We use D' to denote the set of essential vectors, and write $T'_i = T_i \cap D'$ and $T_i^* = T_i \setminus T'_i$ for each i . Given $\mathbf{v} \in D$, we use $\text{LOWER}(\mathbf{v})$ to denote the unique essential vector in $\text{BLOCK}(\mathbf{v})$, i.e., $\text{LOWER}(\mathbf{v})$ is the vector obtained by replacing each h_i in \mathbf{v} by ℓ_i .

We let $\min(S(\mathbf{v}))$ denote the smallest index in $S(\mathbf{v})$ and $S'(\mathbf{v})$ denote $S(\mathbf{v}) \setminus \{\min(S(\mathbf{v}))\}$.

Given a vector $\mathbf{v} \in D$ we follow the convention and write $(\mathbf{v}_{-i}, \alpha)$ to denote the vector obtained from \mathbf{v} by replacing its i th entry v_i with α . We write $(\mathbf{v}_{[n]}, \alpha, \alpha')$ to denote the vector obtained from \mathbf{v} by replacing v_{n+1} with α and v_{n+2} with α' . We let $\rho : T_2 \cup T_3 \cup T_4 \rightarrow T_1$ denote the map with $\rho(\mathbf{v}) = (\mathbf{v}_{[n]}, 0, 0)$.

Given two vectors $\mathbf{v}, \mathbf{w} \in T_i$ of the same type, we write $\mathbf{v} \prec \mathbf{w}$ (or say that \mathbf{v} lies below \mathbf{w} , or \mathbf{w} lies above \mathbf{v}) if either $S(\mathbf{v}) \subset S(\mathbf{w})$, or $S(\mathbf{v}) = S(\mathbf{w})$ and $S^+(\mathbf{v}) \subset S^+(\mathbf{w})$. By definition \prec is transitive.

The linear program $LP'(I)$ is presented in Figure 7.1 which has the same objective function and variables $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ as $LP(I)$. We refer to $u(\mathbf{v})$ and $\mathbf{q}(\mathbf{v})$ as the utility and allocation variables of $\mathbf{v} \in D$, respectively. For convenience, we write $(u(\cdot), \mathbf{q}(\cdot))$ to denote a solution to $LP'(I)$, and call $u(\cdot) : D \rightarrow \mathbb{R}_{\geq 0}$ a utility function. Constraints in Part 1 of $LP'(I)$ concerns variables of type-1 vectors; Part 2 concerns type-2 and type-1 vectors; Part 3 concerns type-3 and type-1 vectors; Part 4 concerns type-4, 3 and 1 vectors.

It is easy to check that $LP'(I)$ is a relaxation of $LP(I)$. The goal is to understand its optimal solutions.

Maximize $\sum_{\mathbf{v} \in D} \Pr[\mathbf{v}] \cdot \left(\sum_{i \in [n+2]} v_i \cdot q_i(\mathbf{v}) - u(\mathbf{v}) \right)$ subject to the following constraints:

Part 0. Same constraints on $u(\mathbf{v})$ and $\mathbf{q}(\mathbf{v})$ as in $LP(I)$:

$$u(\mathbf{v}) \geq 0, \quad q_i(\mathbf{v}) \geq 0, \quad \text{and} \quad \sum_{j \in [n+2]} q_j(\mathbf{v}) \leq 1, \quad \text{for } \mathbf{v} \in D \text{ and } i \in [n+2].$$

Part 1. Constraints on type-1 vectors.

$$u(\mathbf{v}) \geq \sum_{i \in S(\mathbf{v})} d_i \cdot q_i(\mathbf{a}), \quad \text{for } \mathbf{v} \in T'_1; \quad (7.27)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \tau_i \cdot q_i(\mathbf{v}), \quad \text{for } \mathbf{v} \in T_1, i \in S^+(\mathbf{v}) \text{ and } \mathbf{w} = (\mathbf{v}_{-i}, \ell_i); \quad (7.28)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \geq \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{w}), \quad \text{for } \mathbf{v} \in T_1, \mathbf{w} = \text{LOWER}(\mathbf{v}); \quad (7.29)$$

$$u(\mathbf{v}) \geq u(\mathbf{w}), \quad \text{for } \mathbf{v} \in T_1, i \in S(\mathbf{v}), \mathbf{w} = \text{LOWER}(\mathbf{v}_{-i}, a_i); \quad (7.30)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{j \in [n]} (v_j - w_j) \cdot q_j(\mathbf{v}), \quad \text{for } \mathbf{v} \in T_1, i \in S(\mathbf{v}), \mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i). \quad (7.31)$$

Part 2. Constraints on type-2 vectors:

$$u(\mathbf{v}) \geq u(\rho(\mathbf{v})) \text{ and } u(\mathbf{v}) \geq u(\mathbf{c}_2), \quad \text{for } \mathbf{v} \in T_2; \quad (7.32)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \tau_i \cdot q_i(\mathbf{v}), \quad \text{for } \mathbf{v} \in T_2, i \in S^+(\mathbf{v}), \mathbf{w} = (\mathbf{v}_{-i}, \ell_i); \quad (7.33)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{j \in [n]} (v_j - w_j) \cdot q_j(\mathbf{v}), \quad \text{for } \mathbf{v} \in T_2, i \in S(\mathbf{v}), \mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i). \quad (7.34)$$

Part 3: Constraints on type-3 vectors:

$$u(\mathbf{v}) \geq u(\rho(\mathbf{v})) \text{ and } u(\mathbf{v}) \geq u(\mathbf{c}_3), \quad \text{for } \mathbf{v} \in T_3; \quad (7.35)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \tau_i \cdot q_i(\mathbf{v}), \quad \text{for } \mathbf{v} \in T_3, i \in S^+(\mathbf{v}), \mathbf{w} = (\mathbf{v}_{-i}, \ell_i); \quad (7.36)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{j \in [n]} (v_j - w_j) \cdot q_j(\mathbf{v}), \text{ for } \mathbf{v} \in T_3, i \in S(\mathbf{v}), \mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i). \quad (7.37)$$

Part 4: Constraints on type-4 vectors: $u(\mathbf{c}_4) \geq u(\mathbf{c}_3)$ and $u(\mathbf{c}_4) - u(\mathbf{c}_3) \leq s \cdot q_{n+1}(\mathbf{c}_4)$ and

$$u(\mathbf{v}) \geq u(\rho(\mathbf{v})) \text{ and } u(\mathbf{v}) \geq u(\mathbf{c}_4), \quad \text{for } \mathbf{v} \in T_4; \quad (7.38)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \tau_i \cdot q_i(\mathbf{v}), \quad \text{for } \mathbf{v} \in T_4, i \in S^+(\mathbf{v}), \mathbf{w} = (\mathbf{v}_{-i}, \ell_i); \quad (7.39)$$

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{j \in [n]} (v_j - w_j) \cdot q_j(\mathbf{v}), \text{ for } \mathbf{v} \in T_4, i \in S(\mathbf{v}), \mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i). \quad (7.40)$$

Figure 7.1: Relaxed Linear Program $\text{LP}'(I)$

7.6.1.1 Properties of a Small Linear Program

We start with the following lemma on $\mathbf{q}(\mathbf{c}_2)$, $\mathbf{q}(\mathbf{c}_3)$ and $\mathbf{q}(\mathbf{c}_4)$ in any optimal solution to $\text{LP}'(I)$.

Lemma 7.6.1. *If $(u(\cdot), \mathbf{q}(\cdot))$ is an optimal solution to $\text{LP}'(I)$, then it satisfies*

$$q_{n+1}(\mathbf{c}_2) = 1, \quad q_{n+2}(\mathbf{c}_3) = 1, \quad q_{n+1}(\mathbf{c}_4) = \frac{u(\mathbf{c}_4) - u(\mathbf{c}_3)}{s} \quad \text{and} \quad q_{n+2}(\mathbf{c}_4) = 1 - q_{n+1}(\mathbf{c}_4);$$

all other entries of the three vectors $\mathbf{q}(\mathbf{c}_2)$, $\mathbf{q}(\mathbf{c}_3)$ and $\mathbf{q}(\mathbf{c}_4)$ are 0.

Proof. No constraint in $\text{LP}'(I)$ involves $\mathbf{q}(\mathbf{c}_2)$ or $\mathbf{q}(\mathbf{c}_3)$ other than those in Part 0.

For $\mathbf{q}(\mathbf{c}_4)$, in addition to Part 0, there is a constraint in Part 4 that involves $q_{n+1}(\mathbf{c}_4)$: $s \cdot q_{n+1}(\mathbf{c}_4) \geq u(\mathbf{c}_4) - u(\mathbf{c}_3)$. (Note that we have $u(\mathbf{c}_4) \geq u(\mathbf{c}_3)$ by (7.38) in Part 4.)

The lemma then follows from the objective function and that $t \gg s \gg 1$. \square

Let $\hat{D} = T_2 \cup T_3 \cup T_4 \setminus \{\mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4\}$. Vectors $\mathbf{q}(\mathbf{v})$ for $\mathbf{v} \in \hat{D}$ are much more involved. Given a utility function $u : D \rightarrow \mathbb{R}_{\geq 0}$, we define for each $\mathbf{v} \in \hat{D}$ the following small linear

program $\text{LP}(\mathbf{v} : u)$ over $n + 2$ variables $\mathbf{q} = (q_1, \dots, q_{n+2})$:

$$\text{Maximize } \sum_{j \in [n+2]} v_j \cdot q_j - u(\mathbf{v}) \quad \text{subject to} \quad (7.41)$$

$$q_i \geq 0 \quad \text{and} \quad \sum_{j \in [n+2]} q_j \leq 1, \quad \text{for } i \in [n+2]; \quad (7.42)$$

$$\tau_i \cdot q_i \geq u(\mathbf{v}) - u(\mathbf{w}), \quad \text{for } i \in S^+(\mathbf{v}) \text{ and } \mathbf{w} = (\mathbf{v}_{-i}, \ell_i); \quad (7.43)$$

$$\sum_{j \in [n]} (v_j - w_j) \cdot q_j \geq u(\mathbf{v}) - u(\mathbf{w}), \text{ for } i \in S(\mathbf{v}) \text{ and } \mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i). \quad (7.44)$$

Note that $\text{LP}(\mathbf{v} : u)$ uses utilities of \mathbf{v} and \mathbf{w} in blocks nearby \mathbf{v} given by u (so the RHS of the constraints $u(\mathbf{v}) - u(\mathbf{w})$ and $u(\mathbf{v})$ in the objective function are all constants instead of variables), and that q_{n+1} , q_{n+2} and q_i , $i \in [n] \setminus S(\mathbf{v})$, do not appear in constraints of $\text{LP}(\mathbf{v} : u)$ other than (7.42) and the objective function.

Comparing $\text{LP}'(I)$ and $\text{LP}(\mathbf{v} : u)$ gives us the following lemma.

Lemma 7.6.2. *Given a utility function $u(\cdot) : D \rightarrow \mathbb{R}_{\geq 0}$ and $\mathbf{v} \in \hat{D}$, $\mathbf{q}(\mathbf{v})$ satisfies all constraints in $\text{LP}'(\mathbf{v})$ that involve $\mathbf{q}(\mathbf{v})$ iff it is a feasible solution to $\text{LP}(\mathbf{v} : u)$. Moreover, if $(u(\cdot), \mathbf{q}(\cdot))$ is an optimal solution to $\text{LP}'(I)$, then $\mathbf{q}(\mathbf{v})$ must be an optimal solution to $\text{LP}(\mathbf{v} : u)$ for all $\mathbf{v} \in \hat{D}$.*

Proof. The first part is trivial since we included in $\text{LP}(\mathbf{v} : u)$ every constraint in $\text{LP}'(I)$ that involves $\mathbf{q}(\mathbf{v})$.

The second part follows directly from the first part, since the objective function of $\text{LP}(\mathbf{v} : u)$ is exactly $\text{REV}(\mathbf{v})$, the revenue at \mathbf{v} (and we also know that $\Pr[\mathbf{v}] > 0$ for all $\mathbf{v} \in D$). \square

Next we prove a few properties of optimal solutions to $\text{LP}(\mathbf{v} : u)$.

Lemma 7.6.3. *Suppose that $\text{LP}(\mathbf{v} : u)$ is feasible for some utility function $u : D \rightarrow \mathbb{R}_{\geq 0}$ and $\mathbf{v} \in \hat{D}$. Then any optimal solution $\mathbf{q} = (q_1, \dots, q_{n+2})$ to $\text{LP}(\mathbf{v} : u)$ satisfies $q_i(\mathbf{v}) = 0$ for all $i \in [n] \setminus S(\mathbf{v})$ and entries of \mathbf{q} sum to 1. Moreover, we also have $q_{n+2}(\mathbf{v}) = 0$ if $\mathbf{v} \in T_2$, and $q_{n+1}(\mathbf{v}) = 0$ if $\mathbf{v} \in T_3 \cup T_4$.*

Proof. If any of the q_i 's listed above is positive, replacing q_i by 0 and adding q_i to q_{n+1} if $\mathbf{v} \in T_2$ or adding q_i to q_{n+2} if $\mathbf{v} \in T_3 \cup T_4$ would result in a strictly better feasible solution.

If the entries of \mathbf{q} sum to $1 - c$, for some $c > 0$, adding c to either q_{n+1} or q_{n+2} would result in a strictly better feasible solution. \square

In the proof sometimes we need to compare optimal solutions to $\text{LP}(\mathbf{v} : u)$ vs $\text{LP}(\mathbf{v} : u')$ for two utility functions u and u' that are entry-wise close to each other. The following lemma comes in handy.

Lemma 7.6.4. *Assume $\text{LP}(\mathbf{v} : u)$ and $\text{LP}(\mathbf{v} : u')$ are feasible for some $\mathbf{v} \in \hat{D}$ and utilities $u, u' : D \rightarrow \mathbb{R}_{\geq 0}$. Let OPT and OPT' denote optimal values of $\text{LP}(\mathbf{v} : u)$ and $\text{LP}(\mathbf{v} : u')$, respectively. Let $\epsilon > 0$. Then*

1. *If $\mathbf{v} \in T_2$ and $|u(\mathbf{w}) - u'(\mathbf{w})| \leq \epsilon$ for all $\mathbf{w} \in T_2$, then $|\text{OPT} - \text{OPT}'| = O(n\epsilon s/\beta)$.*
2. *If $\mathbf{v} \in T_3$ (or T_4) and $|u(\mathbf{w}) - u'(\mathbf{w})| \leq \epsilon$ for all $\mathbf{w} \in T_3$ (or T_4), then $|\text{OPT} - \text{OPT}'| = O(n\epsilon t/\beta)$.*

Proof. We prove that $\text{OPT}' \geq \text{OPT} - O(n\epsilon s/\beta)$ when $\mathbf{v} \in T_2$. All other cases can be proved similarly.

For this purpose, let \mathbf{q} and \mathbf{q}' denote an optimal solution to $\text{LP}(\mathbf{v} : u)$ and $\text{LP}(\mathbf{v} : u')$, respectively. We consider the following two cases:

Case 1: $q_{n+1} \geq 4n\epsilon/\beta$. Let \mathbf{q}^* denote the following nonnegative vector obtained from \mathbf{q} :

$$q_{n+1}^* = q_{n+1} - |S(\mathbf{v})| \cdot \frac{4\epsilon}{\beta} \quad \text{and} \quad q_i^* = q_i + \frac{4\epsilon}{\beta}, \quad \text{for each } i \in S(\mathbf{v}).$$

It is a feasible solution to $\text{LP}(\mathbf{v} : u')$, given (7.24) and (7.25). Thus, $\text{OPT}' \geq \text{OPT} - O(n\epsilon s/\beta)$.

Case 2: $q_{n+1} < 4n\epsilon/\beta$. This case is more involved. By Lemma 7.6.3 we have $q_{n+2} = q'_{n+2} = 0$. Let

$$c = \max_{i \in [n]} \{q_i - q'_i\}.$$

If $c \leq 8n\epsilon/\beta$, then we immediately have (using $q'_{n+1} \geq 0$)

$$\text{OPT}' \geq \text{OPT} - s \cdot (4n\epsilon/\beta) - n \cdot c \cdot O(1) \geq \text{OPT} - O(n\epsilon s/\beta),$$

since we assumed that $s \gg n$ in (7.26). Otherwise ($c > 8n\epsilon/\beta$), we let $k \in S(\mathbf{v})$ denote an index that achieves the maximum ($k \in S(\mathbf{v})$ since $q_i = q'_i = 0$ for all $i \in [n] \setminus S(\mathbf{v})$) by

Lemma 7.6.3):

$$q_k - q'_k = c > 8n\epsilon/\beta.$$

As $\sum_{i \in S(\mathbf{v})} q_i = 1 - q_{n+1} > 1 - c$ and $\sum_{i \in S(\mathbf{v})} q'_i \leq 1$ we have $q_i \geq q'_i - (n+1)c$ for all $i \in S(\mathbf{v})$. Now let \mathbf{q}^* denote the vector obtained from \mathbf{q} by replacing

$$q_k^* = q_k - (|S(\mathbf{v})| - 1) \cdot \frac{4\epsilon}{\beta} \quad \text{and} \quad q_i^* = q_i + \frac{4\epsilon}{\beta}, \quad \text{for all other } i \in S(\mathbf{v}).$$

One can then verify that \mathbf{q}^* is a feasible solution to $\text{LP}'(\mathbf{v})$. The only nontrivial case in verifying this is to show that the following constraint

$$\sum_{j \in [n]} (v_j - w_j) \cdot q_j^* \geq u'(\mathbf{v}) - u'(\mathbf{w}),$$

holds for any $\mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-k}, a_k)$. To prove this, we note that

$$\sum_{j \in [n]} (v_j - w_j) \cdot q_j^* - \sum_{j \in [n]} (v_j - w_j) \cdot q'_j \geq (v_k - a_k) \cdot \frac{c}{2} - n \cdot O(\beta) \cdot O(nc) = \Omega(c) > 0.$$

As a result, we have $\text{OPT}' \geq \text{OPT} - O(n\epsilon/\beta) \cdot O(n\beta) = \text{OPT} - O(n^2\epsilon)$.

The lemma then follows by combining the two cases and the fact that $s/\beta \gg n$. \square

7.6.1.2 Condition on Utilities of Type-2 Vectors

We show that utilities of type-2 vectors in any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $\text{LP}'(I)$ must satisfy:

CONDITION-TYPE-2: Each type-2 vector $\mathbf{v} \in T_2$ has utility

$$u(\mathbf{v}) = \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_2)\}.$$

Recall that $\rho(\mathbf{v}) = (\mathbf{v}_{-(n+1)}, 0)$ for type-2 vectors. By (7.32) of $\text{LP}'(I)$ in Part 2, $u(\mathbf{v})$ is at least as large as the RHS. So CONDITION-TYPE-2 requires that it is tight for every $\mathbf{v} \in T_2$ in an optimal solution.

We now prove CONDITION-TYPE-2.

Lemma 7.6.5. *Given (7.24), (7.25) and (7.26), any optimal solution to $\text{LP}'(I)$ satisfies CONDITION-TYPE-2.*

Proof. Let $(u(\cdot), \mathbf{q}(\cdot))$ denote an optimal solution to $\text{LP}'(I)$. Let R denote the set of $\mathbf{v} \in T_2$ with

$$u(\mathbf{v}) > \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_2)\}.$$

Note that $\mathbf{c}_2 \notin R$ by the definition of R . Assume for contradiction that R is nonempty.

Our plan is to derive a solution $(u'(\cdot), \mathbf{q}'(\cdot))$ from $(u(\cdot), \mathbf{q}(\cdot))$, by modifying utilities and allocations of type-2 vectors only. We then get a contradiction by showing that $(u'(\cdot), \mathbf{q}'(\cdot))$ is feasible and has a strictly higher revenue than $(u(\cdot), \mathbf{q}(\cdot))$. (Because we only modify utilities and allocations of type-2 vectors, for the feasibility it suffices to verify constraints of $\text{LP}'(I)$ in Part 2.) We use $\text{REV}(\mathbf{v})$ and $\text{REV}'(\mathbf{v})$ to denote the revenue from \mathbf{v} in the old and new solutions. By Lemma 7.6.3 $\text{REV}(\mathbf{v})$ is the value of $\text{LP}(\mathbf{v} : u)$ for $\mathbf{v} \in \hat{D}$.

To define the new solution $(u'(\cdot), \mathbf{q}'(\cdot))$, let $\epsilon > 0$ denote the following parameter:

$$\epsilon = \min \left\{ \min_{\mathbf{v} \in R} \{u(\mathbf{v}) - \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_2)\}\}, \min_{\mathbf{v} \in \hat{D}} \{\text{positive entry in } \mathbf{q}(\mathbf{v})\} \right\}.$$

For each $\mathbf{v} \in T_2$, set $u'(\mathbf{v}) = u(\mathbf{v})$ if $\mathbf{v} \notin R$ and $u'(\mathbf{v}) = u(\mathbf{v}) - \epsilon$ if $\mathbf{v} \in R$. All other entries of u' remain the same as in u . Note that u' still satisfies (7.32) in Part 2. Given $u'(\cdot)$, we set $\mathbf{q}'(\mathbf{v})$ for each $\mathbf{v} \in T_2 \setminus \{\mathbf{c}_2\}$ to be an optimal solution to the linear program $\text{LP}(\mathbf{v} : u')$ (though it is not clear for now if $\text{LP}(\mathbf{v} : u')$ is still feasible or not; we will show that this is indeed the case for every $\mathbf{v} \in T_2 \setminus \{\mathbf{c}_2\}$) and all other allocations remain the same as those in $\mathbf{q}(\cdot)$. This finishes the description of $(u'(\cdot), \mathbf{q}'(\cdot))$.

By Lemma 7.6.2, to show that $(u'(\cdot), \mathbf{q}'(\cdot))$ is well-defined and feasible it suffices to show that $\text{LP}(\mathbf{v} : u')$ is feasible for all $\mathbf{v} \in T_2 \setminus \{\mathbf{c}_2\}$ (because $(u'(\cdot), \mathbf{q}'(\cdot))$ satisfies trivially all constraints of $\text{LP}'(I)$ except (7.33) and (7.34) in Part 2). To see this is the case we fix such a \mathbf{v} . If $\mathbf{v} \in R$ (and $u'(\mathbf{v}) = u(\mathbf{v}) - \epsilon$), every feasible solution to $\text{LP}(\mathbf{v} : u)$ is also feasible to $\text{LP}(\mathbf{v} : u')$. As a result, $\text{LP}(\mathbf{v} : u')$ is feasible as well. Furthermore, we also have $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + \epsilon$ since $u'(\mathbf{v}) = u(\mathbf{v}) - \epsilon$, for any $\mathbf{v} \in R$.

If $\mathbf{v} \notin R$, then either $u'(\mathbf{v}) = u(\mathbf{v}) = u(\mathbf{c}_2)$, or $u'(\mathbf{v}) = u(\mathbf{v}) = u(\rho(\mathbf{v}))$. For the former case, setting $q_{n+1} = 1$ and $q_i = 0$ for all other i is a feasible solution to $\text{LP}(\mathbf{v} : u')$, since $u'(\mathbf{w}) \geq u(\mathbf{c}_2)$ for all $\mathbf{w} \in T_2$. For the latter case, $\mathbf{q} = \mathbf{q}(\rho(\mathbf{v}))$ is a feasible solution to $\text{LP}(\mathbf{v} : u')$ since constraints on $\mathbf{q}(\rho(\mathbf{v}))$ in $\text{LP}'(I)$ are at least as strong as those on \mathbf{q} in $\text{LP}(\mathbf{v} : u)$ using $u'(\mathbf{v}) = u(\rho(\mathbf{v}))$ and $u'(\mathbf{w}) \geq u(\rho(\mathbf{w}))$ for $\mathbf{w} \in T_2$. More specifically, (7.43)

of $\text{LP}(\mathbf{v} : u')$ follows from (7.28) of Part 1 in $\text{LP}'(I)$ over $\mathbf{q}(\rho(\mathbf{v}))$; (7.44) follows from (7.31) of Part 1 in $\text{LP}'(I)$ over $\mathbf{q}(\rho(\mathbf{v}))$. We conclude that $(u'(\cdot), \mathbf{q}'(\cdot))$ is well-defined, feasible to $\text{LP}'(I)$.

