Inefficiency of Ordinal One-Sided Matching Mechanisms

Vasileios Stamatis

Examination committee:
Aris Filos - Ratsikas, University of Edinburgh.
Dimitris Fotakis, National Technical University of Athens.
Evangelos Markakis, Athens University of Economics & Business

Supervisors:
Aris Filos-Ratsikas, Lecturer,
School of Informatics,
University of Edinburgh.

Dimitris Fotakis, Professor,
School of Electrical & Computer Engineering,
National Technical University of Athens.

Μεταπτυχιακό Πρόγραμμα «Αλγόριθμοι, Λογική και Διακριτά Μαθηματικά» – 2016

ABSTRACT

This thesis examines Ordinal One-Sided Matching Mechanisms, that allocate indivisible resources. We focus on two important mechanisms: Random Priority and Probabilistic Serial. Random Priority is safe from player manipulation (truthful), but often leads to suboptimal fair and efficient results, while Probabilistic Serial produces efficient results, but is vulnerable to manipulation.

We study several concepts such as social welfare, efficiency, truthfulness and analyze the performance of these mechanisms. We study the Approximation Ratio, which quantifies how close the outcome of a mechanism is to the optimal solution in terms of social welfare; the Price of Anarchy, which measures how the players' strategic behavior affects the efficiency of the mechanism; the Incentive Ratio, which expresses the maximum benefit a player can gain by manipulating her preferences compared to her true ones.

Finally, we present experimental results and discuss a new metric to assess how agents' strategic behavior affects the mechanism's social welfare relative to its outcome if everyone were telling the truth.

ΣΥΝΟΨΗ

Αυτή η διπλωματική εργασία εξετάζει μηχανισμούς μονομερούς αντιστοίχισης με βάση τις διατακτικές προτιμήσεις παικτών, οι οποίοι κατανέμουν αδιαίρετους πόρους. Επικεντρωνόμαστε σε δύο σημαντικούς μηχανισμούς: τον Random Priority και τον Probabilistic Serial. O Random Priority είναι ασφαλής από χειραγώγιση από τους παίκτες (truthful), αλλά συχνά οδηγεί σε υποβέλτιστα δίκαια αποτελέσματα, ενώ ο Probabilistic Serial παράγει δίκαια αποτελέσματα, αλλά είναι ευάλωτος σε χειραγώγηση.

Εξετάζουμε αρκετές έννοιες, όπως την κοινωνική ευημερία, την αποδοτικότητα, την ασφάλεια στην χειραγώγιση και αναλύουμε την απόδοση αυτών των μηχανισμών. Μελετούμε το Approximation Ratio, το οποίο ποσοτικοποιεί πόσο κοντά βρίσκεται το αποτέλεσμα ενός μηχανισμού στην ιδανική λύση από την άποψη της κοινωνικής ευημερίας· το Τίμημα της Αναρχίας (Price of Anarchy), που μετρά πώς η στρατηγική συμπεριφορά των παικτών επηρεάζει την αποδοτικότητα του μηχανισμού· το Incentive Ratio, το οποίο εκφράζει το μέγιστο όφελος που μπορεί να κερδίσει ένας παίκτης με την αλλαγή των προτιμήσεων σε σύγκριση με την αληθινή του αναφορά.

Τέλος, παραθέτουμε πειραματικά αποτελέσματα και συζητούμε μια νέα μετρική, για να εκτιμηθεί πώς η στρατηγική συμπεριφορά των πρακτόρων επηρεάζει το κοινωνικό όφελος του μηχανισμού, σε σχέση με την έκβαση του, αν όλοι έλεγαν την αλήθεια.

CONTENTS

CHAPTER 1

INTRODUCTION

1.1 Introduction

Artificial intelligence research fundamentally focuses on the study and design of autonomous agents capable of making logical decisions. These autonomous agents find applications in a wide range of areas including, but not limited to, online systems, customer support systems, healthcare, robotics, transportation and traffic control, information management and electronic commerce.

The field of multiagent systems research, draws on a wide range of techniques, from mathematics, economics and computer science to create autonomous agents that can navigate intelligently and achieve specific goals when interacting with potentially uncertain environments. A key concept in modeling decision-making, particularly in multiagent systems, is the handling of individual preferences. Agents often find themselves in decision-making situations and require computational models of preferences to reason over various choices. These preferences may guide agent behavior when seeking collective decisions in social choice problems such as voting, or when making rational decisions while competing or cooperating with other intelligent agents.

In the efficient allocation of indivisible resources among self-interested agents who compete for limited resources, individual preferences are the crucial factor in assessing the quality of the allocation and achieving favourable social outcomes. Techniques from game theory and economics are used in multi-agent resource allocation to ensure various desirable properties based on agents' preferences. Mechanism design approaches, such as matching mechanisms and auctions, are often used to ensure fairness and efficiency while preventing agents from manipulating outcomes.

However, in many scenarios, such as assigning student housing, distributing faculty workloads, assigning student courses, and allocating organs to patients, monetary transactions are prohibited. These systems rely on self-reported preferences, but agents may not always report their true preferences. For example, parents may misrepresent their preferences in school placement systems in order to secure better placements for their children. To address this problem, the fields of mechanism design and matching theory have developed algorithms that encourage truthful preference reporting.

Truthful reporting of preferences is key to ensuring other desirable properties such as efficiency and fairness. Without ensuring truthfulness, a matching mechanism can only guarantee efficient and fair allocations with respect to the, perhaps untruthful, preferences, and thus fails to satisfy the desirable economic properties with respect to the true underlying preferences.

It is therefore crucial the construction allocation mechanisms that provide incentives for self-interested agents to report their preferences truthfully, while ensuring economic efficiency and fairness.

1.2 Outline of the Thesis

The focus of this thesis is a review of the literatute of ordinal mechanisms for the One - Sided Matching Problem - the problem of allocating indivisible items to agents in a way that each agent receives exactly one item. This problem is central to various applications such as student course allocation, housing assignments, and organ donation. The central challenge lies in designing mechanisms that respect agents' preferences while ensuring fair and efficient outcomes without the use of money.

This thesis begins with a bibliographic review of the One - Sided Matching Problem and continues with presenting existing mechanisms designed for the One - Sided Matching Problem. These mechanisms include well-known mechanisms such as the Random Priority and Probabilistic Serial.

By further studying these mechanism we encounter economic inefficiencies by them. For example, the Random Priority mechanism is strategy-proof but can result in suboptimal allocations, while the Probabilistic Serial mechanism, which often yields more efficient and fair outcomes, is susceptible to manipulation.

Social Welfare — a prominent measure of collective utility widely used in the literature, captures the overall satisfaction of all agents under an allocation mechanism and as natural, it is desired to be as high as possible. The inefficiency of these mechanisms, arrising due to strategic behavior or inherent design limitations, can be captured by Social Welfare. Further in this thesis, we will see work done, providing theretical bounds for these inefficiencies, while we will see experiments, showing how these mechanisms behave in the average case.

In the last Chapter of this thesis, we provide possible future directions on this problem. We suggest a new measure that provides a deeper understanding of how strategic behavior can sometimes yield better outcomes than truthful reporting. We state some thoughts concerning possible bounds, and an algorithm that produces such examples.

This thesis is structured as follows:

- **Chapter 1**: Introduction to the thesis.
- **Chapter 2**: This chapter explains the notation and basic concepts used in the thesis.
- **Chapter 3**: This chapter reviews the One-Sided Matching Problem..
- **Chapter 4**: This chapter studies the Approximation Ratio of Ordinal One - Sided Matching Mechanisms.
- **Chapter 5**: This chapter studies the Price of Anarchy of Ordinal One - Sided Matching Mechanisms.
- **Chapter 6**: This chapter studies the Incentive Ratio of the Probabilistic Serial Mechanism.
- **Chapter 7**: This chapter introduces the Price of Truthfulness and its impact.

CHAPTER 2

PRELIMINARIES

We standardize our notation throughout this thesis based on the work in [8], [10] and [5] ensuring consistency with the notational conventions established therein.

2.1 Basic Definitions

In this section, we present an overview of the key definitions and fundamental concepts that form the basis of this thesis. Throught this thesis we let $N = \{1, \dots, n\}$ to be a finite set of agents and $A = \{1, \dots, m\}$ be a finite set of indivisible items. We assume that $n = m$, meaning that the number of items are exactly the same as the number of agents, unless stated otherwise.

2.1.1 Allocations

An **allocation** is a matching of agents to items, that is, an assignment of items to agents where each agent gets assigned exactly one item. We can view an allocation μ as a permutation vector $(\mu_1, \mu_2 \dots, \mu_n)$ where μ_i is the unique item matched with agent i. Let O be the set of all allocations.

2.1.2 Valuation Functions

Each agent i has a (private) valuation function $u_i : A \to \mathbb{R}$ mapping items to real numbers. The two standard ways to fix the canonical representation of u_i in the literature are **unit-range**, where $\max_j u_i(j) = 1$ and $\min_j u_i(j) = 0$, and **unit-sum**, that is, $\sum_j u_i(j) = 1$ with $u_i(j) \geq 0$ for all i, j. Equivalently, we can consider valuation functions as *valuation vectors* $u_i = (u_{i1}, u_{i2}, \dots, u_{in})$ and let V be the set of all valuation vectors of an agent. Let $\mathbf{u} = (u_1, u_2, \dots, u_n)$ denote a typical *valuation profile* and let V^n be the set of all valuation profiles with n agents.

2.1.3 Mechanisms

A *direct revelation mechanism* without money is a function $M : V^n \to O$ mapping *reported* valuation profiles to matchings.

2.1.4 Strategies

We consider *strategic agents* who might have incentives to misreport their valuations. We define $\mathbf{s} = (s_1, s_2, \dots, s_n)$ to be a pure strategy profile, where s_i is the *reported* valuation vector of agent i. We will use $\mathbf{s_{-i}}$ to denote the strategy profile without the ith coordinate and hence $\mathbf{s} = (s_i, \mathbf{s_{-i}})$ is an alternative way to denote a strategy profile.

2.1.5 Ordinal & Cardinal Preferences

A mechanism M is **ordinal** if for any i strategy profile $\mathbf{s}, \mathbf{s}'$ such that for all agents i and for all items j, ℓ, $s_{ij} < s_{i\ell} \Leftrightarrow s'_{ij} < s'_{i\ell}$, it holds that $M(\mathbf{s}) = M(\mathbf{s}')$. A mechanism for which the above does not necessarily hold is **cardinal**.

Equivalently, the strategy space of ordinal mechanisms is the set of all permutations of n items instead of the space of valuation functions V^n. A strategy s_i of agent i is a **preference ordering** of items $(a_1, a_2, \ldots, a_n)$ where $a_\ell \succ a_k$ for $\ell < k$. We will write $j \succ_i j'$ to denote that agent i prefers item j to item j' according to her true valuation function and $j \succ_{s_i} j'$ to denote that she prefers item j to item j' according to her strategy s_i.

When not confusing, we abuse the notation slightly and let u_i denote the truthtelling strategy of agent i, even when the mechanism is ordinal. Note that agents can be indifferent between items and hence the preference order can be a *weak ordering*.

2.1.6 Randomized Mechanisms

For a randomized mechanism, we define M to be a random map $M : V^n \to O$. Let $M_i(\mathbf{s})$ denote the restriction of the outcome of the mechanism to the i'th coordinate, which is the item assigned to agent i by the mechanism. For randomized mechanisms, we let $p_{ij}^{M,\mathbf{s}} = \Pr[M_i(\mathbf{s}) = j]$ and $p_i^{M,\mathbf{s}} = (p_{i1}^{M,\mathbf{s}}, \ldots, p_{in}^{M,\mathbf{s}})$. When it is clear from the context, we drop one or both of the superscripts from the terms $p_{ij}^{M,\mathbf{s}}$. The utility of an agent from the outcome of a deterministic mechanism M on input strategy profile **s** is simply $u_i(M_i(\mathbf{s}))$. For randomized mechanisms, an agent's utility is $E[u_i(M_i(\mathbf{s}))] = \sum_{j=1}^{n} p_{ij}^{M,\mathbf{s}} u_{ij}$.

2.2 Concepts from Economics

2.2.1 Efficiency, Truthfulness & Fairness

When giving away resources, two natural goals are to allocate these resources efficiently and fairly. Achieving these goals requires asking people about their preferences, so a third goal is to allocate in a way that is truthful (incentivizes people to tell us their true preferences).

Can we achieve all three goals simultaneously? Achieving all three goals - efficiency, fairness and truthfulness - simultaneously in resource allocation is generally not possible due to fundamental trade-offs in mechanism design. Theoretical results such as the Gibbard-Satterthwaite theorem and others show that when there are three or more agents or options, no deterministic mechanism can satisfy all three properties without being dictatorial. This impossibility forces mechanism designers to make compromises: between efficiency and truthfulness on the one hand, and fairness on the other; between achieving fairness and truthfulness on the one hand, and some efficiency on the other. Of course, some randomised mechanisms achieve the right trade-offs by slightly relaxing one or the other of these goals, but a harmonious solution that achieves all three at once remains elusive. The reader, might be interested in [22] and [23] for a more inclusive discusion.

Pareto Efficiency and Truthfulness

A deterministic mechanism M is said to be **Pareto Efficient** (PE) if for every preference profile $\succ$, there is no allocation that is as good (according to $\succ$) as $M(\succ)$ for every agent, and strictly better for some agent. In other words, there is no other allocation of the items that can make at least one agent better off without making any other agent worse off.

A deterministic mechanism is **Truthful** (or *strategyproof* or *incentive-compatible*) if, for any preference profile $\succ$ and agent i, the object that i receives from reporting $\succ_i$ is at least as good (according to $\succ_i$) as the object that i receives from reporting any other preference. In, other words, agents do not have incentives to misreport their valuations - telling the truth is a dominant strategy.

Utility-wise, a mechanism $M : V^n \to O$ is **truthfull**, if for every agent i and all $\mathbf{u} = (u_i, u_{-i}) \in V^n$ and $u_i' \in V$ it holds that:

$$u_i(M_i(\mathbf{u})) \geq u_i(M_i(u_i', u_{-i}))$$

Where u_{-i} denotes the valuation profile **u** without the i-th coordinate.

Random Mechanisms & Lotteries

In order to achieve fairness, we often use **random mechanisms**. Formally defined previously, a random mechanism is a function that maps each preference profile to a distribution over allocations.

One way to construct a random mechanism is to use a **lottery** over deterministic mechanisms. A lottery, is a method of randomization used to determine outcomes. It involves assigning probabilities to different possible allocations and then using a random process to select one of these allocations based on the assigned probabilities.

For example, suppose there are two agents and two objects a, b. A simple lottery mechanism might allocate a to agent 1 and a to agent 2 with probability $1/2$, and allocate b to agent 1 and a to agent 2 with probability $1/2$. The lottery ensures that each agent has an equal chance of receiving their preferred object if both prefer the same object.

A random mechanism is **truthful-in-expectation** (or just truthful) if for each agent i and all $\mathbf{u} \in V^n$ and $u'_i \in V$ it holds that:

$$E[u_i(M_i(\mathbf{u}))] \geq E[u_i(M_i(u'_i, u_{-i}))]$$

This condition ensures that, in expectation, an agent cannot benefit by misreporting their preferences.

2.2.2 Ex Ante & Ex Post Truthfulness and Efficiency

Things here start to become a bit complicated, as the randomness affects both the truthfulness. There are two natural approaches to defining what it means for a random mechanism M to satisfy property: it can satisfy this property **ex post** (after the randomness of the mechanism is resolved), or **ex ante** (before the randomness is resolved).

Depending on the property satisfied, the ex ante requirement can be either stronger or weaker than its ex post counterpart. Bogomolnaia and Moulin in [1] proved that **ordinal efficiency** - that is, no other feasible allocation that all agents weakly prefer and at least one agent strictly prefers, based solely on their ordinal preferences, cannot be achieved at all if we also want truthfulness (either ex ante or ex post) and a weak fairness criteria (equal treatment of equals).

Ex Ante Efficiency

A mechanism is ex ante efficient if, before the random allocation is resolved, there is no other mechanism that could provide a better expected outcome for every agent. Formally, a mechanism M is ex ante Pareto efficient if, for every preference profile $\succ$, there is no distribution over allocations that is weakly $\succ$ - preferred by every agent to $M(\succ)$, and strictly $\succ$ - preferred by some agent.

Ex Ante Truthfulness

A mechanism is ex ante truthful if, before the random allocation is resolved, agents have no incentive to misreport their preferences to improve their expected outcome. Formally, a mechanism M is ex ante truthful if, for any preference profile $\succ$, each agent i weakly $\succ$ - prefers $M(\succ)$to the distribution over outcomes that would result if i reported any other preference.

Ex Post Efficiency

A mechanism is ex post efficient if, after the random allocation is resolved, the outcome is Pareto efficient; that is, no agent can be made better off without making another agent worse off. Formally, a mechanism M is ex post Pareto efficient if it can be expressed as a lottery over deterministic mechanisms, each of which is Pareto efficient.

Ex Post Truthfulness

A mechanism is ex post truthful if, after the random allocation is resolved, agents have no incentive to misreport their preferences given the realized outcome. Formally, a mechanism M is ex post truthful if it can be expressed as a lottery over deterministic mechanisms, each of which is truthful.

Ex Post Truthfulness is Stronger than Ex Ante Truthfulness:

If a mechanism is ex post truthful, it is also ex ante truthful. This is because if agents cannot benefit from misreporting their preferences even after knowing the random seed, they certainly cannot benefit if they don't know the seed. However, the converse is not necessarily true. An ex ante truthful mechanism might not be ex post truthful.

Ex Ante Efficiency is Stronger than Ex Post Efficiency

If a mechanism is ex ante efficient, it is also ex post efficient. This is because an ex ante efficient mechanism ensures that no better expected allocation exists, which implies that each realized allocation must be Pareto efficient. However, the converse is not necessarily true. An ex post efficient mechanism might not be ex ante efficient, as it may not maximize expected utility.

Ex Post vs. Ex Ante Truthfulness Example:

Consider a mechanism where an agent can receive their first choice with probability $1/2$ and their second choice with probability $1/2$.

- **Ex Ante:** This mechanism is truthful because the agent maximizes their expected utility by reporting their true preferences.
- **Ex Post:** To see if it's ex post truthful, we would need to check if the agent would still prefer to report truthfully even after the random outcome is known. Some ex ante truthful mechanisms do not hold up under ex post scrutiny.

Ex Post vs. Ex Ante Efficiency Example:

Consider four agents and four objects with specific preference profiles.

- **Ex Ante:** A mechanism that randomizes in a way that, before the allocation, all agents expect to be as well off as possible, considering all possible outcomes.
- **Ex Post:** A mechanism that, after the allocation, no agent can be made better off without making another agent worse off. Random serial dictatorship (RSD) is an example that is ex post efficient but might not always be ex ante efficient.

2.2.3 Fairness and Envy-Freeness

Fairness

Fairness in the context of resource allocation generally means that the resources are distributed in a manner that is equitable and just, considering the preferences and needs of all involved parties. Fairness can be operationalized in various ways depending on the specific context and the criteria used.

Equal Treatment of Equals

Equal treatment of equals is a fairness criterion. It dictates that agents who have identical preferences should be treated the same way in the allocation process. This principle ensures that no arbitrary distinctions are made between agents with the same preferences, thereby promoting fairness and equity.

In formal terms, equal treatment of equals can be defined as follows:

If two agents i and j have the same preference profile, then they should receive allocations that are statistically indistinguishable. In other words, the allocation mechanism should not favor one agent over another when their preferences are identical. This principle helps in maintaining fairness by ensuring that the allocation process does not introduce bias or favoritism.

Envy-Freeness

Envy-freeness is a more stringent criterion for fairness. It ensures that no player prefers another player's share over their own. The formal definition from the document is:

- **Envy relation:** Player p_i envies player p_j (denoted as $p_i \succ p_j$) if the valuation v_i of p_i for their own share X_i is less than their valuation for p_j's share X_j, i.e., $v_i(X_i) < v_i(X_j)$.
- **Envy-free relation:** Player p_j is not envied by player p_i (denoted as $p_i \not\succ p_j$) if $v_i(X_i) \geq v_i(X_j)$.

An allocation is envy-free if no player envies another player.

There are two types of envy-freeness:

- **Ex-Ante Envy-Free:** Before the outcomes are determined, each agent's expected allocation is at least as good as the expected allocation of any other agent.
- **Ex-Post Envy-Free:** After the outcomes are determined, no agent prefers another agent's allocation to their own.

2.2.4 Anonymity & Neutrality

An **anonymous** mechanism has the property that agents with the same valuation functions must have the same probability of receiving each item.

More formally, a mechanism is **anonymous** if for any valuation profile $(u_1, \dots, u_n)$, every agent i, and any permutation $\sigma : N \to N$ it holds that:

$$M_i(\mathbf{u}) = M_{\sigma(i)}(u_{\sigma(1)}, \dots, u_{\sigma(n)})$$

A **neutral** mechanism, is a mechanism mechanism that is invariant to the indices of the items. Formally, a **neutral** mechanism has the property that for any valuation profile $(u_1, \dots, u_n)$, every agent i, and any permutation $\sigma : N \to N$ it holds that:

$$M_i(\mathbf{u}) = \sigma^{-1}(M_i(\sigma(u_1), \dots, \sigma(u_n)))$$

2.3 Nash Equilibrium

The idea of an equilibrium in strategic interaction is pretty old and goes back to Antoine Augustin Cournot, who first introduced his version of it in 1838 when he was studying competition in oligopolies. However, it was John Forbes Nash Jr. who, in the 1950s, extended this concept beyond its initial application to a more general framework, demonstrating that it applies to any strategic situation involving a finite number of players and strategies. In his Ph.D. thesis ([24]), Nash proved that every finite game, regardless of the number of players, has at least one equilibrium in which no player has an incentive to unilaterally change their strategy if all others keep theirs unchanged, now known as a Nash equilibrium. Nash's ingenious work showed that a Nash equilibrium must exist in mixed strategies for any finite set of strategies, using a difficult-to-conceive function and applying Brouwer's Fixed-point theorem. His work not only won him a Nobel Prize in 1994 but also cemented the Nash equilibrium as one of the most important tools for analyzing competitive scenarios in economics.

An **equilibrium** refers to a strategy profile where no agent has an incentive to unilaterally deviate from their chosen strategy, as doing so would not increase their payoff.

In this thesis, we address results on several standard equilibrium concepts: pure Nash, mixed Nash, correlated, coarse correlated, and Bayes-Nash Equilibria. The first four of these concepts correspond to complete information, meaning all agents are fully informed of the structure of the game and the preferences of all others. The Bayesian setting generalizes these by considering incomplete information, where the valuations of the agents come from known distributions, and each agent knows only its own valuation and the distribution of the valuations of the other agents.

We begin by formally defining these equilibrium concepts:

Given a mechanism M, a strategy profile $\mathbf{s}$ is a **pure Nash equilibrium** if, for every agent i, the utility u_i that agent i receives from the outcome under strategy $\mathbf{s}$ satisfies:

$$u_i(M_i(\mathbf{s})) \geq u_i(M_i(s'_i, s_{-i}))$$

for all agents i and all pure deviations s'_i.