The only thing left to show that the expected revenue from $(u'(\cdot), \mathbf{q}'(\cdot))$ is strictly higher. By the definition of $(u'(\cdot), \mathbf{q}'(\cdot))$, we have $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v})$ for all \mathbf{v} other than those in $T_2 \setminus \{\mathbf{c}_2\}$ since each such \mathbf{v} receives the same allocation and utility as in $(u(\cdot), \mathbf{q}(\cdot))$. By Lemma 7.6.4, we also have

$$\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - O(n\epsilon s/\beta), \quad \text{for all } \mathbf{v} \in T_2 \setminus \{\mathbf{c}_2\}.$$

Moreover, if $\mathbf{v} \in T_2 \setminus R$ and there is no $\mathbf{w} \in R$ below \mathbf{v} (or $\mathbf{w} \prec \mathbf{v}$) then $\text{LP}(\mathbf{v} : u')$ is exactly the same as $\text{LP}(\mathbf{v} : u)$ so $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v})$. This inspires us to define $R' \subseteq R$ as the *bottom* of R : $\mathbf{v} \in R'$ if there is no other vector in R below \mathbf{v} . (Since R is nonempty, R' is nonempty as well.) For each $\mathbf{v} \in R'$, we claim that $\text{REV}'(\mathbf{v})$ from the new solution indeed has a much bigger advantage over $\text{REV}(\mathbf{v})$:

$$\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + \Omega(\epsilon s). \quad (7.45)$$

To prove (7.45), we first show that $q_i(\mathbf{v}) > 0$ for some $i \in S(\mathbf{v})$. For this, setting $\mathbf{w} = \text{LOWER}(\mathbf{v}_{-j}, a_j)$ for some $j \in S(\mathbf{v})$ in (7.34) of Part 2 in $\text{LP}'(I)$ (note that $\mathbf{v} \neq \mathbf{c}_2$ implies $S(\mathbf{v}) \neq \emptyset$), we have

$$\sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}) \geq u(\mathbf{v}) - u(\mathbf{w}).$$

It follows from $\mathbf{v} \in R' \subseteq R$ that

$$u(\mathbf{v}) > \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_2)\} \quad \text{and} \quad u(\mathbf{w}) = \max \{u(\rho(\mathbf{w})), u(\mathbf{c}_2)\}.$$

By (7.31) of Part 1 in $\text{LP}'(I)$, we have $u(\rho(\mathbf{v})) \geq u(\rho(\mathbf{w}))$. It follows that $u(\mathbf{v}) > u(\mathbf{w})$ and thus, $q_i(\mathbf{v}) > 0$ for some $i \in S(\mathbf{v})$. Let k be an index in $S(\mathbf{v})$ with $q_k(\mathbf{v}) > 0$. As a result, the following vector \mathbf{q}^* (which is nonnegative because of our choice of ϵ):

$$q_{n+1}^* = q_{n+1}(\mathbf{v}) + \frac{\epsilon}{2} \quad \text{and} \quad q_k^* = q_k(\mathbf{v}) - \frac{\epsilon}{2},$$

and $q_i^* = q_i(\mathbf{v})$ for all other i , must be a feasible solution to $\text{LP}(\mathbf{v} : u')$. (7.45) then follows from $s \gg 1$.

We say a type-2 vector is above R' if it is above one of the vector in R' . Combining all cases together, to show that revenue from $(u'(\cdot), \mathbf{q}'(\cdot))$ is strictly higher than that from $(u(\cdot), \mathbf{q}(\cdot))$, it suffices to show that

$$\Pr[\text{vectors in } R'] \cdot \Omega(\epsilon s) \gg \Pr[(\text{type-2}) \text{ vectors above } R'] \cdot O(n\epsilon s/\beta). \quad (7.46)$$

This follows from our choices of p and r in (7.24). Taking any $\mathbf{v} \in R'$, we have the following bound:

$$\begin{aligned} \Pr[\text{vectors above } \mathbf{v}] &= \Pr[\text{vectors } \mathbf{w} \succ \mathbf{v}, S(\mathbf{w}) = S(\mathbf{v})] \\ &\quad + \Pr[\text{vectors } \mathbf{w} \succ \mathbf{v}, S(\mathbf{v}) \subset S(\mathbf{w})] \\ &= \left(O\left(\frac{nr}{p}\right) + O\left(\frac{np^{|S(\mathbf{v})|+1}}{r^{|S(\mathbf{v})|}\right) \right) \cdot \Pr[\mathbf{v}] \\ &= O\left(\frac{nr}{p} + np \cdot \left(\frac{p}{r}\right)^n\right) \cdot \Pr[\mathbf{v}] \ll \frac{\beta}{n} \cdot \Pr[\mathbf{v}]. \end{aligned}$$

Then (7.46) follows from a union bound. This finishes the proof of the lemma. \square

The arguments used in Lemma 7.6.5 imply the following property. Suppose $(u(\mathbf{v}), \mathbf{q}(\mathbf{v})) : \mathbf{v} \in T_1$) satisfies all the constraints of $\text{LP}'(I)$ in Parts 0 and 1. Given any nonnegative number u_2 , we can extend it to T_2 by setting $u(\mathbf{c}_2) = u_2$ and $u(\mathbf{v}) = \max\{u(\rho(\mathbf{v})), u_2\}$ for each other \mathbf{v} in T_2 , and then setting $\mathbf{q}(\mathbf{c}_2)$ according to Lemma 7.6.1 and $\mathbf{q}(\mathbf{v})$ to be an optimal solution to $\text{LP}(\mathbf{v} : u)$ for each other $\mathbf{v} \in T_2$. It is easy to show, by a similar argument as used in Lemma 7.6.5, that $\text{LP}(\mathbf{v} : u)$ is feasible, and $(u(\mathbf{v}), \mathbf{q}(\mathbf{v})) : \mathbf{v} \in T_1 \cup T_2$) now satisfies all the constraints of $\text{LP}'(I)$ in Parts 0, 1 and 2.

7.6.1.3 Conditions on Utilities of Type-4 Vectors

Next we show that utilities of type-4 vectors satisfy the following condition:

CONDITION-TYPE-4: Each type-4 vector $\mathbf{v} \in T_4$ has utility

$$u(\mathbf{v}) = \max\{u(\rho(\mathbf{v})), u(\mathbf{c}_4)\}.$$

Lemma 7.6.6. *Given (7.24), (7.25) and (7.26), any optimal solution to $\text{LP}'(I)$ satisfies CONDITION-TYPE-4.*

Proof. Let $(u(\cdot), \mathbf{q}(\cdot))$ be an optimal solution, and let R be the set of $\mathbf{v} \in T_4$ with $u(\mathbf{v}) > \{u(\rho(\mathbf{v})), u(\mathbf{c}_4)\}$ (so we have $\mathbf{c}_4 \notin R$). Assume for contradiction that R is nonempty. Our plan is to derive $(u'(\cdot), \mathbf{q}'(\cdot))$ from $(u(\cdot), \mathbf{q}(\cdot))$ by modifying utilities and allocations of vectors in $T_4 \setminus \{\mathbf{c}_4\}$ only. We reach a contradiction by showing that the new solution $(u'(\cdot), \mathbf{q}'(\cdot))$ is feasible and has a strictly higher revenue than $(u(\cdot), \mathbf{q}(\cdot))$.

To define the new solution $(u'(\cdot), \mathbf{q}'(\cdot))$, let $\epsilon > 0$ denote the following parameter:

$$\epsilon = \min \left\{ \min_{\mathbf{v} \in R} \left(u(\mathbf{v}) - \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_3)\} \right), \min_{\mathbf{v} \in D} \{ \text{positive entry in } \mathbf{q}(\mathbf{v}) \} \right\}.$$

First for each $\mathbf{v} \in T_4$ we set $u'(\mathbf{v}) = u(\mathbf{v})$ if $\mathbf{v} \notin R$, and $u'(\mathbf{v}) = u(\mathbf{v}) - \epsilon$ if $\mathbf{v} \in R$; all other entries of u' are the same as those in u . Note that $u'(\cdot)$ still satisfies (7.38) in Part 4 of $\text{LP}'(I)$. Given $u'(\cdot)$ we set $\mathbf{q}'(\mathbf{v})$ for each $\mathbf{v} \in T_4 \setminus \{\mathbf{c}_4\}$ to be an optimal solution to the linear program $\text{LP}(\mathbf{v} : u')$. With an argument similar to that used in Lemma 7.6.5, $\text{LP}(\mathbf{v} : u')$ is feasible (if $\mathbf{v} \in R$, $\mathbf{q}(\mathbf{v})$ is feasible; otherwise $\mathbf{q}(\rho(\mathbf{v}))$ is feasible).

Given that $(u'(\cdot), \mathbf{q}'(\mathbf{v}))$ is well-defined and feasible, we next show that its expected revenue is strictly higher than that of $(u(\cdot), \mathbf{q}(\cdot))$. We follow the approach as in the proof of Lemma 7.6.5. Let R' be the bottom of R : R' contains $\mathbf{w} \in R$ if no other vector in R lies below \mathbf{w} . For each $\mathbf{v} \in T_4 \setminus R'$ with $\mathbf{w} \prec \mathbf{v}$ for some $\mathbf{w} \in R'$, we apply $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - O(\text{net}/\beta)$ by Lemma 7.6.4. For each $\mathbf{v} \in T_4 \setminus R'$ that is not above any vector in R' , we have $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v})$. Finally, for each $\mathbf{v} \in R'$, the same proof of (7.45) in Lemma 7.6.5 gives that $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + \Omega(\epsilon t)$.

Combining all the cases and following the same argument used in Lemma 7.6.5, we have

$$\Pr[\text{vectors in } R'] \cdot \Omega(\epsilon t) \gg \Pr[(\text{type-4}) \text{ vectors above } R'] \cdot O(\text{net}/\beta).$$

This finishes the proof of the lemma. □

The arguments used in Lemma 7.6.6 also imply the following fact. Suppose $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1 \cup T_2 \cup T_3)$ satisfies all the constraints of $\text{LP}'(I)$ in Parts 0, 1, 2 and 3. Given any nonnegative number u_4 with

$$u(\mathbf{c}_3) \leq u_4 \leq u(\mathbf{c}_3) + s,$$

we can extend it to T_4 by setting $u(\mathbf{c}_4) = u_4$ and $u(\mathbf{v}) = \max\{u(\rho(\mathbf{v})), u_4\}$ for other \mathbf{v} in T_4 , and setting $\mathbf{q}(\mathbf{c}_4)$ according to Lemma 7.6.1 and $\mathbf{q}(\mathbf{v})$ to be an optimal solution to $\text{LP}(\mathbf{v} : u)$

for other $\mathbf{v} \in T_4$. By similar arguments used in Lemma 7.6.6, $\text{LP}(\mathbf{v} : u)$ is feasible, and $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ now is feasible to $\text{LP}'(I)$.

7.6.1.4 Condition on Utilities of Type-3 Vectors

A similar condition holds for utilities of type-3 vectors in any optimal solution to $\text{LP}'(I)$:

CONDITION-TYPE-3: Each type-3 vector $\mathbf{v} \in T_3$ has utility

$$u(\mathbf{v}) = \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_3)\}.$$

Lemma 7.6.7. *Given (7.24), (7.25) and (7.26), any optimal solution to $\text{LP}'(I)$ satisfies CONDITION-TYPE-3.*

Proof. Assume for contradiction that $(u(\cdot), \mathbf{q}(\cdot))$ is an optimal solution to $\text{LP}'(I)$ that violates CONDITION-TYPE-3. Let R denote the nonempty set of $\mathbf{v} \in T_3$ with $u(\mathbf{v}) > \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_3)\}$ (so $\mathbf{c}_3 \notin R$).

To reach a contradiction, we derive from $(u(\cdot), \mathbf{q}(\cdot))$ a new solution $(u'(\cdot), \mathbf{q}'(\cdot))$ by modifying utilities and allocations of $\mathbf{v} \in T_3 \setminus \{\mathbf{c}_3\}$ only. (All constraints are satisfied trivially except those in Part 3; note that only $u(\mathbf{c}_3)$ appears in Part 4 but it remains the same in $u'(\cdot)$.) We then show that $(u'(\cdot), \mathbf{q}'(\cdot))$ is better.

We define $(u'(\cdot), \mathbf{q}'(\cdot))$ from $(u(\cdot), \mathbf{q}(\cdot))$ as follows. Let $\epsilon > 0$ denote the following parameter:

$$\epsilon = \min \left\{ \min_{\mathbf{v} \in R} \left(u(\mathbf{v}) - \max \{u(\rho(\mathbf{v})), u(\mathbf{c}_3)\} \right), \min_{\mathbf{v} \in D} \{ \text{positive entry in } \mathbf{q}(\mathbf{v}) \} \right\}.$$

For each $\mathbf{v} \in T_3$ we set $u'(\mathbf{v}) = u(\mathbf{v})$ if $\mathbf{v} \notin R$ and $u'(\mathbf{v}) = u(\mathbf{v}) - \epsilon$ if $\mathbf{v} \in R$; all other entries remain the same. Note that the new u' satisfies (7.35) in Part 3 of $\text{LP}'(I)$. Then for each $\mathbf{v} \in T_3 \setminus \{\mathbf{c}_3\}$, we set $\mathbf{q}'(\mathbf{v})$ to be an optimal solution to $\text{LP}(\mathbf{v} : u')$. With an argument similar to the one used in the proof of Lemma 7.6.5, $\text{LP}(\mathbf{v} : u')$ is feasible (if $\mathbf{v} \in R$, $\mathbf{q}(\mathbf{v})$ is feasible; otherwise, $\mathbf{q}(\rho(\mathbf{v}))$ is feasible). All other entries of $\mathbf{q}'(\cdot)$ remain the same. It is clear now that $(u'(\cdot), \mathbf{q}'(\cdot))$ is a feasible solution to $\text{LP}'(I)$.

We compare the expected revenues from $(u(\cdot), \mathbf{q}(\cdot))$ and $(u'(\cdot), \mathbf{q}'(\cdot))$ and show that the latter is higher. Let R' denote the *bottom* of R : R' contains $\mathbf{v} \in R$ if no other vector in R lies

below \mathbf{v} . For each $\mathbf{v} \in T_3 \setminus R'$ above a vector in R' , we apply $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - O(n\epsilon t/\beta)$ by Lemma 7.6.4. For each vector $\mathbf{v} \in T_3 \setminus R'$ that is not above any vector in R' , we have $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v})$. Finally, for each $\mathbf{v} \in R'$, we can show that $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + \Omega(\epsilon t)$ with an argument similar to that in the proof of Lemma 7.6.5.

Combining all these bounds together and following the same argument used in Lemma 7.6.5, we have

$$\Pr[\text{vectors in } R'] \cdot \Omega(\epsilon t) \gg \Pr[(\text{type-3}) \text{ vectors above } R'] \cdot O(n\epsilon t/\beta).$$

This finishes the proof of the lemma. \square

The arguments used in Lemma 7.6.6 imply the following property. Suppose $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1 \cup T_2)$ satisfies all the constraints of $\text{LP}'(I)$ in Parts 0, 1 and 2. Given a nonnegative number u_3 , we can extend it to T_3 by setting $u(\mathbf{c}_3) = u_3$ and $u(\mathbf{v}) = \max\{u(\rho(\mathbf{v})), u_3\}$ for other \mathbf{v} in T_3 , and setting $\mathbf{q}(\mathbf{c}_3)$ according to Lemma 7.6.1 and $\mathbf{q}(\mathbf{v})$ to be an optimal solution to $\text{LP}(\mathbf{v} : u)$ for other $\mathbf{v} \in T_3$. By similar arguments used in Lemma 7.6.7, $\text{LP}(\mathbf{v} : u)$ is feasible, and $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1 \cup T_2 \cup T_3)$ now satisfies Parts 0, 1, 2 and 3.

7.6.1.5 Expected Revenue from Type-2, 3 and 4 Vectors

Before working on type-1 vectors, which is the most challenging part, we summarize our progress. Let $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ be a partial solution that satisfies all constraints of $\text{LP}'(I)$ in Parts 0 and 1. Given $u_2, u_3, u_4 \geq 0$ that satisfy $u_3 \leq u_4 \leq u_3 + s$, let $\text{Ext}(u(\cdot), \mathbf{q}(\cdot); u_2, u_3, u_4)$ denote the following solutions set $\{u'(\mathbf{v}), \mathbf{q}'(\mathbf{v}) : \mathbf{v} \in D\}$ to $\text{LP}'(I)$:

1. $u'(\mathbf{v}) = u(\mathbf{v})$ and $\mathbf{q}'(\mathbf{v}) = \mathbf{q}(\mathbf{v})$ for all $\mathbf{v} \in T_1$;
2. $u'(\mathbf{c}_2) = u_2$, $u'(\mathbf{c}_3) = u_3$, and $u'(\mathbf{c}_4) = u_4$; $\mathbf{q}'(\mathbf{c}_2) = \mathbf{e}_{n+1}$ and $\mathbf{q}'(\mathbf{c}_3) = \mathbf{e}_{n+2}$;
3. All entries of $\mathbf{q}'(\mathbf{c}_4)$ are 0 except $q_{n+1}(\mathbf{c}_4) = (u_4 - u_3)/s$ and $q_{n+2}(\mathbf{c}_4) = 1 - (u_4 - u_3)/s$;
4. For each $i \in \{2, 3, 4\}$: $\mathbf{v} \in T_i \setminus \{\mathbf{c}_i\}$, $u'(\mathbf{v}) = \max\{u(\rho(\mathbf{v})), u_i\}$ and $\mathbf{q}'(\mathbf{v})$ is an optimal solution to $\text{LP}(\mathbf{v} : u')$.

By discussions at the end of Sections 7.6.1.2, 7.6.1.3, and 7.6.1.4, $\text{Ext}(u(\cdot), \mathbf{q}(\cdot) : u_2, u_3, u_4)$ is well-defined (and nonempty). The next two lemmas summarize our progress so far.

Lemma 7.6.8. *Suppose that $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ satisfies all constraints of $LP'(I)$ in Parts 0 and 1. Given any $u_2, u_3, u_4 \geq 0$, where $u_3 \leq u_4 \leq u_3 + s$, solutions in $\text{Ext}(u(\cdot), \mathbf{q}(\cdot) : u_2, u_3, u_4)$ are feasible to $LP'(I)$ and for each $i = 1, 2, 3, 4$, they all share the same expected revenue from type- i vectors.*

Lemma 7.6.9. *Any optimal solution $(u'(\cdot), \mathbf{q}'(\cdot))$ to $LP'(I)$ must belong to $\text{Ext}(u(\cdot), \mathbf{q}(\cdot) : u_2, u_3, u_4)$ where we set $u_i = u'(\mathbf{c}_i)$ for $i = 2, 3, 4$ and $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ to be the restriction of $(u'(\cdot), \mathbf{q}'(\cdot))$ on T_1 .*

Let $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ and $(u'(\mathbf{v}), \mathbf{q}'(\mathbf{v}) : \mathbf{v} \in T_1)$ denote two partial solutions that satisfy Parts 0 and 1 of $LP'(I)$. The next lemma shows that if $|u(\mathbf{v}) - u'(\mathbf{v})|$ is small for all $\mathbf{v} \in T_1$, then expected revenues of $\text{Ext}(u(\cdot), \mathbf{q}(\cdot) : u_2, u_3, u_4)$ and $\text{Ext}(u'(\cdot), \mathbf{q}'(\cdot) : u_2, u_3, u_4)$ from type-2, 3, 4 vectors are also close.

Lemma 7.6.10. *Suppose $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ and $(u'(\mathbf{v}), \mathbf{q}'(\mathbf{v}) : \mathbf{v} \in T_1)$ satisfy all constraints of $LP'(I)$ in Parts 0 and 1 and $|u(\mathbf{v}) - u'(\mathbf{v})| \leq \epsilon$ for all $\mathbf{v} \in T_1$. Let $u_2, u_3, u_4, u'_2, u'_3, u'_4 \geq 0$ with $u_3 \leq u_4 \leq u_3 + s$, $u'_3 \leq u'_4 \leq u'_3 + s$, and $|u_i - u'_i| \leq \epsilon$ for $i = 2, 3, 4$. Then we have*

$$|\text{REV}_i - \text{REV}'_i| \leq O\left(\frac{\delta^{i-1} n \epsilon s}{\beta}\right)$$

where REV_i and REV'_i denote revenues from type- i vectors in solutions of $\text{Ext}(u(\cdot), \mathbf{q}(\cdot) : u_2, u_3, u_4)$ and solutions of $\text{Ext}(u'(\cdot), \mathbf{q}'(\cdot) : u'_2, u'_3, u'_4)$, respectively.

Proof. We focus on $|\text{REV}_4 - \text{REV}'_4|$. The same argument applies to type-3 and 4 vectors. For convenience, we abuse the notation a little bit and still write $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ and $(u'(\mathbf{v}), \mathbf{q}'(\mathbf{v}) : \mathbf{v} \in D)$ to denote two full feasible solutions of $LP'(I)$ after extension. By definition we have

$$|\text{REV}_4 - \text{REV}'_4| = \left| \sum_{\mathbf{v} \in T_4} \text{Pr}[\mathbf{v}] \cdot (\text{REV}(\mathbf{v}) - \text{REV}'(\mathbf{v})) \right|.$$

It is clear that $|\text{REV}(\mathbf{c}_4) - \text{REV}'(\mathbf{c}_4)| \leq O(\epsilon t/s)$. For other $\mathbf{v} \in T_4$, $\mathbf{q}(\mathbf{v})$ is an optimal solution to $LP(\mathbf{v} : u)$ and $\mathbf{q}'(\mathbf{v})$ is an optimal solution to $LP(\mathbf{v} : u')$, both of which are feasible. It follows from Lemma 7.6.4 and

$$|u(\mathbf{w}) - u'(\mathbf{w})| = |\max\{u(\rho(\mathbf{w})), u_4\} - \max\{u'(\rho(\mathbf{w})), u'_4\}| \leq \epsilon, \quad \text{for all } \mathbf{w} \in T_4,$$

that $|\text{REV}(\mathbf{v}) - \text{REV}(\mathbf{v}')| \leq O(n\epsilon t/\beta)$. Since $\sum_{\mathbf{v} \in T_4} \Pr[\mathbf{v}] < \delta^3$ we have $|\text{REV}_4 - \text{REV}'_4| \leq O(\delta^3 n\epsilon t/\beta)$. This finishes the proof of the lemma. \square

7.6.1.6 Condition over Type-1 Vectors

Finally we show that any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $\text{LP}'(I)$ satisfies the following condition:

CONDITION-TYPE-1: For each type-1 essential vector $\mathbf{v} \in T'_1$ and $\mathbf{v} \neq \mathbf{a}$, we have

$$u(\mathbf{v}) = \sum_{i \in S(\mathbf{v})} d_i \cdot q_i(\mathbf{a}).$$

For each $\mathbf{v} \in T'_1$ and $\mathbf{v} \neq \mathbf{a}$, letting $k = \min(S(\mathbf{v}))$ and $S'(\mathbf{v}) = S(\mathbf{v}) \setminus \{k\}$, we have

$$q_i(\mathbf{v}) = q_i(\mathbf{a}), \quad \text{for all } i \in S'(\mathbf{v}), \quad \text{and} \quad q_k(\mathbf{v}) = 1 - \sum_{i \in S'(\mathbf{v})} q_i(\mathbf{a}),$$

while all other entries of $\mathbf{q}(\mathbf{v})$ are 0. Moreover, for each nonessential type-1 vector $\mathbf{v} \in T_1^*$, letting $\mathbf{w} = \text{LOWER}(\mathbf{v})$, we have $\mathbf{q}(\mathbf{v}) = \mathbf{q}(\mathbf{w})$ and

$$u(\mathbf{v}) = u(\mathbf{w}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{w}) = \sum_{i \in S(\mathbf{v})} d_i \cdot q_i(\mathbf{a}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{w}).$$

Note that CONDITION-TYPE-1 does not require $\sum_{i \in [n]} q_i(\mathbf{a}) = 1$. Actually we will only get to impose this condition later in Section 7.6.2.1 after proper choices of a_i 's.

The following three simple lemmas concern solutions that satisfy CONDITION-TYPE-1.

Lemma 7.6.11. *Assume that $(u(\cdot), \mathbf{q}(\cdot))$ satisfies CONDITION-TYPE-1. If two type-1 vectors \mathbf{v} and \mathbf{w} satisfy $S(\mathbf{w}) \subseteq S(\mathbf{v})$, then $q_j(\mathbf{w}) \geq q_j(\mathbf{v})$ for all $j \in S(\mathbf{w})$.*

Lemma 7.6.12. *Assume that $(u(\cdot), \mathbf{q}(\cdot))$ satisfies CONDITION-TYPE-1. Then we have $\text{REV}(\mathbf{v}) = \text{REV}(\mathbf{v}')$ for any two type-1 vectors \mathbf{v} and \mathbf{v}' in the same block.*

Proof. Let $\mathbf{w} = \text{LOWER}(\mathbf{v}) = \text{LOWER}(\mathbf{v}')$. Then by CONDITION-TYPE-1, $\text{REV}(\mathbf{v})$ is equal to

$$\sum_{i \in [n+2]} v_i \cdot q_i(\mathbf{v}) - u(\mathbf{v}) = \sum_{i \in S(\mathbf{w})} v_i \cdot q_i(\mathbf{w}) - u(\mathbf{w}) - \sum_{i \in S^+(\mathbf{v})} \tau_i \cdot q_i(\mathbf{w}) = \sum_{i \in S(\mathbf{w})} \ell_i \cdot q_i(\mathbf{w}) - u(\mathbf{w}),$$

which does not depend on \mathbf{v} but only $\mathbf{w} = \text{LOWER}(\mathbf{v})$. The lemma then follows. \square

Lemma 7.6.13. *Let \mathbf{q} denote an $(n + 2)$ -dimensional nonnegative vector that sums to at most 1. Then there is a unique $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ that satisfies $\mathbf{q}(\mathbf{a}) = \mathbf{q}$ and CONDITION-TYPE-1.*

Moreover, $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ satisfies all constraints of $LP'(I)$ in Parts 0 and 1.

Proof. Part 0, (7.27) and (7.30) are trivial. For (7.28), given $\mathbf{v} \in T_1$, $i \in S^+(\mathbf{v})$, $\mathbf{w} = (\mathbf{v}_{-i}, \ell_i)$, we have

$$u(\mathbf{v}) - u(\mathbf{w}) = \tau_i \cdot q_i(\text{LOWER}(\mathbf{v})) = \tau_i \cdot q_i(\mathbf{v}),$$

by CONDITION-TYPE-1. For (7.29), letting $\mathbf{w} = \text{LOWER}(\mathbf{v})$, we have

$$u(\mathbf{v}) - u(\mathbf{w}) = \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{w}).$$

For (7.31), given $\mathbf{v} \in T_1$, $i \in S(\mathbf{v})$, $\mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i)$, $\mathbf{v}' = \text{LOWER}(\mathbf{v})$ and $\mathbf{w}' = \text{LOWER}(\mathbf{w})$, we have

$$\begin{aligned} u(\mathbf{v}) - u(\mathbf{w}) &= u(\mathbf{v}) - u(\mathbf{v}') + u(\mathbf{v}') - u(\mathbf{w}') + u(\mathbf{w}') - u(\mathbf{w}) \\ &= \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{v}') + d_i \cdot q_i(\mathbf{a}) - \sum_{j \in S^+(\mathbf{w})} \tau_j \cdot q_j(\mathbf{w}'). \end{aligned}$$

Applying Lemma 7.6.11 on \mathbf{v} and \mathbf{w}' (also $\mathbf{q}(\mathbf{v}) = \mathbf{q}(\mathbf{v}')$ and $q_i(\mathbf{a}) \leq q_i(\mathbf{v})$ for $i \in S(\mathbf{v})$), we have

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{v}) + d_i \cdot q_i(\mathbf{v}) - \sum_{j \in S^+(\mathbf{w})} \tau_j \cdot q_j(\mathbf{v}) = \sum_{j \in S(\mathbf{v})} (v_j - w_j) \cdot q_j(\mathbf{v}).$$

This covers all constraints in Parts 0 and 1, and the lemma is proven. \square

Now we prove CONDITION-TYPE-1.

Lemma 7.6.14. *Given (7.24), (7.25) and (7.26), any optimal solution to $LP'(I)$ satisfies CONDITION-TYPE-1.*

Proof. Let $(u(\cdot), \mathbf{q}(\cdot))$ be an optimal solution to $LP'(I)$. Our plan is the following. We first derive a solution $(u^*(\cdot), \mathbf{q}^*(\cdot))$ from $(u(\cdot), \mathbf{q}(\cdot))$, and show that it is feasible to $LP'(I)$. Then we compare expected revenues from them and show that for $(u(\cdot), \mathbf{q}(\cdot))$ to be optimal as assumed, it must satisfy CONDITION-TYPE-1.

Using Lemma 7.6.13, let $(u'(\mathbf{v}), \mathbf{q}'(\mathbf{v}) : \mathbf{v} \in T_1)$ denote the unique partial solution that satisfies $\mathbf{q}'(\mathbf{a}) = \mathbf{q}(\mathbf{a})$ and CONDITION-TYPE-1. Using Lemma 7.6.13 again $(u'(\mathbf{v}), \mathbf{q}'(\mathbf{v}) : \mathbf{v} \in T_1)$ satisfies all constraints of $LP'(I)$ in Parts 0 and 1. By Lemma 7.6.8, we have that $\text{Ext}(u'(\cdot), \mathbf{q}'(\cdot); u(\mathbf{c}_2), u(\mathbf{c}_3), u(\mathbf{c}_4))$ is a well-defined (nonempty) set of feasible solutions to $LP'(I)$ (here $u(\mathbf{c}_3) \leq u(\mathbf{c}_4) \leq u(\mathbf{c}_3) + s$ as $(u(\cdot), \mathbf{q}(\cdot))$ is feasible). Now we use $(u^*(\mathbf{v}), \mathbf{q}^*(\mathbf{v}) : \mathbf{v} \in D)$ to denote a full feasible solution to $LP'(I)$ in $\text{Ext}(u'(\cdot), \mathbf{q}'(\cdot); u(\mathbf{c}_2), u(\mathbf{c}_3), u(\mathbf{c}_4))$. Now we compare expected revenues of $(u^*(\cdot), \mathbf{q}^*(\cdot))$ and $(u(\cdot), \mathbf{q}(\cdot))$.

For this purpose, let REV_i and REV_i^* denote expected revenues of $(u(\cdot), \mathbf{q}(\cdot))$ and $(u^*(\cdot), \mathbf{q}^*(\cdot))$ from type- i vectors, and let REV and REV' denote their overall expected revenues. Let

$$\epsilon = \max_{\mathbf{v} \in T_1} |u(\mathbf{v}) - u^*(\mathbf{v})|.$$

Then by Lemma 7.6.9 and Lemma 7.6.10 we have

$$\begin{aligned} |(\text{REV}_2 + \text{REV}_3 + \text{REV}_4) - (\text{REV}_2^* + \text{REV}_3^* + \text{REV}_4^*)| &\leq O\left(\frac{\delta n \epsilon s + \delta^2 n \epsilon t + \delta^3 n \epsilon t}{\beta}\right) \\ &= O\left(\frac{\delta n s}{\beta}\right) \cdot \sum_{\mathbf{v} \in T_1} |u(\mathbf{v}) - u^*(\mathbf{v})|, \end{aligned}$$

where we used $s \gg \delta t$ from (7.26) and $\sum_{\mathbf{v}} |u(\mathbf{v}) - u^*(\mathbf{v})|$ as a trivial upper bound for ϵ . By our choice of δ , we have $\delta n s / \beta = o(r^{n+1})$. We also have $\Pr[\mathbf{v}] \geq r^n (1 - \delta)(1 - \delta^2) = \Omega(r^n)$ for all $\mathbf{v} \in T_1$. As a result,

$$\begin{aligned} \text{REV} - \text{REV}^* &\leq \text{REV}_1 - \text{REV}_1^* + |(\text{REV}_2 + \text{REV}_3 + \text{REV}_4) - (\text{REV}_2^* + \text{REV}_3^* + \text{REV}_4^*)| \\ &\leq \sum_{\mathbf{v} \in T_1} \Pr[\mathbf{v}] \cdot (\text{REV}(\mathbf{v}) - \text{REV}^*(\mathbf{v})) + o(r^{n+1}) \cdot \sum_{\mathbf{v} \in T_1} |u(\mathbf{v}) - u^*(\mathbf{v})| \\ &= \sum_{\mathbf{v} \in T_1} \Pr[\mathbf{v}] \cdot \left(\sum_{i \in [n+2]} v_i \cdot (q_i(\mathbf{v}) - q_i^*(\mathbf{v})) + (1 + \zeta_{\mathbf{v}}) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) \right) \\ &:= \sum_{\mathbf{v} \in T_1} \Pr[\mathbf{v}] \cdot \text{DIFF}(\mathbf{v}), \end{aligned}$$

for some $\zeta_{\mathbf{v}}$ with $|\zeta_{\mathbf{v}}| = o(r)$ for all $\mathbf{v} \in T_1$. For convenience we use $\text{DIFF}(\mathbf{v})$ to denote each term for \mathbf{v} .

We bound $\text{DIFF}(\mathbf{v})$ of nonessential type-1 vectors first. Fix a $\mathbf{v} \in T_1^*$. We write $\mathbf{w} = \text{LOWER}(\mathbf{v}) \in T_1'$ and $\mathbf{w}_i = \text{LOWER}(\mathbf{v}_{-i}, a_i) \in T_1'$ for each $i \in S(\mathbf{v})$. We have for each

$i \in S(\mathbf{v})$:

$$u(\mathbf{v}) - u(\mathbf{w}_i) \leq (v_i - a_i) \cdot q_i(\mathbf{v}) + \sum_{i \neq j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{v}) = d_i \cdot q_i(\mathbf{v}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{v}).$$

Applying CONDITION-TYPE-1 on $u^*(\cdot)$, we also have

$$\begin{aligned} u(\mathbf{v}) - u(\mathbf{w}_i) &= u(\mathbf{v}) - u^*(\mathbf{v}) + u^*(\mathbf{v}) - u^*(\mathbf{w}) + u^*(\mathbf{w}) - u^*(\mathbf{w}_i) + u^*(\mathbf{w}_i) - u(\mathbf{w}_i) \\ &= (u(\mathbf{v}) - u^*(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j^*(\mathbf{w}) + d_i \cdot q_i(\mathbf{a}) + (u^*(\mathbf{w}_i) - u(\mathbf{w}_i)). \end{aligned}$$

Combining these two together (and plugging in $\mathbf{q}^*(\mathbf{w}) = \mathbf{q}^*(\mathbf{v})$), we have

$$d_i \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j^*(\mathbf{v}) - q_j(\mathbf{v})) \leq (u^*(\mathbf{v}) - u(\mathbf{v})) - (u^*(\mathbf{w}_i) - u(\mathbf{w}_i)).$$

Let $k = \min(S(\mathbf{v}))$ ($S(\mathbf{v}) \neq \emptyset$ since $\mathbf{v} \in T'_1$) and $S'(\mathbf{v}) = S(\mathbf{v}) \setminus \{k\}$. We consider two cases.

Case 1: $k = \min(S(\mathbf{v})) \notin S^+(\mathbf{v})$. Then we have $q_j^*(\mathbf{v}) = q_j(\mathbf{a})$ for all $j \in S^+(\mathbf{v})$ and thus,

$$d_i \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) \leq (u^*(\mathbf{v}) - u(\mathbf{v})) - (u^*(\mathbf{w}_i) - u(\mathbf{w}_i)). \quad (7.47)$$

Given $q_k^*(\mathbf{v}) = 1 - \sum_{i \in S'(\mathbf{v})} q_i(\mathbf{a})$ and that v_k is the (strictly) largest entry in \mathbf{v} , we have

$$\begin{aligned} \sum_{i \in [n+2]} v_i \cdot q_i^*(\mathbf{v}) &= v_k \left(1 - \sum_{i \in S'(\mathbf{v})} q_i(\mathbf{a}) \right) + \sum_{i \in S'(\mathbf{v})} v_i \cdot q_i(\mathbf{a}) = v_k - \sum_{i \in S'(\mathbf{v})} (v_k - v_i) \cdot q_i(\mathbf{a}), \\ \sum_{i \in [n+2]} v_i \cdot q_i(\mathbf{v}) &\leq v_k \left(1 - \sum_{i \in S'(\mathbf{v})} q_i(\mathbf{v}) \right) + \sum_{i \in S'(\mathbf{v})} v_i \cdot q_i(\mathbf{v}) = v_k - \sum_{i \in S'(\mathbf{v})} (v_k - v_i) \cdot q_i(\mathbf{v}). \end{aligned}$$

Combining these two we get

$$\sum_{i \in [n+2]} v_i \cdot (q_i(\mathbf{v}) - q_i^*(\mathbf{v})) \leq \sum_{i \in S'(\mathbf{v})} (v_k - v_i) \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})). \quad (7.48)$$

Since $\tau_i = O(\beta) = O(1/2^n) \ll d_i \approx 1$, there exists a unique tuple $(\gamma_i : i \in S'(\mathbf{v}))$ such that

$$\sum_{i \in S'(\mathbf{v})} (v_k - v_i) \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) = \sum_{i \in S'(\mathbf{v})} \gamma_i \left(d_i \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) \right) \quad (7.49)$$

This is because $(\gamma_i : i \in S'(\mathbf{v}))$ is the unique solution to a linear system with diagonal entries being d_i or $d_i + \tau_i$ and off-diagonal entries being 0 or τ_j for some $j \in S^+(\mathbf{v})$. Furthermore, given $\tau_j = O(\beta)$ and $\beta \leq v_k - v_i \leq 3n\beta$, we claim that $0 < \gamma_i = O(n\beta)$. To see this, we first prove that $|\gamma_i| \leq 6n\beta$ for all i . Assume for contradiction that $|\gamma_i| = \max_j |\gamma_j| > 6n\beta$. Then we have

$$3n\beta \geq |v_k - v_i| \geq |d_i \gamma_i| - n \cdot O(\beta) \cdot |\gamma_i| > (3/4) \cdot |\gamma_i|,$$

a contradiction. Next, assume for contradiction that $\gamma_i \leq 0$ for some i . Then we have

$$\beta \leq v_k - v_i \leq n \cdot O(\beta) \cdot O(n\beta),$$

contradicting with $\beta = 1/2^n$. It follows from these properties of γ_i 's that

$$\begin{aligned} \sum_{i \in [n+2]} v_i \cdot (q_i(\mathbf{v}) - q_i^*(\mathbf{v})) &\leq \sum_{i \in S'(\mathbf{v})} \gamma_i \cdot ((u^*(\mathbf{v}) - u(\mathbf{v})) - (u^*(\mathbf{w}_i) - u(\mathbf{w}_i))) \\ &= \gamma_{\mathbf{v}} \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) + \sum_{i \in S'(\mathbf{v})} \gamma_{\mathbf{v},i} \cdot (u^*(\mathbf{w}_i) - u(\mathbf{w}_i)), \end{aligned} \quad (7.50)$$

for some $\gamma_{\mathbf{v}}$ and $\gamma_{\mathbf{v},i}$ that satisfy $|\gamma_{\mathbf{v}}| = O(n^2\beta)$ and $|\gamma_{\mathbf{v},i}| = O(n\beta)$ for all $i \in S'(\mathbf{v})$.

Case 2: $k = \min(\mathbf{v}) \in S^+(\mathbf{v})$. Then we have for each $i \in S(\mathbf{v})$:

$$\begin{aligned} d_i \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \tau_k \cdot (q_k^*(\mathbf{v}) - q_k(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v}) \setminus \{k\}} \tau_j \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) \\ \leq (u^*(\mathbf{v}) - u(\mathbf{v})) - (u^*(\mathbf{w}_i) - u(\mathbf{w}_i)). \end{aligned} \quad (7.51)$$

For clarity we use LHS_i to denote the left hand side of the inequality above for each $i \in S(\mathbf{v})$.

Then there exists a unique tuple $(\gamma_i : i \in S(\mathbf{v}))$ such that

$$\sum_{i \in S'(\mathbf{v})} ((v_k - v_i) + \gamma_k) \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \gamma_k \cdot (q_k^*(\mathbf{v}) - q_k(\mathbf{v})) = \sum_{i \in S'(\mathbf{v})} \gamma_i \cdot \text{LHS}_i.$$

This is because $(\gamma_i : i \in S(\mathbf{v}))$ is the unique solution to a linear system with diagonal entries being either d_i or $d_i + \tau_i$ for $i \neq k$ and -1 for k and off-diagonal entries being either 0 or τ_j in general and -1 for the column that corresponds to k . Similarly we have $0 < \gamma_i \leq O(n\beta)$ for all $i \in S(\mathbf{v})$. This gives us a connection between the left hand side above and what we

care about since

$$\begin{aligned}
& \sum_{i \in S'(\mathbf{v})} ((v_k - v_i) + \gamma_k) \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \gamma_k \cdot (q_k^*(\mathbf{v}) - q_k(\mathbf{v})) \\
&= \sum_{i \in S'(\mathbf{v})} (v_k - v_i) \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \gamma_k - \gamma_k \left(q_k(\mathbf{v}) + \sum_{i \in S'(\mathbf{v})} q_i(\mathbf{v}) \right) \\
&\geq \sum_{i \in S'(\mathbf{v})} (v_k - v_i) \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) \geq \sum_{i \in [n+2]} v_i \cdot (q_i(\mathbf{v}) - q_i^*(\mathbf{v})), \tag{7.52}
\end{aligned}$$

where the last inequality follows from (7.48). So (7.50) also holds in this case for some $\gamma_{\mathbf{v}}$ and $\gamma_{\mathbf{v},i}$ with absolute values bounded from above by $O(n^2\beta)$ and $O(n\beta)$, respectively.

To summarize our progress so far, we have shown for each nonessential type-1 vector $\mathbf{v} \in T_1^*$:

$$\begin{aligned}
\text{DIFF}(\mathbf{v}) &\leq (1 + \zeta_{\mathbf{v}} + \gamma_{\mathbf{v}})(u^*(\mathbf{v}) - u(\mathbf{v})) \\
&\quad + \sum_{i \in S'(\mathbf{v})} \gamma_{\mathbf{v},i} \cdot (u^*(\text{LOWER}(\mathbf{v}_{-i}, a_i)) - u(\text{LOWER}(\mathbf{v}_{-i}, a_i))).
\end{aligned}$$

Therefore, we have

$$\sum_{\mathbf{v} \in T_1^*} \Pr[\mathbf{v}] \cdot \text{DIFF}(\mathbf{v}) \leq \sum_{\mathbf{v} \in T_1^*} \Pr[\mathbf{v}] \cdot (1 + \gamma'_{\mathbf{v}}) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) + \sum_{\mathbf{v} \in T_1'} \Pr[\mathbf{v}] \cdot \xi_{\mathbf{v}} \cdot (u^*(\mathbf{v}) - u(\mathbf{v})), \tag{7.53}$$

for some $\gamma'_{\mathbf{v}}$ and $\xi_{\mathbf{v}}$ with $|\gamma'_{\mathbf{v}}| = O(n^2\beta)$ (since $|\zeta_{\mathbf{v}}| = o(r)$) and $|\xi_{\mathbf{v}}| \leq O(n^2p\beta)$. For the latter, we used the fact that for any $\mathbf{v} \in T_1'$ the total probability of all vectors in blocks strictly above $\text{BLOCK}(\mathbf{v})$ is at most $\Omega(np)$ -fraction of that of \mathbf{v} . We continue to simplify the first part of the RHS above.

Let $\mathbf{w} = \text{LOWER}(\mathbf{v})$ for some nonessential vector $\mathbf{v} \in T_1^*$. We have

$$u^*(\mathbf{v}) = u^*(\mathbf{w}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j^*(\mathbf{w}) \quad \text{and} \quad u(\mathbf{v}) \geq u(\mathbf{w}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{w})$$

by CONDITION-TYPE-1 and (7.29) in Part 1 of $\text{LP}'(I)$. As a result, we have

$$u^*(\mathbf{v}) - u(\mathbf{v}) \leq u^*(\mathbf{w}) - u(\mathbf{w}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j^*(\mathbf{w}) - q_j(\mathbf{w})).$$

Fix an essential vector $\mathbf{w} \in T'_1$ and let $B = \text{BLOCK}(\mathbf{w}) \setminus \{\mathbf{w}\}$. Then we have

$$\begin{aligned} & \sum_{\mathbf{v} \in B} \Pr[\mathbf{v}] \cdot (1 + \gamma'_v) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) \\ & \leq \sum_{\mathbf{v} \in B} \Pr[\mathbf{v}] \cdot (1 + \gamma'_v) \cdot \left(u^*(\mathbf{w}) - u(\mathbf{w}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j^*(\mathbf{w}) - q_j(\mathbf{w})) \right) \\ & = \Pr[\mathbf{w}] \cdot \alpha_{\mathbf{w}} \cdot (u^*(\mathbf{w}) - u(\mathbf{w})) + \Pr[\mathbf{w}] \sum_{j \in [n+2]} \alpha_{\mathbf{w},j} \cdot (q_j^*(\mathbf{w}) - q_j(\mathbf{w})), \end{aligned}$$

for some $\alpha_{\mathbf{w}}$ and $\alpha_{\mathbf{w},j}$ with absolute values bounded by $|\alpha_{\mathbf{w}}| = O(nr/p)$ and $|\alpha_{\mathbf{w},j}| = O(nr\beta/p)$.