Now, let $\mathbf{q}$ be a distribution over strategies. Also, for any distribution Δ let Δ_{-i} denote the marginal distribution without the ith index. Then a strategy profile $\mathbf{q}$ is called a:

- In a **mixed Nash equilibrium**, agents randomize over strategies. Formally,
$$\mathbf{q} = \times_i q_i, E_{\mathbf{s}\sim\mathbf{q}}[u_i(M_i(\mathbf{s}))] \geq E_{s_{-i}\sim q_{-i}}[u_i(M_i(s'_i, s_{-i}))],$$
- In a **correlated equilibrium**, players coordinate their strategies based on a shared probability distribution. Formally,
$$E_{\mathbf{s}\sim\mathbf{q}}[u_i(M_i(\mathbf{s}))|s_i] \geq E_{\mathbf{s}\sim\mathbf{q}}[u_i(M_i(s'_i, s_{-i}))|s_i],$$
- In a **coarse correlated equilibrium**, players follow a recommendation drawn from a shared probability distribution. Formally,
$$E_{\mathbf{s}\sim\mathbf{q}}[u_i(M_i(\mathbf{s}))] \geq E_{\mathbf{s}\sim\mathbf{q}}[u_i(M_i(s'_i, \mathbf{s_{-i}}))],$$
- A**Bayes-Nash equilibrium** occurs in games with incomplete information, there, agents optimize their strategies based on their private information. Formally, for a distribution Δ_u where each $(\Delta_u)_i$ is independent, when $\mathbf{u} \sim \Delta_u$, then $\mathbf{q}(\mathbf{u}) = \times_i q_i(u_i)$ and for all u_i in the support of $(\Delta_u)_i$,
$$E_{\mathbf{u_{-i}},\mathbf{s}\sim\mathbf{q}(\mathbf{u})}[u_i(M_i(\mathbf{s}))] \geq E_{\mathbf{u_{-i}},\mathbf{s_{-i}}\sim\mathbf{q_{-i}}(\mathbf{u_{-i}})}[u_i(M_i(s'_i, \mathbf{s_{-i}}))].$$

Where the inequalities above hold for all agents i, and (pure) deviating strategies s'_i. Note that for randomized mechanisms, these definitions are with respect to an expectation over the random choices of the mechanism.

It is know that pure Nash equilibria are contained within mixed Nash equilibria, which are contained within correlated equilibria, which in turn are contained within coarse correlated equilibria. Similarly, in the context of incomplete information, pure Nash equilibria are contained within mixed Nash equilibria, which are contained within Bayes-Nash equilibria.

2.3.1 Price of Anarchy & Price of Stability

Price of Anarchy

The Price of Anarchy (PoA) was introduced to measure the inefficiency in systems due to self-interested strategic behavior by it's agents. It is a measure that quantifies the loss in Social Welfare that occurs when agents act non-cooperatively by comparing the worst-case Nash equilibrium (where no individual has an incentive to deviate) with the socially optimal outcome.

Price of Anarchy, was introduced in the late 1990s by Elias Koutsoupias and Christos Papadimitriou (in [25]) in the context of congestion games, where they sought to understand how the selfish behavior of drivers choosing their routes in a traffic network could lead to suboptimal traffic flows. This concept has since become a fundamental tool in analyzing various strategic settings, including allocation problems.

The Price of Anarchy is particularly useful in multiple equilibrium settings, as it quantifies the potential inefficiency of non-cooperative behaviour by bounding the performance loss. In our case, (as in routing games) in resource allocation games, PoA measures how close the worst-case Nash equilibrium approximates the optimal solution. Remarkably, it often shows that even the least efficient equilibrium maintains a predictable level of efficiency, typically bounded by a constant.

Formally, let $S^M_{\mathbf{u}}$ denote the set of all pure Nash Equilibria of mechanism M under truthful valuation profile $\mathbf{u}$. We define the Price of Anarchy to be:

$$PoA(M) = \sup_{\mathbf{u}\in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s}\in S^M_{\mathbf{u}}} SW_M(\mathbf{u}, \mathbf{s})}$$

where $SW_M(\mathbf{u}, \mathbf{s}) = \sum_{i=1}^{n} E[u_i(M_i(\mathbf{s}))]$ is the expected **Social Welfare** of mechanism M on strategy profile **s** under true valuation profile **u**, and $SW_{OPT}(\mathbf{u}) = \max_{\mu \in O} \sum_{i=1}^{n} u_i(\mu_i)$ is the Social Welfare of the optimal matching. Let $OPT(\mathbf{u})$ be the optimal matching on profile **u** and let $OPT_i(\mathbf{u})$ be the restriction to the i^{th} coordinate. The Price of Anarchy for other Equilirbium Concepts is defined similarly.

Price of Stability

The Price of Stability (PoS) is another standard measure in the literature that measures the inefficiency of a system due to strategic behaviour, similar to the Price of Anarchy (PoA). However, while the PoA looks at the worst-case equilibrium, PoS focuses on the best-case scenario. It is defined as the ratio of the optimal Social Welfare to the Social Welfare at the most efficient Nash equilibrium.

Historically, the PoS was developed to provide a more optimistic counterpart to the PoA. This concept has been particularly useful in network design games, where it assesses how close the best stable solution is to the optimal.

For example, in network design problems, the PoS is often bounded, meaning that even under strategic behaviour, the best equilibrium is not far from the optimal solution. This has important implications for designing systems that are robust to individual incentives and can still achieve high efficiency.

Formally, let $S_{\mathbf{u}}^{M}$ denote the set of all pure Nash Equilibria of mechanism M under truthful valuation profile **u**. We define the Price of Stability to be:

$$PoS(M) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\max_{\mathbf{s} \in S_{\mathbf{u}}^{M}} SW_M(\mathbf{u}, \mathbf{s})}$$

2.3.2 Approximation Ratio

The Approximation Ratio is a measure of performance (efficiency) loss that relates the result obtained by a particular mechanism or algorithm to the optimal solution. Generally, this comparison is made when exact optimisation is rather impractical or strategically challenging.

The Approximation Ratio has its roots in approximation algorithms, which were developed as a response to the complexity of solving NP-hard problems exactly. In this setting, algorithms are designed to find solutions close to the optimum within a provable bound. This idea naturally extended to game theory and mechanism design, where the goal is often to design strategies or mechanisms that perform well even in the face of selfish behaviour by agents.

The formal definition of the Approximation Ratio for a mechanism M is:

$$ar(M) = \inf_{\mathbf{u} \in V^n} \frac{\sum_{i=1}^{n} u_i(M_i(\mathbf{u}))}{\max_{\mu \in O} \sum_{i=1}^{n} u_i(\mu_i)} = \inf_{\mathbf{u} \in V^n} \frac{SW_M(\mathbf{u})}{SW_{OPT}(\mathbf{u})}$$

Here, $SW_M(\mathbf{u})$ represents the Social Welfare achieved by the mechanism M on a given valuation profile **u**, while $SW_{OPT}(\mathbf{u})$ represents the optimal Social Welfare.

2.3.3 Incentive Ratio

The Incentive Ratio is a measure, in mechanism design, that quantifies the maximum potential gain an agent can have by deviating from being truthful. Specifically, it evaluates how much an agent's utility can be increased by strategic manipulation compared to the utility they would receive by reporting truthfully.

Formally, let the utility of agent i when they report truthfully be denoted as $u_i(M_i(u))$ and the utility when they manipulate their preferences as $u_i(M_i(u'))$, the Incentive Ratio $r_i(M)$ is equal to:

$$r_i(M) = \sup_{u, u'} \frac{u_i(M_i(u'))}{u_i(M_i(u))},$$

where u is the true preference profile and u' is the "manipulated" preference profile.

The goal is often to minimise this ratio, so that the incentive to manipulate is as small as possible. A lower Incentive Ratio indicates a mechanism, where agents have little or no benefit from misreporting their preferences.

CHAPTER 3

ONE-SIDED MATCHING

The reader can also read [11] and [17].

3.1 The Stable Marriage Problem

David Gale and Lloyd Shapley's 1962 article, "College Admissions and the Stability of Marriage" in American Mathematical Monthly (see [18]), introduced the stable marriage problem and the deferred acceptance algorithm, which have become cornerstone methodologies in matching theory and has significantly influenced economic theory and practical applications in various domains.

The model Gale and Shapley presented is simple. A number of boys and girls have preferences for each other and would like to be matched. The question Gale and Shapley were especially interested in was whether there is a “stable” way to match each boy with a girl so that no unmatched pair can later find out that they can both do better by matching each other. They found that there indeed is such a stable matching, and they presented an algorithm that achieves this objective. Versions of this algorithm are used today to match hospitals with residents and students with public schools in New York City and Boston.

While two-sided matching markets, where Gale and Shapley focused, in which two distinct groups of agents (such as students and schools or residents and hospitals) seek to form mutually beneficial pairings, have been extensively studied, one-sided matching markets where agents are paired based on individual preferences without the explicit need for mutual agreement, has gained prominence for its applicability in various real-world scenarios.

3.2 One - Sided Matching

One-sided matching markets often involve scenarios where individuals or items need to be allocated efficiently without explicit pairing agreements from both sides. The challenge in one-sided matching is to design mechanisms that respect individual preferences while ensuring overall efficiency and fairness.

Examples include:

- **Kidney Exchange Programs**: Inspired by the TTC algorithm, kidney exchange programs facilitate organ transplants by matching donors and recipients based on compatibility, thereby increasing the number of successful transplants and saving lives.
- **Public Housing Allocation**: One-sided matching mechanisms are used to allocate public housing to applicants based on their preferences and eligibility, ensuring a fair and efficient distribution of housing resources.

One of the most influential models in one-sided matching is the "housing market" model introduced by Lloyd Shapley and Herbert Scarf in [15] . In this model, each agent is initially endowed with a single indivisible item (such as a house) and seeks to trade it for another agents' item, aiming to maximize their utility based on individual preferences. The core question is whether there exists a stable allocation where no group of agents can collectively deviate to improve their situation.

3.2.1 The House Allocation Problem

Formally, a **house allocation problem** (or *one-sided matching allocation problem*, Hylland & Zeckhauser [16]) consists of a set I of n agents, a set H of n indivisible items (*houses*) and preferences of agents over items. In many cases we consider, an initial **endownment** for each agent - the items (in this context, houses) that each agent initially possesses, which is then subject to potential reallocation through the matching.

The problem asks for a matching $\mu : I \to H$ of each agent to a distinct item so as to ensure desirable normative properties, such as:

- *Pareto efficiency*: There is no other matching that makes at least one agent better off without making any other agent worse off.
- *Individual rationality*: Each agent gets a house that they consider at least as good as their initial endowment (the house they initially possess).
- *Envy-freeness*: No one feels envious because they think someone else got a better house than they did.

In simpler words, the question becomes, whether it is possible to find a matching where no group of agents can come together, reallocate the houses among themselves, and make at least one member of the group better off without making any other member worse off.

The **core** of a housing market, that is, a **house allocation problem** with initial **endownments**, are the matchings of a housing market, where no group of agents can come together, reallocate houses among themselves, and make at least one member better off without making any member worse off.

So, an even more minimal question, equivalent to the initial, comes down to wheter there exists an allocation of houses to agents that is in the core.

Lloyd Shapley and Herbert Scarf in 1974 (see [15]) demonstrated the existence of such an allocation and proposed the Top Trading Cycles (TTC) algorithm to achieve it.

3.2.2 Gale's Top Trading Cycles (TTC) Algorithm

The TTC algorithm works by allowing agents to form trading cycles where each agent in a cycle exchanges their item with another agent's item. The process continues until no more beneficial trades can be made, ensuring that the final allocation is stable and Pareto efficient. This algorithm has been instrumental in practical applications, such as kidney exchange programs, where patients and donors are matched to maximize the compatibility of kidney transplants, saving thousands of lives globally.

- Step 1: Each agent "points" to the owner of his favourite house. Since there is a finite number of agents, there is at least one cycle of agents pointing to each other. Each agent in a cycle is assigned the house of the agent he is pointing to and is removed from the market with his assignment. When there is at least one remaining agent, proceed to the next step.
- Step k: Each remaining agent points to the owner of her favorite house among the remaining houses. Every agent in a cycle is assigned the house of the agent she is pointing to and removed from the market with his assignment. If there is at least one remaining agent, proceed with the next step.

Alvin Roth and Andrew Postlewaite later showed that the TTC algorithm has some additional remarkable properties when preferences are strict.

Theorem 3.1. (Theorem 2 in [12]) The outcome of the TTC algorithm is the unique matching in the core of each housing market. Moreover, this matching is the unique competitive allocation.

Theorem 3.2. (Theorem 1 in [13]): The core (as a direct mechanism) is strategy-proof.

In housing markets, as opposed to other one-sided matching problems without **endowments**, there exists a plausible mechanism where truth-telling is a dominant strategy for all agents. Indeed, the core is the only mechanism that ensures strategy-proofness.

Theorem 3.3. (Theorem 1 in [19]): Core is the only mechanism that is Pareto efficient, individually rational, and strategy-proof.

3.2.3 Still Ordinal, but without endowments?

One-sided matching markets can be classified along two dimensions: the nature of the utility functions of agents and whether agents have initial endowments or not. As stated in the previous chapter, utility functions may either have only an ordinal component, i.e., they are modeled via preference relations, or they may have a cardinal component as well. Thus we get four possibilities, which are summarized below, together with the most well-known mechanisms for each.

- (Ordinal, No Endowments): Random Priority and Probabilistic Serial
- (Ordinal, Endowments): Top Trading Cycle
- (Cardinal, No Endowments): Hylland-Zeckhauser
- (Cardinal, Endowments): ε-Approximate ADHZ

The two types of utility functions described above have their individual pros and cons, and neither dominates the other. Whereas the former are easier to elicit, the latter are more expressive, enabling an agent to not only report if she prefers one item to another but also by how much, thereby producing higher quality allocations.

In this thesis, we are especially interested in the Random Priority and Probabilistic Serial mechanisms, which each have unique advantages and disadvantages without one being superior. Random Priority is truthful and easy to implement, ensuring agents have no incentive to misreport their preferences, and it guarantees ex-post Pareto efficiency, providing efficient allocations after preferences are realized. However, it may not always result in fair allocations.

On the other hand, Probabilistic Serial ensures ordinal efficiency and envy-freeness, leading to fairer and more balanced distributions of items, but it is not truthful, meaning agents might benefit from misrepresenting their preferences. These differing strengths and weaknesses highlight the inherent trade-offs in allocation mechanism design.

The Priority Mechanism

As a first step towards the Random Priority, the following deterministic mechanism, called **Priority** (also called *Serial Dictatorship*).

The Priority Mechanism: The mechanism, picks an ordering π of the n agents and in the i^{th} iteration, lets the agent $\pi(i)$ pick her most preferred item among the currently available items and declares the chosen item unavailable, while initially all items are available.

The priority mechanism is obviously not envy-free: if several agents prefer the same item the most, the one earliest in the ordering π will get it.

Theorem 3.4. The priority mechanism is truthful and the allocation produced by it is Pareto efficient.

Proof. First, for truthfulness: In each iteration, the active agent has the opportunity of obtaining the best available item, according to her preference list. Therefore, misreporting preferences can only lead to a suboptimal allocation.

Next, for Pareto efficiency: Let μ be the allocation produced by the priority mechanism and assume, for contradiction, assume that μ' is an allocation that dominates μ formally, $\mu'(\pi(i)) \succeq_{\pi(i)} \mu(\pi(i))$ i.e. for each agent i, $u_i(\mu'(\pi(i))) \geq u_i(\mu(\pi(i)))$. Let i be the first agent that the allocation μ' dominates μ. Clearly, for all agents j before i, agent $\pi(j)$ is assigned the same item under μ and μ'. Therefore, in the i^{th} iteration, agent $\pi(i)$ has available item $\mu'(\pi(i))$. Since $\pi(i)$ picks the best available item, for agent $\mu(\pi(i)) \succeq_{\pi(i)} \mu'(\pi(i))$, leading to a contradiction. □

The priority mechanism suffers from the obvious drawback of not being fair, since agents at the top of the list π have the opportunity of choosing their favorite items while those at the bottom get the left-overs.

The Random Priority corrects this.

Random Priority: Random Priority, also called Random Serial Dictatorship, iterates over all $n!$ orderings of the n agents. For each ordering π, it runs the priority mechanism and when an agent chooses an item, it assigns a $\frac{1}{n!}$ share of the item to the agent.

Theorem 3.5. Random Priority is truthful and ex-post Pareto efficient.

Proof. As we saw earlier, in the previous Theorem, in each of the $n!$ iterations, truthfulness is the dominant strategy of an agent. Thus, Random Priority is truthful. We next argue that the random allocation output by RP can be decomposed into a convex combination of perfect matchings of agents to items such that the allocation made by each perfect matching is Pareto efficient. This is obvious if we choose the $n!$ perfect matchings corresponding to the $n!$ orderings of agents. Clearly, two different orderings may yield the same perfect matching, therefore the convex combination may be over fewer than $n!$ perfect matchings. Therefore, RP is ex post Pareto efficient. □

Interestingly, the ex-post Pareto efficiency of the priority mechanism doesn't lead to ordinal efficiency of Random Priority. Consider the following example:

Example 3.6. *Consider agents $A = \{A_1, A_2, A_3, A_4\}$, items $I = \{I_1, I_2, I_3, I_4\}$ and the preferences of the agents as follows:*

$$A_1 : I_1 \succ I_2 \succ I_3 \succ I_4$$
$$A_2 : I_1 \succ I_2 \succ I_3 \succ I_4$$
$$A_3 : I_2 \succ I_1 \succ I_3 \succ I_4$$
$$A_4 : I_2 \succ I_1 \succ I_3 \succ I_4$$

Then RP will return the following random allocation:

Agent / Item	I_1	I_2	I_3	I_4
A_1	$\frac{5}{12}$	$\frac{1}{12}$	$\frac{5}{12}$	$\frac{1}{12}$
A_2	$\frac{5}{12}$	$\frac{1}{12}$	$\frac{5}{12}$	$\frac{1}{12}$
A_3	$\frac{1}{12}$	$\frac{5}{12}$	$\frac{1}{12}$	$\frac{5}{12}$
A_4	$\frac{1}{12}$	$\frac{5}{12}$	$\frac{1}{12}$	$\frac{5}{12}$

Table 3.1: Random allocation produced by RP

However, it is stochastically dominated by the following random allocation:

Agent / Item	I_1	I_2	I_3	I_4
A_1	$\frac{1}{2}$	0	$\frac{1}{2}$	0
A_2	$\frac{1}{2}$	0	$\frac{1}{2}$	0
A_3	0	$\frac{1}{2}$	0	$\frac{1}{2}$
A_4	0	$\frac{1}{2}$	0	$\frac{1}{2}$

Table 3.2: Stochastically dominant allocation

Even worse, RP is not envy-free, despite the fact that the reason for generalizing from priority to RP was to introduce fairness. Consider the following example:

Example 3.7. *Consider agents $A = \{A_1, A_2, A_3\}$, items $I = \{I_1, I_2, I_3\}$ and the preferences of the agents as follows:*

$$A_1 : I_1 \succ I_2 \succ I_3$$
$$A_2 : I_2 \succ I_1 \succ I_3$$
$$A_3 : I_2 \succ I_3 \succ I_1$$

Then RP will return the following allocation:

Agent / Item	I_1	I_2	I_3
A_1	$\frac{5}{6}$	0	$\frac{1}{6}$
A_2	$\frac{1}{6}$	$\frac{3}{6}$	$\frac{2}{6}$
A_3	0	$\frac{3}{6}$	$\frac{3}{6}$

Table 3.3: Random allocation produced by RP

Since the total allocation of the two items I_1 and I_2 to agents A_1 and A_2 is $\frac{5}{6}$ and $\frac{4}{6}$, respectively, agent A_2's allocation does not stochastically dominate agent A_1's allocation.

While in the Priority mechanism, agents are assigned items sequentially based on a pre-determined order. In the Random Priority (RP) mechanism, this order is determined by a single random permutation of the agents. Each agent, following this random order, selects their most preferred available item.

It's easy to see that RP is efficient, as it does not need to evaluate multiple permutations — only one is randomly chosen and used for the entire process. However, it sacrifices certain aspects of fairness and ex-ante Pareto efficiency. While truthful, the randomness in the ordering may lead to outcomes that are not always fair or efficient.

For example, suppose we have three agents A_1, A_2, A_3 and three items I_1, I_2, I_3. Suppose it's known that A_1 and A_2 both prefer I_1 the most, while A_3 is indefferent between I_2 and I_3. In Random Priority, depending on the order, A_1 or A_2 might end up with I_1 and the other might end up with their second or third choice, leading to a significant difference in their utility.

If the mechanism were designed with knowledge of these preferences in advance, however, it could find a way of allocating items in such a way, to make both A_1 and A_2 better off on average, perhaps by giving them partial shares of I_1 and I_2, thereby increasing their expected utility.

This discussion, however, highlights the inherent trade-offs between fairness, truthfulness, and Pareto efficiency. Despite the appeal of RP for its strategy-proofness and potential fairness, it turns out that no mechanism can simultaneously satisfy all these properties. The search for a mechanism that is both fair and truthful and also ensures ex-ante Pareto efficiency is ultimately fruitless.

It turns out that, there is no "fair" mechanism that is strategy-proof and ex-ante Pareto efficient.

Theorem 3.8. (Theorem 2 in [1]) There is no mechanism that satisfies ordinal efficiency, thruthfulness and equal treatment of equals (i.e. agents with same preferences should receive the same random consumption) for more than three agents.

This impossibility forces designers to make critical choices, often prioritizing one or two properties at the expense of the others, depending on the specific needs and constraints of the application at hand.

Probabilistic Serial

With respect to the four properties studied above for RP, PS behaves in exactly the opposite manner. It is time-efficient, ordinally efficient, and envy-free but not truthful. The Probabilistic Serial (PS) mechanism was introduced by Bogomolnaia and Moulin in 2001 (see [1] as an alternative to Random Priority.

Suppose that each good can be divided into probability shares, each share representing a portion of the good. The PS (Probabilistic Serial) mechanism continuously allocates these shares so that over the course of an hour, each agent receives exactly one unit of probability share across all goods.

Initially, each agent receives the probability of their most preferred good. If several agents prefer the same good, the good is consumed faster, at a rate proportional to the number of agents. Once an agent's most preferred good is fully consumed, they move on to their next preferred good that still has shares available.

Since each agent has a complete ranking of all goods, by the end of the hour each agent will have received exactly one unit of allocation and all goods will be fully distributed. This continuous allocation process can be broken down into discrete steps by calculating the exact moments when each good is fully allocated. The final allocations are expressed in rational numbers, and the whole process is done in polynomial time.

At this point, we will see an example of how the PS mechanism works.

Example 3.9. *Let agents* $A = \{A_1, A_2, A_3\}$ *and items* $I = \{I_1, I_2, I_3\}$ *and the preferences of the agents as follows:*

$$A_1 : I_1 \succ I_2 \succ I_3$$
$$A_2 : I_1 \succ I_3 \succ I_2$$
$$A_3 : I_2 \succ I_1 \succ I_3$$

PS will return the following allocation.

Agent / Item	I_1	I_2	I_3
A_1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$
A_2	$\frac{1}{2}$	0	$\frac{1}{2}$
A_3	0	$\frac{3}{4}$	$\frac{1}{4}$

Table 3.4: Allocation produced by PS

Now, regarding truthfulness, fairness and efficiency, we gather the following results:

Lemma 3.10. The allocation computed by PS is envy-free.

Proof. At each point in the algorithm, each agent is obtaining probability share of her most favorite good. Therefore, at any time in the algorithm, agent i cannot prefer agent j's current allocation to her own. Hence PS is envy-free. ☐

For showing that the random allocation computed by PS is ordinally efficient, we will appeal to the following property of stochastic dominance which follows directly from its definition. Assume that x and y are two allocations made to agent i having equal total probability shares; let $t \leq 1$ be this total. Assume $x \succeq_i^{sd} y$. Let $\alpha < t$ and remove α amount of the least desirable probability shares from each of x and y to obtain x' and y', respectively. Then $x' \succeq_i^{sd} y'$.