Combining all these inequalities together, we have

$$\begin{aligned} & \sum_{\mathbf{v} \in T_1} \Pr[\mathbf{v}] \cdot \text{DIFF}(\mathbf{v}) \\ & \leq \sum_{\mathbf{v} \in T'_1} \Pr[\mathbf{v}] \cdot \left(\sum_{j \in [n+2]} v_j \cdot (q_j(\mathbf{v}) - q_j^*(\mathbf{v})) + (1 + \zeta_{\mathbf{v}}) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) \right) \\ & \quad + \sum_{\mathbf{v} \in T'_1} \Pr[\mathbf{v}] \cdot \left(\alpha_{\mathbf{v}} \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) + \sum_{j \in [n+2]} \alpha_{\mathbf{v},j} \cdot (q_j^*(\mathbf{v}) - q_j(\mathbf{v})) \right) \\ & \quad + \sum_{\mathbf{v} \in T'_1} \Pr[\mathbf{v}] \cdot \xi_{\mathbf{v}} \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) \\ & = \sum_{\mathbf{v} \in T'_1} \Pr[\mathbf{v}] \cdot \left((1 + \zeta_{\mathbf{v}} + \alpha_{\mathbf{v}} + \xi_{\mathbf{v}}) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) + \sum_{j \in [n+2]} (v_j - \alpha_{\mathbf{v},j}) \cdot (q_j(\mathbf{v}) - q_j^*(\mathbf{v})) \right). \end{aligned}$$

Recall that $|\zeta_{\mathbf{v}}| = o(r)$ and $|\xi_{\mathbf{v}}| \leq O(n^2p\beta)$. We have $1 + \zeta_{\mathbf{v}} + \alpha_{\mathbf{v}} + \xi_{\mathbf{v}} = 1 \pm o(1)$. Fix an essential $\mathbf{v} \in T'_1$. We have $v_j - \alpha_{\mathbf{v},j} \approx 2$ for $j \in S(\mathbf{v})$, and $k = \min(S(\mathbf{v}))$ still has the (strictly) largest coefficient $v_k - \alpha_{\mathbf{v},k}$ since $|\alpha_{\mathbf{v},j}| = O(nr\beta/p) \ll \beta$. As a result, we have (recall that $S'(\mathbf{v}) = S(\mathbf{v}) \setminus \{k\}$)

$$\begin{aligned} \sum_{j \in [n+2]} (v_j - \alpha_{\mathbf{v},j}) \cdot q_j^*(\mathbf{v}) &= (v_k - \alpha_{\mathbf{v},k}) \cdot \left(1 - \sum_{j \in S'(\mathbf{v})} q_j(\mathbf{a}) \right) + \sum_{j \in S'(\mathbf{v})} (v_j - \alpha_{\mathbf{v},j}) \cdot q_j(\mathbf{a}), \\ \sum_{j \in [n+2]} (v_j - \alpha_{\mathbf{v},j}) \cdot q_j(\mathbf{v}) &\leq (v_k - \alpha_{\mathbf{v},k}) \cdot \left(1 - \sum_{j \in S'(\mathbf{v})} q_j(\mathbf{v}) \right) + \sum_{j \in S'(\mathbf{v})} (v_j - \alpha_{\mathbf{v},j}) \cdot q_j(\mathbf{v}). \end{aligned} \tag{7.54}$$

Let $\phi_{\mathbf{v},j} = v_k - v_j - \alpha_{\mathbf{v},k} + \alpha_{\mathbf{v},j}$ for each $j \in S'(\mathbf{v})$. Then we have $\Omega(\beta) \leq \phi_{\mathbf{v},j} \leq O(n\beta)$ and

$$\sum_{j \in [n+2]} (v_j - \alpha_{\mathbf{v},j}) \cdot (q_j(\mathbf{v}) - q_j^*(\mathbf{v})) \leq \sum_{j \in S'(\mathbf{v})} \phi_{\mathbf{v},j} \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})).$$

Plugging this in, we have

$$\sum_{\mathbf{v} \in T_1} \Pr[\mathbf{v}] \cdot \text{DIFF}(\mathbf{v}) \leq \sum_{\mathbf{v} \in T_1'} \Pr[\mathbf{v}] \cdot \left((1 \pm o(1)) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) + \sum_{j \in S'(\mathbf{v})} \phi_{\mathbf{v},j} \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) \right)$$

We also have $u^*(\mathbf{v}) - u^*(\mathbf{v}_{-j}, a_j) = d_j \cdot q_j(\mathbf{a})$ for $\mathbf{v} \in T_1'$ and each $j \in S(\mathbf{v})$, and

$$u(\mathbf{v}) - u(\mathbf{v}_{-j}, a_j) \leq d_j \cdot q_j(\mathbf{v})$$

by (7.31) of $\text{LP}'(I)$. As a result, we have

$$\sum_{j \in S'(\mathbf{v})} \phi_{\mathbf{v},j} \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) \leq \sum_{j \in S'(\mathbf{v})} \frac{\phi_{\mathbf{v},j}}{d_j} \cdot ((u^*(\mathbf{v}) - u(\mathbf{v})) - (u^*(\mathbf{v}_{-j}, a_j) - u(\mathbf{v}_{-j}, a_j))). \quad (7.55)$$

Plugging it back, we have

$$\begin{aligned} \sum_{\mathbf{v} \in T_1} \Pr[\mathbf{v}] \cdot \text{DIFF}(\mathbf{v}) &\leq \sum_{\mathbf{v} \in T_1'} \Pr[\mathbf{v}] \cdot (1 \pm o(1)) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})) \\ &\quad + \sum_{\mathbf{v} \in T_1'} \Pr[\mathbf{v}] \sum_{j \in S'(\mathbf{v})} \frac{\phi_{\mathbf{v},j}}{d_j} \cdot ((u^*(\mathbf{v}) - u(\mathbf{v})) - (u^*(\mathbf{v}_{-j}, a_j) - u(\mathbf{v}_{-j}, a_j))) \\ &\leq \sum_{\mathbf{v} \in T_1'} \Pr[\mathbf{v}] \cdot (1 + \delta_{\mathbf{v}}) \cdot (u^*(\mathbf{v}) - u(\mathbf{v})), \end{aligned}$$

for some $\delta_{\mathbf{v}}$ with absolute value bounded from above by $|\delta_{\mathbf{v}}| \leq o(1) + O(n^2\beta) + O(n\beta \cdot np) = o(1)$.

Since $u(\mathbf{v}) \geq u^*(\mathbf{v})$ for all $\mathbf{v} \in T_1'$ (due to (7.27) of $\text{LP}'(I)$), we must have $u(\mathbf{v}) = u^*(\mathbf{v})$ for all $\mathbf{v} \in T_1'$ by the optimality of $(u(\cdot), \mathbf{q}(\cdot))$. This proved the part of **CONDITION-TYPE-1** on $\mathbf{q}(\mathbf{v})$ of essential vectors.

Combining this with (7.31) of $\text{LP}'(I)$, we have for each $\mathbf{v} \in T_1'$, $i \in S(\mathbf{v})$ and $\mathbf{w} = (\mathbf{v}_{-i}, a_i)$:

$$d_i \cdot q_i(\mathbf{a}) = u(\mathbf{v}) - u(\mathbf{w}) \leq d_i \cdot q_i(\mathbf{v})$$

and thus, $q_i(\mathbf{v}) \geq q_i(\mathbf{a})$. On the other hand, it follows from the optimality of $(u(\cdot), \mathbf{q}(\cdot))$ that both (7.55) and (7.54) must be tight. This implies that $\mathbf{q}(\mathbf{v}) = \mathbf{q}^*(\mathbf{v})$ for all essential vectors $\mathbf{v} \in T'_1$.

For a nonessential type-1 vector $\mathbf{v} \in T_1^*$, letting $\mathbf{w} = \text{LOWER}(\mathbf{v})$, (7.29) in Part 1 of $\text{LP}'(I)$ implies that

$$u(\mathbf{v}) \geq u(\mathbf{w}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{w}) = u^*(\mathbf{v}),$$

because we have proved that $u(\mathbf{w}) = u^*(\mathbf{w})$ and $\mathbf{q}(\mathbf{w}) = \mathbf{q}^*(\mathbf{w})$ (since \mathbf{w} is essential). Then $u(\mathbf{v}) = u^*(\mathbf{v})$ follows from the tightness of (7.53).

Finally, for each nonessential vector $\mathbf{v} \in T_1^*$, we consider two cases ($k = \min(S(\mathbf{v}))$):

Case 1: $k \notin S^+(\mathbf{v})$. $\mathbf{q}(\mathbf{v}) = \mathbf{q}^*(\mathbf{v})$ follows from the tightness of (7.47) and (7.48). (7.47) yields that

$$d_i \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) = 0$$

for all $i \in S'(\mathbf{v})$ (note that we actually do not use $i = k$ in (7.49)). These equations together imply that $q_i(\mathbf{v}) = q_i(\mathbf{a})$ for all $i \in S'(\mathbf{v})$. $q_k(\mathbf{v}) = q_k^*(\mathbf{v})$ follows from the tightness of (7.48).

Case 2: $k \in S^+(\mathbf{v})$. The tightness of (7.52) implies that

$$q_k(\mathbf{v}) = 1 - \sum_{j \in S'(\mathbf{v})} q_j(\mathbf{v}). \quad (7.56)$$

The tightness of (7.51) implies that

$$d_i \cdot (q_i(\mathbf{a}) - q_i(\mathbf{v})) + \tau_k \cdot (q_k^*(\mathbf{v}) - q_k(\mathbf{v})) + \sum_{j \in S^+(\mathbf{v}) \setminus \{k\}} \tau_j \cdot (q_j(\mathbf{a}) - q_j(\mathbf{v})) = 0$$

for all $i \in S'(\mathbf{v})$. Plugging in $q_k^*(\mathbf{v}) = 1 - \sum_{j \in S'(\mathbf{v})} q_j(\mathbf{a})$ and (7.56), we must have $q_i(\mathbf{a}) = q_i(\mathbf{v})$ for all $i \in S'(\mathbf{v})$ and thus, $q_k(\mathbf{v}) = q_k^*(\mathbf{v})$ by (7.56). It follows that $\mathbf{q}(\mathbf{v}) = \mathbf{q}^*(\mathbf{v})$.

This finishes the proof of the lemma. \square

7.6.1.7 Characterization of Optimal Solutions

Let \mathbf{q} be a nonnegative $(n+2)$ -dimensional vector that sums to at most 1, and $u_2, u_3, u_4 \geq 0$ that satisfy $u_3 \leq u_4 \leq u_3 + s$. We use $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ to denote the following set of solutions to $\text{LP}'(I)$:

Let $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ be the unique partial solution that satisfies both $\mathbf{q}(\mathbf{a}) = \mathbf{q}$ and CONDITION-TYPE-1. By Lemma 7.6.13, $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in T_1)$ satisfies all constraints in Parts 0 and 1 of $LP'(I)$. Then we set $\text{Ext}(\mathbf{q}, u_2, u_3, u_4) = \text{Ext}(u(\cdot), \mathbf{q}(\cdot); u_2, u_3, u_4)$.

We record the following lemma.

Lemma 7.6.15. *Given any nonnegative vector \mathbf{q} that sums to at most 1, and $u_2, u_3, u_4 \geq 0$ with $u_3 \leq u_4 \leq u_3 + s$, $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ is a nonempty set of feasible solutions to $LP'(I)$.*

Our characterization of optimal solutions to $LP'(I)$ is summarized in the theorem below.

Theorem 15. *Any optimal solution $(u(\mathbf{v}), \mathbf{q}(\mathbf{v}) : \mathbf{v} \in D)$ to $LP'(I)$ belongs to $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$, where $\mathbf{q} = \mathbf{q}(\mathbf{a})$ and $u_i = u(\mathbf{c}_i)$ for each $i = 2, 3, 4$.*

7.6.2 Choices of Parameters and their Consequences

Now we pin down the rest of parameters: a_i, s, h_i, t , and see how they affect optimal solutions of $LP'(I)$.

7.6.2.1 Setting a_i 's

First, we set a_i 's (see (7.58) below) such that they satisfy (7.25), i.e. $|a_i - 1| = O(np)$, and the expected revenue from type-1 vectors in any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $LP'(I)$ is of the following form

$$\text{CONST} + c \cdot \sum_{i \in [n]} q_i(\mathbf{a}), \quad (7.57)$$

for some $c \approx 1$. By Theorem 15, $(u(\cdot), \mathbf{q}(\cdot))$ lies in $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ for some nonnegative vector \mathbf{q} that sums to at most 1, and some $u_2, u_3, u_4 \geq 0$ that satisfy $u_3 \leq u_4 \leq u_3 + s$. Given that $(u(\cdot), \mathbf{q}(\cdot))$ satisfies CONDITION-TYPE-1, expected revenue from type-1 vectors only depends on $\mathbf{q}(\mathbf{a}) = \mathbf{q}$. We next calculate the expected revenue from type-1 vectors given \mathbf{q} , and the choices of a_i 's will become clear.

First, we have $\text{REV}(\mathbf{a}) = \sum_{i \in [n]} a_i \cdot q_i$ (since $u(\mathbf{a}) = 0$). Given that $(u(\cdot), \mathbf{q}(\cdot))$ satisfies CONDITION-TYPE-1, each essential type-1 vector $\mathbf{v} \in T_1'$ and $\mathbf{v} \neq \mathbf{a}$ has revenue (letting

$$k = \min(S(\mathbf{v}))$$

$$\text{REV}(\mathbf{v}) = \sum_{i \in S'(\mathbf{v})} \ell_i \cdot q_i + \ell_k \cdot \left(1 - \sum_{i \in S'(\mathbf{v})} q_i\right) - \sum_{i \in S(\mathbf{v})} d_i \cdot q_i = \ell_k - \sum_{i \in S'(\mathbf{v})} (\ell_k - a_i) \cdot q_i - d_k \cdot q_k.$$

Given Lemma 7.6.12, the block B that contains $\mathbf{v} \in T'_1$ and $\mathbf{v} \neq \mathbf{a}$ overall contributes

$$\Pr[B] \cdot \text{REV}(\mathbf{v}) = \Pr[B] \cdot \left(\ell_k - \sum_{i \in S'(\mathbf{v})} (\ell_k - a_i) \cdot q_i - (\ell_k - a_k) \cdot q_k \right).$$

It is clear now that expected revenue from type-1 vectors is an affine linear form of q_i 's, $i \in [n]$.

Let c_i denote the coefficient of each q_i in the expected revenue from type-1 vectors. Then \mathbf{a} contributes $\Pr[\mathbf{a}] \cdot a_i$ to c_i (note that $\Pr[\mathbf{a}] \approx 1 - np$ which as we will see is the dominating term in c_i). A block B that contains $\mathbf{v} \in T'_1$ and $\mathbf{v} \neq \mathbf{a}$ contributes 0 if $i \notin S(\mathbf{v})$;

$$- \Pr[B] \cdot (\ell_i - a_i) \quad \text{if } i = \min(S(\mathbf{v})); \text{ and} \quad - \Pr[B] \cdot (\ell_{\min(S(\mathbf{v}))} - a_i) \quad \text{if } i \in S'(\mathbf{v}).$$

More specifically, the total probability of type-1 blocks B and $\mathbf{v} \in B$ with $i = \min(S(\mathbf{v}))$ is

$$(1 - \delta) \cdot (1 - \delta^2) \cdot (1 - p - r)^{i-1} \cdot (p + r);$$

for each $k < i$, the total probability of type-1 blocks B with $i \in S(\mathbf{v})$ and $\min(S(\mathbf{v})) = k$ is

$$(1 - \delta) \cdot (1 - \delta^2) \cdot (1 - p - r)^{k-1} \cdot (p + r) \cdot (p + r).$$

As a result, we have the following explicit expression for c_i (setting $\psi = (1 - \delta)(1 - \delta^2)$):

$$\psi \left((1 - p - r)^n a_i - \sum_{k < i} (1 - p - r)^{k-1} (p + r)^2 (\ell_k - a_i) - (1 - p - r)^{i-1} (p + r) (\ell_i - a_i) \right)$$

To meet both goals, i.e., $c_1 = \dots = c_n \approx 1$ and $|a_i - 1| \leq O(np)$, we set

$$a_i = \frac{1 + \sum_{k < i} (1 - p - r)^{k-1} \cdot (p + r)^2 \cdot \ell_k + (1 - p - r)^{i-1} \cdot (p + r) \cdot \ell_i}{(1 - p - r)^n + \sum_{k < i} (1 - p - r)^{k-1} \cdot (p + r)^2 + (1 - p - r)^{i-1} \cdot (p + r)}. \quad (7.58)$$

It is easy to verify that a_i 's satisfy $1 < a_i \leq 1 + O(np)$. The length of binary representations of each a_i is polynomial in n and a_i 's can be computed efficiently, given p, r and ℓ_i 's as in (7.24) and (7.25).

We summarize the consequence of our choices of a_i 's in the following lemma:

Lemma 7.6.16. *Given choices of a_i 's in (7.58), revenue from type-1 vectors in any feasible solution to $LP'(I)$ that satisfies CONDITION-TYPE-1 is of the form in (7.57) with $c = (1 - \delta)(1 - \delta^2) \approx 1$.*

It is now time to prove that $\mathbf{q}(\mathbf{a})$ sums to 1 in any optimal solution to $LP'(I)$.

Lemma 7.6.17. *Given our choices of a_i 's in (7.58) any optimal solution to $LP'(I)$ satisfies $\sum_{i \in [n]} q_i(\mathbf{a}) = 1$.*

Proof. Assume for contradiction that $\{u(\cdot), \mathbf{q}(\cdot)\}$ is optimal but $\mathbf{q}(\mathbf{a})$ does not satisfy $\sum_{i \in [n]} q_i(\mathbf{a}) = 1$. Let \mathbf{q}' be the vector obtained from $\mathbf{q}(\mathbf{a})$ as follows: If $\sum_{i \in [n+2]} q_i(\mathbf{a}) < 1$, we replace its first entry by

$$q'_1 = q_1(\mathbf{a}) + \epsilon, \quad \text{where } \epsilon = 1 - \sum_{i \in [n+2]} q_i(\mathbf{a}) > 0;$$

otherwise, letting $\epsilon = q_{n+1}(\mathbf{a}) + q_{n+2}(\mathbf{a}) > 0$, we set

$$q'_1 = q_1(\mathbf{a}) + \epsilon \quad \text{and} \quad q'_{n+1} = q'_{n+2} = 0.$$

Let $\{u'(\cdot), \mathbf{q}'(\cdot)\}$ denote a feasible solution from $\text{Ext}(\mathbf{q}', u(\mathbf{c}_2), u(\mathbf{c}_3))$. It follows from Lemma 7.6.16 that expected revenue from type-1 vectors goes up by $\Omega(\epsilon)$ in $\{u'(\cdot), \mathbf{q}'(\cdot)\}$. However, by CONDITION-TYPE-1,

$$|u(\mathbf{v}) - u'(\mathbf{v})| \leq O(\epsilon), \quad \text{for all } \mathbf{v} \in T_1.$$

By Lemma 7.6.10, expected revenue from type-2, 3 and 4 vectors goes down in $\{u'(\cdot), \mathbf{q}'(\cdot)\}$ by at most $O(\delta n \epsilon s / \beta) + O(\delta^2 n \epsilon t / \beta) \ll \epsilon$. This contradicts with the assumption that $\{u(\cdot), \mathbf{q}(\cdot)\}$ is optimal. \square

Given Lemma 7.6.17 we will from now on restrict \mathbf{q} to be a nonnegative n -dimensional vector that sums to exactly 1 in $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$. We also use $\text{REV}(\mathbf{q}, u_2, u_3, u_4)$ to denote the expected revenue of solutions in $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$. All parameters of I have been chosen except s , h_i 's and t .

7.6.2.2 Setting s

Our goal in this section is to show that, by setting

$$s = 2 + \frac{1}{(n - 0.5)p} = \Theta\left(\frac{1}{np}\right). \quad (7.59)$$

any optimal solution to $\text{LP}'(I)$ from $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ must satisfy

$$d_1 \cdot q_1(\mathbf{a}) = d_2 \cdot q_2(\mathbf{a}) = \cdots = d_n \cdot q_n(\mathbf{a}). \quad (7.60)$$

Note that we have chosen both a_i and ℓ_i , and so is $d_i = \ell_i - a_i$. (7.60) then uniquely determines $\mathbf{q}(\mathbf{a})$ in any optimal solution, as by Lemma 7.6.17, $\mathbf{q}(\mathbf{a})$ must sum to 1. (7.60) also implies that $\mathbf{q}(\mathbf{a})$ is indeed very close to the uniform distribution over $[n]$ since $d_i \approx 1$ (more precisely, $|d_i - 1| = O(np + n\beta)$).

In the rest of Section 7.6.2.2 we use \mathbf{v}_i for each $i \in [n]$ to denote the type-2 vector with $v_{i,i} = \ell_i$, $v_{i,j} = a_j$ for other $j \in [n]$, $v_{i,n+1} = s$ and $v_{i,n+2} = 0$. To prove (7.60), we start with the following lemma:

Lemma 7.6.18. *Let \mathbf{q} be any nonnegative n -dimensional vector that sums to 1, and $u'_2 = \min_{i \in [n]} d_i \cdot q_i$. If $u_2, u_3, u_4 \geq 0$ satisfy $u_3 \leq u_4 \leq u_3 + s$ and $u_2 \neq u'_2$, then $\text{REV}(\mathbf{q}, u_2, u_3, u_4) < \text{REV}(\mathbf{q}, u'_2, u_3, u_4)$.*

Proof. Let $(u(\cdot), \mathbf{q}(\cdot))$ be a feasible solution in $\text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ and $(u'(\cdot), \mathbf{q}'(\cdot))$ be a feasible solution in $\text{Ext}(\mathbf{q}, u'_2, u_3, u_4)$. Below we compare their revenues REV_2 and REV'_2 from type-2 vectors since it is clear that $\text{REV}_i = \text{REV}'_i$ for all $i \in \{1, 3, 4\}$. We consider two cases: $u_2 < u'_2$ or $u_2 > u'_2$.

Case 1: $u_2 < u'_2$. Let $\epsilon = u'_2 - u_2 > 0$. We compare revenues from type-2 vectors one by one. For \mathbf{c}_2 , $\text{REV}'(\mathbf{c}_2) = \text{REV}(\mathbf{c}_2) - \epsilon$. For $\mathbf{v} \in T_2$ other than \mathbf{c}_2 , $\mathbf{q}(\mathbf{v})$ and $\mathbf{q}'(\mathbf{v})$ are optimal solutions to $\text{LP}(\mathbf{v} : u)$ and $\text{LP}(\mathbf{v} : u')$, respectively. Given $|u(\mathbf{w}) - u(\mathbf{w}')| \leq \epsilon$ for all $\mathbf{w} \in T_2$, we have by Lemma 7.6.4

$$|\text{REV}'(\mathbf{v}) - \text{REV}(\mathbf{v})| = O(n\epsilon s/\beta).$$

We compare $\text{REV}'(\mathbf{v}_i)$ and $\text{REV}(\mathbf{v}_i)$ more carefully. Since $u_2 < u'_2$, by CONDITION-TYPE-2 and the definition of u'_2 , we have $u(\mathbf{v}_i) = u'(\mathbf{v}_i)$. Constraints of $\text{LP}(\mathbf{v}_i : u)$ are

$$q_i \geq 0, \quad \sum_{j \in [n+2]} q_j \leq 1, \quad \text{and} \quad d_i \cdot q_i \geq u(\mathbf{v}_i) - u_2$$

and constraints in $\text{LP}(\mathbf{v}_i : u')$ are

$$q_i \geq 0, \quad \sum_{j \in [n+2]} q_j \leq 1, \quad \text{and} \quad d_i \cdot q_i \geq u'(\mathbf{v}_i) - u'_2.$$

It follows that $\mathbf{q}(\mathbf{v})$ has $q_i(\mathbf{v}) = (u(\mathbf{v}_i) - u_2)/d_i$ and puts the rest of probability on $q_{n+1}(\mathbf{v})$, while $\mathbf{q}'(\mathbf{v})$ has $q'_i(\mathbf{v}) = (u'(\mathbf{v}_i) - u'_2)/d_i = q_i(\mathbf{v}) - (\epsilon/d_i)$ and puts the rest on $q'_{n+1}(\mathbf{v})$. As a result we have

$$\text{REV}'(\mathbf{v}_i) = \text{REV}(\mathbf{v}_i) + \frac{\epsilon}{d_i} \cdot (s - \ell_i) - \epsilon, \quad \text{for each } i \in [n].$$

To summarize, we have

$$\text{REV}'_2 - \text{REV}_2 \geq \sum_{i \in [n]} \Pr[\mathbf{v}_i] \cdot \left(\frac{\epsilon}{d_i} \cdot (s - \ell_i) - \epsilon \right) - \Pr[\mathbf{c}_2] \cdot \epsilon - O(nr\delta) \cdot O\left(\frac{n\epsilon s}{\beta}\right).$$

Plugging in that $\Pr[\mathbf{c}_2] \leq \delta$, $1/d_i \geq 1 - O(n\beta)$, $s - \ell_i \geq s - 2 - O(n\beta)$, and

$$\Pr[\mathbf{v}_i] = p \cdot (1 - p - r)^{n-1} \cdot \delta \cdot (1 - \delta^2) \geq p\delta \cdot (1 - O(np)),$$

we have

$$\begin{aligned} \sum_{i \in [n]} \Pr[\mathbf{v}_i] \cdot \frac{\epsilon}{d_i} \cdot (s - \ell_i) &\geq n \cdot p\delta \cdot (1 - O(np)) \cdot \epsilon \cdot (1 - O(n\beta)) \cdot \left(\frac{1}{(n - 0.5)p} - O(n\beta) \right) \\ &\geq \frac{n\delta\epsilon}{n - 0.5} \cdot (1 - O(n\beta)). \end{aligned}$$

As a result, we have $\text{REV}'_2 - \text{REV}_2 > 0$ given our choices of p , r and β .

Case 2: $u_2 > u'_2$. Let $\epsilon = u_2 - u'_2 > 0$. In this case, we have $\text{REV}'(\mathbf{c}_2) = \text{REV}(\mathbf{c}_2) + \epsilon$ and similarly, $|\text{REV}'(\mathbf{v}) - \text{REV}(\mathbf{v})| = O(n\epsilon s/\beta)$ for all other type-2 vectors $\mathbf{v} \in T_2$. For each $i \in [n]$, a similar analysis of $\text{LP}(\mathbf{v}_i : u)$ and $\text{LP}(\mathbf{v}_i : u')$ as in Case 1 implies that

$$q_i(\mathbf{v}) = \frac{u(\mathbf{v}_i) - u_2}{d_i} \quad \text{and} \quad q'_i(\mathbf{v}) = \frac{u'(\mathbf{v}_i) - u'_2}{d_i},$$

and both vectors have the rest of probability allocated on their $(n + 1)$ th entries.

Let I denote the nonempty set of i that has the minimum $d_i \cdot q_i$ among all indices in $[n]$. It then follows from the definition of u'_2 and the assumption of $u_2 > u'_2$ that $u(\mathbf{v}_i) = u_2$ and $u(\mathbf{v}'_i) = u'_2$ for each $i \in I$ and thus, $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v}) + \epsilon$ for each $i \in I$. For each $i \notin I$, we have

$$u'(\mathbf{v}_i) - u'_2 \leq u(\mathbf{v}_i) - u_2 + \epsilon.$$

This follows by considering both cases of $u(\rho(\mathbf{v}_i)) \leq u_2$ or $u(\rho(\mathbf{v}_i)) > u_2$. As a result, for each $i \notin I$,

$$\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - \frac{\epsilon}{d_i} \cdot (s - \ell_i) + \epsilon.$$

Combining them together, we have

$$\text{REV}'_2 - \text{REV}_2 \geq \Pr[\mathbf{c}_2] \cdot \epsilon + \sum_{i \in I} \Pr[\mathbf{v}_i] \cdot \epsilon - \sum_{i \notin I} \Pr[\mathbf{v}_i] \cdot \left(\frac{\epsilon}{d_i} \cdot (s - \ell_i) - \epsilon \right) - O(nr\delta) \cdot O\left(\frac{n\epsilon s}{\beta}\right).$$

Plugging in $\Pr[\mathbf{c}_2] \geq \delta(1 - O(np))$ and

$$\sum_{i \notin I} \Pr[\mathbf{v}_i] \cdot \frac{\epsilon}{d_i} \cdot (s - \ell_i) \leq (n-1) \cdot p\delta \cdot \epsilon \cdot (1 + O(n\beta)) \cdot \frac{1}{(n-0.5)p} = \delta\epsilon \cdot \frac{n-1}{n-0.5} \cdot (1 + O(n\beta)),$$

we have $\text{REV}'_2 - \text{REV}_2 > 0$. This finishes the proof of the lemma. \square

We are now ready to prove the main lemma of this section.

Lemma 7.6.19. *Any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $\text{LP}'(I)$ satisfies $d_1 \cdot q_1(\mathbf{a}) = \dots = d_n \cdot q_n(\mathbf{a})$.*

Proof. Let $(u(\cdot), \mathbf{q}(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ be an optimal solution to $\text{LP}'(I)$, where \mathbf{q} is an n -dimensional nonnegative vector that sums to 1 and $u_2, u_3, u_4 \geq 0$ with $u_3 \leq u_4 \leq u_3 + s$. By Lemma 7.6.18, we have

$$u_2 = \min_{i \in [n]} d_i \cdot q_i.$$

Assume for contradiction that \mathbf{q} does not satisfy $d_1 \cdot q_1 = \dots = d_n \cdot q_n$. We use $K \subset [n]$ to denote the set of indices k with $d_k \cdot q_k = \min_i d_i \cdot q_i$, and $t \in [n]$ denote an index with $d_t \cdot q_t > \min_i d_i \cdot q_i$. Then we replace \mathbf{q} by \mathbf{q}' , where $q'_k = q_k + (\epsilon/d_k)$ for each $k \in K$ and $q'_t = q_t - \sum_{k \in K} (\epsilon/d_k)$, for a sufficiently small $\epsilon > 0$ such that \mathbf{q}' remains nonnegative and indices $k \in K$ still have the smallest $d_k \cdot q'_k = d_k \cdot q_k + \epsilon$ in \mathbf{q}' . We also replace u_2 by $u'_2 = u_2 + \epsilon$. Let $(u'(\cdot), \mathbf{q}'(\cdot)) \in \text{Ext}(\mathbf{q}', u'_2, u_3, u_4)$ be a feasible solution. Then we reach a contradiction by showing that the revenue of $(u'(\cdot), \mathbf{q}'(\cdot))$ is strictly higher than that of $(u(\cdot), \mathbf{q}(\cdot))$.

First it is clear that $\text{REV}'_1 = \text{REV}_1$ since both \mathbf{q} and \mathbf{q}' sum to 1. By Lemma 7.6.10, we have

$$|(\text{REV}_3 + \text{REV}_4) - (\text{REV}'_3 + \text{REV}'_4)| \leq O\left(\frac{\delta^2 n^2 \epsilon t}{\beta}\right) + O\left(\frac{\delta^3 n^2 \epsilon t}{\beta}\right) = O\left(\frac{\delta^2 n^2 \epsilon t}{\beta}\right),$$

where we used the loose bound of $|u(\mathbf{w}) - u(\mathbf{w}')| \leq O(n\epsilon)$ for all $\mathbf{w} \in T_1$. The RHS above is negligible as we will see due to δ^2 . Recall that $\delta = 1/2^{n^6}$ and $t = 2^{\Theta(n^5)}$. It remains to compare REV_2 and REV'_2 . For all $\mathbf{v} \in T_2$ other than \mathbf{c}_2 and \mathbf{v}_i , $i \in [n]$, by Lemma 7.6.4: $|\text{REV}'(\mathbf{v}) - \text{REV}(\mathbf{v})| = O(n \cdot n\epsilon \cdot s/\beta) = O(n^2\epsilon s/\beta)$. On the other hand, we have $\text{REV}'(\mathbf{c}_2) = \text{REV}(\mathbf{c}_2) - \epsilon$. For each $i \in [n]$, it follows from $\text{LP}(\mathbf{v}_i : u)$ that

$$\text{REV}(\mathbf{v}_i) = \ell_i \cdot q_i(\mathbf{v}) + s \cdot (1 - q_i(\mathbf{v})) - d_i \cdot q_i(\mathbf{a}) = s - (s - \ell_i) \cdot \frac{d_i \cdot q_i - u_2}{d_i} - d_i \cdot q_i.$$

A similar expression holds for $\text{REV}'(\mathbf{v}_i)$ (replacing q_i by q'_i and u_2 by u'_2). As a result,

$$\begin{aligned} & \sum_{i \in [n]} \Pr[\mathbf{v}_i] \cdot (\text{REV}'(\mathbf{v}_i) - \text{REV}(\mathbf{v}_i)) \\ &= \delta \cdot (1 - \delta^2) \cdot p \cdot (1 - p - r)^{n-1} \cdot \left(\sum_{i \in [n]} \frac{s - \ell_i}{d_i} \cdot \epsilon + (s - a_t) \sum_{k \in K} \frac{\epsilon}{d_k} - \sum_{k \in K} (s - a_k) \cdot \frac{\epsilon}{d_k} \right) \\ &\geq \delta \cdot (1 - \delta^2) \cdot p \cdot (1 - p - r)^{n-1} \cdot \left(\frac{n}{(n - 0.5)p} \cdot (1 - O(n\beta)) \cdot \epsilon - O(np \cdot n\epsilon) \right) \\ &= \frac{n\delta\epsilon}{n - 0.5} \cdot (1 - O(n\beta)) - O(n^2 p^2 \delta\epsilon), \end{aligned}$$

where we used $|a_i - 1| = O(np)$. Combining all these bounds together, we have

$$\text{REV}'_2 - \text{REV}_2 \geq \frac{n\delta\epsilon}{n - 0.5} \cdot (1 - O(n\beta)) - O(n^2 p^2 \delta\epsilon) - \epsilon \cdot \delta - O(nr\delta) \cdot O(n^2\epsilon s/\beta) \gg O\left(\frac{\delta^2 n^2 \epsilon t}{\beta}\right),$$

given our choices of parameters. This contradicts with the optimality of $(u(\cdot), \mathbf{q}(\cdot))$. \square

Given that $\mathbf{q}(\mathbf{a})$ is close to a uniform distribution, we record a lemma that will be useful later.

Lemma 7.6.20. *Let $\mathbf{v}, \mathbf{v}' \in D$ denote two valuation vectors that differ at the i th entry only, for some $i \in [n]$, and $v'_i > v_i$. Then we have $u(\mathbf{v}') \geq u(\mathbf{v})$ in any optimal solution to $\text{LP}'(I)$.*

Proof. It suffices to prove the lemma for two type-1 vectors $\mathbf{v}, \mathbf{v}' \in T_1$ (due to CONDITION-TYPE-2, 3, and 4). The case when $v_i = \ell_i$ and $v'_i = h_i$ follows directly from CONDITION-TYPE-1. The case when $v_i = a_i$ and $v_i = \ell_i$ follows from Lemma 7.6.19, that $\mathbf{q}(\mathbf{a})$ is close to a uniform distribution. In particular, we have

$$u(\text{LOWER}(\mathbf{v}')) = u(\text{LOWER}(\mathbf{v})) + d_i \cdot q_i(\mathbf{a}) \approx u(\text{LOWER}(\mathbf{v})) + (1/n),$$

while both $u(\mathbf{v}') - u(\text{LOWER}(\mathbf{v}'))$ and $u(\mathbf{v}) - u(\text{LOWER}(\mathbf{v}))$ are much smaller than $1/n$ since $\tau_i = O(n\beta)$ for all i and $\beta = 1/2^n$. This finishes the proof of the lemma. \square

7.6.2.3 Setting h_i 's and t

Before giving our choices of h_i 's and t , we introduce the problem COMP, and show that it is #P-hard. Here an input (\mathcal{G}, H, M) of COMP consists of a tuple $\mathcal{G} = (g_2, \dots, g_n)$ of $n - 1$ integers between 1 and $N = 2^n$, a subset $H \subset [2 : n]$ of size $|H| = m = \lceil n/2 \rceil$, and an integer M between 1 and $\binom{n-1}{m}$. For convenience, we write $\text{Sum}(T) = \sum_{i \in T} g_i$ for $T \subseteq [2 : n]$. We use t^* to denote the M -th largest integer in the multiset

$$\left\{ \text{Sum}(T) : T \subset [2 : n] \text{ and } |T| = m \right\}. \quad (7.61)$$

The problem is then to decide whether $\text{Sum}(H) > t^*$ or $\text{Sum}(H) \leq t^*$, i.e. compare $\text{Sum}(H)$ to the M -th largest integer in (7.61). We first show that COMP is #P-hard.

Lemma 7.6.21. *COMP is #P-hard.*

Proof. We reduce from a related problem called LEX-RANK, which was shown to be #P-hard in [Daskalakis *et al.*, 2014a]. In LEX-RANK, the input consists of a collection $C = \{c_1, \dots, c_n\}$ of positive integers, a subset $S \subseteq [n]$, and a positive integer k . Order the subsets T of $[n]$ of cardinality $|S|$ according to their sums, $\text{Sum}_C(T) = \sum_{i \in T} c_i$, from smallest to largest, with subsets that have equal sums ordered lexicographically; that is, we have $T <_C T'$ if and only if $\text{Sum}_C(T) < \text{Sum}_C(T')$, or $\text{Sum}_C(T) = \text{Sum}_C(T')$ and the largest element in the symmetric difference $T \Delta T'$ belongs to T' . The LEX-RANK problem is to determine for a given input (C, S, k) whether the rank of S in this ordering (among subsets of cardinality $|S|$) is at most k .

Let (C, S, k) be an instance of LEX-RANK. Let $c'_i = 2^{2^n} \cdot c_i + 2^i$, for all i , and let $C' = \{c'_1, \dots, c'_n\}$. Clearly, any two subsets of C' have unequal sums and furthermore, $T <_C T'$ iff $\text{Sum}_{C'}(T) < \text{Sum}_{C'}(T')$, for all $T, T' \subseteq [n]$. In the new instance (C', S, k) , the rank of a set S (among sets of the same cardinality) is the same as its rank in the old (C, S, k) , and the rank of S is at most k iff $\text{Sum}_{C'}(S)$ is at most the M -th largest sum, where $M = \binom{n}{|S|} - k + 1$. Thus, the LEX-RANK problem in the new instance is equivalent

to the COMP problem, except that in the latter problem we also require that $|S| = \lceil n/2 \rceil$ and that all input integers are at most 2^n .

Let B be the maximum number of bits of the integers in C' ; note, $B \geq 2n$. Add $2B - n - 1$ new elements to the set C' to form the new set \mathcal{G} ; $B - |S|$ of the new elements have value $n2^{B+1}$, and the rest have value 1. Let H be the set that consists of S and the new elements with value $n2^{B+1}$. Thus, \mathcal{G} has $2B - 1 = n' - 1$ elements, S has size $B = n'/2$ and all the integers are between 1 and $2^{2B} = 2^{n'}$. Let $M = \binom{n}{|S|} - k + 1$, as above. The instance (\mathcal{G}, H, M) of COMP now satisfies the required constraints. If we order the subsets of cardinality $B = |H|$ from largest sum to smallest, the first $\binom{n}{|S|}$ subsets will each consist of the $B - |S|$ new elements with the large value of $n2^{B+1}$ and then a subset of cardinality $|S|$ of the original elements, ordered according to their sum. Therefore, $\text{Sum}_{\mathcal{G}}(H)$ is at most the M -th largest sum in the instance (\mathcal{G}, H, M) of COMP iff $\text{Sum}_{C'}(S)$ is at most the M -th largest sum in (C', S, k) , i.e., iff the rank of S is at most k in the original instance (C, S, k) of LEX-RANK. \square

We embed COMP in I . Let (\mathcal{G}, H, M) be an instance of COMP where $\mathcal{G} = (g_2, \dots, g_n)$ is a sequence of $n - 1$ integers between 1 and $N = 2^n$, $H \subset [2 : n]$ with $|H| = m = \lceil n/2 \rceil$, and M is an integer between 1 and $\binom{n-1}{m}$. Here are our choices of τ_i 's and then $h_i = \ell_i + \tau_i$. Recall that we promised in (7.25), (7.26) that

$$\beta \leq \tau_i \leq \left(1 + \frac{1}{N^2}\right) \beta, \quad t = O\left(\frac{\beta}{r^{m+1}m}\right) \quad \text{and} \quad t = \Omega\left(\frac{\beta}{r^{m+1}m2^n}\right). \quad (7.62)$$

By our choices and a_i 's and ℓ_i 's, $d_1 = \max_{j \in [n]} d_j$. Set $\tau_i = \tau'_i + \beta$ for each $i \in [n]$ with $\tau'_1 = \beta/N^2$ and

$$\tau'_i = \frac{\beta}{N^2} \cdot \frac{d_1 - d_i}{d_1} + g_i \cdot \frac{d_i \beta}{N^4} = O\left(\frac{n\beta^2}{N^2}\right), \quad \text{for each } i > 1.$$

Recall g_i is from \mathcal{G} . As $\beta = 1/N$, (7.62) on τ_i is satisfied. The choice of t needs to be done more carefully.

Let R denote the set of $\mathbf{v} \in T_3$ satisfying $|S(\mathbf{v})| = m + 1$ and $|S^+(\mathbf{v})| = m$, and let R' denote the set of $\mathbf{v} \in T_3$ with $|S(\mathbf{v})| = |S^+(\mathbf{v})| = m + 1$. Let R^* denote the set of $\mathbf{v} \in R'$ with $1 \in S^+(\mathbf{v})$. Let h denote the probability $\Pr[\mathbf{v}]$ of each vector $\mathbf{v} \in R'$ (note that they

all share the same probability $\Pr[\mathbf{v}]$:

$$h = (1 - \delta) \cdot \delta^2 \cdot r^{m+1} \cdot (1 - p - r)^{n-m} \approx \delta^2 r^{m+1}.$$

We are now ready to set t using M from the instance of COMP as follows:

$$t = 2 + \frac{\beta \delta^2}{h(m+1)(M - (1/2))},$$

which clearly satisfies the promise on t in (7.62).

Fix a type-3 vector $\mathbf{v} \in R^*$, and let $\mathbf{w} = \rho(\mathbf{v}) \in T_1$. Let $(u(\cdot), \mathbf{q}(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ be a feasible solution to $\text{LP}'(I)$ for some nonnegative \mathbf{q} that sums to 1 and $u_2, u_3, u_4 \geq 0$ that satisfy

$$u_2 = d_1 \cdot q_1 = \dots = d_n \cdot q_n = \Theta(1/n) \quad \text{and} \quad u_3 \leq u_4 \leq u_3 + s.$$

To see the connection between the two problems, we calculate $u(\mathbf{w})$. As $\min(S(\mathbf{w})) = \min(S(\mathbf{v})) = 1$,

$$\begin{aligned} u(\mathbf{w}) &= \sum_{i \in S(\mathbf{v})} d_i \cdot q_i + \tau_1 \cdot \left(1 - \sum_{i \in S'(\mathbf{v})} q_i\right) + \sum_{i \in S'(\mathbf{v})} \tau_i \cdot q_i \\ &= (m+1) \cdot u_2 + \tau_1 - \sum_{i \in S'(\mathbf{v})} (\tau'_1 - \tau'_i) \cdot \frac{u_2}{d_i} \\ &= (m+1) \cdot u_2 + \tau_1 - \sum_{i \in S'(\mathbf{v})} \left(\frac{\beta d_i}{N^2 d_1} - g_i \cdot \frac{d_i \beta}{N^4} \right) \cdot \frac{u_2}{d_i} \\ &= C + \frac{\beta u_2}{N^4} \sum_{i \in S'(\mathbf{v})} g_i, \end{aligned}$$

where we write the constant C (independent of the choice of $\mathbf{v} \in R^*$) as

$$C = (m+1) \cdot u_2 + \tau_1 - \frac{m \beta u_2}{N^2 d_1}.$$

This suggests a natural one-to-one correspondence: $T \mapsto \mathbf{v} \in R^*$ with $S(\mathbf{v}) = \{1\} \cup T$, between

$$\left\{ T : T \subset [2 : n] \text{ and } |T| = m \right\}$$

and R^* with respect to which the order over $\text{Sum}(T)$ is the same as that over $u(\rho(\mathbf{v}))$. Moreover, since τ'_1 is much larger than τ'_i with $i > 1$, other $\mathbf{v} \in R'$ have strictly smaller utility $u(\rho(\mathbf{v}))$ than those in R^* .

To see this, note that for each $\mathbf{v} \in R^*$, we have

$$u(\rho(\mathbf{v})) \geq (m+1) \cdot u_2 + \beta + \tau'_1 - \sum_{i \in S'(\mathbf{v})} \tau'_1 \cdot q_i = (m+1) \cdot u_2 + \beta + \Omega(\tau'_1).$$

On the other hand, let $k = \min(S(\mathbf{v}')) > 1$ for some $\mathbf{v}' \in R' \setminus R^*$. We have

$$u(\rho(\mathbf{v}')) \leq (m+1) \cdot u_2 + \beta + \tau'_k + \sum_{i \in S'(\mathbf{v}')} \tau'_i \cdot q_i = (m+1) \cdot u_2 + \beta + O\left(\max_{i \geq 2} \tau'_i\right).$$

It is also easier to verify that $u(\rho(\mathbf{v}))$ with $\mathbf{v} \in R^*$ are strictly higher than $u(\rho(\mathbf{v}'))$ of $\mathbf{v}' \in R$.

We write u^* to denote the M -th largest element of the multiset $\{u(\rho(\mathbf{v})) : \mathbf{v} \in R^*\}$. Then the next two lemmas together show that $u(\mathbf{c}_3) = u_3$ must be exactly u^* in any optimal solution to $\text{LP}'(I)$.

Lemma 7.6.22. *Any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $\text{LP}'(I)$ must satisfy $u(\mathbf{c}_3) \leq u^*$.*

Proof. This direction is easy. Assume for contradiction that $(u(\cdot), \mathbf{q}(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ is optimal but $u_3 > u^*$. Let $\epsilon > 0$ be sufficiently small such that $u(\mathbf{v}) < u_3$ implies that $u(\mathbf{v}) < u_3 - \epsilon$, for all $\mathbf{v} \in D$.

We show that $(u'(\cdot), \mathbf{q}'(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u'_3, u'_4)$, where $u'_3 = u_3 - \epsilon$ and $u'_4 = u_4 - \epsilon$ (note that we still have $u'_3 \leq u'_4 \leq u'_3 + s$), results in strictly higher expected revenue from type-3 and type-4 vectors, which contradicts with the optimality of $(u(\cdot), \mathbf{q}(\cdot))$. By Lemma 7.6.10, we have $|\text{REV}'_4 - \text{REV}_4| = O(\delta^3 n \epsilon t / \beta)$.

We now bound $\text{REV}'_3 - \text{REV}_3$. To this end, let A denote the set of $\mathbf{v} \in T_3$ with $u(\rho(\mathbf{v})) \geq u_3$ and let B denote the rest of type-3 vectors with $u(\rho(\mathbf{v})) < u_3$ (so $\mathbf{c}_3 \in B$). For each $\mathbf{v} \in B$, we have $u(\mathbf{v}) = u_3$ and $u'(\mathbf{v}) = u'_3$ (by our choice of ϵ). By $\text{LP}(\mathbf{v} : u)$ and $\text{LP}(\mathbf{v} : u')$, we have both $\mathbf{q}(\mathbf{v})$ and $\mathbf{q}'(\mathbf{v})$ put probability 1 on item $n+2$. As a result, we have $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v}) + \epsilon$ for each $\mathbf{v} \in B$. On the other hand, for each $\mathbf{v} \in A$, by Lemma 7.6.4 we have $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - O(n \epsilon t / \beta)$.

We need to take a closer look at vectors $\mathbf{v} \in R^* \cap A$ (which can be empty but by $u_3 > u^*$, $|R^* \cap A|$ is at most $M-1$). To understand $\mathbf{q}(\mathbf{v})$ and $\mathbf{q}'(\mathbf{v})$, we note that all $u(\mathbf{w})$ in $\text{LP}(\mathbf{v} : u)$ are u_3 and all $u'(\mathbf{w})$ in $\text{LP}(\mathbf{v} : u')$ are u'_3 . As a result, we only need to consider

the following constraints in $\text{LP}(\mathbf{v} : u)$

$$q_i \geq 0, \quad \sum_{j \in [n+2]} q_j \leq 1, \quad \text{and} \quad \tau_i \cdot q_i \geq u(\mathbf{v}) - u_3, \quad \text{for } i \in S^+(\mathbf{v}) = S(\mathbf{v}).$$

since all other constraints would be implied. As a result, $q_i(\mathbf{v}) = (u(\mathbf{v}) - u_3)/\tau_i$ for each $i \in S^+(\mathbf{v})$, and $\mathbf{q}(\mathbf{v})$ puts the rest of probability on $q_{n+2}(\mathbf{v})$. Similarly, $q'_i(\mathbf{v}) = (u(\mathbf{v}) - u'_3)/\tau_i$ for each $i \in S^+(\mathbf{v})$ and $\mathbf{q}'(\mathbf{v})$ puts the rest of probability on $q'_{n+2}(\mathbf{v})$. This implies that

$$\text{REV}'(\mathbf{v}) > \text{REV}(\mathbf{v}) - (m+1) \cdot \frac{\epsilon}{\beta} \cdot t, \quad \text{for each } \mathbf{v} \in R^* \cap A.$$

Combining all these inequalities, we have

$$\text{REV}'_3 - \text{REV}_3 \geq \Pr[B] \cdot \epsilon - \Pr[A \setminus R^*] \cdot O(n\epsilon t/\beta) - (M-1) \cdot h \cdot (m+1) \cdot \frac{\epsilon}{\beta} \cdot t.$$

Plugging in $\Pr[B] \geq \Pr[\mathbf{c}_3] \geq \delta^2(1 - O(np))$ (since $\mathbf{c}_3 \in B$) and $\Pr[A \setminus R^*] < 3^n \cdot \delta^2 \cdot p^{m+2}$ since $A \setminus R^*$ only has vectors $\mathbf{v} \in T_3$ with $|S(\mathbf{v})| \geq m+2$, we have

$$\begin{aligned} & \text{REV}'_3 - \text{REV}_3 \\ & \geq \delta^2 \epsilon (1 - O(np)) - O\left(\frac{3^n \delta^2 p^{m+2} n \epsilon t}{\beta}\right) + (M-1)h(m+1) \frac{\epsilon}{\beta} \cdot \frac{\beta \delta^2 (1 + o(r^m))}{h(m+1)(M - (1/2))} \\ & = \delta^2 \epsilon \cdot \left(\frac{1}{2M-1} - O(np) - o(r^m) - O\left(\frac{3^n p^{m+2} n}{r^{m+1} m}\right)\right) \gg O(\delta^3 n \epsilon t/\beta), \end{aligned}$$

where the inequalities follow from choices of p, r and δ in (7.24). This finishes the proof. \square

Lemma 7.6.23. *Any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $\text{LP}'(I)$ must satisfy $u(\mathbf{c}_3) \geq u^*$.*

Proof. This direction is more difficult. Assume that $(u(\cdot), \mathbf{q}(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u_3, u_4)$ is an optimal solution but $u_3 < u^*$. Let $\epsilon > 0$ be a sufficiently small positive number such that $u(\mathbf{v}) > u_3$ implies $u(\mathbf{v}) > u_3 + \epsilon$ for all $\mathbf{v} \in D$. Our plan is to show that $(u'(\cdot), \mathbf{q}'(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u'_3, u'_4)$, where $u'_3 = u_3 + \epsilon$ and $u'_4 = u_4 + \epsilon$, results in strictly higher expected revenue, a contradiction.