Lemma 3.11. The random allocation computed by PS is ordinally efficient.

Proof. During the run of PS on the given instance, let $t_0 = 0, t_1, \ldots, t_m = 1$ be the times at which some agent exhausts the good she is currently being allocated. By induction on k, we will prove that at time t_k, the partial allocation computed by PS is ordinally efficient among all allocations which give t_k amount of probability shares to each agent.

The induction basis, for $k = 0$, is obvious since the empty allocation is vacuously ordinally efficient. Let A_k denote the allocation at time t_k and let A_k^i denote the allocation made to agent i under A_k. Assume the induction hypothesis, namely that the assertion holds for k, i.e., A_k is ordinally efficient.

For the induction step, we need to show that A_{k+1} is ordinally efficient. Suppose not and let it be stochastically dominated by random allocation P. Let $\alpha = t_{k+1} - t_k$. For each agent i, remove α amount of the least desirable probability shares from P_i to obtain P_i'. Since $P_i \succeq_i^{sd} A_{k+1}^i$, by the property stated above, $P_i' \succeq_i^{sd} A_k^i$. By the induction hypothesis, $A_k^i \succeq_i^{sd} P_i'$ as well. Therefore, $P_i' = A_k^i$. In the time period between t_k and t_{k+1}, each agent obtains α units of probability shares of her most preferred good remaining. Therefore, $A_{k+1}^i \succeq_i^{sd} P_i$, leading to a contradiction. ☐

To prove that a mechanism is strategyproof, we would need to show that x_i stochastically dominates x_i' where x and x' are the allocations when i reports preferences truthfully and misreports, respectively. The following example shows that this does not hold for PS.

Example 3.12. *Let's consider the agents with their preferences from the previous example and assume that agent A_3 misreports her preference list as $I_1 \succ I_2 \succ I_3$, then PS will return the following allocation.*

Agent / Item	I_1	I_2	I_3
A_1	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{6}$
A_2	$\frac{1}{3}$	0	$\frac{2}{3}$
A_3	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{6}$

Table 3.5: Allocation produced by PS when A_1 misreports

Therefore by lying, agent A_3 obtains $\frac{5}{6}$ units of her two most prefered items, instead of $\frac{3}{4}$ units.

Finally, we provide an example in which RP and PS are Pareto incomparable in the sense that different agents prefer different allocations.

Example 3.13. *Consider agents $A = \{A_1, A_2, A_3, A_4\}$, items $I = \{I_1, I_2, I_3, I_4\}$ and the preferences of the agents as follows:*

$$\begin{aligned}
A_1 &: I_1 \succ I_2 \succ I_3 \succ I_4 \\
A_2 &: I_1 \succ I_2 \succ I_4 \succ I_3 \\
A_3 &: I_2 \succ I_1 \succ I_3 \succ I_4 \\
A_4 &: I_3 \succ I_4 \succ I_1 \succ I_2
\end{aligned}$$

Then RP will return the following random allocation:

Agent / Item	I_1	I_2	I_3	I_4
A_1	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{12}$	$\frac{1}{4}$
A_2	$\frac{1}{2}$	$\frac{1}{6}$	0	$\frac{1}{3}$
A_3	0	$\frac{2}{3}$	$\frac{1}{12}$	$\frac{1}{4}$
A_4	0	0	$\frac{5}{6}$	$\frac{1}{6}$

Table 3.6: Random allocation produced by RP

PS will return the following allocation.

Agent / Item	I_1	I_2	I_3	I_4
A_1	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{9}$	$\frac{2}{9}$
A_2	$\frac{1}{2}$	$\frac{1}{6}$	0	$\frac{1}{3}$
A_3	0	$\frac{2}{3}$	$\frac{1}{9}$	$\frac{2}{9}$
A_4	0	0	$\frac{7}{9}$	$\frac{2}{9}$

Table 3.7: Allocation produced by PS

Agents A_1 and A_3 prefer the PS allocation, agent A_4 prefers the RP allocation and agent A_2 is indifferent.

Even worse, RP is not envy-free, despite the fact that the reason for generalizing from the Priority Mechanism to RP was to introduce fairness, as discussed in Example 3.7. Although the priority mechanism runs efficiently, RP would require exponential time to compute the ex-ante random allocation by evaluating all possible permutations of agents, making it impractical for all but very small values of n. However, if an integral matching of agents to goods is desired, then picking one ordering of agents at random, as in RP, is clearly time-efficient.

3.3 Summary

We summarize the main points of this chapter.

3.3.1 Probabilistic Serial (PS)

Ordinal Efficiency

- **PS is ordinally efficient.**
  - The PS mechanism simultaneously allocates fractions of goods based on agents' rankings, ensuring that no other allocation could make someone better off without making someone else worse off according to their preference order.

Ex-Ante Pareto Efficiency

- **PS achieves ex-ante Pareto efficiency.**
  - Before the randomization is resolved, the expected allocations under PS cannot be improved upon without harming at least one agent's expected utility.

Ex-Post Pareto Efficiency

- **PS does not guarantee ex-post Pareto efficiency.**
  - After the randomization, the realized allocation may not be Pareto efficient; there might exist another allocation that could make some agents better off without making others worse off.

Truthfulness

- **PS is not strategy-proof (not truthful).**
  - Agents may have an incentive to misreport their preferences to achieve a more favorable outcome.

Fairness (Equal Treatment of Equals)

- **PS satisfies ETE.**

3.3.2 Random Priority (RP)

Ordinal Efficiency

- **RP is not ordinally efficient.**
  - RP assigns a random order to agents who then pick their top available choices sequentially. This process can lead to allocations where a different feasible allocation would be unanimously preferred by all agents based on their rankings.

Ex-Ante Pareto Efficiency

- **RP does not guarantee ex-ante Pareto efficiency.**
  - Before the randomization, there may exist alternative expected allocations that could make some agents better off without harming others, based on expected utilities.

Ex-Post Pareto Efficiency

- **RP achieves ex-post Pareto efficiency.**
  - Once the random order is realized and allocations are made, no agent can be made better off without making another agent worse off in the realized outcome.

Truthfulness

- **RP is strategy-proof (truthful).**
  - It is in each agent's best interest to report their true preferences because misreporting cannot lead to a better outcome.

Fairness (Equal Treatment of Equals)

- **RP does not satisfy ETE**

Property	PS Mechanism	RP Mechanism
Ordinal Efficiency	Yes	No
Ex-Ante Pareto Efficiency	Yes	No
Ex-Post Pareto Efficiency	No	Yes
Truthful (Strategy-Proof)	No	Yes
Fairness (ETE)	Yes	No

3.2.4 Conclusion

RP and PS have their obvious advantages and drawbacks, as summarized in the previous table. In the following chapters we will see many, some would say beautiful, results about these two mechanisms. We will try to understand agents' incentives to manipulate PS, better understand the fairness issues of RP, see experiments on how these mechanisms behave in real life, while also proposing a new measure of inefficiency and showing that sometimes letting agents lie might actually be beneficial! Brace yourselves for the journey ahead! :)

CHAPTER 4

APPROXIMATE SOCIAL WELFARE BOUNDS

This chapter is based on the work of Filos-Ratsikas et al. in [10], with some adjustments and clarifications for ease of reading.

4.1 Synopsis

In this chapter, we follow the work of Filos-Ratsikas et al. in [10] and study the problem of maximizing Social Welfare in One-Sided Matching mechanisms, where agents have unrestricted preferences over items. Each agent assigns a value to items, representing their preference (these values can be interpreted as von Neumann-Morgenstern utilities).

Our concentration focuses on truthful mechanisms. A typical example is the Random Priority mechanism, in which agents are chosen randomly and serially choose their most prefered items. As we saw in Chapter 3, while Random Priority is truthful and ex-post Pareto efficient, it doen't satisfy ex-ante Pareto efficiency - there exists a matching that makes all agents at least as satisfied, with one strictly better off in expectation. Economic inefficiencies in turn translate to a lower Social Welfare, which measures the total satisfaction of agents. Previous research has focused on mechanisms that satisfy some efficiency criteria, rather than mechanisms that try to maximise Social Welfare while remaining truthful.

In this paper, the authors study Random Priority and prove its Approximation Ratio to be $\Theta(\frac{1}{\sqrt{n}})$ - best possible among all truthful-in-expectation mechanisms for the problem. Furthermore, they show that all ordinal mechanisms have an upper bound of $O(\frac{1}{\sqrt{n}})$, and hence making Random Priority is the optimal truthful and ordinal mechanism for this problem.

4.2 Introduction

Mechanism design without money focuses on creating systems that allocate resources or make decisions based on the preferences of participants, in order to ensure that they are truthful. The great challenge is to design mechanisms that not only encourage truthful behaviour, but also produce outcomes that are close to the optimal solution.

The **Approximation Ratio**, as seen in Chapter 2. , originates from approximation algorithms, designed to solve NP-hard problems approximately. In this context, algorithms aim to find solutions close to the optimum within a region. This concept extends to game theory and mechanism design, where the goal often includes developing mechanisms that perform effectively despite agents' selfish behavior.

In "Approximate Mechanism Design without Money", Procaccia and Tennenholtz ([26]) explore the creation of strategyproof mechanisms for various optimization problems. They focus on achieving favorable Approximation Ratios while ensuring participants have no incentive to misreport their preferences.

The authors' primary contribution is demonstrating that approximation can be a viable approach to achieve strategyproofness without monetary incentives. They show that it is possible to design mechanisms that ensure truthful reporting and produce outcomes reasonably close to the optimal.

Most prior work has concentrated on designing mechanisms that achieve efficiency for different Pareto efficiency notions, with less focus on whether truthful mechanisms achieve high levels of Social Welfare. This question is crucial, as the Approximation Ratio provides a systematic way to compare mechanisms or demonstrate their limitations.

In this Chapter, we investigate the behavior of the Random Priority mechanism under the lens of Social Welfare, defined as the sum of agents' valuations for the items they receive in the mechanism's outcome. While, the problem of SW maximazation of RP has been studied before for other variants of SW (see [27]) there is strong evidence that consider the SW with underlying cardinal preferences of the agents, is a rather natural and widely embraced, perspective.

We assume the valuation functions are unit-range, meaning agents are not penalized for liking more than one item. The results can, however, be extended to the unit-sum setting. We evaluate the performance of a mechanism by its Approximation Ratio—the worst-case ratio between the (expected) Social Welfare of the mechanism and the welfare of the optimal allocation, across all valuation profiles. The reader can see **Chapter 2.** for a reminding of the definition of the Approximation Ratio.

4.3 Anonymity & Random Priority

We begin our analysis, by concentrating in the anonymity of the mechanisms. The following lemma states that upper bounds on the Approximation Ratio of mechanisms can be proven by focusing on anonymous mechanisms. An anonymous mechanism, like RP, treats all agents equally, meaning the outcome depends solely on their preferences, not their identities. When considering upper bounds on Approximation Ratio, we are typically interested in the worst-case scenario. Since anonymous mechanisms treat all agents the same, they tend to naturally cover a wide range of possible scenarios, including the worst-case. This means that if we can prove a bound for an anonymous mechanism, that bound is valid for non-anonymous mechanisms as well.

Lemma 4.1. For any mechanism M, there exists an anonymous mechanism M' such that $ar(M') \geq ar(M)$. In addition, if M is truthful (or truthful-in-expectation) then it holds that M' is truthful (or truthful-in-expection) too.

Proof. Let M' be a mechanism that takes a valuation profile **u** as input and performs a uniformly at random permutation on the set of agents, before applying M. By randomizing the order of the agents uniformly, M' ensures that all agents have an equal chance of being in any position in the order. Thus, M' is anonymous.

Furthermore, since M' only permutes the agents randomly before applying M, the set of possible outputs of M' is the same as the set of possible outputs of M, just with permuted agent identities. Thus, the Social Welfare achieved by M' is the same as that of M for any permutation of the agents. Therefore, the Approximation Ratio of M' cannot be worse than that of M.

For the same reason, if M is truthful (or truthful-in-expectation), the expected utility for any agent reporting truthfully is at least as high as if they misreport, given the random permutation is independent of the reports.

Since the permutation step in M' is independent of the agents' reports and M is truthful-in-expectation, M' maintains the truthfulness. The random permutation doesn't provide any additional advantage or disadvantage to any agent in terms of strategizing their reports, so M' remains truthful-in-expectation. □

We know that Random Priority fixes an ordering of the agents uniformly at random and then lets them pick their most preferred items from the set of available items based on this ordering. Random Priority is truthful-in-expectation, ordinal, anonymous and neutral - the final allocation is invariant to the relabeling or reordering of the items. The following lemma shows that the RP mechanism, by randomly determining the order in which agents select items, is capable of producing any allocation, including the optimal.

Lemma 4.2. For any valuation profile **u**, the optimal allocation on **u** is a possible outcome of Random Priority.

Proof. Assume that in the optimal allocation, no agent is matched with their most preferred item. Then there must be a cycle of agents $i_1, ..., i_k$ such that:

- Agent i_1 is matched with an item that is most preferred by agent i_k.
- Agent i_2 is matched with an item that is most preferred by agent i_1.
- Agent i_k is matched with an item that is most preferred by agent i_{k-1}.

By swapping items along this cycle, each agent would receive their most preferred item. This makes all agents in the cycle better off.

Since the initial allocation was assumed to be optimal, it cannot be improved. The existence of such a cycle implies that the initial allocation was not optimal, which leads to a contradiction.

Thus, there must exist at least one agent j in the optimal allocation who is matched with their most preferred item j. Consider the valuation profile **u**. Remove the agent j and their most preferred item from the profile to form a smaller valuation profile **u**$'$.

The optimal allocation for **u**$'$ remains optimal for **u** minus the removed agent and item. By applying the same argument recursively, we can show that for the reduced profile **u**$'$, there must be at least one agent matched with their most preferred item in the optimal allocation.

By continuously applying this argument, starting from any profile **u** and reducing it step by step, we demonstrate that at every step, there exists at least one agent who is matched with their most preferred item. Thus, this holds for the original profile **u** as well. □

We consider, for our analysis, a special class of valuation functions C_ϵ we refer to as *quasi-combinatorial valuation functions* were it captures all valuation functions where the valuations of each agent for every item are ϵ-close to 1 or close to 0.

$$C_\epsilon = \{u \in V | u(M) \subset [0, \epsilon) \cup (1-\epsilon, 1]\} \subseteq V$$

where $u(M)$ is the image of the valuation function u. The reason we are interested in C_ϵ is that or any valuation profile $\mathbf{u}$, there exists a valuation profile $\mathbf{u}'$ in C_ϵ^n that achieves an Approximation Ratio at least as good as $\mathbf{u}$.

Lemma 4.3. Let M be an ordinal, anonymous and neutral randomized mechanism for unit-range representation and let $\epsilon > 0$. Then

$$ar(M) = \inf_{\mathbf{u} \in C_\epsilon^n} \frac{E[\sum_{i=1}^n u_i(M(\mathbf{u}))]}{\sum_{i=1}^n u_i(\mu_i^*)}.$$

Proof. Since M is anonymous and neutral, meaning that it's outcome does not depend on the identities of the agents nor the items. We can assume that the optimal matching is μ^* where μ^* is the matching with $\mu_i^* = i$ for every agent i.

We define, for any valuation profile $\mathbf{u}$, $g(\mathbf{u})$ to be:

$$g(\mathbf{u}) = \frac{E[\sum_{i=1}^n u_i(M(\mathbf{u}))]}{\sum_{i=1}^n u_i(\mu_i^*)}.$$

Hence, the Approximation Ratio can be written as $ar(M) = \inf_{\mathbf{u} \in V^n} g(\mathbf{u})$.

Since, $C_\epsilon^n \subseteq V^n$. The lemma is equivalent with the following claim:

For all valuation profiles $\mathbf{u} \in V^n$, there exists a $\mathbf{u}' \in C_\epsilon^n$ such that $g(\mathbf{u}') \leq g(\mathbf{u})$

This will be done via induction to the number of valuations within the interval $[\epsilon, 1-\epsilon]$.

- **Base Case:** When $\sum_{i=1}^n \#\{u_i(M) \cap [\epsilon, 1-\epsilon]\} = 0$, all valuations u_i are either below ϵ or above $1-\epsilon$. Thus, $u_i \in C_\epsilon^n$ for all agents, so $\mathbf{u} = \mathbf{u}' \in C_\epsilon^n$ and trivially, $g(\mathbf{u}) = g(\mathbf{u}')$
- **Inductive Hypothesis:** Assume that the claim is true for any profile $\mathbf{u}$ such that $\sum_{i=1}^n \#\{u_i(M) \cap [\epsilon, 1-\epsilon]\} \leq k$.
- **Inductive Step:** Consider a profile $\mathbf{u}$ such that $\sum_{i=1}^n \#\{u_i(M) \cap [\epsilon, 1-\epsilon]\} > k$. Then, there exists an agent such that $\#\{u_i(M) \cap [\epsilon, 1-\epsilon]\} > 0$. Then, there exist l and r in $[\epsilon, 1-\epsilon]$ that both are in $u_i(M)$.
  Let, w.l.o.g. $l \leq r$ and $\bar{l}$ be the largest value in $[0, \epsilon)$ within $u_i(M)$ and $\bar{r}$ the smallest value in $(1-\epsilon, 1]$ withing $u_i(M)$. Both of this numbers exist since $\{0, 1\} \subseteq u_i(M)$. Let, $\tilde{l} = \frac{\bar{l}+\epsilon}{2}$ and $\tilde{r} = \frac{\bar{r}+1-\epsilon}{2}$.

For $x \in [\tilde{l} - l, \tilde{r} - r]$, define a modiefied valuation u_i^x by increasing (or decreasing if $x < 0$) $u_i(j)$ by x for items j with valuations in $[\epsilon, 1-\epsilon]$.

We will prove the claim by induction in $\sum_{i=1}^n \#\{u_i(M) \cap [\epsilon, 1-\epsilon]\}$. Let $(u_i^x, \mathbf{u}_{-i})$ be the valuation profile where all agents have the same valuation functions as in $\mathbf{u}$ except for agent i, who has valuation function u_i^x.

Define $f(x) = g((u_i^x, \mathbf{u}_{-i}))$. Since our mechanism M is ordinal, by definition of g, f can only be a fractional linear function, defined in the whole interval $[\tilde{l}-l, \tilde{r}-r]$. Thus, it can only be monotonically increasing, decreasing or constant. Let $\tilde{\mathbf{u}}$ be:

- $\tilde{\mathbf{u}} = (u_i^{\tilde{l}-l}, \mathbf{u}_{-i})$ if f is monotonically increasing.
- $\tilde{\mathbf{u}} = (u_i^{\tilde{r}-r}, \mathbf{u}_{-i})$, otherwise.

Then $g(\tilde{\mathbf{u}}) \leq g(\mathbf{u})$ and $\sum_{i=1}^n \#\{\tilde{u}_i(M) \cap [\epsilon, 1-\epsilon]\} < \sum_{i=1}^n \#\{u_i(M) \cap [\epsilon, 1-\epsilon]\} \leq k$.

Hence, by applying the induction hypothesis on $\tilde{\mathbf{u}}$, there exists a profile $\mathbf{u}'' \in C_\epsilon^n$ such that $g(\mathbf{u}'') \leq g(\tilde{\mathbf{u}}) \leq g(\mathbf{u})$ meaning that $\mathbf{u}''$ provides an equal or better Approximation Ratio and this completes the proof. □

The lemma implies that the mechanism M behaves optimally over the entire set V^n can be captured by its performance on C_ϵ^n, formalizing the intuition that because the mechanism is ordinal, the worst-case Approximation Ratio is encountered on extreme valuation profiles. Critical for proving the following theorem, in the next page.

4.4 Lower Bound

In this section, we prove the lower bound of the Approximation Ratio of Random Priority.

Lemma 4.4. For unit-range representation, the Approximation Ratio of Random Priority is $\Omega\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Since RP is a neutral and anonymous mechanism, because of the previous lemma 4.3 it suffices to establish a lower bound for quasi-combinatorial valuation profiles.

Let $w^*(\mathbf{u})$ denote the Social Welfare of the optimal matching (maximum weight matching) on valuation profile $\mathbf{u}$, it is true that $w^*(\mathbf{u}) \leq n$ and there exists $k \in \mathbf{N}$ and $\epsilon \leq \frac{1}{n^3}$ such that

$$|k - w^*(\mathbf{u})| \leq n \cdot \epsilon \leq \frac{1}{n^2},$$

Random Priority can achieve an expected Social Welfare of 1 trivially, since for any permutation the first agent will be matched to her most preferred item, we can assume that $k \geq \sqrt{n}$, otherwise , if $k < \sqrt{n}$

$$ar(RP) = \inf_{\mathbf{u} \in C_\epsilon^n} \frac{E[\sum_{i=1}^n u_i(M(\mathbf{u}))]}{w^*(\mathbf{u})} \geq \frac{1}{\sqrt{n}}$$

and we are done.

The optimal matching μ^* assigns k items to agents with $u_i(\mu_i) \in (1-\epsilon, 1]$. Without loss of generality, due to RPs anonymity and neutrality, we can assume that these agents are $\{1, \dots, k\}$ and that each agent j is assigned to the item with their own index in this optimal matching, $\mu_j^* = j$. Thus $u_{j \in \{1,\dots,k\}}(j) \in (1-\epsilon, 1]$ and $u_{j \in \{k+1,\dots,n}\}(j) \in [0, \epsilon)$.

Let, l denote a (current) run of RP. Let $l \in \{0, \dots, n-1\}$ be any of the n rounds. We odefine the following sets:

$$\begin{aligned}
U_l &= \{j \in \{1, \dots, n\} : \text{ agent } j \text{ has not been selected before round } l\} \\
G_l &= \{j \in U_l : u_j(j) \in (1-\epsilon, 1] \text{ and item } j \text{ is still unmatched}\} \\
B_l &= \{j \in U_l : u_j(j) \in [0, \epsilon) \text{ or item } j \text{ has already been matched to some agent}\}
\end{aligned}$$

- U_l : contains all the agents still waiting to be matched by round l of the RP.
- G_l : the "good" subset of U_l that includes all agents who have high valuations (in $(1-\epsilon, 1]$) for their assigned items in the optimal matching and which are still available.
- B_l : the "bad" subset of U_l that includes all agents who have low valuations (in $[0, \epsilon)$) for their assigned items in the optimal matching and which are still available.

The probability that an agent $i \in G_l$ is picked in round l of RP is $\frac{|G_l|}{|G_l|+|B_l|}$, where the probability that an agent $i \in B_l$ is picked is $\frac{|B_l|}{|G_l|+|B_l|}$. When an agent in G_l is picked, they contribute at least $1-\epsilon$ to the Social Welfare, while when an agent from B_l is picked, they contribute less than ϵ.

Thus, round l's expected contribution is $\frac{|G_l|}{|G_l|+|B_l|} \cdot (1-\epsilon) + \frac{|B_l|}{|G_l|+|B_l|} \cdot \epsilon = \frac{|G_l|}{|G_l|+|B_l|} - \epsilon \cdot \frac{|G_l|-|B_l|}{|G_l|+|B_l|} \geq \frac{|G_l|}{|G_l|+|B_l|} - \epsilon$.

We will now see how $|G_l|$ and $|B_l|$ change in each round l. Suppose that an agent i from G_l gets picked and matched with item j.

- If $j \neq i$ and agent $j \in G_l$, then in round $l+1$: $|G_{l+1}| = |G_l| - 2$ and $B_{l+1} = |B_l| + 1$. This is because agent j, lost its optimal item and now moves to $|B_{l+1}|$.
- If $j = i$ or if agent $j \in B_l$, then in round $l+1$: $|G_{l+1}| = |G_l| - 1$ and $B_{l+1} = |B_l|$.