By Lemma 7.6.10, we have $|\text{REV}'_4 - \text{REV}_4| = O(\delta^3 n \epsilon t/\beta)$. Next we compare REV'_3 and REV_3 . For this purpose we define A as the set of $\mathbf{v} \in T_3$ with $u(\rho(\mathbf{v})) > u_3$ and B as the rest of $\mathbf{v} \in T_3$ with $u(\rho(\mathbf{v})) \leq u_3$ (so $\mathbf{c}_3 \in B$). By an argument similar to the previous lemma, we have $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v}) - \epsilon$ for all $\mathbf{v} \in B$. For $\mathbf{v} \in A$, we have $u(\mathbf{v}) = u'(\mathbf{v})$ by

our choice of ϵ . Since $u'_3 = u_3 + \epsilon$, we have $u'(\mathbf{w}) \geq u(\mathbf{w})$ for each \mathbf{w} in $\text{LP}(\mathbf{v} : u)$ and thus, constraints in $\text{LP}(\mathbf{v} : u)$ are at least as strong as those in $\text{LP}(\mathbf{v} : u')$. As a result, we have $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - \epsilon$ for every $\mathbf{v} \in A$.

Let C denote the subset of $\mathbf{v} \in A$ that satisfies 1) \mathbf{v} is below a vector in R^* and 2) every type-3 vector below \mathbf{v} has $u(\mathbf{w}) = u_3$ (so $\mathbf{w} \in B$). Note that C can be empty. Fix a $\mathbf{v} \in C$ when it is nonempty. We have $u'(\mathbf{v}) = u(\mathbf{v}) > u'_3 = u_3 + \epsilon$ and $u'(\mathbf{w}) = u'_3 = u(\mathbf{w}) + \epsilon$, for all type-3 vectors \mathbf{w} below \mathbf{v} . As a result, every constraint (other than those on \mathbf{q} only) in $\text{LP}(\mathbf{v} : u)$ has its RHS larger than the corresponding RHS of $\text{LP}(\mathbf{v} : u')$ by ϵ . We claim that $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + \Omega(t\epsilon)$. To see this, let \mathbf{q}^* denote the vector derived from $\mathbf{q}(\mathbf{v})$ as follows: $q_i^* = q_i(\mathbf{v}) - (\epsilon/2)$ for some $i \in S(\mathbf{v})$ and $q_{n+2}^* = q_{n+2}(\mathbf{v}) + (\epsilon/2)$; all other entries remain the same. It is clear that \mathbf{q}^* is nonnegative (since $d_i \cdot q_i(\mathbf{v}) \geq u(\mathbf{v}) - u_3 > \epsilon$) and is also a feasible solution to $\text{LP}(\mathbf{v} : u')$. It follows that $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + \Omega(t\epsilon)$.

To finish the proof, we consider the following two cases: **Case 1:** $C \neq \emptyset$. Then (taking the worst case that $|C| = 1$ and the vector is in R) we have

$$\text{REV}'_3 - \text{REV}_3 \geq \delta^2 \cdot r^m \cdot p \cdot \Omega(t\epsilon) - \delta^2 \cdot \epsilon > \delta^2 \epsilon \cdot \left(\frac{r^m p \beta \delta^2}{h(m+1)(M-0.5)} - 1 \right) \gg O(\delta^3 \text{net}/\beta),$$

where the second to the last inequality follows from $p/r = 2^{n^2} \gg mM/\beta$.

Case 2: $C = \emptyset$. Then every $\mathbf{v} \in R^* \cap A$ satisfies that all vectors \mathbf{w} below \mathbf{v} have $u(\mathbf{w}) = u_3$, and for each $\mathbf{v} \in R^* \cap A$, $\text{LP}(\mathbf{v} : u)$ boils down to the following constraints:

$$q_i \geq 0, \quad \sum_{i \in [n+2]} q_i \leq 1, \quad \text{and} \quad \tau_i \cdot q_i \geq u(\mathbf{v}) - u_3, \quad \text{for all } i \in S^+(\mathbf{v}) = S(\mathbf{v}),$$

since all other constraints would be trivially implied. As a result, we have $q_i(\mathbf{v}) = (u(\mathbf{v}) - u_3)/\tau_i$ for each $i \in S(\mathbf{v})$ and $q_{n+2}(\mathbf{v})$ takes the rest of probability. Similarly, we have $q'_i(\mathbf{v}) = (u(\mathbf{v}) - u'_3)/\tau_i$ for each $i \in S(\mathbf{v})$ and $q'_{n+2}(\mathbf{v})$ takes the rest of probability. Plugging in $u'_3 = u_3 + \epsilon$, we have

$$\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) + (m+1) \cdot \frac{\epsilon}{\max_{i \in [n]} \tau_i} \cdot \left(t - \max_{i \in [n]} h_i \right) - \epsilon.$$

Given that $u_3 < u^*$, we have $|A \cap R^*| \geq M$. Combining all bounds together, we have

$$\begin{aligned} \text{REV}'_3 - \text{REV}_3 &\geq Mh \cdot \frac{(m+1)\epsilon}{\beta(1 + O(1/N^2))} \cdot (t - (2 + 3n\beta)) - \delta^2 \epsilon \\ &= \delta^2 \epsilon \left(\frac{M(1 - O(1/N^2))}{M - 0.5} - 1 \right) > 0 \end{aligned}$$

and is $\gg O(\delta^3 net/\beta)$. This finishes the proof of the lemma. \square

Before we pin down $u(\mathbf{c}_4)$, recall that the second part of the input (\mathcal{G}, H, M) is a set $H \subset [2 : n]$ of size m . Let \mathbf{v}_H denote the vector in R^* with $S^+(\mathbf{v}_H) = S(\mathbf{v}_H) = \{1\} \cup H$. Given that $u_3 = u^*$, we have 1) if $\text{Sum}(H) > t^*$ then $u(\mathbf{v}_H) > u(\mathbf{c}_3)$; and 2) if $\text{Sum}(H) \leq t^*$ then $u(\mathbf{v}_H) = u(\mathbf{c}_3)$, in any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $\text{LP}'(I)$. It also follows from $\text{LP}(\mathbf{v} : u)$ that 1) if $\text{Sum}(H) > t^*$ then $q_{n+2}(\mathbf{v}_H) < 1$, and 2) if $\text{Sum}(H) \leq t^*$ then $q_{n+2}(\mathbf{v}_H) = 1$, in any optimal solution. We summarize it below.

Corollary 7.6.1. *If $\text{Sum}(H) > t^*$, then $q_{n+2}(\mathbf{v}_H) < 1$ in every optimal solution to $\text{LP}'(I)$.*

If $\text{Sum}(H) \leq t^$, then $q_{n+2}(\mathbf{v}_H) = 1$ in every optimal solution to $\text{LP}'(I)$.*

Finally we show that $u(\mathbf{c}_4) = u(\mathbf{c}_3) = u^*$ in any optimal solution to $\text{LP}'(I)$.

Lemma 7.6.24. *Any optimal solution to $\text{LP}'(I)$ must satisfy $u(\mathbf{c}_4) = u(\mathbf{c}_3) = u^*$.*

Proof. As for Lemma 7.6.22, suppose that $(u(\cdot), \mathbf{q}(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u^*, u_4)$ is optimal but $u_4 > u^*$ (and $u_4 \leq u^* + s$ for it to be feasible). Let $\epsilon > 0$ be a sufficiently small positive number, such that $u_4 - \epsilon \geq u^*$ and $u(\mathbf{v}) < u_4$ implies that $u(\mathbf{v}) < u_4 - \epsilon$, for every $\mathbf{v} \in D$. Our goal is then to show that $(u'(\cdot), \mathbf{q}'(\cdot)) \in \text{Ext}(\mathbf{q}, u_2, u^*, u'_4)$ is strictly better, where $u'_4 = u_4 - \epsilon$, a contradiction.

It suffices to compare REV'_4 and REV_4 since $\text{REV}'_i = \text{REV}_i$ for $i = 1, 2, 3$.

For \mathbf{c}_4 we have $\text{REV}'(\mathbf{c}_4) \geq \text{REV}(\mathbf{c}_4) + \Omega(\epsilon t/s)$. Let A be the set of $\mathbf{v} \in T_4 \setminus \{\mathbf{c}_4\}$ with $u(\rho(\mathbf{v})) \geq u_4$ and B be the rest of $\mathbf{v} \in T_4 \setminus \{\mathbf{c}_4\}$ with $u(\rho(\mathbf{v})) < u_4$. Following the same argument used in Lemma 7.6.22, we have $\text{REV}'(\mathbf{v}) = \text{REV}(\mathbf{v}) + \epsilon$ for each $\mathbf{v} \in B$, and $\text{REV}'(\mathbf{v}) \geq \text{REV}(\mathbf{v}) - O(n\epsilon t/\beta)$ for each $\mathbf{v} \in A$.

These bounds are strong enough for the current lemma. Given u^* we have $\Pr[A] \leq 3^n \cdot \delta^3 \cdot r^{m+1}$. So

$$\text{REV}'_4 - \text{REV}_4 \geq \Pr[\mathbf{c}_4] \cdot \Omega(\epsilon t/s) - \Pr[A] \cdot O(n\epsilon t/\beta) = \Omega(\delta^3 \epsilon t/s) - O(3^n \delta^3 r^{m+1} n\epsilon t/\beta) > 0.$$

This finishes the proof of the lemma. \square

7.6.3 Returning to the Standard Linear Program

Let (\mathcal{G}, H, M) be an input instance of COMP and I be the input instance of the optimal mechanism design problem (or the lottery problem) constructed from (\mathcal{G}, H, M) in Section 7.6.1 and 7.6.2.

We show that any optimal solution to $LP'(I)$ is a feasible solution to the standard $LP(I)$.

Lemma 7.6.25. *Any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $LP'(I)$ is a feasible solution to $LP(I)$.*

Before proving Lemma 7.6.25, we use it to prove Theorem 14.

Proof of Theorem 14 Assuming Lemma 7.6.25. By Lemma 7.6.25, we claim $(u(\cdot), \mathbf{q}(\cdot))$ is an optimal solution to $LP(I)$ if and only if it is an optimal solution to $LP'(I)$. To see this, let OPT and OPT' denote the optimal values of $LP(I)$ and $LP'(I)$, respectively. As $LP'(I)$ is a relaxation of $LP(I)$, we have $\text{OPT}' \geq \text{OPT}$. Lemma 7.6.25, on the other hand, implies $\text{OPT} \geq \text{OPT}'$. So $\text{OPT} = \text{OPT}'$, and from here the claim follows easily.

Suppose that $\mathcal{A}(\cdot, \cdot)$ satisfies both properties stated in Theorem 14. Then it follows from the connection between the optimal mechanism design problem and $LP(I)$ (see Section 7.2.1) that $(\bar{\mathcal{A}}(I, \mathbf{v}) : \mathbf{v} \in D)$ is an optimal solution to $LP(I)$ and thus, an optimal solution to $LP'(I)$.

It follows from Corollary 7.6.1 that 1) when $\text{Sum}(H) > t^*$, $\mathcal{A}(I, \mathbf{v}_H)$ assigns an item other than $n + 2$ or no item to the buyer with a positive probability; 2) when $\text{Sum}(H) \leq t^*$, $\mathcal{A}(I, \mathbf{v}_H)$ always assigns item $n + 2$ with probability 1. Given I and \mathbf{v}_H , the problem of deciding which case it is belongs to NP, because \mathcal{A} is a randomized algorithm that always terminates in polynomial time by assumption.

The theorem follows directly from the #P-hardness of COMP proved in Lemma 7.6.21. \square

We prove Lemma 7.6.25 in the rest of the section. It suffices to show that any optimal solution $(u(\cdot), \mathbf{q}(\cdot))$ to $LP'(I)$ satisfies (7.5) for all ordered pairs (\mathbf{v}, \mathbf{w}) in D .

7.6.3.1 Reducing to (\mathbf{v}, \mathbf{w}) with $S(\mathbf{w}) \subseteq S(\mathbf{v})$

First we handle the special case when $\mathbf{v} = \mathbf{a}$.

Let $(u(\cdot), \mathbf{q}(\cdot))$ be an optimal solution to $LP'(I)$ and $\mathbf{w} \in T_1$. Then by CONDITION-TYPE-1:

$$u(\mathbf{w}) \geq \sum_{i \in S(\mathbf{w})} (w_i - a_i) \cdot q_i(\mathbf{a}).$$

We extend ρ by setting $\rho(\mathbf{w}) = \mathbf{w}$ if $\mathbf{w} \in T_1$. Then $u(\mathbf{w}) \geq u(\rho(\mathbf{w}))$ for all $\mathbf{w} \in D$. Thus,

$$u(\mathbf{w}) \geq u(\rho(\mathbf{w})) \geq \sum_{i \in S(\rho(\mathbf{w}))} (w_i - a_i) \cdot q_i(\mathbf{a}) = \sum_{i \in S(\mathbf{w})} (w_i - a_i) \cdot q_i(\mathbf{a}).$$

As $u(\mathbf{a}) = 0$, this implies (7.5) on (\mathbf{a}, \mathbf{w}) for all $\mathbf{w} \in D$. We assume $\mathbf{v} \neq \mathbf{a}$ in (\mathbf{v}, \mathbf{w}) from now on.

Now we claim that it suffices to prove (7.5) for (\mathbf{v}, \mathbf{w}) that satisfies $S(\mathbf{w}) \subseteq S(\mathbf{v})$ (though \mathbf{v} and \mathbf{w} here may belong to different blocks). Suppose that we have proved (7.5) over (\mathbf{v}, \mathbf{w}) with $S(\mathbf{w}) \subseteq S(\mathbf{v})$. Given any general pair (\mathbf{v}, \mathbf{w}) with $\mathbf{v} \neq \mathbf{a}$ (otherwise it is done), we use \mathbf{w}' to denote the vector obtained from \mathbf{w} by replacing every w_i , $i \in S(\mathbf{w}) \setminus S(\mathbf{v})$, by a_i . Then clearly we have $S(\mathbf{w}') \subseteq S(\mathbf{v})$. Because (7.5) holds for $(\mathbf{v}, \mathbf{w}')$, then by monotonicity of $u(\cdot)$ (Lemma 7.6.20), we have

$$\begin{aligned} u(\mathbf{v}) - u(\mathbf{w}) &\leq u(\mathbf{v}) - u(\mathbf{w}') \leq \sum_{i \in [n]} (v_i - w'_i) \cdot q_i(\mathbf{v}) + \sum_{i \in \{n+1, n+2\}} (v_i - w'_i) \cdot q_i(\mathbf{v}) \\ &= \sum_{i \in [n]} (v_i - w_i) \cdot q_i(\mathbf{v}) + \sum_{i \in \{n+1, n+2\}} (v_i - w_i) \cdot q_i(\mathbf{v}). \end{aligned}$$

The last equation follows from two observations: $w_{n+1} = w'_{n+1}$ and $w_{n+2} = w'_{n+2}$; for every $i \in [n]$ but $i \notin S(\mathbf{v})$, $q_i(\mathbf{v}) = 0$ (due to CONDITION-TYPE-1, Lemma 7.6.1 and Lemma 7.6.3).

From now on we consider pairs (\mathbf{v}, \mathbf{w}) that satisfy $S(\mathbf{w}) \subseteq S(\mathbf{v})$.

7.6.3.2 Both \mathbf{v} and \mathbf{w} are Type-1

We start with the case when \mathbf{v} and \mathbf{w} are both type-1 vectors (and satisfy $S(\mathbf{w}) \subseteq S(\mathbf{v})$).

Note that (7.5) means that \mathbf{w} does not envy the lottery of \mathbf{v} . As \mathbf{v} buys the same lottery as $\text{LOWER}(\mathbf{v})$ (by CONDITION-TYPE-1), we may assume without loss of generality

that $\mathbf{v} \in T_1'$ and $S(\mathbf{w}) \subset S(\mathbf{v})$. Then

$$\begin{aligned} u(\mathbf{v}) - u(\mathbf{w}) &= \sum_{i \in S(\mathbf{v}) \setminus S(\mathbf{w})} d_i \cdot q_i(\mathbf{a}) - \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot q_i(\text{LOWER}(\mathbf{w})) \\ &\leq \sum_{i \in S(\mathbf{v}) \setminus S(\mathbf{w})} d_i \cdot q_i(\mathbf{v}) - \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot q_i(\mathbf{v}), \end{aligned}$$

where the last inequality follows from $q_i(\mathbf{v}) \geq q_i(\mathbf{a})$ for all $i \in S(\mathbf{v})$, and $q_i(\text{LOWER}(\mathbf{w})) \geq q_i(\mathbf{v})$ for all $i \in S(\mathbf{w})$ by Lemma 7.6.11.

7.6.3.3 Both \mathbf{v} and \mathbf{w} are Type-2

Next we prove (7.5) for pairs (\mathbf{v}, \mathbf{w}) of type-2 vectors that satisfy $S(\mathbf{w}) \subseteq S(\mathbf{v})$.

The special case of $|S(\mathbf{v})| \leq 1$ is easy to check. Let \mathbf{v}_i be the type-2 vector with $S(\mathbf{v}_i) = \{i\}$ and its i th entry being ℓ_i and let \mathbf{v}'_i denote the type-2 vector with $S(\mathbf{v}'_i) = \{i\}$ and its i th entry being h_i . The constraint (7.5) over $(\mathbf{v}, \mathbf{w}) = (\mathbf{v}_i, \mathbf{c}_2)$, $(\mathbf{v}'_i, \mathbf{c}_2)$, or $(\mathbf{v}'_i, \mathbf{v}_i)$ is part of $\text{LP}'(I)$; for $(\mathbf{v}, \mathbf{w}) = (\mathbf{v}_i, \mathbf{v}'_i)$, (7.5) follows trivially from the fact that $u(\mathbf{v}'_i) \geq u(\mathbf{v}_i)$ (by CONDITION-TYPE-2), and $q_i(\mathbf{v}_i) = 0$. To see the latter, note that by Lemma 7.6.19 we have $u(\mathbf{v}_i) = u(\mathbf{c}_2)$ and thus, an optimal solution to $\text{LP}(\mathbf{v}_i : u)$ must have $q_{n+1} = 1$.

For type-2 (\mathbf{v}, \mathbf{w}) with $|S(\mathbf{v})| \geq 2$, we need to understand $\mathbf{q}(\mathbf{v})$ better. We prove the following lemma regarding $\mathbf{v} \in T_2 \cup T_3 \cup T_4 \setminus \{\mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4\}$ that satisfies certain conditions.

Lemma 7.6.26. *Let $(u(\cdot), \mathbf{q}(\cdot))$ be an optimal solution to $\text{LP}'(I)$ and $\mathbf{v} \in T_2 \cup T_3 \cup T_4 \setminus \{\mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4\}$. Assume that $u(\mathbf{v}) = u(\rho(\mathbf{v}))$ and $u(\mathbf{w}) = u(\rho(\mathbf{w}))$ for every \mathbf{w} that appears in $\text{LP}(\mathbf{v} : u)$. Then $\text{LP}(\mathbf{v} : u)$ has the following unique optimal solution \mathbf{q} : (letting $k = \min(S(\mathbf{v}))$ and $S'(\mathbf{v}) = S(\mathbf{v}) \setminus \{k\}$)*

- *If $k \notin S^+(\mathbf{v})$, $q_i = q_i(\mathbf{a})$ for all $i \in S(\mathbf{v})$, and \mathbf{q} puts the rest of probability $1 - \sum_{i \in S(\mathbf{v})} q_i(\mathbf{a})$ (if any) on q_{n+1} if $\mathbf{v} \in T_2$ or q_{n+2} if $\mathbf{v} \in T_3 \cup T_4$; all other entries of \mathbf{q} are 0.*
- *If $k \in S^+(\mathbf{v})$, $q_i(\mathbf{v}) = q_i(\mathbf{a})$ for all $i \in S'(\mathbf{v})$ and $q_k(\mathbf{v}) = 1 - \sum_{i \in S'(\mathbf{v})} q_i(\mathbf{a})$; all other entries of \mathbf{q} are 0. In this case we have $\mathbf{q}(\mathbf{v}) = \mathbf{q}(\rho(\mathbf{v}))$.*

Proof. We relax $\text{LP}(\mathbf{v} : u)$: its second batch of constraints is now over $i \in S(\mathbf{v})$ and $\mathbf{w}_i = \text{LOWER}(\mathbf{v}_{-i}, a_i)$ only. Denote this linear program by $\text{LP}^*(\mathbf{v} : u)$:

$$\begin{aligned} & \text{Maximize } \sum_{j \in [n+2]} v_j \cdot q_j - u(\mathbf{v}) \text{ subject to} \\ & q_i \geq 0 \text{ and } \sum_{j \in [n+2]} q_j \leq 1, \quad \text{for } i \in [n+2]; \\ & \tau_i \cdot q_i \geq u(\mathbf{v}) - u(\mathbf{w}), \quad \text{for } i \in S^+(\mathbf{v}) \text{ and } \mathbf{w} = (\mathbf{v}_{-i}, \ell_i); \\ & \sum_{j \in [n]} (v_j - w_j) \cdot q_j \geq u(\mathbf{v}) - u(\mathbf{w}), \quad \text{for } i \in S(\mathbf{v}) \text{ and } \mathbf{w} = \text{LOWER}(\mathbf{v}_{-i}, a_i). \end{aligned}$$

We start with the case when $k = \min(S(\mathbf{v})) \notin S^+(\mathbf{v})$. The first batch of constraints yields $q_i \geq q_i(\mathbf{a})$, for all $i \in S^+(\mathbf{v})$, where we used $u(\mathbf{v}) = u(\rho(\mathbf{v}))$, $u(\mathbf{w}) = u(\rho(\mathbf{w}))$, and **CONDITION-TYPE-1**. For each $i \in S(\mathbf{v})$, the second batch requires

$$(v_i - a_i) \cdot q_i + \sum_{j \in S(\mathbf{w})} (v_j - w_j) \cdot q_j \geq d_i \cdot q_i(\mathbf{a}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{a}).$$

Rearranging terms results in for each $i \in S(\mathbf{v})$:

$$d_i \cdot (q_i - q_i(\mathbf{a})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j - q_j(\mathbf{a})) \geq 0.$$

These are the only constraints in $\text{LP}^*(\mathbf{v} : u)$ other than those on \mathbf{q} itself. We now show that $\text{LP}^*(\mathbf{v} : u)$ has a unique optimal solution \mathbf{q} with $q_i = q_i(\mathbf{a})$ for all $i \in S(\mathbf{v})$, and \mathbf{q} allocates all the rest of probability on q_{n+1} or q_{n+2} , depending on whether $\mathbf{v} \in T_2$ or $\mathbf{v} \in T_3 \cup T_4$.

Assume for contradiction that $q_\ell < q_\ell(\mathbf{a})$ for some $\ell \in S(\mathbf{v})$ (this is actually without loss of generality since if $q_i \geq q_i(\mathbf{a})$ for all $i \in S(\mathbf{v})$, then to be optimal \mathbf{q} must be the vector described above). Take ℓ to be an index in $S(\mathbf{v})$ that maximizes $q_\ell(\mathbf{a}) - q_\ell$, denoted by $\epsilon > 0$. For the second constraint on ℓ we must have $q_t - q_t(\mathbf{a}) \geq \Omega(\epsilon/(n\beta))$ for some $t \in S(\mathbf{v})$. Let \mathbf{q}' denote the following vector derived from \mathbf{q} : $q'_i = q_i + \epsilon$ for all $i \in S(\mathbf{v})$ and $i \neq t$; $q'_t = q_t - 2n\epsilon$; q'_{n+1} or q'_{n+2} takes the rest of probability. Then \mathbf{q}' is feasible and strictly better than \mathbf{q} . For feasibility, the only nontrivial constraint to check is the second one on t :

$$d_t \cdot (q'_t - q_t(\mathbf{a})) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j - q_j(\mathbf{a})) \geq \Omega(\epsilon/(n\beta)) - n \cdot O(\beta) \cdot \epsilon > 0.$$

Given that \mathbf{q} described above is the unique optimal solution to $\text{LP}^*(\mathbf{v} : u)$, it is easy to verify that \mathbf{q} is indeed a feasible solution to $\text{LP}(\mathbf{v} : u)$. Taking a $\mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-i}, a_i)$ for some $i \in S(\mathbf{v})$, we have

$$\begin{aligned} u(\mathbf{v}) - u(\mathbf{w}) &= d_i \cdot q_i(\mathbf{a}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\rho(\mathbf{v})) - \sum_{j \in S^+(\mathbf{w})} \tau_j \cdot q_j(\rho(\mathbf{w})) \\ &\leq d_i \cdot q_i(\mathbf{a}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{a}) - \sum_{j \in S^+(\mathbf{w})} \tau_j \cdot q_j(\mathbf{a}) \\ &= \sum_{j \in S(\mathbf{v})} (v_j - w_j) \cdot q_j(\mathbf{a}) = \sum_{j \in S(\mathbf{v})} (v_j - w_j) \cdot q_j. \end{aligned}$$

This finishes the proof of the case when $k = \min(S(\mathbf{v})) \notin S^+(\mathbf{v})$.

We consider the case when $k = \min(\mathbf{v}) \in S^+(\mathbf{v})$. Let $\mathbf{v}' = \rho(\mathbf{v})$. The first batch requires $q_i \geq q_i(\mathbf{v}')$ for all $i \in S^+(\mathbf{v})$ (including k). For each $i \in S^-(\mathbf{v})$, the second batch of $\text{LP}^*(\mathbf{v} : u)$ requires

$$(v_i - a_i) \cdot q_i + \sum_{j \in S(\mathbf{w})} (v_j - w_j) \cdot q_j \geq d_i \cdot q_i(\mathbf{a}) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot q_j(\mathbf{v}').$$

Since $i \in S^-(\mathbf{v})$ and $i \neq k$, we have $q_i(\mathbf{a}) = q_i(\mathbf{v}')$. Rearranging terms, we get for each $i \in S^-(\mathbf{v})$:

$$d_i \cdot (q_i - q_i(\mathbf{v}')) + \sum_{j \in S^+(\mathbf{v})} \tau_j \cdot (q_j - q_j(\mathbf{v}')) \geq 0.$$

It turns out that $\mathbf{q} = \mathbf{q}(\mathbf{v}')$ is the unique *feasible* solution to these constraints (as $\mathbf{q}(\mathbf{v}')$ sums to 1, \mathbf{q} sums to at most 1, and $d_i \gg \tau_j$). Hence, we have $\mathbf{q}(\mathbf{v}) = \mathbf{q}(\mathbf{v}')$ (as $\text{LP}(\mathbf{v} : u)$ is feasible and $\mathbf{q} = \mathbf{q}(\mathbf{v}')$ is the only feasible solution to its relaxation $\text{LP}^*(\mathbf{v} : u)$). This finishes the proof for the case when $k \in S^+(\mathbf{v})$. \square

We summarize below the following property of $\mathbf{q}(\mathbf{v})$ for all $\mathbf{v} \in T_2$ that will be useful later:

Lemma 7.6.27. *For all $\mathbf{v} \in T_2$ and $i \in S(\mathbf{v})$, $q_i(\mathbf{v}) \leq q_i(\rho(\mathbf{v}))$. Moreover, $q_{n+1}(\mathbf{v}) = 1 - \sum_{i \in S(\mathbf{v})} q_i(\mathbf{v})$.*

Proof. Recall \mathbf{v}_i and \mathbf{v}'_i at the beginning of Section 7.6.3.3. For \mathbf{c}_2 and \mathbf{v}_i we have $q_{n+1}(\mathbf{c}_2) = q_{n+1}(\mathbf{v}_i) = 1$; for \mathbf{v}'_i we have $q_i(\mathbf{v}'_i) = 1 = q_i(\rho(\mathbf{v}'_i))$. The rest of $\mathbf{v} \in T_2$ follows from Lemma 7.6.26. \square

Now let (\mathbf{v}, \mathbf{w}) be a pair of type-2 vectors with $S(\mathbf{w}) \subseteq S(\mathbf{v})$ and $|S(\mathbf{v})| \geq 2$ (so Lemma 7.6.26 applies to \mathbf{v} and we know exactly what $\mathbf{q}(\mathbf{v})$ is). The rest of the proof is similar to that for type-1 vectors.

Using $u(\mathbf{v}) = u(\rho(\mathbf{v}))$ and $u(\mathbf{w}) \geq u(\rho(\mathbf{w}))$, we have

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{i \in S(\mathbf{v}) \setminus S(\mathbf{w})} d_i \cdot q_i(\mathbf{a}) + \sum_{i \in S^+(\mathbf{v})} \tau_i \cdot q_i(\rho(\mathbf{v})) - \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot q_i(\rho(\mathbf{w})).$$

For the case when $k = \min(\mathbf{v}) \notin S^+(\mathbf{v})$, we have $q_i(\mathbf{v}) = q_i(\mathbf{a})$ for all $i \in S(\mathbf{v})$. We have

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{i \in S(\mathbf{v}) \setminus S(\mathbf{w})} d_i \cdot q_i(\mathbf{a}) + \sum_{i \in S^+(\mathbf{v})} \tau_i \cdot q_i(\mathbf{a}) - \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot q_i(\mathbf{a}) = \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}),$$

where we used $q_i(\rho(\mathbf{v})) = q_i(\mathbf{a})$ for all $i \neq \min(S(\mathbf{v}))$ and $q_i(\rho(\mathbf{w})) \geq q_i(\mathbf{a})$ for all $i \in S(\mathbf{w})$.

For the case when $k = \min(\mathbf{v}) \in S^+(\mathbf{v})$, we have $\mathbf{q}(\mathbf{v}) = \mathbf{q}(\rho(\mathbf{v}))$. Then

$$u(\mathbf{v}) - u(\mathbf{w}) \leq \sum_{i \in S(\mathbf{v}) \setminus S(\mathbf{w})} d_i \cdot q_i(\mathbf{v}) + \sum_{i \in S^+(\mathbf{v})} \tau_i \cdot q_i(\mathbf{v}) - \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot q_i(\mathbf{v}) = \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}),$$

where we used $q_i(\rho(\mathbf{w})) \geq q_i(\rho(\mathbf{v})) = q_i(\mathbf{v})$ for all $i \in S(\rho(\mathbf{w})) = S(\mathbf{w})$ by Lemma 7.6.11.

This finishes the proof of (7.5) over all pairs (\mathbf{v}, \mathbf{w}) of type-2 vectors.

7.6.3.4 Both \mathbf{v} and \mathbf{w} are Type-3

Now we turn to pairs (\mathbf{v}, \mathbf{w}) of type-3 vectors that satisfy $S(\mathbf{w}) \subseteq S(\mathbf{v})$.

When $|S(\mathbf{v})| \geq m+3$, we note that by Lemma 7.6.22 and 7.6.23, \mathbf{v} satisfies the condition of Lemma 7.6.26 which completely characterizes $\mathbf{q}(\mathbf{v})$. The same argument above for type-2 vectors with $|S(\mathbf{v})| \geq 2$ can be used to prove (7.5) for type-3 (\mathbf{v}, \mathbf{w}) with $S(\mathbf{w}) \subseteq S(\mathbf{v})$ and $|S(\mathbf{v})| \geq m+3$.

Next we check the case when $|S(\mathbf{v})| \leq m+1$. The case when $u(\mathbf{v}) = u(\mathbf{c}_3)$ is simple as $q_{n+2}(\mathbf{v}) = 1$ (note that this includes $\mathbf{v} = \mathbf{c}_3$). As a result, we have (using $u(\mathbf{w}) \geq u(\mathbf{c}_3)$ by CONDITION-TYPE-3)

$$u(\mathbf{v}) - u(\mathbf{w}) \leq 0 = \sum_{i \in [n+2]} (v_i - w_i) \cdot q_i(\mathbf{v}).$$

For the case when $u(\mathbf{v}) > u(\mathbf{c}_3)$ and $|S(\mathbf{v})| \leq m+1$, by Lemma 7.6.22 and 7.6.23 we must have $\mathbf{v} \in R^*$. $\mathbf{q}(\mathbf{v})$ is then an optimal solution to the following (relaxed) linear

program (from $\text{LP}(\mathbf{v} : u)$):

$$q_i \geq 0, \quad \sum_{i \in [n+2]} q_i \leq 1, \quad \tau_i \cdot q_i \geq u(\mathbf{v}) - u(\mathbf{c}_3), \quad \text{for } i \in S^+(\mathbf{v}),$$

as all other constraints in $\text{LP}(\mathbf{v} : u)$ would be implied. This implies that $q_i(\mathbf{v}) = (u(\mathbf{v}) - u(\mathbf{c}_3))/\tau_i$ for all $i \in S(\mathbf{v})$ and $q_{n+2}(\mathbf{v})$ takes the rest of probability. We now prove (7.5) on (\mathbf{v}, \mathbf{w}) . Using $S(\mathbf{w}) \subseteq S(\mathbf{v})$ and $\mathbf{w} \neq \mathbf{v}$ (so $u(\mathbf{w}) = u(\mathbf{c}_3)$), there must be an index $t \in S(\mathbf{v})$ such that $w_t < v_t$. As a result we have

$$\sum_{i \in [n+2]} (v_i - w_i) \cdot q_i(\mathbf{v}) = \sum_{i \in S(\mathbf{v})} (h_i - w_i) \cdot q_i(\mathbf{v}) \geq \tau_t \cdot q_t(\mathbf{v}) = u(\mathbf{v}) - u(\mathbf{c}_3) = u(\mathbf{v}) - u(\mathbf{w}).$$

The only case left for type-3 (\mathbf{v}, \mathbf{w}) is when $|S(\mathbf{v})| = m + 2$. We need the next lemma about its $\mathbf{q}(\mathbf{v})$.

Lemma 7.6.28. *For each $\mathbf{v} \in T_3$ with $|S(\mathbf{v})| = m + 2$, $\mathbf{q}(\mathbf{v})$ satisfies $q_i(\mathbf{v}) = q_i(\rho(\mathbf{v}))$ for each $i \in S^+(\mathbf{v})$, $q_i(\mathbf{v}) \leq q_i(\mathbf{a})$ for each $i \in S^-(\mathbf{v})$, and $\mathbf{q}(\mathbf{v})$ puts the rest of probability on $q_{n+2}(\mathbf{v})$.*

Proof. By $\text{LP}(\mathbf{v} : u)$, $q_i(\mathbf{v})$ for each $i \in S^+(\mathbf{v})$ must satisfy

$$\tau_i \cdot q_i(\mathbf{v}) \geq u(\mathbf{v}) - u(\mathbf{v}_{-i}, \ell_i) = u(\rho(\mathbf{v})) - u(\rho(\mathbf{v}_{-i}, \ell_i)) = \tau_i \cdot q_i(\rho(\mathbf{v})), \quad (7.63)$$

since we have $u(\mathbf{w}) = u(\rho(\mathbf{w}))$ for $\mathbf{w} \in T_3$ with $|S(\mathbf{w})| \geq m + 2$. Let \mathbf{q} be the vector with $q_i = q_i(\rho(\mathbf{v}))$ for all $i \in S^+(\mathbf{v})$ and $q_i = q_i(\mathbf{a})$ for all $i \in S^-(\mathbf{v})$. Let $c = \max_i (q_i(\mathbf{v}) - q_i)$, and assume for contradiction that $c > 0$. Let $t \in S(\mathbf{v})$ denote an index with $q_t(\mathbf{v}) = q_t + c$. We consider two cases below.

Case 1: One of the constraints in the second batch of $\text{LP}(\mathbf{v} : u)$ with $i = t$ is tight, i.e., there is a type-3 vector $\mathbf{w} \in \text{BLOCK}(\mathbf{v}_{-t}, a_t)$ such that

$$\sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}) = u(\mathbf{v}) - u(\mathbf{w}).$$

Since $u(\mathbf{v}) = u(\rho(\mathbf{v}))$ and $u(\mathbf{w}) \geq u(\rho(\mathbf{w}))$, we have

$$\sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}) \leq u(\rho(\mathbf{v})) - u(\rho(\mathbf{w})) = d_t \cdot q_t(\mathbf{a}) + \sum_{i \in S^+(\mathbf{v})} \tau_i \cdot q_i(\rho(\mathbf{v})) - \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot q_i(\rho(\mathbf{w})).$$

Plugging in $q_t(\mathbf{a}) \leq q_t$, $q_i(\rho(\mathbf{v})) = q_i$ for $i \in S^+(\mathbf{v})$, and $q_i(\rho(\mathbf{w})) \geq q_i(\rho(\mathbf{v})) \geq q_i$ for $i \in S(\mathbf{w})$,

$$d_t \cdot (q_t(\mathbf{v}) - q_t) + \sum_{i \in S^+(\mathbf{v})} \tau_i \cdot (q_i(\mathbf{v}) - q_i) \leq \sum_{i \in S^+(\mathbf{w})} \tau_i \cdot (q_i(\mathbf{v}) - q_i) \leq n \cdot O(\beta) \cdot c.$$

Given that $d_t \approx 1 \gg O(n\beta)$, there must exist an $i \in S^+(\mathbf{v})$ such that $q_i(\mathbf{v}) - q_i < 0$, contradicting (7.63).

Case 2: All constraints in the second batch with $i = t$ are loose. For this case we lower $q_t(\mathbf{v})$ by ϵ , for some sufficiently small $\epsilon > 0$, increase $q_i(\mathbf{v})$ by $\epsilon/(2n)$ for other $i \in S(\mathbf{v})$, and move the rest of (at least $\epsilon/2$) probability to $q_{n+2}(\mathbf{v})$. This gives a feasible solution that is strictly better than $\mathbf{q}(\mathbf{v})$, a contradiction. \square

We summarize below the following property of $\mathbf{q}(\mathbf{v})$ for all $\mathbf{v} \in T_3$ that will be useful later:

Lemma 7.6.29. *For all $\mathbf{v} \in T_3$ and $i \in S(\mathbf{v})$, $q_i(\mathbf{v}) \leq q_i(\rho(\mathbf{v}))$. Moreover, $q_{n+1}(\mathbf{v}) = 1 - \sum_{i \in S(\mathbf{v})} q_i(\mathbf{v})$.*

Proof. The case of $u(\mathbf{v}) = u(\mathbf{c}_3)$ is trivial. The case of $u(\mathbf{v}) > u(\mathbf{c}_3)$ and $|S(\mathbf{v})| \leq m + 1$ follows from

$$q_i(\mathbf{v}) = (u(\mathbf{v}) - u(\mathbf{c}_3))/\tau_i \leq (u(\rho(\mathbf{v})) - u(\rho(\mathbf{v}_{-i}, l_i)))/\tau_i = q_i(\rho(\mathbf{v})).$$

The rest of $\mathbf{v} \in T_4$ follows from either Lemma 7.6.26 or Lemma 7.6.28. \square

We now return to prove (7.5) for pairs (\mathbf{v}, \mathbf{w}) of type-3 vectors with $S(\mathbf{w}) \subseteq S(\mathbf{v})$ and $|S(\mathbf{v})| = m + 2$. The only nontrivial case here is when \mathbf{w} also has $|S(\mathbf{w})| = m + 2$. For other cases, we have

1. $|S(\mathbf{w})| = m + 1$: Trivial since the constraint is indeed part of $\text{LP}'(I)$;
2. $|S(\mathbf{w})| < m + 1$: Let \mathbf{w}^* denote a type-3 vector in R such that $\mathbf{w} \prec \mathbf{w}^* \prec \mathbf{v}$ and $w_i^* = h_i$ for all $i \in S(\mathbf{w})$. Then we have $u(\mathbf{w}^*) = u(\mathbf{w}) = u(\mathbf{c}_3)$. It follows from (7.5) over $(\mathbf{v}, \mathbf{w}^*)$ that

$$u(\mathbf{v}) - u(\mathbf{w}) = u(\mathbf{v}) - u(\mathbf{w}^*) \leq \sum_{i \in S(\mathbf{v})} (v_i - w_i^*) \cdot q_i(\mathbf{v}) \leq \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}),$$

where the last inequality follows from $w_i \leq w_i^*$ for all i .

When $|S(\mathbf{w})| = m + 2$, we have $S(\mathbf{v}) = S(\mathbf{w})$. Then $u(\mathbf{v}) - u(\mathbf{w}) = u(\rho(\mathbf{v})) - u(\rho(\mathbf{w}))$ and

$$\begin{aligned} u(\rho(\mathbf{v})) - u(\rho(\mathbf{w})) &= \sum_{i \in S^+(\mathbf{v}) \setminus S^+(\mathbf{w})} \tau_i \cdot q_i(\rho(\mathbf{v})) - \sum_{i \in S^+(\mathbf{w}) \setminus S^+(\mathbf{v})} \tau_i \cdot q_i(\rho(\mathbf{v})) \\ &\leq \sum_{i \in S^+(\mathbf{v}) \setminus S^+(\mathbf{w})} \tau_i \cdot q_i(\mathbf{v}) - \sum_{i \in S^+(\mathbf{w}) \setminus S^+(\mathbf{v})} \tau_i \cdot q_i(\mathbf{v}) \\ &= \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}), \end{aligned}$$

since $q_i(\mathbf{v}) = q_i(\rho(\mathbf{v}))$ for all $i \in S^+(\mathbf{v})$ and $q_i(\mathbf{v}) \leq q_i(\mathbf{a}) \leq q_i(\rho(\mathbf{v}))$ for all $i \in S^-(\mathbf{v})$.

This finishes the proof of (7.5) over pairs of type-3 vectors.

7.6.3.5 Both \mathbf{v} and \mathbf{w} are Type-4

For each $\mathbf{v} \in T_4$, let $\Phi(\mathbf{v}) = (\mathbf{v}_{-(n+1)}, 0)$. So Φ is a one-to-one correspondence between type-4 and type-3 vectors. As $u(\mathbf{c}_4) = u(\mathbf{c}_3)$, we have $u(\mathbf{v}) = u(\Phi(\mathbf{v}))$ for all $\mathbf{v} \in T_4$. This suggests the following lemma.

Lemma 7.6.30. *Let $(u(\cdot), \mathbf{q}(\cdot))$ be an optimal solution to $\text{LP}'(I)$ and $\mathbf{v} \in T_4$. Then $(u(\cdot), \mathbf{q}(\cdot))$ remains to be an optimal solution to $\text{LP}'(I)$ after replacing $\mathbf{q}(\Phi(\mathbf{v}))$ by $\mathbf{q}(\mathbf{v})$.*

Proof. The statement is trivial for $\mathbf{v} = \mathbf{c}_4$ since $q_{n+2}(\mathbf{c}_3) = q_{n+2}(\mathbf{c}_4) = 1$.

For $\mathbf{v} \neq \mathbf{c}_4$, note that $\text{LP}(\mathbf{v} : u)$ is essentially the same as $\text{LP}(\Phi(\mathbf{v}) : u)$, with the only subtle difference being that the coefficient of q_{n+1} is s in $\text{LP}(\mathbf{v} : u)$ but 0 in $\text{LP}(\Phi(\mathbf{v}) : u)$. However, neither $\mathbf{q}(\Phi(\mathbf{v}))$ nor $\mathbf{q}(\mathbf{v})$ can put any probability on q_{n+1} . The lemma then follows. \square

To prove (7.5) on a pair (\mathbf{v}, \mathbf{w}) of type-4 vectors we simply replace $\mathbf{q}(\Phi(\mathbf{v}))$ by $\mathbf{q}(\mathbf{v})$ to get a new optimal solution by Lemma 7.6.30, and (7.5) must hold on $(\Phi(\mathbf{v}), \Phi(\mathbf{w}))$ in the new solution (since we have proved (7.5) between type-3 vectors in any optimal solution). This then implies (7.5) on (\mathbf{v}, \mathbf{w}) in the original solution.

7.6.3.6 Pairs with Different Types

Finally we prove (7.5) for pairs (\mathbf{v}, \mathbf{w}) of vectors with $S(\mathbf{w}) \subseteq S(\mathbf{v})$ and of different types.

The following lemma helps us further reduce cases that need to be considered.

Lemma 7.6.31. *Assume that $\mathbf{v}, \mathbf{v}' \in D$ differ at the i th entry only, for some $i \in \{n+1, n+2\}$, and $v'_i > v_i$. Then we have $u(\mathbf{v}') \geq u(\mathbf{v})$ in any optimal solution to $\text{LP}'(I)$.*

Proof. The case when $\mathbf{v} \in T_1$ follows directly from CONDITION-TYPE-2 and CONDITION-TYPE-3.

The case when $\mathbf{v} \in T_3$ and $i = n+1$ follows from $u(\mathbf{c}_3) = u(\mathbf{c}_4)$.

The case when $\mathbf{v} \in T_2$ and $i = n+2$ follows from the fact that $u(\mathbf{c}_3) > u(\mathbf{c}_2)$. \square

It suffices to prove (7.5) for (\mathbf{v}, \mathbf{w}) that satisfies $v_{n+1} \geq w_{n+2}$ and $v_{n+2} \geq w_{n+2}$. To see this, we let \mathbf{w}' denote the vector obtained from \mathbf{w} by replacing w_i by $\min(w_i, v_i)$, $i \in \{n+1, n+2\}$. Then $u(\mathbf{w}) \geq u(\mathbf{w}')$ by Lemma 7.6.31 and $(\mathbf{v}, \mathbf{w}')$ satisfies $v_{n+1} \geq w'_{n+1}$ and $v_{n+2} \geq w'_{n+2}$. Assuming that (7.5) holds for $(\mathbf{v}, \mathbf{w}')$,

$$u(\mathbf{v}) - u(\mathbf{w}) \leq u(\mathbf{v}) - u(\mathbf{w}') \leq \sum_{i \in [n+2]} (v_i - w'_i) q_i(\mathbf{v}) = \sum_{i \in [n]} (v_i - w_i) q_i(\mathbf{v}) + \sum_{i \in \{n+1, n+2\}} (v_i - w'_i) q_i(\mathbf{v})$$

The RHS is indeed the same as $\sum_i (v_i - w_i) \cdot q_i(\mathbf{v})$. This is because for either $i \in \{n+1, n+2\}$, $w_i \neq w'_i$ would imply that $v_i = 0$ and thus, $q_i(\mathbf{v}) = 0$.

Now we need to consider the following cases of types of (\mathbf{v}, \mathbf{w}) : (2, 1), (3, 1), (4, 1), (4, 2) and (4, 3). We start with the case when \mathbf{v} is type-2 and \mathbf{w} is type-1.

We consider two cases: $u(\mathbf{v}) = u(\mathbf{c}_2)$ or $u(\mathbf{v}) > u(\mathbf{c}_2)$. For the former, $q_{n+1}(\mathbf{v}) = 1$ and thus,

$$u(\mathbf{v}) - u(\mathbf{w}) \leq u(\mathbf{c}_2) \ll s = \sum_{i \in [n+2]} (v_i - w_i) \cdot q_i(\mathbf{v}).$$

For the latter, $u(\mathbf{v}) = u(\rho(\mathbf{v}))$. By Lemma 7.6.27, let $\gamma_i = q_i(\rho(\mathbf{v})) - q_i(\mathbf{v}) \geq 0$ for each $i \in S(\mathbf{v})$. Then

$$\begin{aligned} u(\mathbf{v}) - u(\mathbf{w}) &= u(\rho(\mathbf{v})) - u(\mathbf{w}) \leq \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\rho(\mathbf{v})) \\ &= \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}) + \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot \gamma_i \\ &\leq \sum_{i \in S(\mathbf{v})} (v_i - w_i) \cdot q_i(\mathbf{v}) + s \cdot \sum_{i \in S(\mathbf{v})} \gamma_i \\ &= \sum_{i \in [n+2]} (v_i - w_i) \cdot q_i(\mathbf{v}), \end{aligned}$$

where the last equation used $q_{n+1}(\mathbf{v}) = 1 - \sum_{i \in S(\mathbf{v})} q_i(\mathbf{v})$ from Lemma 7.6.27.