In any case, $|G_{l+1}| \geq |G_l| + 2$ and $|B_{l+1}| \leq |B_l| + 1$. Intuitively, the selected agent might take away some item from a "good" agent and downgrade it into a bad agent.
In summary, in each round l of RP, we can assume the size of B_l increases by at most 1 and the size of G_l decreases by at most 2. Assuming that $|G_0| = k$, $|B_0| = n - k$ and that $|G_l| > 0$ for $l \leq \lfloor k/2 \rfloor$, we get:

$$
\begin{aligned}
E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right] &\geq \sum_{l=0}^{n}\left(\frac{|G_l|}{|G_l|+|B_l|}-\epsilon\right)\\
&\geq \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\left(\frac{|G_l|}{|G_l|+|B_l|}-\epsilon\right)\\
&\geq \left(\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\frac{k-2l}{n-k+l+k-2l}\right)-n\epsilon\\
&= \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\frac{k-2l}{n-l}-n\epsilon
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right]}{w^*(\mathbf{u})} &\geq \frac{\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\frac{k-2l}{n-l}-n\epsilon}{k+\frac{1}{n^2}} \geq \frac{\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\frac{k-2l}{n-l}-n\epsilon}{2k}\\
&= \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\frac{1-\frac{2l}{k}}{2(n-l)}-\frac{n\epsilon}{2k} > \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor}\frac{1-\frac{2l}{k}}{2n}-\frac{n\epsilon}{2k} \geq \frac{k-11}{8n}-\frac{n\epsilon}{2k}.
\end{aligned}
$$

Since, $ar(RP) = \inf_{\mathbf{u}\in C_\epsilon^n} \frac{E[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))]}{w^*(\mathbf{u})}$, we need to minimize $\frac{k-11}{8n} - \frac{n\epsilon}{2k}$. Since, $k \geq \sqrt{n}$, the bound is minimum when $k = \sqrt{n}$. Thus,

$$
ar(RP) = \inf_{\mathbf{u}\in C_\epsilon^n} \frac{E[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))]}{w^*(\mathbf{u})} \geq \frac{\sqrt{n}-11}{8n} - \frac{n\epsilon}{2\sqrt{n}}.
$$

We can choose ϵ so that the Approximation Ratio is at least $\frac{1}{20\sqrt{n}}$ for $n \geq 400$ and for $n \leq 400$, the bound holds trivially since Random Priority matches at least one agent with its most preferred item. $\square$

4.5 Upper Bound

In this section, we prove an upper bound of the Approximation Ratio of *any* ordinal mechanism for the problem, as well us an upper bound of any truthful-in-expectation mechanism (ordinal or not).

Lemma 4.5. Let M be any ordinal mechanism for unit-range representation. The Approximation Ratio of M is $O\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Let n be the number of agents, and k an integer such that $k = \lfloor\sqrt{n}\rfloor$. Consider a valuation profile $\mathbf{u} = (u_1, u_2, \ldots, u_n)$ where:

$$u_i(j) = \begin{cases} 1 - \frac{j-1}{n}, & \text{if } 1 \le j \le i, \\ \frac{n-j}{n^2}, & \text{otherwise,} \end{cases} \quad \forall i \in \{1, \ldots, k\}$$

$$u_i(j) = \begin{cases} 1, & \text{if } j = 1, \\ \frac{n-j}{n^2}, & \text{otherwise,} \end{cases} \quad \forall i \in \{k+1, \ldots, n\}$$

Without loss of generality, by Lemma 4.1 we can assume that M is anonymous. Also, by the valuation profile, we can see that for any agent i, the valuation $u_i(j)$ assigned to item i is greater that $u_i(j')$ for any $j < j'$.

Hence, $\mathbf{u}$ is *ordered* and thus M on $\mathbf{u}$ is anonymous and ordinal. This implies that every agent has an equal probability of being matched with any given item. Formally, the probability that an agent and therefore, no agent has an advantage over another, and the mechanism's output is entirely based on randomness with respect to the agents' preferences. Meaning that the only fair outcome for M is a uniformly random matching.

From the above mentioned, the expected welfare of the mechanism on valuation profile $\mathbf{u}$ will be

$$E[\sum_{i=1}^{n} u_i(M(\mathbf{u}))] = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{n} u_i(j) \le \frac{1}{n}\left[\sum_{i=1}^{k}\left(i + \frac{n-i}{n}\right) + \sum_{i=k+1}^{n}\left(1 + \frac{n-1}{n}\right)\right] \le 4 + \frac{1}{2\sqrt{n}} \le 5,$$

where in the above expression, we upper bound each term $\frac{n-j}{n^2}$ by $\frac{1}{n}$ and each term $1 - \frac{j}{n}$ by 1.

On the other hand, the Social Welfare of the maximum weight matching is

$$w^*(\mathbf{u}) = \sum_{i=1}^{k}\left(1 - \frac{i-1}{n}\right) + \sum_{i=k+1}^{n}\frac{n-i}{n^2} \ge \sum_{i=1}^{k}\left(1 - \frac{i-1}{n}\right) \ge k - 1 \ge \frac{\sqrt{n}}{4}.$$

Where the final inequality holds for $n \ge 4$, the Approximation Ratio then is at most $\frac{20}{\sqrt{n}}$ for $n \ge 4$, and the bound holds trivially for $n < 4$. $\square$

Lemma 4.6. Let M be a truthful-in-expectation mechanism for unit-range representation. Then $ar(M) = O\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Using Lemma 4.1, we can once again assume that M is anonymous. Consider $k \geq 2$ and let $\mathbf{u} = (u_1, u_2, \ldots, u_n)$ be the valuation profile defined as follows:

For each $i \in \{1, \ldots, k+1\}$:

$$u_i(j) = \begin{cases} 1, & \text{if } j = i, \\ \frac{2}{k} - \frac{j}{n}, & \text{if } 1 \leq j \leq k+1 \text{ and } j \neq i, \\ \frac{n-j}{n^2}, & \text{otherwise.} \end{cases}$$

For each $i \in \{k+2, \ldots, n\}$:

$$u_i(j) = \begin{cases} 1, & \text{if } j = 1, \\ \frac{2}{k} - \frac{j}{n}, & \text{if } 2 \leq j \leq k+1, \\ \frac{n-j}{n^2}, & \text{otherwise.} \end{cases}$$

Now let $\mathbf{u^i} = (u_i', u_{-i})$ for all $i = 2, \ldots, k+1$, be the valuation profile where all agents besides i have the same valuations as in $\mathbf{u}$ and $u_i' = u_{k+2}$. Notice that when agent i in valuation profile $\mathbf{u^i}$, reports u_i instead of u_i', the resulting valuation profile is $\mathbf{u}$. Since M is anonymous and $u_i' = u_1 = u_{k+2} = \ldots = u_n$, agent i receives at most a uniform lottery among these agents on valuation profile $\mathbf{u^i}$. Therefore, it follows that

$$\begin{aligned} E[u_i'(M(\mathbf{u^i}))] &\leq \frac{1}{n-k+1} + \sum_{j=2}^{k+1} \frac{1}{n-k+1}\left(\frac{2}{k} - \frac{j}{n}\right) + \sum_{j=k+2}^{n} \frac{1}{n-k+1} \cdot \frac{n-j}{n^2} \\ &\leq \frac{4}{n-k+1} \end{aligned}$$

For $i = 2, \ldots, k+1$, consider the valuation profile $\mathbf{u^i} = (u_i', u_{-i})$, where all agents except agent i have the same valuations as in the original profile $\mathbf{u}$, and the valuation function u_i' for agent i is identical to that of agent $k+2$. In this setting, when agent i, under the valuation profile $\mathbf{u^i}$, truthfully reports their original valuation u_i instead of misreporting u_i', the resulting valuation profile will be exactly $\mathbf{u}$.

Now, since the mechanism M is anonymous, it treats all agents equally, irrespective of their identities. The anonymity of the mechanism implies that the allocation probabilities depend only on the reported preferences and not on the specific identities of the agents. Given that $u_i' = u_1 = u_{k+2} = \ldots = u_n$, agent i is indistinguishable from agents $1, k+2$, and up to agent n in terms of their reported preferences. Therefore, in the valuation profile $\mathbf{u^i}$, the mechanism cannot favor agent i over any of these other agents who share the same reported valuation function u_i'.

As a result, agent i will receive an item through a process that is equivalent to a uniform lottery among all these agents with identical valuation functions. In other words, agent i has no better chance of obtaining a more preferred item than any of the other agents with the same reported preferences. Consequently, the expected outcome for agent i in this scenario is at most as favorable as if they were part of a uniform random draw among these agents.

This uniformity in treatment ensures fairness but also limits the influence of any individual agent on the outcome when their reported valuation matches that of a large group. Thus, it holds that:

$$\mathbb{E}[u_i'(M(\mathbf{u^i})_i)] \leq \frac{1}{n-k+1} \sum_{j=1}^{n-k+1} u_i'(M(\mathbf{u^i})_j)$$

where the expectation reflects the uniform lottery's outcome, ensuring that agent i cannot improve their expected utility by deviating from reporting u_i' truthfully.

For all $i = 2, \ldots, k+1$, let p_i be the probability that $(\mathbf{u}) = i$. Then, it holds that

$$E[u_i'(M(\mathbf{u}))] \geq p_i\left(\frac{2}{k} - \frac{i}{n}\right) \geq p_i\left(\frac{2}{k} - \frac{k+1}{n}\right)$$

Hence,

$$p_i\left(\frac{2}{k} - \frac{k+1}{n}\right) \leq E[u_i'(M(\mathbf{u}))] \leq E[u_i'(M(\mathbf{u^i}))] \leq \frac{4}{n-k+1}$$

$$\text{Thus, } p_i \leq \frac{4}{n-k+1} \cdot \frac{kn}{2n-k(k+1)} \leq \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$$

Let $p = \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$. Our goal is to establish an upper bound on the expected Social Welfare achieved by the mechanism M on the valuation profile $\mathbf{u}$.

First, consider the contribution to Social Welfare from item $j = 1$. This contribution is trivially upper bounded by 1. Similarly, for each item $j = k+2, \ldots, n$, the contribution to Social Welfare is upper bounded by $\frac{1}{n}$. Thus, the total contribution from item $j = 1$ and items $j = k+2, \ldots, n$ is collectively upper bounded by 2.

Next, we examine the contribution to Social Welfare from items $j = 2, \ldots, k+1$. Define the random variables

$$X_j = \begin{cases} 1, & \text{if } M(\mathbf{u})_j = j, \\ \frac{2}{k} - \frac{j}{n}, & \text{otherwise.} \end{cases}$$

The total contribution from items $j = 2, \ldots, k+1$ is then given by $\sum_{j=2}^{k+1} X_j$. Therefore, we have:

$$\mathbb{E}\left[\sum_{j=2}^{k+1} X_j\right] = \sum_{j=2}^{k+1} \mathbb{E}[X_j] \leq \sum_{j=2}^{k+1} \left(p + \frac{2}{k} - \frac{j}{n}\right) \leq kp + 2.$$

In summary, the expected Social Welfare of the mechanism M is at most $4 + pk$, while the Social Welfare of the optimal matching is $k + 1 + \sum_{i=k+2}^{n} \frac{n-i}{n^2}$, which is at least k. Given that $p = \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$, the Approximation Ratio of M is:

$$ar(M) \leq \frac{4 + pk}{k} = \frac{4}{k} + \frac{4}{n-k} \cdot \frac{kn}{2n - (k+1)^2}.$$

Let $k = \lfloor \sqrt{n} \rfloor - 1$ and note that $\sqrt{n} - 2 \leq k \leq \sqrt{n} - 1$. Then,

$$\begin{aligned} ar(M) &\leq \frac{4}{k} + \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2} \leq \frac{4}{\sqrt{n}-2} + \frac{4}{n-\sqrt{n}+1} \cdot \frac{(\sqrt{n}-1)n}{2n-(\sqrt{n})^2} \\ &\leq \frac{4}{\sqrt{n}-2} + \frac{4}{\sqrt{n}} \leq \frac{12}{\sqrt{n}} + \frac{4}{\sqrt{n}} = \frac{16}{\sqrt{n}}. \end{aligned}$$

The last inequality holds for $n \geq 9$, and for $n < 9$, the bound holds trivially. This completes the proof. $\square$

We are now prepared to present the central result of this work:

Theorem 4.7. The approximation ratio of Random Priority is $\Theta\left(\frac{1}{\sqrt{n}}\right)$. Moreover, Random Priority is asymptotically the best truthful-in-expectation mechanism as well as the best ordinal (not necessarily truthful-in-expectation) mechanism for this problem.

Proof. By Lemma 4.4 and Lemma 4.5, the approximation ratio of RP is $\Theta\left(\frac{1}{\sqrt{n}}\right)$. Lemma 4.5 shows that no truthful-in-expectation mechanism can outperform RP in the asymptotic limit. $\square$

4.6 Beyond Unit-range Representation & Extensions

4.6.1 Unit-sum valuation functions

In this section we, prove Theorem 4.5 for the unit-sum representation setting, by the following three lemmas.

As seen in Chapter 2. , in the unit-sum representation, the sum of an agent's valuations for all available items is normalized to 1. Formally, if an agent i has a valuation function $u_i : M \to \mathbb{R}_+$, where M is the set of items, the sum of the valuations over all items satisfies:

$$\sum_{j \in M} u_i(j) = 1$$

Additionally, it is required that $u_i(j) \geq 0$ for all items j. This ensures that the agent's total preference weight is evenly distributed across all items.

Lower Bound

Lemma 4.8. For unit-sum representation, the Approximation Ratio of RP is $\Omega\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Let n denote the number of agents and consider a valuation profile $\mathbf{u}$ such that for each agent i, the sum of their valuations for all items equals 1 and let c be the constant in the bound from Lemma 4.4. Without loss of generality, suppose first that $w^*(\mathbf{u}) < \frac{4\sqrt{n}}{c}$.

We see that Random Priority guarantees an expected Social Welfare of 1. First, pick an agent i and notice that in the l^{th} - round of RP,the probability that the agent gets picked $\frac{l}{n}$, hence the probability of the agent getting one of its l-most preferred items is at least $\frac{l}{n}$.

Let u_i^l be agent i's valuation for its l'th most preferred item; agent i's expected utility for the 1^{st}- round is then at least $\frac{u_i^1}{n}$. For the 2^{nd} - round, in the worst case, agent i's most preferred item has already been matched to a different agent and so the expected utility of the agent for the first two rounds is at least $\frac{u_i^1}{n} + \frac{u_i^2}{n}$. By the same argument, agent i's expected utility after n rounds is at least $\sum_{i=1}^{n} \frac{u_i^l}{n} = \frac{1}{n}$. Since this holds for each of the n agents, the expected Social Welfare is at least 1.

If $w^*(\mathbf{u}) \geq \frac{4\sqrt{n}}{c}$ it suffices to transform the valuation profile $\mathbf{u}$ into a unit-range valuation profile $\mathbf{u}''$. By Lemma 4.2, the optimal allocation can be achieved by a run of Random Priority, so we know that in the optimal allocation at most 1 agent will be matched with its least preferred item. Now consider the valuation profile $\mathbf{u}'$ where each agent i's valuation for its least preferred item is set to 0 and the rest of the valuations are as in $\mathbf{u}$. Since the ordinal preferences of agents remain unchanged, Random Priority will perform no better on $\mathbf{u}'$ than on $\mathbf{u}$ performs worse on this valuation profile and because of Lemma 4.2, $w^*(\mathbf{u}') \geq w^*(\mathbf{u}) - \frac{1}{n}$.

Next, consider the valuation profile $\mathbf{u}''$, defined as:

$$\mathbf{u}'' = \begin{pmatrix} \mathbf{u}' & \mathbf{1} \\ \mathbf{o}^T & 1 \end{pmatrix}$$

where $\mathbf{o} = (\mathbf{o}_j)_{j \in [n]} \in \mathbf{R}^n$ is a vector where each is given by $\mathbf{o}_j = \frac{j-1}{n^5}$.

This, results in a valuation profile $\mathbf{u}''$ has $n+1$ agents and items. Agents $1, \ldots, n$ retain the same valuations for items $1, \ldots, n$ as in $\mathbf{u}'$. Additionally, every agent assigns a valuation of 1 for item $n+1$ and agent $n+1$ only has a significant valuation for item $n+1$.

Our constructed $\mathbf{u}''$ is a unit-range valuation profile.

Furthermore, $w^*(\mathbf{u}'') \geq w^*(\mathbf{u}')+1$ and $E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right] \geq E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}'))\right] \geq E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}''))\right] - 2$ and $w^*(\mathbf{u}) \leq w^*(\mathbf{u}') + \frac{1}{n} \leq w^*(\mathbf{u}'') + \frac{1}{n} - 1 \leq w^*(\mathbf{u}'')$

Hence,

$$\begin{aligned}
\frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right]}{w^*(\mathbf{u})} &\geq \frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}'))\right]}{w^*(\mathbf{u}') + \frac{1}{n}} \geq \frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}''))\right] - 2}{w^*(\mathbf{u}'') + \frac{1}{n} - 1} \\
&\geq \frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}''))\right]}{w^*(\mathbf{u}'')} - \frac{2}{w^*(\mathbf{u}'')} \geq \frac{c}{\sqrt{n}} - \frac{2}{w^*(\mathbf{u})} \\
&\geq \frac{c}{\sqrt{n}} - \frac{2}{\frac{4\sqrt{n}}{c}} = \frac{c}{\sqrt{n}} - \frac{c}{2\sqrt{n}} = \frac{c}{2\sqrt{n}}.
\end{aligned}$$

Thus,

$$ar(RP) = \inf_{\mathbf{u} \in C_\epsilon^n} \frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right]}{w^*(\mathbf{u})} = \Omega(\frac{1}{\sqrt{n}})$$

□

Upper Bounds

Lemma 4.9. Let M be an ordinal mechanism for unit-sum representation. The Approximation Ratio of M is $O\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Let n be the number of agents, such that, without loss of generality, $\sqrt{n} = k \in \mathbb{N}$. Since, we are trying to find an upper bound for the Approximation Ratio, from 4.1 we can assume, without loss of generality that M is anonymous. Consider the following valuation profile **u**, for all $i \in \{1, \ldots, \sqrt{n}\}$:

$$u_i(j) = \begin{cases} 1 - \sum_{j \neq i} u_i(j), & \text{if } j = i, j \leq \sqrt{n} \\ \frac{n-j}{10n^5}, & \text{otherwise} \end{cases}$$

$$u_{i+l\cdot\sqrt{n}}(j) = \begin{cases} 1 - \sum_{j \neq i} u_i(j), & \text{if } j = i, j \leq \sqrt{n} \\ \frac{1}{\sqrt{n}} - \frac{j}{10n^2}, & \text{if } j \neq i, j \leq \sqrt{n} \\ \frac{n-j}{10n^5}, & \text{otherwise} \end{cases} \quad, \ l \in \{1, \ldots, \sqrt{n} - 1\}$$

More intuitevely, the valuation profile **u** is constructed so that for each $i \in \{1, \ldots, \sqrt{n}\}$, the valuation function of agent i, induces the same ordering as the valuation function of agent $i + l \cdot \sqrt{n}$ for any $l \in \{1, \ldots, \sqrt{n} - 1\}$. Because of the anonymity of M, each agent $i = 1, \ldots, \sqrt{n}$, can at most expect to get a uniform lottery over all the items with each of the other $\sqrt{n} - 1$ agents that have the same ordering of valuations. For agents $i \in \{\sqrt{n} + 1, \ldots, n\}$, the contribution to the Social Welfare from items 1, ..., $\sqrt{n}$ is at most 2 since their valuations for these items are bounded by $\frac{2}{\sqrt{n}}$. Similarly, their contribution to the Social Welfare from items $\sqrt{n} + 1, \ldots, n$ is similarly bounded by 1.

Therefore, we get the following upper bound on the expected Social Welfare:

$$\sum_{i=1}^{\sqrt{n}} E\left[u_i(M(\mathbf{u}))\right] + \sum_{i=\sqrt{n}+1}^{n} E\left[u_i(M(\mathbf{u}))\right] \leq \sum_{i=1}^{\sqrt{n}} \frac{1}{\sqrt{n}} + 3 = 4,$$

while the Social Welfare of the optimal allocation is at least $\frac{\sqrt{n}-1}{10n^3}$. From this, we get $ar(M) \leq \frac{8}{\sqrt{n}} = O(\frac{1}{\sqrt{n}})$. □

Finally, the upper bound for any truthful-in-expectation mechanism is given by the following lemma.

Lemma 4.10. Let M be a truthful-in-expectation mechanism for unit-sum representation. The Approximation Ratio of M is $O\left(\frac{1}{\sqrt{n}}\right)$.

Proof. The intuition behind the lemma lies in the fact that the valuation profile used in the proof of Lemma 4.5 can be adjusted so that the sum of valuations across each row of the valuation matrices equals one. To achieve this, consider the following modified valuation profile:

$$u_i(j) = \begin{cases} 1 - \sum_{j \neq i} u_i(j), & \text{for } j = i \\ \frac{2}{10k} - \frac{j}{10n}, & \text{for } 1 \leq j \leq k+1, j \neq i \\ \frac{n-j}{10n^2}, & \text{otherwise} \end{cases} \qquad \forall i \in \{1, \ldots, k+1\}$$

$$u_i(j) = \begin{cases} 1 - \sum_{j \neq 1} u_i(j), & \text{for } j = 1 \\ \frac{2}{10k} - \frac{j}{10n}, & \text{for } 1 < j \leq k+1 \\ \frac{n-j}{10n^2}, & \text{otherwise} \end{cases} \qquad \forall i \in \{k+2, \ldots, n\}$$

This profile is a direct adaptation of the one used in Lemma 4.5, with the key difference that all entries are divided by 10, except those entries where the valuation is 1, which are now equal to 1 minus the sum of the valuations for the rest of the items. This modification introduces a scaling factor of $\frac{1}{10}$ through the calculations. However, it does not effect the he asymptotic bound established in Lemma 4.5, hence the same proven bound holds. □

4.6.2 Allowing ties

The results of this paper, extend even when we have ties, meaning when agents value at least two items the same. Random Priority though, needs a way to decide which item an agent should pick.

To handle ties, we introduce small perturbations of the agents' valuations. Such perturbations are taken to break the ties, according to some tie-breaking rule. In such cases, the outcome of the RP mechanism — what items the agents actually end up with — would still be the same as if we had applied RP directly on the original profile with the tie-breaking rule put in place.

The reason lies in the fact that, since these perturbations do not change the ranking of an agent over different items to a great extent, the assignment probabilities (i.e., the probability of any agent receiving a certain item) stay the same. In this way, the performance and guarantees — like the approximation ratio — given by the RP mechanism remain unchanged even when ties are present.

4.6.3 Unit - range* - valuation functions

All the results apply to the extension of the unit-range representation where 0 is not required to be in the image of the function, that is, $\max_j u_i(j) = 1$ and for all j, $u_i(j) \in [0, 1]$. This captures cases where an agent might slightly prefer one item to another, without having to completely discard some item - no item needs to have its valuation exactly 0. For example, the agent might value each of the items pretty much similarly and vary the preferences slightly between the items. This is different from the strict unit-range model, which requires at least one item to be valued at 0, indicating complete indifference to that particular item.

Since all the unit-range valuation profiles are valid for this representation too, the upper bounds are trivial. The lower bound is not — the reason a lower bound doesn't hold directly is that the transformation of a $[0, 1]$ profile to a unit-range profile can reduce the optimal Social Welfare by a factor of up to 1 unit. The following proof adjusts this by measuring welfare before and after the transformation, so the $\Omega(\frac{1}{\sqrt{n}})$ bound still holds with only minor loss.

Corollary 4.11. The Approximation Ratio of Random Priority, for the unit-range* valuation functions is $\Omega\left(\frac{1}{\sqrt{n}}\right)$.

Proof. Let $\mathbf{u}$ be any $[0, 1]$ valuation profile, where each agent i assigns a value $u_i(j) \in [0, 1]$ to each item j. Let c be the constant appearing in the lower bound of Lemma 4.4. We will relate the Approximation Ratio of Random Priority on this generalized profile $\mathbf{u}$ to the established ratio for the unit-range profile.

To achieve this, we transform $\mathbf{u}$ into a profile $\mathbf{u}'$. For each agent i, we modify the valuation profile by setting the valuation of the least-preferred item to 0, if it is not already 0. Formally, let $\mathbf{u}'$ be defined as:

$$u'_i(j) = \begin{cases} 0 & \text{, if } j = \text{argmin}_k u_i(k) \\ u_i(j) & \text{, otherwise} \end{cases}$$

This transformation ensures that $\mathbf{u}'$ is a valid unit-range profile, meaning $u'_i(j) \in [0, 1]$ and there exists at least one item j such that $u'_i(j) = 0$.

Since the transformation only decreases the valuation of each agent's least-preferred item, the expected Social Welfare of Random Priority under $\mathbf{u}'$ cannot exceed that under $\mathbf{u}$. Thus:

$$E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}'))\right] \leq E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right].$$

Now, since the optimal welfare $w^*(\mathbf{u}')$ for the modified profile $\mathbf{u}'$ is at most 1 unit less than the optimal welfare $w^*(\mathbf{u})$ for the original profile, as the optimal matching on $\mathbf{u}$ may have assigned the least-preferred item to exactly one agent:

$$w^*(\mathbf{u}') \geq w^*(\mathbf{u}) - 1.$$

Given that $\mathbf{u}'$ is a unit-range profile, we apply Lemma 4.4:

$$\frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}'))\right]}{w^*(\mathbf{u}')} \geq \frac{c}{\sqrt{n}}.$$

And since $w^*(\mathbf{u}') \geq w^*(\mathbf{u}) - 1$:

$$\frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right]}{w^*(\mathbf{u})} \geq \frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}'))\right]}{w^*(\mathbf{u}') + 1}.$$

Since $w^*(\mathbf{u}') \geq w^*(\mathbf{u}) - 1$, it follows that:

$$\frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right]}{w^*(\mathbf{u})} \geq \frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}'))\right]}{2w^*(\mathbf{u}')}.$$

Combining the above inequalities, we get:

$$\frac{E\left[\sum_{i=1}^{n} u_i(RP(\mathbf{u}))\right]}{w^*(\mathbf{u})} \geq \frac{c}{2\sqrt{n}}.$$

Thus, the Approximation Ratio of Random Priority for any unit - range* valuation profile **u** is $\Omega\left(\frac{1}{\sqrt{n}}\right)$.

□

4.6.4 An Improved Approximation

As we saw, Theorem 4.5 states that RP is the best truthful-in-expectation mechanism for the problem, when *only* considering the asymptotic behavior of mechanisms. The authors, consider non-asymptotic behavior by studying the case when $n = 3$ and present an non-ordinal mechanism that achieves better bounds than any ordinal mechanism, when the representation of the valuation functions is unit-range.

We begin by looking the Approximation Ratio of RP for $n = 3$. It is easy to see that the Optimal obtainable Social Welfare by RP in the unit-range setting is 3 - each agent get's their most desired item, valued at 1. Now, for the Worst-Case Social Welfare. Consider the following profile:

$$\mathbf{u} = \begin{pmatrix} 1 & 1-\epsilon & 0 \\ 1 & \epsilon & 0 \\ 1 & \epsilon & 0 \end{pmatrix}$$

The best matching that RP can achieve in this profile is for agent 2 to get item 1, agent 1, item 2 and agent 3, the item 3 the Social Welfare is the $2 - \epsilon \to 2$ when ϵ is small.

The **Hybrid Mechanism** builds upon the cubic lottery, a non-ordinal mechanism proven to be truthful-in-expectation (see [28]) which matches each agent with their most prefered item with a certain probability.

The cubic lottery for 1 agent can be explained easily, as they get:

- The most preferred item with probability $\frac{6-2\alpha^3}{8}$,
- The second-most preferred item with probability $\frac{1+3\alpha^2}{8}$,
- The remaining item with probability $1 - (\frac{6-2\alpha^3}{8} + \frac{1+3\alpha^2}{8})$.

  Where, α is the agent's valuation for their second-most preferred item.

The Hybrid Mechanism is, hence, the following:

Steps in HM

1. Uniformly at random choose a permutation σ of the agents.
2. Agent $\sigma(1)$ is matched with an item based on the cubic lottery.
3. Agent $\sigma(2)$ selects their favorite item from the remaining items.

4. Agent $\sigma(3)$ is assigned the final remaining item.

The Hybrid Mechanism achieves an Approximation Ratio of 0.699, which is better than the $\frac{2}{3}$ achieved by any ordinal mechanism, including RP. The authors prove that by solving a non-linear optimization problem. The proof is omited but can by found in [10]. Although for $n = 3$ this method works well, for larger n we face computational challenges. But the possible results by Extending the Hybrid Mechanism for larger n's can be valuable.

CHAPTER 5

PRICE OF ANARCHY BOUNDS

This chapter is based on the work of George Christodoulou, Aris Filos-Ratsikas, Soren Kristoffer Stiil Frederiksen, Paul W. Goldberg, Jie Zhang and Jinshan Zhang, in their paper "Social Welfare in One-Sided Matching Mechanisms" presented at the Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016, with some adjustments and clarifications for ease of reading.

5.1 Synopsis

In this paper, the authors provide a lower bound of $\Omega(\sqrt{n})$ for the Price of Anarchy across all types of mechanisms for the problem - including both ordinal and cardinal ones.

They focus on the Probabilistic Serial and Random Priority mechanisms, which obtain a matching upper bound. The analysis is extended to deterministic mechanisms where we will see that they exhibit significantly worse performance in terms of their bounds.

The authors also study how these mechanisms behave when agents have incomplete information. They demonstrate that even in such cases, the PoA remains within the same bounds.

Finally, the authors examine the Price of Stability, which measures inefficiency at the best possible equilibrium. They show that the PoS has a similar lower bound of $\Omega(\sqrt{n})$ in both ordinal and deterministic mechanisms.

5.2 Introduction

In Chapter 4, we studied the performance of ordinal one-sided matching mechanisms through their Approximation Ratio under unit-range valuations. We focused on the Random Priority (RP) mechanism due to its truthfulness, but it is known that the RP does not satisfy ex-ante Pareto efficiency and envy-freeness.

The truthfulness of RP gives a simple route to the analysis of the mechanism's Social Welfare inefficiency. Other ordinal mechanisms, however, do not offer this property and if they do, this does not necessarily mean that misreporting cannot obtain an outcome with higher Social Welfare.

In settings without money, valuations are represented using a canonical way. Common approaches are the *unit-range* and *unit-sum* representantions. While in Chapter 4. we used the *unit-range* representation mainly, easily extending the result to the *unit-sum* representantion. In this chapter, it will be the other way around, fixing our representation to *unit-sum*. As we imply cardinal utilities below the ordinal preferences, each agent's allocation induces a certain utility. Forming a game at hand, by Nash's theorem, every game has a Nash Equilibrium. The natural question becomes how to measure the mechanism's inefficiency.

In Chapter 5, we shift our focus, from bounding the inneficiency of truthtelling, to studying mechanism's behaviour in Equilibria. While the results of Chapter 4 prove to be fundamental, as they establish further results for studying the, other than truthtelling, equilibria of RP, we also focus on Probabilistic Serial, PS is known for it's ex-ante Pareto efficient and envy-free properties, but also because it is not truthful, making it interesting to study it's performance in equilibria.

The Price of Anarchy (PoA), as seen in Chapter 2. , was introduced by Koutsoupias and Papadimitriou in 1999 and is a celebrated measure quantifying the efficiency loss at equilibrium due to selfish behavior compared to the optimal outcome. Initially used in network congestion problems, PoA is now widely applied to analyze inefficiencies from strategic behavior in various settings, including resource allocation and mechanism design. The reader can see Chapter 2. for a reminding of the definition of the Price of Anarchy.

In this chapter, our analysis moves from the direct comparison of Social Welfare under truthful reporting, to examining Social Welfare at Nash Equilibria. Hence, the bridge from Chapter 4 to Chapter 5 is guided by incentive compatibility.

5.3 Upper Bounds

In this section, we establish the (pure) Price of Anarchy upper bound guarantees for the Probabilistic Serial and Random Priority mechanisms. Together with the results from the next section, we will be able to provide strict bounds.

5.3.1 Random Priority Bounds

As we have seen in Chapter 4. , Filos-Ratsikas et al. in [10] proved that the Social Welfare in any truthtelling equilibrium is an $\Omega(\frac{1}{\sqrt{n}})$-fraction of the maximum Social Welfare. Although Random Priority may have other equilibria, establishing the Price of Anarchy bound relies on the observation that, any non-truthful strategy does not affect the allocation of other agents and consequently, does not affect Social Welfare.

We assume that valuations are distinct - that is there are no ties. If ties emerge, as seen in Chapter 4. , small pertubations on the valuation functions can do the trick, with a negliglible cost of Social Welfare.

Lemma 5.1. If valuations are distinct, the Social Welfare is the same in all mixed Nash Equilibria of Random Priority.

Proof. Let i be an agent and B be a subset of the items. Let **s** be a mixed Nash equilibrium where there is a positive probability that agent i will be selected to choose an item when B is the set of remaining items. Since agent i has distinct valuations for the items, their strategy should always rank their most preferred item in B at the top of the preference list for that set.

Now, for two items j and j', suppose that there is no set of items B that may be offered to i with positive probability where either j or j' is the optimal choice for i. In this case, agent i may rank j and j' in any order (can report $j \succ_i j'$ or $j' \succ_i j$). However, this report has no impact on the other agents and it does not influence their Social Welfare. □

Theorem 4.5 from Chapter 4. , states that the the approximation ratio of Random Priority is $\Theta\left(\frac{1}{\sqrt{n}}\right)$. Together with Lemma 5.3.1 we can easily see the following result.

Theorem 5.2. For distinct valuations, the Price of Anarchy of Random Priority is $\Theta(\sqrt{n})$.

When dealing with ties, the same upper bound holds.

Theorem 5.3. The Price of Anarchy of Random Priority is $O(\sqrt{n})$.

Proof. We know from Theorem 4.5 from Chapter 4. . that the Social Welfare of Random Priority when agents are truthful is within $O(\sqrt{n})$ of the social optimum. The Social Welfare of a (mixed) Nash equilibrium **q** cannot be worse than the worst pure profile from q that occurs with positive probability, so let **s** be such a pure profile. We will say that agent i misranks items j and j' if $j \succ_i j'$, but $j' \succ_{s_i} j$.

If an agent misranks two items for which she has distinct values, it is because she has 0 probability in **s** to receive either item. So we can change **s** so that no items are misranked, without affecting the Social Welfare or the allocation. For items that the agent values equally (which are then not misranked) we can apply arbitrarily small perturbations to make them distinct.

Profile **s** is thus consistent with rankings of items according to perturbed values and is truthful with respect to these values, which, being arbitrarily close to the true ones, have optimum Social Welfare *arbitrarily* close to the true optimal Social Welfare. □

5.3.2 Probabilistic Serial Bounds

We explain the Probabilistic Serial Mechanism in Chapter 2. , as a reminder, the mechanism works as follows: Each item can be viewed as an infinitely divisible item that all agents can consume at unit speed during the unit time interval $[0, 1]$. Initially each agent consumes her most preferred item (or one of her most preferred items in case of ties) until the item is entirely consumed. Then, the agent moves on to consume the item on top of her preference list, among items that have not yet been entirely consumed. The mechanism terminates when all items have been entirely consumed. The fraction p_{ij} of item j consumed by agent i is interpreted as the probability that agent i will be matched with item j under the mechanism.

Aziz et.al. in [6] prove the the Probabilistic Serial has Pure Nash Equilibria. Since, we are trying to upper bound the Price of Anarchy of PS, it suffices to consider only Pure Nash Equilibria. We procceed by proving two lemmas,

We start with the following two lemmas.

Let $t_j(\mathbf{s})$ be the time when item j is entirely consumed on profile **s** under $PS(\mathbf{s})$.

Lemma 5.4. Let **s** be any strategy profile and let s_i^* be any strategy where agent i places item j on top of her preference list. i.e. $j \succ_{s_i^*} \ell$ for all items $\ell \neq j$. Then it holds that $t_j(s_i^*, \mathbf{s_{-i}}) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$.

Proof. Let, for ease of notation, $\mathbf{s}^* = (s_i^*, \mathbf{s_{-i}})$ be any strategy where agent i places item j on top of her preference list. If $j \succ_{s_i} \ell$ for all $\ell \neq j$ and since all other agents' reports are fixed, $t_j(\mathbf{s}^*) = t_j(\mathbf{s})$ and the statement of the lemma holds.

Hence, we assume that there exists some item $j' \neq j$ such that $j' \succ_{s_i} j$. If agent i is the only one consuming item j for the duration of the mechanism, then $t_j(\mathbf{s}^*) = 1$ and we are done. Hence, assume that at least one other agent consumes item j at some point, and let τ be the time when the first agent besides agent i starts consuming item j in $\mathbf{s}^*$. Agent i consumes j for more time, thus, $t_j(\mathbf{s}^*) > \tau$, therefore if $\tau \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ then $t_j(\mathbf{s}^*) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ and we are done.

Assume that $\tau < \frac{1}{4} \cdot t_j(\mathbf{s})$. Observe that in the interval $[\tau, t_j(\mathbf{s}^*)]$, since there are at least two agents consuming item j, agent i can consume at most half of what remains of it. Overall, agent i's consumption is at most $\frac{1}{2} + \frac{1}{4}t_j(\mathbf{s})$ so at least $1 - [\frac{1}{2} + \frac{1}{4}t_j(\mathbf{s})] = \frac{1}{2} - \frac{1}{4}t_j(\mathbf{s})$, of the item j will be consumed by the rest of the agents.

Now consider all agents other than i in profile **s** and let $\alpha \in (0, 1]$ be the the amount of item j that they have consumed by time $t_j(\mathbf{s})$. The total consumption speed of an item is non-decreasing in time which means in particular that for any $0 \leq \beta \leq 1$, agents other than i need at least $\beta t_j(\mathbf{s})$ time to consume $\alpha \cdot \beta$ in profile **s**. Since agent i starts consuming item j at time 0 in $\mathbf{s}^*$ and all other agents' strategies are the same in both profiles, **s** and $\mathbf{s}^*$, it holds that every agent $k \neq i$ starts consuming item j in $\mathbf{s}^*$ no sooner than she does in **s**. This means that in profile $\mathbf{s}^*$, agents other than i will need more time to consume $\beta \cdot \alpha$ and in particular they will need at least $\beta t_j(\mathbf{s})$ time, so $t_j(\mathbf{s}^*) \geq \beta t_j(\mathbf{s})$.

However, we saw earlier that they will consume at least $\frac{1}{2} - \frac{1}{4}t_j(\mathbf{s})$, and thus $\alpha \cdot \beta \geq \frac{1}{2} - \frac{1}{4}t_j(\mathbf{s})$. Letting $\beta \geq \frac{1}{\alpha}\left(\frac{1}{2} - \frac{1}{4} \cdot t_j(\mathbf{s})\right)$ we get:

$$t_j(\mathbf{s}^*) \geq \beta t_j(\mathbf{s}) \geq \frac{1}{\alpha}\left(\frac{1}{2} - \frac{1}{4} \cdot t_j(\mathbf{s})\right) t_j(\mathbf{s}) \geq \left(\frac{1}{2} - \frac{1}{4} \cdot t_j(\mathbf{s})\right) t_j(\mathbf{s}) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$$

□

We can now, lower bound the utility of an agent at any Pure Nash Equilibrium (PNE).

Lemma 5.5. Let $\mathbf{u} = (u_1, \ldots, u_n)$ be the profile of true agent valuations and let $\mathbf{s}$ be a pure Nash equilibrium. For any agent i and any item j it holds that the utility of agent i at $\mathbf{s}$ is at least $\frac{1}{4} \cdot t_j(\mathbf{s}) \cdot u_{ij}$.

Proof. Let $\mathbf{s}' = (s_i', \mathbf{s_{-i}})$ be the strategy profile resulting from agent i deviating from $\mathbf{s}$ to a new strategy s_i' where s_i' ranks item j above all other items $\ell \neq j$. Since $\mathbf{s}$ is a PNE, it holds that $u_i(PS_i(\mathbf{s})) \geq u_i(PS_i(\mathbf{s}')) \geq t_j(\mathbf{s}') \cdot u_{ij}$, where the last inequality holds because the utility of agent i from their deviation to s_i' must be at least as much as the utility they derive from the consumption of item j, since, by deviating to the strategy s_i', agent i ensures that item j is ranked highest. This guarantees that their utility from $\mathbf{s}'$ is at least her utility from consuming item j alone for the time $t_j(\mathbf{s}')$. By Lemma 5.4, it holds that $t_j(\mathbf{s}') \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ and hence $u_i(PS_i(\mathbf{s})) \geq \frac{1}{4} \cdot t_j(\mathbf{s}) \cdot u_{ij}$. □

The intuition behind these proofs is that, in a PNE, an agent's utility cannot be significantly lower than the utility they would receive if they were consuming the item matched to them in the optimal allocation from the start, until the item is fully consumed.

Now, we can upper bound the PoA of PS.

Theorem 5.6. The pure Price of Anarchy of Probabilistic Serial is $O(\sqrt{n})$.

Proof. Let $\mathbf{u} = (u_1, \ldots, u_n) \in \mathbb{R}_+^{n \times n}$ be the profile of true agent valuations and let $\mathbf{s}$ be a pure Nash equilibrium. It's easy to see that in the Probabilistic Serial mechanism, by reporting truthfully, each agent i is assured of a fair share of the utility, no matter what the other agents do. This fair share translates to at least $\frac{1}{n}$ of the total available utility for each agent, making truthfulness a safe strategy and ensuring that the Social Welfare of the mechanism is at least 1. To see this, first consider time $t = \frac{1}{n}$ and observe that during the interval $[0, \frac{1}{n}]$, agent i is consuming her favorite item (w.l.o.g. say a_1) and hence $p_{ia_1} \geq \frac{1}{n}$. Next, consider time $t' = \frac{2}{n}$ and observe that during the interval $[0, \frac{2}{n}]$, agent i is consuming one or both of her two favorite items (a_1 and a_2) and hence $p_{ia_1} + p_{ia_2} \geq \frac{2}{n}$. Consequently, for any k, it holds that $\sum_{j=1}^{k} p_{ia_j} \geq \frac{k}{n}$. This implies that regardless of other agents' strategies, agent i can achieve a utility $\sum_{j=1}^{n} p_{ia_j} \cdot u_{ij} \geq \frac{1}{n} \sum_{j=1}^{n} u_{ij}$. Since $\mathbf{s}$ is a pure Nash equilibrium, it holds that $u_i(PS_i(\mathbf{s})) \geq \frac{1}{n} \cdot \sum_{j=1}^{n} u_{ij}$ as well. Thus, for the Social Welfare of the mechanism, summing over all agents, we get that $SW_{PS}(\mathbf{u}, \mathbf{s}) \geq \sum_{i=1}^{n} \frac{1}{n} \cdot \sum_{j=1}^{n} u_{ij} = 1$. If $SW_{OPT}(\mathbf{u}) \leq \sqrt{n}$, we get that:

$$PoA(PS) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_{PS}(\mathbf{u}, \mathbf{s})} \leq \frac{\sqrt{n}}{1} = O(\sqrt{n})$$

So we are done. Let us assume that $SW_{OPT}(\mathbf{u}) > \sqrt{n}$.

Recall that PS is neutral, meaning that it's allocation process does not favor any specific item over others based on their identity. Thus, we can assume that $t_j(\mathbf{s}) \leq t_{j'}(\mathbf{s})$ for $j < j'$ without loss of generality.

By any time $0 \leq t \leq 1$, the total fraction of items consumed by all agents combined is exactly $t \cdot n$. This is because there are n items, and the consumption rate is uniform. Since $t_j(s)$ represents the time at which the j^{th} item is entirely consumed, the mass of items consumed by this time must be at least j. Thus, $t_j(\mathbf{s}) \cdot n \geq j$, or, $t_j(\mathbf{s}) \geq \frac{j}{n}$.

For each item j let i_j be the agent who receives item j in the optimal allocation and for ease of notation, let w_{i_j} denotes the valuation of agent i_j for item j, i.e. $w_{i_j} = u_{i_j j}$. Now by Lemma 5.5, since $\mathbf{s}$ is a pure Nash equilibrium, it holds that

$$u_{i_j}(PS(\mathbf{s})) \geq \frac{1}{4} \cdot t_j(\mathbf{s}) \cdot w_{i_j} \geq \frac{1}{4} \cdot \frac{j}{n} \cdot w_{i_j} \quad \text{and} \quad SW_{PS}(\mathbf{u}, \mathbf{s}) = \sum_{i=1}^{n} u_i(PS(\mathbf{s})) = \sum_{j=1}^{n} u_{i_j}(PS(\mathbf{s})) \geq \frac{1}{4} \sum_{j=1}^{n} \frac{j}{n} \cdot w_{i_j}.$$

Since, by definition, $SW_{OPT}(\mathbf{u}) = \sum_{j=1}^{n} w_{i_j}$, the Price of Anarchy is then:

$$PoA(PS) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_{PS}(\mathbf{u}, \mathbf{s})} \leq \frac{SW_{OPT}(\mathbf{u})}{SW_{PS}(\mathbf{u}, \mathbf{s})} = \frac{\sum_{j=1}^{n} w_{i_j}}{\frac{1}{4} \sum_{j=1}^{n} \frac{j}{n} \cdot w_{i_j}}.$$

To establish an upper bound for the quantity

$$4n \cdot \frac{\sum_{j=1}^{n} w_{i_j}}{\sum_{j=1}^{n} j \cdot w_{i_j}}.$$

we consider the case when the ratio

$$\frac{\sum_{j=1}^{n} w_{i_j}}{\sum_{j=1}^{n} j \cdot w_{i_j}}.$$

is maximized. Let k be an integer such that $k \leq \sum_{j=1}^{n} w_{i_j} \leq k+1$. Given that $w_{i_j} \in [0,1]$ (as valuation functions), it follows that $w_{i_j} = 1$ for $j = 1, \ldots, k$ and $w_{i_j} = 0$, for $j \geq k+2$.