The case when \mathbf{v} is type-3 and \mathbf{w} is type-1 can be proved similarly using Lemma 7.6.29. From this case, the case when \mathbf{v} is type-3 and \mathbf{w} is type-2 follows from Lemma 7.6.31 (we mention it since it is used below).

For the case when \mathbf{v} is type-4 and \mathbf{w} is type-3, we simply replace $\mathbf{q}(\Phi(\mathbf{v}))$ by $\mathbf{q}(\mathbf{v})$ to get a new optimal solution by Lemma 7.6.30. (7.5) on (\mathbf{v}, \mathbf{w}) in the original solution then follows from that on $(\Phi(\mathbf{v}), \mathbf{w})$ in the new solution (note that this is the (3, 2) case we already handled), given that $q_{n+1}(\mathbf{v}) = 0$.

For the case when \mathbf{v} is type-4 and \mathbf{w} is type-2, we again replace $\mathbf{q}(\Phi(\mathbf{v}))$ by $\mathbf{q}(\mathbf{v})$ to get a new optimal solution by Lemma 7.6.30. (7.5) on (\mathbf{v}, \mathbf{w}) in the original solution then follows from that on $(\Phi(\mathbf{v}), \mathbf{w})$ in the new solution, given $q_{n+1}(\mathbf{v}) = 0$. The same argument works for the case when \mathbf{v} is type-4 and \mathbf{w} is type-1.

This finishes the proof of Lemma 7.6.25.

Part III

Conclusions

This page intentionally left blank.

Chapter 8

Conclusions

8.1 Conclusions

In this thesis we studied the Computational Complexity of Equilibrium Computation and Revenue Maximization. Here, we summarize our results and provide future research directions.

8.1.1 Summary of Results

8.1.1.1 Market Equilibria

In the first part of the thesis we studied the computation of Market Equilibria. Prior to our work, all PPAD-hardness proofs were based on reductions fine-tuned to the utility functions under consideration. In an effort to unify these results, we took a step towards a systematic understanding of what features make equilibrium computation hard.

Our first result is a complete characterization of the complexity of finding an ϵ -approximate equilibrium in Arrow-Debreu markets with CES utilities based on the parameters ρ_i of the utilities. In particular, we show (Theorem 3) that it is PPAD-hard to find an ϵ -approximate market equilibrium when the traders are allowed to have CES utilities with parameters $\rho_i < -1$, thus resolving this long-standing open problem.

Building on this result, we introduced the notion of non-monotone utilities, which covers a wide variety of important utility functions. Our main result (Theorem 6) states that for

any family \mathcal{U} of non-monotone utilities, it is PPAD-hard to compute an ϵ -approximate equilibrium for a market with utilities that are drawn from \mathcal{U} or are linear. Our result is the first to provide a general sufficient PPAD-hardness condition for Arrow-Debreu markets. We note however that although all previously known PPAD-hard families of utilities satisfy non-monotonicity they are not subsumed by our result because of the use of linear utilities.

8.1.1.2 Revenue Maximization

In the second part of the thesis we concentrate on the complexity of Revenue Maximization when there is a Unit-demand buyer. We first study Bayesian Unit-Demand Item-Pricing problem with independent distributions. Prior to our work, only approximation algorithms existed for the problem; indeed, the problem was not even known to be in NP. Our first result is a *structural characterization* (Theorem 8) of the optimal prices as solutions to a Linear Program restricted to cells of a partition of the price space. NP-membership follows by guessing the appropriate cell, which can be described with a polynomial number of linear constraints. Our next result (Theorem 9) is a *polynomial time algorithm* for the case that each marginal distribution has support size at most 2. Finally, our main result (Theorem 10) *resolves the computational complexity* of the problem showing that it is NP-hard even for distributions of support size at most 3.

The final contribution of this thesis is on the complexity of Unit-demand Lottery Pricing. We begin by studying the special case of distributions with support-size 2 and the same high value for all items and prove (Theorem 12) that *randomization does not improve revenue* in this case. Therefore, one can use our algorithm for Item Pricing and distributions of support size 2. After providing some simple examples where randomization improves revenue we prove our main result (Theorem 14); unless $P^{NP} = P^{\#P}$, *there is no universal efficient randomized algorithm that computes implicitly an optimal menu*. This concludes the contributions of this thesis.

8.1.2 Future Directions and Open Problems

Our work raises many interesting questions for both Market Equilibrium computation and Revenue Maximization. Here, we present what we believe are the most interesting ones.

8.1.2.1 Market Equilibria

We begin by pointing out that our hardness result for CES markets characterizes the computation of approximate market equilibria, that is, finding prices that approximately clear the market. In general, approximate equilibrium prices, when viewed as a vector, may be far from an actual equilibrium price vector and this makes them undesirable in many cases; economists are more interested in approximating actual equilibrium prices. In our work, we prove that the problem of approximating an actual equilibrium for CES markets is in FIXP. Theorem 3 further implies that the problem is PPAD-hard; by continuity a price-vector close to an actual equilibrium is an approximate market equilibrium. Furthermore, it is known that when $\rho \rightarrow -\infty$, that is all utilities are Leontief functions, it is FIXP-hard to approximate an actual equilibrium. However, the exact complexity for finite ρ remains unknown and we conjecture that it is FIXP-complete to approximate an actual equilibrium of a CES market.

Regarding our general result for arbitrary non-monotone utilities, can we dispense with the linear functions, i.e., is it true that for any family \mathcal{U} of non-monotone utilities, the approximate equilibrium problem is PPAD-hard for markets that use utilities from \mathcal{U} only? The reduction for CES utilities is essentially a fine-tuned version of the reduction for non-monotone utilities that makes use of the properties of CES utilities. Can a similar approach work in general for all non-monotone utilities? Related to this question, what other general features of utilities (if any) make the market equilibrium problem hard? Non-monotonicity is related with markets that have disconnected sets of market equilibria for which currently we do not have any efficient algorithmic methods to deal with. Convexity has been critical essentially in all tractable cases so far, whether the set of market equilibria itself is convex or a convex formulation can be obtained after a change of variables. Can we obtain a complexity dichotomy theorem that allows us to classify any family of utility functions (under standard, generally acceptable, mild assumptions for utilities) into those that can be solved efficiently and those that are intractable (PPAD-hard and/or FIXP-hard)?

On the algorithmic side of the problem, the work of Codenotti et al [Codenotti *et al.*, 2005a] provides convex programs for the Arrow-Debreu setting in the cases that either every trader i has a $0 < \rho_i$ or every trader i has a $-1 \leq \rho_i < 0$. However, the two convex

programs cannot be combined because the two cases have different structure of equilibria. In the case $0 < \rho_i$ the set of equilibria is convex while in the case $-1 \leq \rho_i < 0$ it is log-convex. Currently, no algorithm is known for the case that every trader i has a $-1 \leq \rho_i < 1$ with $\rho_i \neq 0$. We conjecture that this case also admits a polynomial time algorithm but we believe that a new approach is needed to obtain it.

Second, convexity seems to be the main property exploited by all algorithms for finding an (approximate) market equilibrium; with WGS being the largest family of markets with convex set of equilibria for which algorithms are known. Another general property that implies convexity of the set of equilibria is the Weak Axiom of Revealed Preference (WARP) [Mas-Colell *et al.*, 1995] but currently there is no polynomial time algorithm that works for all WARP markets. Is it possible to obtain an efficient algorithm for all WARP markets?

A more ambitious direction is towards obtaining a general sufficient condition that allows for an efficient algorithm. While many families of utilities satisfy WGS or WARP, they do not seem to cover all the efficiently solvable market problems. For example, the family of CES utilities with parameter $-1 \leq \rho < 0$ does not satisfy WGS or WARP but has a convex formulation. Is there a more general property that covers all the markets with connected sets of equilibria and, if so, can this property lead to polynomial time algorithms?

8.1.2.2 Revenue Maximization

Our results in this thesis address the computational complexity for the general cases of both Item Pricing and Lottery Pricing for a Unit-demand buyer by proving that they are NP-complete and #P-hard respectively. In addition, we provide efficient algorithms for a few special cases. Many important questions however remain widely open.

Regarding Item Pricing, our NP-hardness results do not preclude the existence of a PTAS or even an FPTAS (Fully Polynomial Time Approximation Scheme). Actually, by adapting techniques from [Cai and Daskalakis, 2011] we can give an FPTAS for the case when the supports of the distributions are integers in a bounded interval. Thus, the most natural and important open question is the following: *Is there a PTAS or, even better, an FPTAS for unit-demand deterministic pricing?* Addressing this question will require new intuition about the problems and their structure as well as application of tools from other

areas such as geometry and probability theory. A first step towards this direction would be identifying special cases of distributions that allow for a PTAS. Existing work provides a PTAS for Monotone Hazard Rate distributions. Can this result be extended to different distributions? Finally, our hardness proof for the i.i.d. case uses distributions of support polynomial to the number of items. In fact, when the distributions are identical the problem obtains a nice structure. Namely, one has then to decide how many items get each of the prices suggested by our NP-membership. When all the prices come from the support, this observation gives an $O(n^k)$ algorithm for the problem, where k is the size of the support. We conjecture that this can be extended to a polynomial algorithm for identical distributions of constant support size, even when the candidate prices suggested by our NP-membership do not lie in the support.

Regarding the randomized setting, the question of whether a PTAS or an FPTAS exists transfers here, since our hardness proof does not rule out this possibility. Similarly, identifying special cases of the problem that accept polynomial-time exact or approximation algorithms is a major open problem in this setting too. Especially for the i.i.d. case, it is possible that one can obtain exact algorithms using the symmetries that arise in this case. Finally, regarding the menu-size of the problem, obtaining a characterization of the cases that need an exponentially large menu remains an important open problem.

This page intentionally left blank.

Part IV

Bibliography

This page intentionally left blank.

Bibliography

- [Arrow and Debreu, 1954] K.J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954.
- [Arrow *et al.*, 1959] K.J. Arrow, H.D. Block, and L. Hurwicz. On the stability of the competitive equilibrium, II. *Econometrica*, 27(1):82–109, 1959.
- [Arrow *et al.*, 1961] K.J. Arrow, H. Chenery, B. Minhas, and R. Solow. Capital-labor substitution and economic efficiency. *Review of Economics and Statistics*, 63:225–250, 1961.
- [Babaioff *et al.*, 2014] M. Babaioff, N. Immorlica, B. Lucier, and S.M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. *SIGecom Exchange*, 13(2):31–35, 2014.
- [Briest *et al.*, 2010] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing randomized allocations. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 585–597, 2010.
- [Briest, 2008] P. Briest. Uniform budgets and the envy-free pricing problem. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 808–819, 2008.
- [Cai and Daskalakis, 2011] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 522–531, 2011.

- [Chawla *et al.*, 2007] S. Chawla, J.D. Hartline, and R.D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 243–251, 2007.
- [Chawla *et al.*, 2010a] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing*, pages 311–320, 2010.
- [Chawla *et al.*, 2010b] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 149–158, 2010.
- [Chen and Deng, 2008] X. Chen and X. Deng. Matching algorithmic bounds for finding a Brouwer fixed point. *Journal of the ACM*, 55(3):1–26, 2008.
- [Chen and Teng, 2007] X. Chen and S.-H. Teng. Paths beyond local search: A tight bound for randomized fixed-point computation. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 124–134, 2007.
- [Chen and Teng, 2009] X. Chen and S.-H. Teng. Spending is not easier than trading: On the computational equivalence of Fisher and Arrow-Debreu equilibria. In *Proceedings of the 20th International Symposium on Algorithms and Computation*, pages 647–656, 2009.
- [Chen and Teng, 2011] X. Chen and S.-H. Teng. A complexity view of markets with social influence. In *Proceedings of the 2nd Symposium on Innovations in Computer Science*, 2011.
- [Chen *et al.*, 2004] N. Chen, X. Deng, X. Sun, and A.C.-C. Yao. Fisher equilibrium price with a class of concave utility functions. In *Proceedings of the 12th Annual European Symposium on Algorithms*, pages 169–179, 2004.
- [Chen *et al.*, 2006] X. Chen, L.-S. Huang, and S.-H. Teng. Market equilibria with hybrid linear-Leontief utilities. In *Proceedings of the 2nd Workshop on Internet and Network Economics*, pages 274–285, 2006.

- [Chen *et al.*, 2008] X. Chen, X. Sun, and S.-H. Teng. Quantum separation of local search and fixed point computation. In *Proceedings of the 14th Annual International Computing and Combinatorics Conference*, pages 170–179, 2008.
- [Chen *et al.*, 2009a] X. Chen, D. Dai, Y. Du, and S.-H. Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 273–282, 2009.
- [Chen *et al.*, 2009b] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):1–57, 2009.
- [Chen *et al.*, 2013] Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. The complexity of non-monotone markets. In *In Proceedings of the 45th ACM Symposium on Theory of Computing*, 2013.
- [Chen *et al.*, 2014] X. Chen, I. Diakonikolas, D. Paparas, X. Sun, and M. Yannakakis. The complexity of optimal multidimensional pricing. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms*, 2014.
- [Chen *et al.*, 2015] X. Chen, I. Diakonikolas, A. Orfanou, D. Paparas, X. Sun, and M. Yannakakis. On the complexity of optimal lottery pricing and randomized mechanisms. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Sciences*, 2015.
- [Codenotti and Varadarajan, 2007] B. Codenotti and K. Varadarajan. Computation of market equilibria by convex programming. In E. Tardos, N. Nisan, T. Roughgarden, and V.V. Vazirani, editors, *Algorithmic Game Theory*, pages 135–158. Cambridge Press, 2007.
- [Codenotti *et al.*, 2005a] B. Codenotti, B. McCune, S. Penumatcha, and K. Varadarajan. Market equilibrium for exchange CES economies: Existence, multiplicity, and computation. In *Proceedings of the 25th Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 505–516, 2005.

- [Codenotti *et al.*, 2005b] B. Codenotti, B. McCune, and K. Varadarajan. Market equilibrium via the excess demand function. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 74–83, 2005.
- [Codenotti *et al.*, 2005c] B. Codenotti, S. Pemmaraju, and K. Varadarajan. On the polynomial time computation of equilibria for certain exchange economies. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 72–81, 2005.
- [Codenotti *et al.*, 2006] B. Codenotti, A. Saberi, K. Varadarajan, and Y. Ye. Leontief economies encode nonzero sum two-player games. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 659–667, 2006.
- [Daskalakis *et al.*, 2009] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1), 2009.
- [Daskalakis *et al.*, 2013] C. Daskalakis, A. Deckelbaum, and C. Tzamos. Mechanism design via optimal transport. In *Proceedings of the 14th ACM conference on Electronic commerce*, pages 269–286, 2013.
- [Daskalakis *et al.*, 2014a] C. Daskalakis, A. Deckelbaum, and C. Tzamos. The complexity of optimal mechanism design. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms*, pages 1302–1318, 2014.
- [Daskalakis *et al.*, 2014b] C. Daskalakis, A. Deckelbaum, and C. Tzamos. Strong duality for a multiple-good monopolist. *CoRR abs/1409.4150*, 2014.
- [de La Grandville, 2009] O. de La Grandville. *Economic Growth: A Unified Approach*. Cambridge University Press, 2009.
- [Debreu, 1959] G. Debreu. *Theory of Value*. New York, Wiley, 1959.
- [Deng and Du, 2008] X. Deng and Y. Du. The computation of approximate competitive equilibrium is PPAD-hard. *Information Processing Letters*, 108:369–373, 2008.
- [Deng *et al.*, 2003] X. Deng, C.H. Papadimitriou, and S. Safra. On the complexity of price equilibria. *Journal of Computer and System Sciences*, 67(2):311–324, 2003.

- [Devanur and Kannan, 2008] N.R. Devanur and R. Kannan. Market equilibria in polynomial time for fixed number of goods or agents. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 45–53, 2008.
- [Devanur and Vazirani, 2003] N.R. Devanur and V.V. Vazirani. An improved approximation scheme for computing Arrow-Debreu prices for the linear case. In *Proceedings of the 23rd Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 149–155, 2003.
- [Devanur and Vazirani, 2004] N.R. Devanur and V.V. Vazirani. The spending constraint model for market equilibrium: Algorithmic, existence and uniqueness results. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 519–528, 2004.
- [Devanur *et al.*, 2008] N.R. Devanur, C.H. Papadimitriou, A. Saberi, and V.V. Vazirani. Market equilibrium via a primal-dual algorithm for a convex program. *Journal of the ACM*, 55(5):1–18, 2008.
- [Dickinson, 1954] H.D. Dickinson. A note on dynamic economics. *Review of Economic Studies*, 22(3):169–179, 1954.
- [Dobzinski *et al.*, 2011] Shahar Dobzinski, Hu Fu, and Robert D. Kleinberg. Optimal auctions with correlated bidders are easy. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 129–138, 2011.
- [Eaves, 1985] B.C. Eaves. Finite solution of pure trade markets with Cobb-Douglas utilities. *Mathematical Programming Studies*, 23:226–239, 1985.
- [Eisenberg and Gale, 1959] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [Ellickson, 1994] B. Ellickson. *Competitive Equilibrium: Theory and Applications*. Cambridge University Press, 1994.

- [Etessami and Yannakakis, 2010] K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.
- [Gale, 1976] D. Gale. The linear exchange model. *Journal of Mathematical Economics*, 3:205–209, 1976.
- [Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, 1979.
- [Garg and Kapoor, 2004] R. Garg and S. Kapoor. Auction algorithms for market equilibrium. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 511–518, 2004.
- [Garg *et al.*, 2004] R. Garg, S. Kapoor, and V. Vazirani. An auction-based market equilibrium algorithm for the separable gross substitutability case. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems and 8th International Workshop on Randomization and Computation*, pages 128–138, 2004.
- [Garg *et al.*, 2014] J. Garg, R. Mehta, V. V. Vazirani, and S. Yazdanbond. Leontief exchange markets can solve multivariate polynomial equations yielding fixp and etr hardness. arXiv:1411.5060, 2014.
- [Garg *et al.*, 2016] Jugal Garg, Ruta Mehta, and Vijay V. Vazirani. Dichotomies in equilibrium computation and membership of plc markets in fixp. *Theory of Computing*, 2016.
- [Gjerstad, 1996] S. Gjerstad. Multiple equilibria in exchange economies with homothetic, nearly identical preferences. Technical Report 288, Minnesota – Center for Economic Research, 1996.
- [Guruswami *et al.*, 2005] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1164–1173, 2005.

- [Hart and Nisan, 2012] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 656–656, 2012.
- [Hart and Nisan, 2013] S. Hart and N. Nisan. The menu-size complexity of auctions. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 565–566, 2013.
- [Hicks, 1932] J.R. Hicks. *The Theory of Wages*. MacMillan, London, 1932.
- [Hirsch *et al.*, 1989] M.D. Hirsch, C.H. Papadimitriou, and S. Vavasis. Exponential lower bounds for finding Brouwer fixed points. *Journal of Complexity*, 5:379–416, 1989.
- [Huang and Teng, 2007] L.-S. Huang and S.-H. Teng. On the approximation and smoothed complexity of Leontief market equilibria. In *Proceedings of the 1st International Frontiers of Algorithmics Workshop*, pages 96–107, 2007.
- [Jain and Mahdian, 2005] K. Jain and M. Mahdian. Computing equilibria in a Fisher market with linear single constraint production units. In *Proceedings of the 1st International Workshop on Internet and Network Economics*, pages 788–792, 2005.
- [Jain and Varadarajan, 2006] K. Jain and K. Varadarajan. Equilibria for economies with production: Constant-returns technologies and production planning constraints. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithm*, pages 688–697, 2006.
- [Jain *et al.*, 2003] K. Jain, M. Mahdian, and A. Saberi. Approximating market equilibria. In *Proceedings of the 6th International Workshop on Approximation Algorithms*, pages 98–108, 2003.
- [Jain *et al.*, 2005] K. Jain, V.V. Vazirani, and Y. Ye. Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 63–71, 2005.
- [Jain, 2007] K. Jain. A polynomial time algorithm for computing an Arrow-Debreu market equilibrium for linear utilities. *SIAM Journal on Computing*, 37:303–318, 2007.

- [Janovskaya, 1968] E.B. Janovskaya. Equilibrium points in polymatrix games. *Latvian Mathematical Collection*, 1968.
- [Kakutani, 1941] S. Kakutani. A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal*, 8:457–459, 1941.
- [Li and Yao, 2013] X. Li and A.C.-C. Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.
- [Manelli and Vincent, 2006] A.M. Manelli and D.R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.
- [Manelli and Vincent, 2007] Alejandro M. Manelli and Daniel R. Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- [Mas-Colell *et al.*, 1995] A. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford Press, 1995.
- [Maxfield, 1997] R.R. Maxfield. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics*, 27(1):23–51, 1997.
- [McAfee and Mc-Millan, 1988] P. McAfee and J. Mc-Millan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theor*, 46(2):335–354, 1988.
- [McAfee *et al.*, 1989] R. Preston McAfee, John McMillan, and Michael D. Whinston. Multi-product monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 104(2):371–383, 1989.
- [Myerson, 1981] R.B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [Nenakov and Primak, 1983] E.I. Nenakov and M.E. Primak. One algorithm for finding solutions of the Arrow-Debreu model. *Kibernetika*, 3:127–128, 1983.

- [Papadimitriou and Pierrakos, 2011] C.H. Papadimitriou and G. Pierrakos. On optimal single-item auctions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 119–128, 2011.
- [Papadimitriou and Wilkens, 2011] C.H. Papadimitriou and C. Wilkens. Economies with non-convex production and complexity equilibria. In *Proc. of the 12th ACM Conf. on Electronic Commerce*, pages 137–146, 2011.
- [Papadimitriou and Yannakakis, 2010] C. H. Papadimitriou and M. Yannakakis. An impossibility theorem for price adjustment mechanisms. *Proc. Natl. Acad. Sci.*, 107(5):1854–1859, 2010.
- [Papadimitriou, 1994] C.H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, pages 498–532, 1994.
- [Pavlov, 2010] G. Pavlov. Optimal Mechanism for Selling Two Goods. UWO Department of Economics Working Papers 20103, University of Western Ontario, Department of Economics, 2010.
- [Robinson, 1933] J. Robinson. *The Economics of Imperfect Competition*. MacMillan, London, 1933.
- [Rubinstein, 2015] A. Rubinstein. Inapproximability of nash equilibrium. In *Proc. 47th Annual ACM Symposium on Theory of Computing*, pages 409–418, 2015.
- [Rubinstein, 2016] Aviad Rubinstein. On the computational complexity of optimal simple mechanisms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 21–28, 2016.
- [Rutherford, 1999] T. Rutherford. Applied general equilibrium modeling with MPSGE as a GAMS subsystem: An overview of the modeling framework and syntax. *Computational Economics*, 14:1–46, 1999.
- [Samuelson, 1947] P.A. Samuelson. *Foundations of Economic Analysis*. Harvard University Press, 1947.

- [Scarf, 1973] H. Scarf. *The Computation of Economic Equilibria*. Yale University Press, 1973.
- [Shoven and Whalley, 1992] J.B. Shoven and J. Whalley. *Applying General Equilibrium*. Cambridge University Press, 1992.
- [Solow, 1956] R.M. Solow. A contribution to the theory of economic growth. *Quarterly Journal of Economics*, 70(1):65–94, 1956.
- [Thanassoulis, 2004] J. Thanassoulis. Haggling over substitutes. *Journal of Economic Theory*, 117(2):217–245, 2004.
- [Vazirani and Yannakakis, 2011] V.V. Vazirani and M. Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM*, 58(3):1–25, 2011.
- [Walras, 1874] L. Walras. *Elements of Pure Economics, or the Theory of Social Wealth*. 1954 English Translation, 1874.
- [Wang and Tang, 2014] Z. Wang and P. Tang. Optimal mechanisms with simple menus. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 227–240, 2014.
- [Wilson, 1996] Robert Butler Wilson. Nonlinear pricing and mechanism design. In H.M. Amman, D.A. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, volume 1, chapter 5, pages 253–293. Elsevier, 1996.
- [Yao, 2015] A.C.-C. Yao. An n -to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109, 2015.
- [Ye, 2007] Y. Ye. Exchange market equilibria with Leontief’s utility: Freedom of pricing leads to rationality. *Theoretical Computer Science*, 378(2):134–142, 2007.
- [Ye, 2008] Y. Ye. A path to the Arrow-Debreu competitive market equilibrium. *Mathematical Programming*, 111:315–348, 2008.