Hence,

$$\frac{\sum_{j=1}^{n} w_{i_j}}{\sum_{j=1}^{n} j \cdot w_{i_j}} \leq \frac{k + w_{i_{k+1}}}{(k+1) \cdot w_{i_{k+1}} + \sum_{j=1}^{k} j \cdot 1} = \frac{k + w_{i_{k+1}}}{(k+1) \cdot w_{i_{k+1}} + \frac{k(k+1)}{2}} = \frac{k + w_{i_{k+1}}}{a w_{i_{k+1}} + b}$$

For $w_{i_{k+1}}$ in $[0,1]$, the above ratio is decreasing. Therefore, the maximum value of $(k + w_{i_{k+1}})/(a w_{i_{k+1}} + b)$ is achieved when $w_{i_{k+1}} = 0$. Consequently, the Price of Anarchy is at most:

$$4n \cdot \frac{\sum_{j=1}^{n} w_{i_j}}{\sum_{j=1}^{n} j \cdot w_{i_j}} \leq 4n \cdot \frac{k + w_{i_{k+1}}}{a w_{i_{k+1}} + b} \leq 4n \cdot \frac{k}{\frac{k(k+1)}{2}} = \frac{8n}{k+1}$$

Thus, the Price of Anarchy is maximized when k is minimized. Since, $SW_{OPT}(\mathbf{u}) = \sum_{j=1}^{n} w_{i_j} \geq k$ and we assumed that $SW_{OPT}(\mathbf{u}) > \sqrt{n}$, we find that $k > \sqrt{n}$. Hence, the ratio is $O(\sqrt{n})$. $\square$

5.4 Lower Bounds

The following result, bounds the Price of Anarchy of any mechanism. Including, ordinal, cardinal, deterministic and randomized mechanisms. Since we're interested in mechanisms with good properties, it's natural to focus on those with pure Nash equilibria, where outcomes are deterministic and no randomness is involved.

Theorem 5.7. The pure Price of Anarchy of any mechanism is $\Omega(\sqrt{n})$.

Proof. Let M be a mechanism, $n = k^2$ for some $k \in \mathbb{N}$ the number of agents and consider the following valuation profile **u** described as follows. There are $\sqrt{n}$ sets of agents and let G_j denote the j-th set. For every $j \in \{1, \ldots, \sqrt{n}\}$ and every agent $i \in G_j$, it holds that $u_{ij} = \frac{1}{n} + \alpha$ and $u_{ik} = \frac{1}{n} - \frac{\alpha}{n-1}$, for $k \neq j$, where $\alpha > 0$ is sufficiently small. Let **s** be a pure Nash equilibrium and for every set G_j, let $i_j = \arg\min_{i \in G_j} p_{ij}^{M,\mathbf{s}}$ (break ties arbitrarily).

Since, the total probability of agents in G_j to get item j is 1 and there are $\sqrt{n}$ agents in G_j: $\sum_{i \in G_j} p_{ij}^{M,\mathbf{s}} = p_{i_1 j}^{M,\mathbf{s}} + p_{i_2 j}^{M,\mathbf{s}} + \ldots + p_{i_{\sqrt{n}} j}^{M,\mathbf{s}} = 1$. Thus, for all $j = 1, \ldots, \sqrt{n}$, it holds that $p_{i_j j}^{M,\mathbf{s}} \leq \frac{1}{\sqrt{n}}$. Observe that for all $j = 1, \ldots, \sqrt{n}$, it holds that $p_{i_j j}^{M,\mathbf{s}}$. Now, let $I = \{i_1, i_2, \ldots, i_{\sqrt{n}}\}$ and consider the valuation profile $\mathbf{u}'$ where:

- For every agent $i \notin I$, $u'_i = u_i$.
- For every agent $i_j \in I$, let $u'_{i_j j} = 1$ and $u'_{i_j k} = 0$ for all $k \neq j$.

The strategy **s** is a pure Nash equilibrium under $\mathbf{u}'$ as well:

- For agents $i \notin I$: the valuations have not changed and hence they have no incentive to deviate.
- Assume now that some agent $i \in I$ whose most preferred item is item j , deviates to some beneficial strategy s'_i. Since agent i only values item j, this would imply that $p_{ij}^{M,(s'_i,\mathbf{s}_{-i})} > p_{ij}^{M,\mathbf{s}}$. However, since agent i values all items other than j equally under u_i and her most preferred item is item j, such a deviation would also be beneficial under profile **u**, contradicting the fact that **s** is a pure Nash equilibrium in profile **u**.

Now consider the expected Social Welfare of M under valuation profile $\mathbf{u}'$ at the pure Nash equilibrium **s**.

$$SW_M(\mathbf{u}', \mathbf{s}) = \sum_{i \in I} u'_i + \sum_{i \notin I} u'_i$$

- For agents not in I :

  $$\sum_{i \notin I} u'_i = \sum_{i \notin I} u_i \leq \sum_{i \notin I} \frac{1}{n} + \alpha = (n - \sqrt{n}) \cdot (\frac{1}{n} + \alpha) < 1 \text{ , after letting } \alpha < \frac{1}{n^3}.$$

- For agents in I:

  Under $\mathbf{u}'$ in the expected Social Welfare contribute only the agents that get their most prefered item j.

  $$\sum_{i \in I} u'_i = \sum_{j=1}^{\sqrt{n}} p_{i_j j}^{M,\mathbf{s}} \cdot u'_{i_j j} \leq \frac{1}{\sqrt{n}} \cdot \sqrt{n} + 1 \leq 2$$

Hence the expected Social Welfare of M is at most 3.

As for the optimal Social Welfare, it is achieved when each agent in I gets their most preferred item. Since there are $\sqrt{n}$ such agents and each values their most preferred item at 1, the optimal Social Welfare is at least $\sqrt{n}$.

$$SW_{OPT}(\mathbf{u}') \geq \sum_{i_j \in I} u'_{i_j j} = \sqrt{n} \cdot 1 = \sqrt{n}$$

Thus:

$$PoA(M) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})} \geq \frac{SW_{OPT}(\mathbf{u'})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})} \geq \frac{SW_{OPT}(\mathbf{u'})}{SW_M(\mathbf{u'}, \mathbf{s})} \geq \frac{\sqrt{n}}{3} = \Omega(\sqrt{n}).$$

□

Below, we see that deterministic mechanisms perform poorly in terms of efficiency at equilibrium.

Theorem 5.8. The Pure Price of Anarchy of any deterministic mechanism is $\Omega(n^2)$.

Proof. Let M be a deterministic mechanism, that always has a Pure Nash Equilibrium.

Let **u** be a valuation profile such that for all agents i, i', it holds that $u_i = u_{i'}$, $u_{i1} = \frac{1}{n} + \frac{1}{n^3}$ and $u_{ij} > u_{ik}$ for $j < k$. Let **s** be a Pure Nash Equilibrium for this profile and assume without loss of generality that $M_i(s) = i$.

Now fix another true valuation profile u' such that $u'_1 = u_1$ and for agents $i \in \{2, \ldots, n\}$, $u'_{i,i-1} = 1 - \epsilon'_{i,i-1}$ and $u_{ij} = \epsilon'_{ij}$ for $j \neq i-1$, where $0 \leq \epsilon'_{ij} \leq \frac{1}{n^3}$, $\sum_{j \neq i-1} \epsilon'_{ij} = {\epsilon'}_{i,i-1}$ and $\epsilon'_{ij} > \epsilon'_{ik}$ if $j < k$ when $j, k \neq i-1$.

Intuitively, in profile $\mathbf{u'}$, each agent $i \in \{2, \ldots, n\}$ has valuation close to 1 for item $i-1$ and small valuations for all other items. Futhermore, she prefers items with smaller indices, except for item $i-1$.

We claim that **s** is a Pure Nash Equilibrium under true valuation profile $\mathbf{u'}$ as well. Assume that it is not. Meaning that some agent i has a benefiting deviation, matching her with an item that she prefers more than i. But then, since the set of items that she prefers more than i in both **u** and $\mathbf{u'}$ is $\{1, \ldots, i\}$, the same deviation would match her with a more preferred item under **u** as well, contradicting the fact that **s** is a Pure Nash Equilibrium for profile **u**.

For the final step, recall the definition of the Price of Anarchy.

$$PoA(M) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})}$$

It holds that:

$$PoA(M) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})} \geq \frac{SW_{OPT}(\mathbf{u'})}{\min_{\mathbf{s} \in S_{\mathbf{u'}}^M} SW_M(\mathbf{u'}, \mathbf{s})} \geq \frac{SW_{OPT}(\mathbf{u'})}{SW_M(\mathbf{u'}, \mathbf{s})}$$

It holds that:

$$SW_{OPT}(\mathbf{u'}) \geq \sum_{i=1}^n \sum_{j=1}^n u'_{ij} = \sum_{j=1}^n u_{1j} + \sum_{i=2}^n \sum_{j=1}^n u'_{ij} = \sum_{j=1}^n u_{1j} + \sum_{i=2}^n u'_{i,i-1} + \sum_{i=2}^n \sum_{j \neq i-1} \epsilon'_{ij} = \sum_{j=1}^n u_{1j} + \sum_{i=2}^n u'_{i,i-1} + \sum_{i=2}^n \epsilon'_{i,i-1}$$

$$= \sum_{j=1}^n u_{1j} + \sum_{i=2}^n (1 - \epsilon'_{i,i-1}) + \sum_{i=2}^n \epsilon'_{i,i-1} = \sum_{j=1}^n u_{1j} + n - 2 - \sum_{i=2}^n \epsilon'_{i,i-1} + \sum_{i=2}^n \epsilon'_{i,i-1} \geq n - 2$$

Whereas,

$$SW_M(\mathbf{u'}, \mathbf{s}) \leq \sum_{i=1}^n \sum_{j=1}^n u'_{ij}$$

The mechanism allocates item i to agent i. Thus,

$$SW_M(\mathbf{u}', \mathbf{s}) \leq \sum_{i=1}^{n} \sum_{j=1}^{n} u'_{ij} = \sum_{i=1}^{n} u'_{i,i} \leq \sum_{i=1}^{n} u'_{i,1}$$

We have established that:

$$SW_{OPT}(\mathbf{u}') \geq \sum_{i=1}^{n} \left(1 - \epsilon'_{i,i-1}\right) = n - 2,$$

While for the mechanism's social welfare $SW_M(\mathbf{u}', \mathbf{s})$, since the mechanism assigns item i to agent i, it holds that:

$$SW_M(\mathbf{u}', \mathbf{s}) = \sum_{i=1}^{n} u'_{ii} \leq 1 + \sum_{i=2}^{n} \epsilon'_{i,i-1}.$$

Given that $\epsilon'_{i,i-1} \leq \frac{1}{n^3}$:

$$SW_M(\mathbf{u}', \mathbf{s}) \leq 1 + (n-1) \cdot \frac{1}{n^3} \leq 2.$$

Therefore:

$$PoA(M) \geq \frac{n-2}{2} = \Omega(n^2).$$

This completes the proof, showing that the Pure Price of Anarchy for any deterministic mechanism is $\Omega(n^2)$.

$\square$

5.5 Unit-range valuation functions

In this section we extend the above bounds to the unit-range valuation setting, that is, $\max_j u_i(j) = 1$ and $\min_j u_i(j) = 0$.

For Random Priority, since the results of Chapter 4. ([10]) hold for this normalization as well, we can apply the same techniques to prove the bounds.

For Probabilistic Serial, observe that Lemma 5.5 holds independently of the representation. Hence, in the proof of Theorem 5.6, it now holds that:

$$SW_{PS}(\mathbf{u}, \mathbf{s}) \geq \frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{n} u_{ij} \geq 1,$$

which is sufficient for bounding the Price of Anarchy when $SW_{OPT}(\mathbf{u}) \leq \sqrt{n}$. Finally, the arguments for the case when $SW_{OPT}(\mathbf{u}) \leq \sqrt{n}$ hold for both representations.

Concerning the lower bounds, we can prove the following theorem on the Price of Anarchy of deterministic mechanisms.

Lemma 5.9. The Price of Anarchy of any deterministic mechanism that always has pure Nash equilibria is $\Omega(n)$ for the unit-range representation.

Proof. (Sketch.) Let u and u' be valuation profiles with the same preference ordering, and s a pure Nash equilibrium under u. If s is not a Nash equilibrium under u', then agent i can deviate to a more preferred item under u'_i. Since the preference ordering is the same, i would have also deviated under u, contradicting the equilibrium. Thus, s must be a Nash equilibrium under u'. □

Through Lemma 5.9, we can prove the following theorem

Theorem 5.10. The Price of Anarchy of any deterministic mechanism that always has pure Nash equilibria is $\Omega(n)$ for the unit-range representation.

Proof. (Sketch.) Consider a deterministic mechanism M with Pure Nash Equilibria. Let u be a valuation profile such that all agents have identical preferences, and assume the equilibrium assigns item i to agent i. Construct a profile where the first half of agents highly value their top item and the second half value the top half of items slightly less. The optimal social welfare is then at least $\frac{n}{2}$, but the mechanism's welfare is at most 2, leading to a Price of Anarchy of $\Omega(n)$. □

5.6 On More General Equilibrium Concepts

In previous sections, we used pure Nash equilibriua to bound the inefficiency of mechanisms. Here, we extend the results to broader equilibrium concepts: coarse correlated equilibrium (for games with complete information) and Bayes-Nash equilibrium (for games with incomplete information). Since concepts like mixed Nash equilibrium are special cases of these, they suffice for our purposes. One can find a formal definition of these Equilirbium Concepts in Chapter 2. .

Bellow we state the extensions of our theorems, omitting the proofs, which can be found in [8],

Extensions to Random Priority follow naturally, as agents, even with probabilistic mixtures over strategies, always choose their most preferred available item when chosen. The order of remaining items does not affect the distribution.

Theorem 5.11. The coarse correlated Price of Anarchy of Random Priority is $O(\sqrt{n})$. The Bayesian Price of Anarchy of Random Priority is $O(\sqrt{n})$.

For the Probabilistic Serial mechanism, the results hold for both coarse correlated and Bayes-Nash equilibria. While tactics like Roughgarden's *smoothness* framework, are used to establish Price of Anarchy bounds, here the authors we employ a simpler approach. Instead of analyzing all possible outcomes, they focus on the items an agent actually receives, which leads to similarly strong results. This makes the analysis more straightforward while achieving comparable bounds. The full proofs can be found in [8],

Theorem 5.12. The coarse correlated Price of Anarchy of Probabilistic Serial is $O(\sqrt{n})$.

For the incomplete information setting, when valuations are drawn from some publically known distributions, we can prove the same upper bound on the Bayesian Price of Anarchy of the mechanism.

Theorem 5.13. The Bayesian Price of Anarchy of Probabilistic Serial is $O(\sqrt{n})$.

5.7 On the Price of Stability

Theorem 5.7 provides a bound the Price of Anarchy of all mechanisms. A more optimistic (and hence stronger when proving lower bounds) measure of efficiency is the *Price of Stability*, formally defined in Chapter 2., Price of Stability is the worst-case ratio over all valuation profiles of the optimal social welfare over the welfare attained at the *best* equilibrium.

In this section, we extend Theorem 5.7 to the Price of Stability of all mechanisms, that satisfy a "proportionality" property.

Definition 5.14 (Stochastic Dominance)**.** Let $a_1 \succ_i a_2 \succ_i \cdots \succ_i a_n$ be the (possibly weak) preference ordering of agent i. A random assignment vector p_i for agent i *stochastically dominates* another random assignment vector q_i if $\sum_{j=1}^{k} p_{ia_j} \geq \sum_{j=1}^{k} q_{ia_j}$, for all $k = 1, 2, \cdots, n$. The notation that we will use for this relation is $p_i \succ_i^{sd} q_i$.

Definition 5.15 (Safe strategy)**.** Let M be a mechanism. A strategy s_i is a *safe strategy* if for any strategy profile s_{-i} of the other players, it holds that $M_i(s_i, s_{-i}) \succ_i^{sd} \left(\frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n}\right)$.

We will say that a mechanism M has a safe strategy if every agent i has a safe strategy s_i in M. Now, we can state the following theorem:

Theorem 5.16. The pure Price of Stability of any mechanism that has a safe strategy is $\Omega(\sqrt{n})$.

Proof. Let M be a mechanism and $I = \{k+1, \ldots, n\}$ be a subset of agents. We define the valuation profile u as follows:

- For each agent $i \in I$, the utility for items $j = 1, \ldots, k$ is $u_{ij} = \frac{1}{k}$ and for the other items, $u_{ij} = 0$.
- For each agent $i \notin I$, we set $u_{ii} = 1$ (i.e., agent i values their own item most) and for all other items, $u_{ij} = 0$ for $j \neq i$.

Now, let s be a PNE for the profile u, and let s'_i denote a safe strategy for agent i.

At equilibrium s, the expected utility of each agent $i \in I$ is given by:

$$E[u_i(s)] = \sum_{j \in [n]} p_{ij}(s_i, s_{-i}) v_{ij}.$$

Because s is a Nash equilibrium and s'_i a safe strategy, we know that:

$$E[u_i(s)] \geq \sum_{j \in [n]} p_{ij}(s'_i, s_{-i}) v_{ij} \geq \frac{1}{n} \sum_{j \in [n]} v_{ij} = \frac{1}{n}.$$

Additionally, since s is a pure Nash equilibrium, for all agents $i \in I$, the probability that agent i gets one of the first k items is at least $\frac{k}{n}$, meaning:

$$\sum_{j=1}^{k} p_{ij} \geq \frac{k}{n}.$$

For the agents not in I, the total probability that these agents receive one of the first k items is:

$$\sum_{i \in N \setminus I} \sum_{j=1}^{k} p_{ij} = k - \sum_{i \in I} \sum_{j=1}^{k} p_{ij}.$$

Since each agent in I has a probability at least $\frac{k}{n}$ of getting one of the first k items, the total contribution from the agents in I is:

$$\sum_{i \in I} \sum_{j=1}^{k} p_{ij} \geq (n-k)\frac{k}{n}.$$

Thus, the contribution to the social welfare from agents not in I is bounded by:

$$k - (n-k)\frac{k}{n} = \frac{k^2}{n}.$$

The total expected social welfare is the sum of the contributions from agents in I and those outside I. For agents outside I, the maximum contribution is at most $\frac{k^2}{n}$. Therefore, the total expected social welfare of the mechanism

M is at most $1 + \frac{k^2}{n}$. The optimal social welfare $SW_{\text{OPT}}(u)$ is at least k hence, letting $k = \sqrt{n}$ the lower bound of the Price of Stability follows.

□

Due to Theorem 5.16, in order to obtain an $\Omega(\sqrt{n})$ bound for a mechanism M, it suffices to prove that M has a safe strategy. In fact, most reasonable mechanisms, including Random Priority and Probabilistic Serial, as well as all ordinal *envy-free* mechanisms satisfy this property.

Lemma 5.17. Let M be an ordinal, envy-free mechanism. Then for any agent i, the truth-telling strategy u_i is a safe strategy.

Proof. Let $\mathbf{s} = (u_i, \mathbf{s_{-i}})$ be the strategy profile where agent i is truthfully reporting, while the remaining agents follow strategies $\mathbf{s_{-i}}$. Since the mechanism M is envy-free and ordinal, it holds that $\sum_{j=1}^{\ell} p_{ij}^{\mathbf{s}} \geq \sum_{j=1}^{\ell} p_{rj}^{\mathbf{s}}$ for all agents $r \in \{1, \dots, n\}$ and for all $\ell \in \{1, \dots, n\}$. Summing these inequalities over all agents $r = 1, 2, \dots, n$, we get:

$$n \sum_{j=1}^{\ell} p_{ij}^{\mathbf{s}} \geq \sum_{j=1}^{\ell} \sum_{r=1}^{n} p_{rj}^{\mathbf{s}} = \ell,$$

which implies that $\sum_{j=1}^{\ell} p_{ij}^{\mathbf{s}} \geq \frac{\ell}{n}$ for all $i \in \{1, \dots, n\}$ and for all $\ell \in \{1, \dots, n\}$. □

Note that since Probabilistic Serial is ordinal and envy-free, by Lemma 5.17, it has a safe strategy and hence Theorem 5.16 applies. It is not hard to see that Random Priority has a safe strategy too.

Lemma 5.18. Random Priority has a safe strategy.

Proof. Random Priority assigns agents a random order uniformly, each agent i has a probability of $1/n$ of being selected first to choose an item, $2/n$ of being selected within the first two, and so on. If an agent ranks their items truthfully, then for every $\ell = 1, \dots, n$, the probability that they receive one of their top ℓ items is $\sum_{i=1}^{\ell} p_{ij} \geq \frac{\ell}{n}$. □

The safe strategy property plays a very important role ensuring the bound holds. For example, the *randomly dictatorial* mechanism, which selects an agent uniformly at random and assigns them their most preferred item, while assigning the remaining items based solely on that agent's preferences, achieves a constant Price of Stability. However, this same mechanism has a Price of Anarchy of $\Omega(n)$.

CHAPTER 6

BOUNDING THE INCENTIVES OF PS

This chapter is based on the work of Zihe Wang, Zhide Wei and Jie Zhang, in their paper "Social Welfare in One-Sided Matching Mechanisms" presented at the Proceedings of the AAAI Conference on Artificial Intelligence 2020, with some adjustments and clarifications for ease of reading.

6.1 Synopsis

As we have seen in previous chapters, Probabilistic Serial is a celebrated mechanism for the One-Sided Matching problem. While PS is ex-ante Pareto efficient and envy-free, it is well known that PS is not truthful. In this Chapter, we present the work of Wang et al. in [5], where they examined the degree to which an agent has an incentive to manipulate the mechanism.

The textitincentive ratio, formally defined in Chapter 2., is a measure, that quantifies the maximum potential gain an agent can have by deviating from being truthful. In this work, we will show that no agent can increase their utility by more than 50% through strategic behavior, in the PS mechanism. This worst-case guarantee holds under conditions where the agent has complete information about others' reports and can compute the best response, even if doing so is computationally hard ([6]).

In addition to this worst-case analysis, the authors conduct an experiment to evaluate the incentives of agents to manipulate PS in the average-case. The results show that the incentive-ratio in the average-case is much smaller than the worst-case bound, with observed values ranging from 1.02 to 1.06, well below 1.5, which by an example, show that is tight. These results highlight the resilience of the Probabilistic Serial mechanism to strategic manipulation, supporting the idea that the PS mechanism is "approximately" truthful in practice, even in small-scale instances. Future research could focus on developing alternative mechanisms that can better balance fairness, efficiency and truthfulness.

6.2 Introduction

In this chapter, we bound the extent to which an agent can increase their utility through strategic manipulation by reporting preferences, not consistent with their true valuations. Chapter 6 follows naturally from Chapter 5, in which we studied the inefficiencies caused by strategic behavior system-wide but not at the individual level.

In Chapter 5, we analyzed the Price of Anarchy (PoA) to understand how individual self-interest in one-sided matching mechanisms can result in inefficiency. We compared mechanism's outcomes under self-interested behavior to those which are socially optimal and thereby obtained upper and lower bounds on the Price of Anarchy. In Chapter 6, we turn attention to the Probabilistic Serial (PS) mechanism, building on the work of Wang et al. (2020) in [5]. Once again, we investigate how agents may manipulate the system by misreporting their preferences, but this time, we bound the ratio between an agent's utility when reporting truthfully and when strategizing.

The Incentive Ratio, formally defined in Chapter 2, is a measure in mechanism design that quantifies the maximum potential gain an agent can achieve by deviating from truthful behavior. Specifically, it evaluates how much an agent's utility can be increased through strategic manipulation compared to the utility they would receive by reporting truthfully. We adopt the *incentive ratio* notion to quantify agents' incentives to deviate from reporting their actual private information. Informally, the Incentive Ratio is the factor representing the largest possible utility gain that an agent can achieve by behaving strategically, assuming all other agents' strategies are fixed.

The main theorem of this work can be stated as follows:

Theorem 6.1. In the Probabilistic Serial mechanism, when the number of agents is no less than the number of items, no agent is able to unilaterally manipulate and increase their utility to more than $\frac{3}{2}$ times the utility they would receive when reporting truthfully.

There are two scenarios where agents might refrain from manipulation in non-incentive-compatible mechanisms. First, when computational complexity makes manipulation difficult, agents may act truthfully. Second, if utility gains from manipulation are minimal and the cost of gathering necessary information is high, agents may prefer truthfulness.

The incentive ratio is defined in a worst-case sense, representing the strongest approximation guarantee for manipulation incentives. However, this bound overlooks the likelihood of extreme cases; the probability of an agent achieving $3/2$ times their utility is often negligible, as our tight bound example is artificially constructed.

The results presented here assume complete information and perfect rationality. That is, agents are assumed to have complete information about other agents' preferences and are able to compute the best response strategy accordingly. If either of these assumptions is missing, the agents' power to manipulate the mechanism would be much smaller than the $\frac{3}{2}$ bound implies. In fact, computing the best response strategy is intractable in general [6].

The One-Sided Matching problem consists of n agents and m divisible items which we denote as $j, j \in \{1, \ldots, m\}$. In general, n and m are not necessarily equal. The results of this paper hold for the case that $n \geq m$, but cannot be extended to the case $n < m$ straightforwardly. W.l.o.g. we can suppose that m is an integer multiple of n, by adding items of no utility to any agent to achieve this. In the Probabilistic Serial mechanism, agents express strict ordinal preferences, $\succ$, over items. In other words, they are not indifferent between any two items. The assumption is to simplify our analysis; otherwise, one needs to equip the Probabilistic Serial mechanism with a tie-breaking rule, like the ones that we saw earlier in this thesis, that make no significant difference in quantifying the incentive ratio.

The expected utility of agent i is $u_i = \sum_j u_{ij} p_{ij}$. Where $0 \leq u_{ij} \leq 1$ denotes the utility derived by agent i on obtaining a unit of item j and p_{ij}, the probability of agent i receiving item j. Since, PS is not truthful. Agents may misreport their ordinal preferences if that results in a better allocation in their perspective. In that case, $u_i(s_i, s_{-i}) < u_i(s'_i, s_{-i})$, where s_i is agent i's true preference (or true strategy), s_{-i} is other agents' preferences, and s'_i is a misreport by agent i.

The *incentive ratio* captures the extent to which utilities can be increased by strategic plays of individuals. The *incentive ratio* of agent i in mechanism M is:

$$r_i(M) = \max_{\mathbf{s}_{-i}} \frac{\max_{\mathbf{s}'_i} u'_i(s'_i, s_{-i})}{u_i(\mathbf{s})}.$$

Where, $\mathbf{s} = (s_i, s_{-i})$, $u_i(\mathbf{s})$ denotes the utility of agent i when they truthfully reports their preferences and $\max_{\mathbf{s}'_i} u'_i(s'_i, s_{-i})$ denotes the largest possible utility of agent i when she unilaterally misreports her preferences. The **incentive ratio of agent** i is then the maximum value of the ratio over all possible inputs of other agents. The **incentive ratio of a mechanism** M is then $\max_i r_i(M)$. Throughout the paper, w.l.o.g., we consider the strategic manipulation of agent 1.

6.2.1 Probabilistic Serial is not Truthful

We have seen in Section 1 that PS is not truthful. The following example from [1], can serve as a reminder why.

Example 6.2. *Let agents $A = \{A_1, A_2, A_3\}$ and items $I = \{I_1, I_2, I_3\}$ and the preferences of the agents as follows:*

$$A_1 : I_2 \succ I_1 \succ I_3$$
$$A_2 : I_1 \succ I_2 \succ I_3$$
$$A_3 : I_1 \succ I_3 \succ I_2$$

PS will return the following allocation.

Agent / Item	I_1	I_2	I_3
A_1	0	$\frac{3}{4}$	$\frac{1}{4}$
A_2	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$
A_3	$\frac{1}{2}$	0	$\frac{1}{2}$

Table 6.1: Allocation produced by PS

Assume that agent A_1 misreport as follows:

$$A_1 : I_1 \succ I_2 \succ I_3$$
$$A_2 : I_1 \succ I_2 \succ I_3$$
$$A_3 : I_1 \succ I_3 \succ I_2$$

PS will return the following allocation.

Agent / Item	I_1	I_2	I_3
A_1	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{6}$
A_2	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{6}$
A_3	$\frac{1}{3}$	0	$\frac{2}{3}$

Table 6.2: Allocation produced by PS

For some underlying cardinal utilities compatible with agent A_1's true ordinal preferences, for example, $u_{A_1}(I_1) = 0.9, u_{A_1}(I_2) = 1, u_{A_1}(I_3) = 0$, $\mathbf{u}_{A_1} = \frac{3}{4}$ in the truthful profile *and $\mathbf{u'}_{A_1} = 0.8$ when A_1 misreports. The intuition behind this is that, both items I_1 and I_2 are important to agent A_1, but item I_1 is more competitive than I_2 as the other two agents place it as their most preferred item.Sso, instead of start eating a less-competitive item I_2, it is better for agent A_1 to start eating item I_1.*

6.3 Bounding the Incentive Ratio

In this section we prove the main theorem of this paper, Theorem 6.3. To prove the $\frac{3}{2}$ incentive ratio upper bound we will need to prove some lemmas, begining by a reduction on the instances we need to consider. We will prove that it is sufficient to consider the instances in which agents' utilities are *dichotomous*. That is, each agent's preferences are either close to 1, or close to 0. This approach is similar to the concept of *quasi-combinatorial valuation functions* we saw in Chapter 3, where agents' valuations are restricted to being ϵ-close to 0 or 1.

Lemma 6.3. Given an agent's truthful valuations of a preference ordering u_i, define the ratio $c = \frac{u_i'}{u_i}$ where u_i' represents the maximum utility that agent i by strategizing. Then, it is **always possible** to construct a corresponding dichotomous valuation v_i, consistent with the truthful preference ordering, such that $c' = \frac{v_i'}{v_i} \geq c$.

Proof. Assume, without loss of generality, that agent 1 prefers item j more than item $j+1$ for all items $\{1, \dots m\}$. Denote as l_j, the length of time that agent 1 spends on eating item j in the truthful profile and l_j' the length of time that agent 1 spends on eating item j in another strategy. By definition,

$$c(u_{1j}) = \frac{u_1'}{u_1} = \frac{u_{11}l_1' + u_{12}l_2' + \cdots + u_{1m}l_m'}{u_{11}l_1 + u_{12}l_2 + \cdots + u_{1m}l_m}.$$

Note that $u_{1j}, \forall j$ is not necessarily either close to 1, or close to 0. We will show that by carefully pushing u_{1j}'s towards 1 and 0, the ratio c is non-decreasing. Let,

$$k =_j \frac{l_1' + l_2' + \cdots + l_j'}{l_1 + l_2 + \cdots + l_j}, \quad c_{\text{max}} = \frac{\sum_{j=1}^{k} l_j'}{\sum_{j=1}^{k} l_j}.$$

Where k The index j (from the list of items) where this cumulative ratio is the highest - the point in the sequence of items where agent 1's relative gain from manipulating their preferences reaches it's peak.

c_{max} is the ratio of the total manipulated utility to the total truthful utility up to the item k, where k is the index that maximizes this ratio - it quantifies the maximum relative benefit that agent 1 can obtain through strategizing, considering the cumulative impact up to item k.

We can rewrite c as

$$c(u_{1j}) = \frac{\sum_{j=1}^{m}(u_{1j} - u_{1,j+1}) \sum_{h=1}^{j} l_h'}{\sum_{j=1}^{m}(u_{1j} - u_{1,j+1}) \sum_{h=1}^{j} l_h}, \ \textit{where} \ u_{1(m+1)} = 0.$$

Now we construct a new preference profile:

$$b_{1j} = \begin{cases} 1 - (j-1)\epsilon, & \text{for } j = 1, \dots, k, \\ (m-j)\epsilon, & \text{for } j = k+1, \dots, m. \end{cases}$$

This new profile is still consistent with the original preference order u_1. As a result, the allocations based on truthful preferences, l_j, remain unchanged. Similarly, using the same strategy, the allocations resulting from any manipulations, l_j', also stay the same.

Moreover,

$$\begin{aligned} c(b_{1j}) &= \frac{\sum_{j=1}^{m}(b_{1j} - b_{1(j+1)}) \sum_{h=1}^{j} l_h'}{\sum_{j=1}^{m}(b_{1j} - b_{1(j+1)}) \sum_{h=1}^{j} l_h} \\ &= \frac{\epsilon \sum_{j \neq k} \sum_{h=1}^{j} l_h' + (1 - (m-2)\epsilon) \sum_{h=1}^{k} l_h'}{\epsilon \sum_{j \neq k} \sum_{h=1}^{j} l_h + (1 - (m-2)\epsilon) \sum_{h=1}^{k} l_h} \xrightarrow{\epsilon \to 0} c_{\text{max}} \end{aligned}$$

□

Intuitively, this lemma shows that due to the ordinal nature of the mechanism, the worst-case incentive ratio arises in scenarios with extreme valuation profiles.

By Lemma 6.3, we can categorize the items $i_{j\in[m]}$ into two groups:

- $\overline{I} = \{i_j \mid u_{1j}$ is close to $1\}$ - those that agent 1 is interested in.
- $\underline{I} = \{i_j \mid u_{1j}$ is close to $0\}$ - those that agent 1 is not interested in.

Without loss of generality, assume that agent 1 is interested in the first k items, so $\overline{I} = \{i_1, \ldots, i_k\}$. In the truthful profile, if agent 1 spends time p_j consuming item j, then their utility is approximately $u_1 \approx \sum_{j=1}^{k} p_j$. It's important to note that agent 1 may not obtain a positive fraction of each item in $\overline{I}$, so some p_j values may be zero. Although agent 1 might also consume some items from $\underline{I}$, their contribution to the total utility is negligible.

Given agents' ordinal preferences, at any moment t, the following lemma compares the amount of each item that is not eaten up yet in two scenarios.

- In the *normal scenario*, all agents eat items according to their reported ordinal preferences as normal.
- In the *pause scenario*, a set of agents is paused from time t for some time while the other agents continue eating normally.

Lemma 6.4. For any item, at any time from moment t until it gets fully consumed, the amount remaining in the *pause scenario* is always at least as large as the amount remaining in the *normal scenario*.

Proof. We will prove this by contradiction. Let t_{inf} be the earliest time at which there exists an item j^* such that the remaining amount of j^* in the pause scenario is less than its remaining amount in the normal scenario.

Consider a small $t_\delta > 0$. At time $t_{\text{inf}} + t_\delta$, for the amount of j^* to be less in the pause scenario, the number of agents consuming j^* in the pause scenario, must be greater than the number of agents consuming j^* in the normal scenario. If this were not the case, the amount of j^* in the pause scenario could not be less than in the normal scenario at this time.

Now, w.l.o.g. let agent 2 be one of the agents who is consuming j^* in the pause scenario but not in the normal scenario at time $t_{\text{inf}} + t_\delta$. Since agents consume items according to their preferences, starting with their most preferred items, the fact that agent 2 is eating j^* in the pause scenario but not in the normal scenario implies that there exists another item j', which is more preferred than j^* in the preference list of agent 2 .

For agent 2 to switch from eating j' to eating j^* in the pause scenario by time $t_{\text{inf}} + t_\delta$, j' must be completely consumed in the pause scenario but still available in the normal scenario at this time. However, this situation contradicts our assumption that t_{inf} is the earliest moment when the remaining amount of any item j^* in the pause scenario is less than in the normal scenario.

Thus, our assumption that such a t_{inf} exists must be incorrect, proving that at no time does the amount of any item j^* in the pause scenario fall below the amount in the normal scenario. ☐

This lemma holds true regardless of how many agents are paused or for how long they are paused. Furthermore, it applies to any input of the PS mechanism, whether the agents are truthful or not.

Now, w.l.o.g. for agent 1, denote T and T' the moment by which all items in $\overline{I}$ are eaten up in the truthful profile u_1 and the strategic profile u_1' respectively and denote as $\tilde{T}$ and $\tilde{T}'$ the moment by which all items in $\overline{I}$ are eaten up while agent 1 is paused all the time - or say, agent 1 is eliminated from the eating process in the truthful profile u_1 and the strategic profile u_1' respectively.

Ignoring the ϵ terms maintaining a strict preference ordering , it's clear that $u_1 = T \leq T' = u_1'$. By Lemma 6.4, we have $T' \leq \tilde{T}'$. Additionally, since agent 1's reports do not affect the eating process once he's is removed, we observe that $\tilde{T} = \tilde{T}'$.

Therefore, to establish our main result $u_1' \leq \frac{3}{2}u_1$, it suffices to show that $\tilde{T} \leq \frac{3}{2}T$. This approach allows us to avoid the need to determine agent 1's best response strategies and instead focus on the additional time required for the other agents to consume the items in the absence of agent 1. This method is used in proving Theorem 6.5 and Case 1 of Theorem 6.6. We prove that our main theorem holds in all of the three possible cases when the number of agents is less or equal to the number of items, according to T.

1. $0 < T < \frac{1}{2}$
2. $\frac{1}{2} \leq T < \frac{2}{3}$
3. $\frac{2}{3} \leq T \leq 1$

Theorem 6.5. 1. When $0 < T < \frac{1}{2}$, the Incentive Ratio is upper bounded by $\frac{3}{2}$.

Proof. We define the set $\overline{I}^* \subseteq \overline{I}$ to be the set of items that agent 1 gets a positive fraction of in the truthful profile. W.l.o.g. let $\overline{I}^* = \{i_{j\in 1,\ldots,k^*}\}.Let,l_j$ be the time agent 1 spends eating the item i_j. Then, i_j is completely eaten by time $\sum_{h=1}^{j} l_h$. The total time agent 1 spends is $T = \sum_{j=1}^{k^*} l_j$, where $0 < T < \frac{1}{2}$.

It takes at least time $\frac{1}{2}$ for two agents to finish and item. At time $\sum_{h=1}^{j} l_h$, there are at least three agents eating $i_j, j = 1, \ldots, k^*$. So, there must be at least three agents eating each item i_j at the moment it is fully consumed.

Consider now the following process for agent 1:

- At time l_1, all agents but those eating item i_1 are paused. Without agent 1, the remaining agents take extra time δ_1 to finish item i_1, where $\delta_1 \leq \frac{l_1}{2}$.
- Repeat this process for all items $i_j \in \bar{I}^*$

The total time for the other agents to finish the items after agent 1 is eliminated is:

$$\tilde{T} \leq \sum_{j=1}^{k^*}(l_j + \delta_j) \leq \sum_{j=1}^{k^*}(l_j + \frac{l_j}{2}) \leq \frac{3}{2}T$$

Therefore, $\tilde{T} \leq \frac{3}{2} \cdot T$, meaning the time increases by at most a factor of $\frac{3}{2}$ when agent 1 is eliminated.

Thereforw, we have that $\tilde{T} \leq \sum_{j=1}^{k^*}(l_j + \delta_j) \leq \sum_{j=1}^{k^*}(l_j + \frac{l_j}{2}) \leq \frac{3}{2}T$. □

The proof of 2. requires a stricter investigation. We state the following theorem, followed by three lemmata, which will constitute the final piece of the proof.

Theorem 6.6. 2. When $\frac{1}{2} \leq T < \frac{2}{3}$, the Incentive Ratio is upper bounded by $\frac{3}{2}$.

Consider any item $i_j \in \overline{I}^*$. If there are at least two other agents consuming the item at the same time as agent 1 at the moment $\sum_{h=1}^{j} l_h$, the theorem follows directly from the proof of Theorem 6.5.

However, we assume the existence of an item for which only one other agent is consuming it simultaneously with agent 1 when it is completely eaten. Such a scenario can occur for only one item, given that it takes a minimum of time $\frac{1}{2}$ for two agents to eat up one item. Therefore, assume that there exists an item that only one other agent is eating it with agent 1 at the moment it is eaten up. Note that there could only exist one such item, as it takes at least time $\frac{1}{2}$ for two agents to eat up one item. Denote this item by $i_{k'}$ and the other agent by agent 2. For ease of notation, let $t_1 = l_1 + \cdots + l_{k'-1}$, $t_2 = l_{k'}$, and $t_3 = l_{k'+1} + \cdots + l_{k^*}$. Then $t_1 + t_2 + t_3 = T$.

Lemma 6.7. $t_1 + t_2 \geq \frac{1}{2}$, $t_3 < \frac{1}{6}$, $t_1 < \frac{1}{3}$, and $t_2 > 2t_3$.

Proof. Since agents 1 and 2 are the only ones consuming item $i_{k'}$, it takes them at least $\frac{1}{2}$ time for them to finish it. Therefore, $t_1 + t_2 \geq \frac{1}{2}$. Additionally, since $T < \frac{2}{3}$, we have $t_3 = T - (t_1 + t_2) < \frac{1}{6}$.

Agent 1 has been consuming item $i_{k'}$ for t_2 time and even if agent 2 started from the beginning, they have consumed it for $t_1 + t_2$ time. This implies $t_2 + (t_1 + t_2) \geq 1$. Given that $t_1 + t_2 \leq T < \frac{2}{3}$, it follows that $t_2 > \frac{1}{3}$ and $t_1 < \frac{1}{3}$. Thus, $t_2 > \frac{1}{3} \geq 2t_3$. □

Let I_1 be the set of items that are consumed by time t_1 in the truthful profile, I_2 the set of items consumed during the interval $(t_1, t_1 + t_2]$, and I_3 the set of items consumed during the interval $(t_1 + t_2, T]$. These three sets include all the items that agent 1 is interested in, i.e., $\overline{I} \subseteq I_1 \cup I_2 \cup I_3$.

Since $t_1 < \frac{1}{3}$, there are at least three agents consuming any item that agent 1 was eating during the interval $[0, t_1]$. Thus, the analysis for this interval is similar to that in Theorem 6.5.

We can conclude that if agent 1 were removed, all items in I_1 would be eaten up within time $\frac{3}{2}t_1$.

For the set I_3, although we do not obtain the same $\frac{3}{2}$ bound straightforwardly, we can derive a slightly looser bound. This bound, combined with other methods for handling I_2, will allow us to establish an overall $\frac{3}{2}$ bound. To proceed, we first show the following lemma.

For the set I_3, we would not obtain the same $\frac{3}{2}$ bound straightforwardly, but are able to obtain a slightly looser bound, which will be used together with some other approaches for handling I_2 to obtain an overall $\frac{3}{2}$ bound. We first show the following lemma.

Lemma 6.8. In the normal scenario, for each item in I_3, at the moment that it is finished, there will be at least two agents other than agents 1 and 2 who will be eating the item.

Proof. For each item $i_j \in I_3$, we distinguish three possible cases:

1. **Case** 1.**:** At the moment $t_1 + t_2$, no agent is eating item i_j. Yet, since $t_3 < \frac{1}{6}$ and i_j is fully consumed within the interval $[t_1 + t_2, t_1 + t_2 + t_3]$, there must be at least six agents eating the item. Amongst these six agents, even if agents 1 and 2 are among them, there are another four agents.
2. **Case** 2.**:** At the moment $t_1 + t_2$, one agent is eating item i_j. In this case, even if this agent is eating item i_j from the beginning, there are is at least $1 - (t_1 + t_2)$ amount of this item remaining, and it will be eaten up before time $t_1 + t_2 + t_3$, by $\frac{1-(t_1+t_2)}{t_3} = 1 + \frac{1-t}{t_3} > 1 + \frac{1-2/3}{1/6} = 3$,we know that there are at least four agents eating the item. Amongst these agents, even if agents 1 and 2 are among them, there are another two agents.
3. **Case** 3.**:** At the moment $t_1 + t_2$, at least two agents are eating item i_j. Since agents 1 and 2 are eating item $i_{k'}$, there must be two different agents, each contributing to the consumption of i_j.

□

If two agents are absent for some time, it will take another two agents the same amount of time to consume the amount of items left over due to their absence. Thus, the following bound, arises from the previous lemma.

Corollary 6.9. After eliminating agent 1 from eating items in I_1, if we eliminate agents 1 and 2 from the moment $\frac{3}{2}t_1 + t_2$, all items in I_3 will be eaten up by at most an extra t_3 time.

Corollary 6.10. After removing agent 1 from consuming the items in I_1, if agents 1 and 2 are both removed starting from time $\frac{3}{2}t_1 + t_2$, then all items in I_3 will be fully consumed within an additional t_3 time.

We now prove Theorem 6.6 by integrating these intermediate results and analyzing item $i_{k'}$ and the set I_3. As outlined in our earlier discussion, we will begin by removing agent 1 from eating item $i_{k'}$. This will delay the time that item $i_{k'}$ will be fully consumed and it will lead to two possible outcomes.

Proof. **Case 1: Extended Eating Process of Item** $i_{k'}$ **:**

We consider a case, where the process of eating the item $i_{k'}$ is extended - in the sense where, some agents who were initially eating items in I_3, may start eating $i_{k'}$ before time $\frac{3}{2}t_1 + t_2 + 2t_3$. We will focus on agents 1 and 2 who have a positive fraction of items in I_3. We define $\{s_1, \cdots, s_j\} \subseteq I_3$, in the truthful profile to be the set of items that either agent 1 or 2 get a positive fraction of. Let c_h and d_h, for $h = 1, \cdots, j$ be the fraction of item s_h that agents 1 and 2 get, respectively. In the strategic profile, agent 1 is eliminated and agent 2 is eating $i_{k'}$. Define z_h as the time when item s_h is fully consumed and let agent 3 to be the first agent who starts eating $i_{k'}$ before time $\frac{3}{2}t_1 + t_2 + 2t_3$. Suppose that at time $\frac{3}{2}t_1 + t_2 + x$, where $0 \leq x \leq 2t_3$, agent 3 starts eating item $i_{k'}$ and this time $\frac{3}{2}t_1 + t_2 + x \in (z_{w-1}, z_w]$ for some $w \leq j$. In the truthful profile, both agents 1 and 2 have not yet started eating items in I_3 by the time $\frac{3}{2}t_1 + t_2$ leading to a delay.

The delay incured is $\frac{1}{2}t_1 + \frac{\sum_{j=1}^{w-1} c_j + \sum_{j=1}^{w-1} d_j}{2}$, where the factor $\frac{1}{2}$ comes from Lemma 6.8, that implies two agents take longer to eat without the third agent.

Now, let agents 2 and 3 will now eat item $i_{k'}$ and pause others. This, will take an adittional time of $\frac{t_2-x}{2}$ where

t_2 is due to the absence of agent 1, x is how much agent 2 has eaten, the $\frac{1}{2}$ is again, due to Lemma 6.8. At time $\frac{3t_1}{2} + t_2 + x + \frac{t_2 - x}{2}$ all agents, except 1, continue eating.

We calculate the amount of time after this moment when the items in I_3 will be fully consumed in the manipulation profile. We categorize this time into three categories.

1. **First Category:** The remaining time for the items in I_3 after the pause is $\frac{t_3 - \sum_{h=1}^{w-1} d_h}{2}$.
2. **Second Category:** If agent 2 starts eating another item $i_q \in I_3$, it adds an additional delay, bounded by $\frac{x}{2} - \frac{3\sum_{h=1}^{w-1} c_h + \sum_{h=1}^{w-1} d_h}{4}$
3. **Third Category:** After the items in I_3 are fully eaten, the time taken for any remaining items is $t_3 - (x - \frac{1}{2} \cdot \sum_{j=1}^{w-1} c_j + d_j$

Summing up the time from all categories, we find that the total time $\tilde{T}$ for all items $I_1 \cup I_2 \cup I_3$ to be eaten is bounded by:

$$\tilde{T} \leq \frac{3}{2} \cdot (t_1 + t_2 + t_3) - \frac{\sum_{h=1}^{w-1} c_h + \sum_{h=1}^{w-1} d_h}{4} \leq 1$$

This shows that the total time in the strategic profile is at most $\frac{3}{2}$ the total time in the truthful profile.

Case 2: No Agent switches back to $i_{k'}$ Before time $\frac{3}{2}t_1 + t_2 + 2t_3$, we assume that agents who where eating items in I_3 continue eating their next items in I_3 instead of switching back to item $i_{k'}$. Agent 1's utility is partitioned into two parts, before and after time $\frac{3}{2}t_1 + t_2 + 2t_3$.

It's known that $u_1'|_{\leq \frac{3}{2}t_1 + t_2 + 2t_3} \leq \frac{3}{2}t_1 + t_2 + 2t_3$

Next, we upper bound $u_1'|_{> \frac{3}{2}t_1 + t_2 + 2t_3}$. After this time, most items are eaten except from, possibly, $i_{k'}$. Due to agent 1's absence, item $i_{k'}$ has t_2 time left but has already been partially eaten by agent 2 for $2t_3$ time. This leaves $t_2 - 2t_3$ of $i_{k'}$.

Agent 1's best strategy is to eat the remaining portion of $i_{k'}$ alongside agent 2, giving agent 1 an adittional utility of at most $\frac{t_2 - 2t_3}{2}$.

Thus,

$$u_1' \leq u_1'|_{\leq \frac{3}{2}t_1 + t_2 + 2t_3} + u_1'|_{> \frac{3}{2}t_1 + t_2 + 2t_3} = \frac{3}{2}t_1 + \frac{3}{2}t_2 + t_3 \leq T$$

Thus, agent 1's total utility from strategizing doesn't exceed $\frac{3}{2}$ times the utility when truthful and the proof is complete. □

About the third and final case, the following claim is trivial:

Theorem 6.11. 3. When $\frac{2}{3} \leq T \leq 1$, the Incentive Ratio is upper bounded by $\frac{3}{2}$.

Since the optimal utility u_1' is upper bound by $\frac{m}{n} \leq 1$. Hence the utility achieved in truthful profile is quite large compared to u_1'. Formally, $r^{PS} = \max \frac{u_1'}{u_1} \leq \frac{1}{\frac{2}{3}} \leq \frac{3}{2}$.

Combining Theorems 6.5, 6.6, 6.11, we complete the proof of our main Theorem.

6.4 A Tight Bound Example

The authors of [5], also provide an example where the Incentive Ratio of an agent (w.l.o.g. agent 1) is exactly $\frac{3}{2}$. Therefore, making the $\frac{3}{2}$ upper bound tight. [5]

Example 6.12. *Let agents $A = \{A_1, A_2, A_3, \ldots, A_n\}$ and items $I = \{i_1, i_2, i_3, \ldots, i_n\}$ and the preferences of the agents as follows:*

$$
\begin{aligned}
A_1 &: i_1 \succ i_2 \succ i_3 \succ \ldots \succ i_{n-1} \succ i_n \\
A_2 &: i_1 \succ i_2 \succ i_3 \succ \ldots \succ i_{n-1} \succ i_n \\
A_3 &: i_2 \succ i_3 \succ i_4 \succ \ldots \succ i_n \succ i_1 \\
\vdots &\quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \quad \vdots \\
A_n &: i_2 \succ i_3 \succ i_4 \succ \ldots \succ i_n \succ i_1
\end{aligned}
$$

Agent A_1 is interested in the first $\frac{n}{2} - 1$ items, i.e., $\overline{I} = \{i_1, \cdots, i_{\frac{n}{2}-1}\}$. Then in the truthful profile, $u_1 = \frac{1}{2}$. Agent 1 will only get half fraction of item 1. By using the strategy $i_2 \succ i_3 \succ \cdots \succ i_{\frac{n}{2}-1} \succ i_1 \succ i_{\frac{n}{2}} \succ \cdots \succ i_n$. Agent A_1 will get $\frac{1}{n-1}$ fraction of item i_2 to $\frac{n}{2} - 1$ and $\frac{1}{4}$ fraction of item A_1. Agent A_1's utility becomes $u_1' = \frac{3}{4}$. So the ratio is $\frac{3}{2}$.

6.5 In the Average - Case

The following section is an analysis of an experiment conducted by the authors of [5], which evaluates the performance of Probabilistic Serial, complementing the analysis of the worst-case bound.

The authors look into the symmetric setting of PS - that is, when the number of agents is equal to the number of items, varying this number from 8 to 20. For each value of n, they generated $10,000$ instances, giving a total of $120,000$ instances. In each instance, the agents' ordinal preferences were generated uniformly at random, and the manipulator's cardinal preferences were made *dichotomous* to maximise the potential utility gain, as suggested by Lemma 6.3.

They varied the number of items the manipulator was interested in, denoted as k, from 2 to 6. For each instance, they enumerate the manipulator's all $k!$ strategies, in order to find out the largest possible utility the agent can obtain. By dividing the largest attained utility by the utility obtained in the truthful profile, we get a ratio to evaluate the agent's utility increment.

The figure 6.1, reveals the following: as the number of items in $\overline{O}$ (the number of items the manipulator is interested in) increases from 2 to 6, the average-case incentive ratio decreases across all values of n, indicating that as an agent becomes interested in more items, the potential for manipulation decreases. The figure also shows that for a fixed number of items in $\overline{O}$, the ratio increases slightly with the number of agents/items n. For example, the ratio is generally higher for $n = 20$ compared to $n = 8$, suggesting a slightly higher potential for manipulation with more agents/items. However, the average-case incentive ratio remains close to 1, ranging from about 1.02 to 1.06, implying that the overall benefit of manipulation is small compared to the theoretical bound of 50%. These results suggest that while there is some potential for manipulation, it is limited, and the PS mechanism is close to being truthful in practical scenarios.

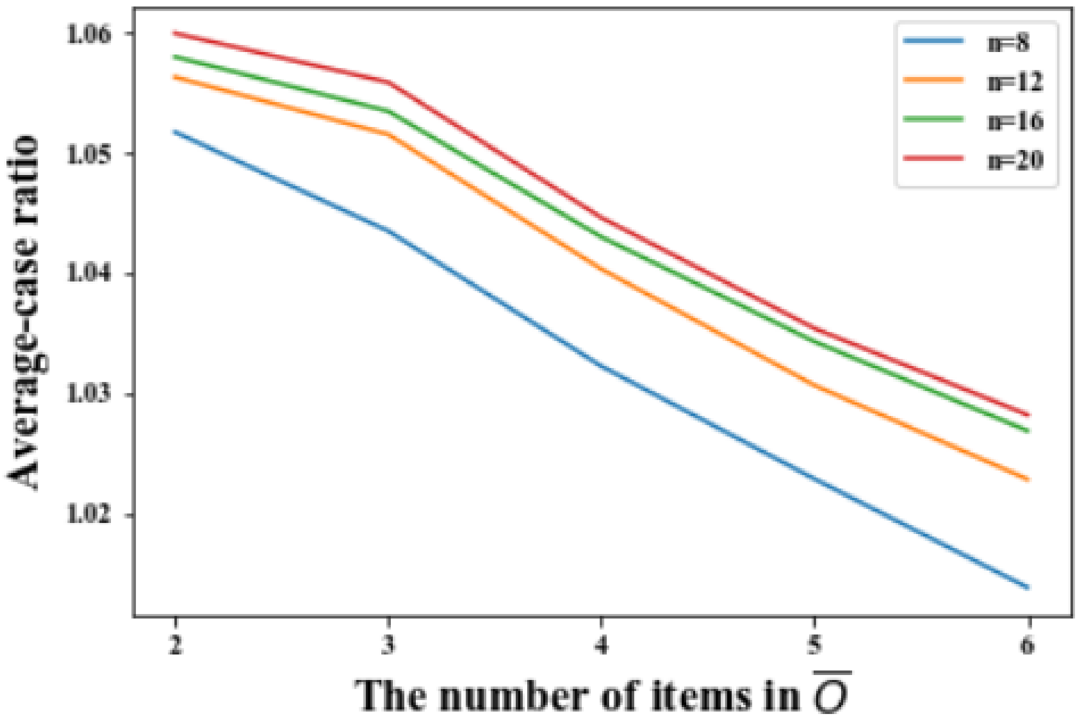

Figure 6.1: Average-case incentive ratio as a function of the number of items $\overline{O}$ that a manipulating agent is interested in. The graph displays different lines for varying numbers of agents/items n (where $n = 8, 12, 16, 20$), image from [5].

CHAPTER 7

THE PRICE OF BEING TRUTHFUL

As we have discussed throughout this thesis, the Probabilistic Serial (PS) mechanism plays an important role in addressing the one-sided matching problem. In Chapter 2, we established that PS is not a truthful mechanism. In Chapter 4, we proved a tight bound of $\Theta(\sqrt{n})$ for the Price of Anarchy of PS. In Chapter 6, we derived a bound of $\frac{3}{2}$ for the Incentive Ratio. However, we demonstrated that cases where PS reaches this $\frac{3}{2}$ bound are highly unusual. In typical scenarios, the Incentive Ratio ranges between 1.02 and 1.06, which is very close to one.

From early on in [8], the authors identify a non-constant lower bound on the Price of Anarchy for the well-known ordinal mechanisms for the problem. However, these findings are crucial for understanding the challenges faced by a social-welfare maximizer.

An interesting direction for future research, as suggested by the authors, would be to explore conditions on the valuation space that could lead to **constant** Price of Anarchy values, or to apply distributional assumptions to quantify the average welfare loss due to selfish behavior.

A key point of interest is understanding how PS performs in terms of Social Welfare when agents report truthfully, especially compared to its performance under equilibrium conditions. Computationally, determining this, seems to be highly complex and appears to be intractable.
Aziz et al. (2015) [6] conducted a series of detailed experiments to better understand the nature and quality of equilibria under the Probabilistic Serial (PS) rule. Given the computational complexity inherent in this analysis, the experiments were limited to scenarios involving small numbers of agents ($n = 2, 3, 4$) and houses ($m = 2, 3, 4$). This is because the size of the search space, which scales as $m!^n$, grows exponentially with the number of agents and houses. For each combination of agents, houses, and preference models, 1000 samples were generated. In total, the experiments required over 40 days of compute time for each model. Due to the need of my PC to write this thesis, we present their results. However, we provide an algorithm providing such "good" Equilibria examples, see in subsection 8.1 and in appendix.

The experiments drew on a variety of preference generation models, including Impartial Culture (IC), Single-Peaked Impartial Culture (SP-IC), the Mallows model, and the Polya-Eggenberger Urn model. These models capture different levels of correlation between agent preferences, from random preferences in IC to correlated preferences in the Urn model. Additionally, real-world data from the PREFLIB AGH Course Selection dataset was used, with students' preferences for courses modeled as agent preferences in the experiments. Each agent's preferences were also associated with a utility value for each house, generated using the Random model. This model assigns a random real number between 0 and 1 to each house for each agent, and these values are then normalized to sum to a constant, representing each agent's total utility for the houses. The Random model was found to be the most manipulable and often led to the worst equilibria.

In the experiments, equilibria were classified into three categories based on social welfare: those where SW remained the same as in the truthful profile, those where SW decreased, and those where SW increased. The percentage change between the social welfare of two profiles, SW_1 and SW_2, was calculated using the formula $\frac{|SW_1 - SW_2|}{SW_1} \times 100\%$. The results revealed several important trends. First, the vast majority of equilibria produced the same social welfare as the truthful profile, with only slight variations. There were generally more equilibria that increased social welfare compared to those that decreased it, though the changes were small, capped at less than 23%. Furthermore, the number of equilibria observed was relatively low, particularly in models with correlated preferences (such as the Urn and AGH models) compared to models with less correlation (such as IC and SP-IC).

From these experiments, the authors suggest, that manipulation under the PS, rarely leads to significant welfare loss. Even when agents strategically misreported their preferences, the resulting equilibria tended to either maintain or slightly improve social welfare compared to the truthful profile. This provides strong empirical support for the robustness of the PS rule in terms of social welfare, even in the presence of strategic behavior. This robustness was observed across all combinations of agents and houses tested, and the trends appeared consistent regardless of the preference models used. The fact that the PS rule performed well in terms of social welfare, even when manipulation was possible, suggests that it is well-suited for use in strategic environments.

The experiments also highlighted the computational challenges involved in computing and verifying equilibria under PS. The large search space required for analyzing all possible misreports and preference profiles makes it computationally difficult to compute pure Nash equilibria (PNE) in general. In fact, the problem of computing a PNE is known to be NP-hard, and verifying whether a given profile is a PNE is coNP-complete. Such computational barriers discourage strategic manipulation: it is hard to find equilibria, so agents do not engage in complex manipulations, particularly in large settings.

The analysis of Aziz et. al. [6] revealed that while the PS rule is not strategy-proof, it offers a level of resistance to manipulation that is both empirically and theoretically significant. The fact that manipulation rarely results in significant welfare loss, combined with the computational difficulty of finding equilibria, makes the PS rule a strong candidate for use in real-world applications involving resource allocation. Despite these positive findings, the study also points to several avenues for future research. One potential direction is to extend the current analysis to cases where indifferences exist in agents' preferences. Additionally, investigating strong Nash equilibria, where groups of agents rather than individual agents deviate from their truthful preferences, could provide further insights into the strategic properties of the PS rule. Another interesting direction would be to explore the dynamics of Nash equilibria in more detail, potentially leading to a deeper understanding of how equilibria evolve over time under the PS rule.

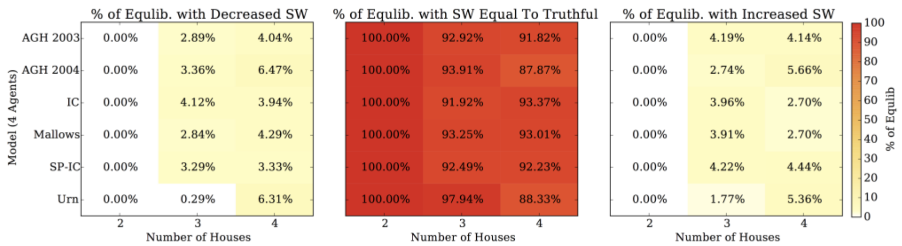

Figure 2: Classification of equilibria for all 1000 samples per setting with four agents ($n = 4$), 2 to 4 houses ($m \in \{2, 3, 4\}$), and preferences drawn from the six models. We can see that the vast majority of the equilibria found across all samples have the same social welfare as the truthful profile. In general, there are roughly the same number of equilibria that increase as those that decrease it.

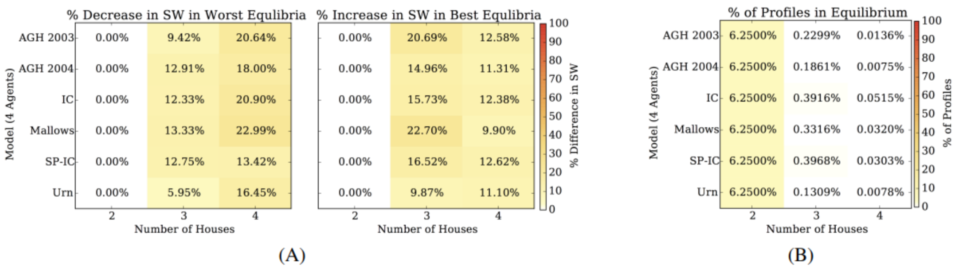

Figure 3: (A) The maximum and minimum percentage increase or decrease in social welfare over all 1000 samples in settings with four agents ($n = 4$), 2 to 4 houses ($m \in \{2, 3, 4\}$), and preferences drawn from the six models. We see that for three houses the gain of the best profile is, in general, slightly more than the loss in the worst profile with respect to the truthful profile; this trend appears to reverse for settings with four houses. (B) The average number of the $m!^4$ profiles that are in equilibria per sample with four agents ($n = 4$), 2 to 4 houses ($m \in \{2, 3, 4\}$). The more uncorrelated models (i.e., IC and SP-IC) admit the highest number of equilibria.

Figure 7.1: Result from the experiment in [6]

7.1 An Example & Further Questions

The analysis of social welfare gain and loss under the Probabilistic Serial Mechanism raises several interesting questions. A measure, similar to the Price of Anarchy, that could quantify of how the social welfare under truthful strategies compares with the worst-case equilibria. Could be of interest, in understanding PS.

$$\sup_{\mathbf{u} \in V^n} \frac{SW(\mathbf{u}_{\text{truthful}})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})}$$

Where, $SW(\mathbf{u}_{\text{truthful}})$ represents the social welfare under the truthful ordinal preferences derived from valuations, while $\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})$ denotes the worst-case social welfare among all pure Nash equilibria $S_{\mathbf{u}}^M$ under the same valuation profile $\mathbf{u}$. The closer it is to 1, it indicates that strategic manipulation does not degrade social welfare at all, meaning the mechanism is robust against strategic behavior.

The experiments in [6] suggest that the number of equilibria with significantly different social welfare values is limited. This opens up the possibility that the bounds of this quatnity could be relatively tight. It is known [22] that Probabilistic Serial and Random Priority converge in large economies. That is as the number of agents and objects increases. Meaning that Probabilistic Serial tends to become truthful. As we saw [5] the incentive ratio for PS is upper bounded by 1.5 when the number of agents is no less than the number of items, while very recent research published earlier this year [30], show that when the items are more than the agents, the upper bound is 2. All these come as a strong indicator that a constant bound on this measure might not be surprising, given the limited deviation in outcomes observed in experiments. Another interesting research direction could be the exploration of any specific types of valuation profiles or distributions of preferences that lead to a higher disparity between truthful and strategic behavior.

7.1.1 Example

Let agents $A = \{A_1, A_2, A_3\}$ and items $I = \{I_1, I_2, I_3\}$. The preferences of the agents are as follows:

$$\begin{aligned} A_1 &: I_2 \succ I_1 \succ I_3 \\ A_2 &: I_3 \succ I_2 \succ I_1 \\ A_3 &: I_3 \succ I_2 \succ I_1 \end{aligned}$$

The utilities of the agents for each item are:

$$\begin{aligned} A_1 &: u(I_1) = 0.51599, \quad u(I_2) = 0.56056, \quad u(I_3) = 0.24756 \\ A_2 &: u(I_1) = 0.03775, \quad u(I_2) = 0.42752, \quad u(I_3) = 0.45960 \\ A_3 &: u(I_1) = 0.58086, \quad u(I_2) = 0.59542, \quad u(I_3) = 0.79550 \end{aligned}$$

The Probabilistic Serial (PS) mechanism returns the following allocation for truthful reporting:

Agent / Item	I_1	I_2	I_3
A_1	$\frac{1}{3}$	$\frac{2}{3}$	0
A_2	$\frac{1}{3}$	$\frac{1}{6}$	$\frac{1}{2}$
A_3	$\frac{1}{3}$	$\frac{1}{6}$	$\frac{1}{2}$

Table 7.1: Allocation produced by PS for truthful reporting

The resulting utilities for each agent in the truthful reporting scenario are:

$$
\begin{aligned}
u(A_1) &= 0.51599 \times \frac{1}{3} + 0.56056 \times \frac{2}{3} + 0.24756 \times 0 = 0.54570 \\
u(A_2) &= 0.03775 \times \frac{1}{3} + 0.42752 \times \frac{1}{6} + 0.45960 \times \frac{1}{2} = 0.31364 \\
u(A_3) &= 0.58086 \times \frac{1}{3} + 0.59542 \times \frac{1}{6} + 0.79550 \times \frac{1}{2} = 0.69061
\end{aligned}
$$

The social welfare by truthful reporting is thus $SW_{Tr.} = 0.54570 + 0.31364 + 0.69061 = 1.54995$.

Assume that agents misreport their preferences as follows:

$$
\begin{aligned}
A_1 &: I_2 \succ I_1 \succ I_3 \\
A_2 &: I_2 \succ I_3 \succ I_1 \\
A_3 &: I_3 \succ I_2 \succ I_1
\end{aligned}
$$

This strategy profile is a Pure Nash Equilibrium. The PS mechanism returns the following allocation:

Agent / Item	I_1	I_2	I_3
A_1	$\frac{1}{2}$	$\frac{1}{2}$	0
A_2	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{1}{4}$
A_3	$\frac{1}{4}$	0	$\frac{3}{4}$

Table 7.2: Allocation produced by PS for the Pure Nash Equilibrium

The resulting utilities for each agent in the Pure Nash Equilibrium scenario are:

$$
\begin{aligned}
u(A_1) &= 0.51599 \times \frac{1}{2} + 0.56056 \times \frac{1}{2} + 0.24756 \times 0 = 0.53828 \\
u(A_2) &= 0.03775 \times \frac{1}{4} + 0.42752 \times \frac{1}{2} + 0.45960 \times \frac{1}{4} = 0.33810 \\
u(A_3) &= 0.58086 \times \frac{1}{4} + 0.59542 \times 0 + 0.79550 \times \frac{3}{4} = 0.74184
\end{aligned}
$$

The social welfare for the Pure Nash Equilibrium is $SW_{PNE} = 0.53828 + 0.33810 + 0.74184 = 1.61822$.

Comparing the social welfare, we observe that $SW_{PNE} = 1.61822$ is higher than $SW_{Tr.} = 1.54995$, indicating that the misreporting strategy leads to a higher overall social welfare.

The example above was constructed by the following algorithm:

Algorithm 1 PS : Be Truthful or not to Be

1. **Generate Random Preference Profiles**
   - Generate uniformly at random, ordinal preference profiles for each player, where each player receives a random permutation of the items. Call this profile Truthful.
2. **Generate Random Valuation Profiles**
   - Based on the ordinal preferences, generate random cardinal preferences, according to the ordinal ones. No ties, are allowed, all must add to the same number.
3. **Apply the Probabilistic Serial (PS) Rule**
   - Use the PS rule to allocate fractions of items to players. Players consume fractions of their most preferred available item until the items are fully allocated. The code for PS was writen by Dominik Peters and is available in Github: [29]
4. **Calculate Social Welfare**
   - Compute the social welfare by summing the products of each player's valuation and their allocated share of the items.
5. **Generate Permutations of Preferences**
   - Generate all possible permutations of the item preferences to explore all possible strategies for each player. The space is now $(m!)^n$ big, where $m =$ # items and $n =$ # agents.
6. **Generate Payoff Matrix**
   - For each combination of strategies (preference profiles), calculate the corresponding allocation using the PS rule, then compute the utilities for all players.
7. **Find Pure Nash Equilibria (PNE)**
   - Check the payoff matrix for combinations where each player is playing a best response. If no player can improve their utility by switching strategies, mark the combination as a PNE.
8. **Compare Social Welfare of PNE and Truthful Strategies**
   - Compare the social welfare of PNE strategies with that of the truthful strategy. If any PNE improves the social welfare, it is identified as a better outcome.
9. **Repeat Process**
   - Repeat the entire process for a fixed number of attempts to explore different random preference profiles and search for better PNE profiles.

BIBLIOGRAPHY

[1] Anna Bogomolnaia and Herve Moulin, *A new solution to the random assignment problem*, Journal of Economic Theory, 100, 2 (2001), 295–328.

[2] Hadi Hosseini, Kate Larson and Robin Cohen, *Investigating the Characteristics of One-Sided Matching Mechanisms Under Various Preferences and Risk Attitudes*, 2017, arXiv:1703.00320.

[3] Joel Watson, *Strategy: An Introduction to Game Theory*, Norton Company, 2013.

[4] Akshay-Kumar Katta and Jay Sethuraman, *A solution to the random assignment problem on the full preference domain*, Journal of Economic Theory, Volume 131, Issue 1, 2006, Pages 231-250, ISSN 0022-0531, doi:10.1016/j.jet.2005.05.001.

[5] Zihe Wang, Zhide Wei and Jie Zhang, *Bounded Incentives in Manipulating the Probabilistic Serial Rule*, 2020, arXiv:2001.10640.

[6] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska and Toby Walsh, *Equilibria Under the Probabilistic Serial Rule*, 2015, arXiv:1502.04888.

[7] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li and Chenhao Wang, *Mechanism Design for Facility Location Problems: A Survey*, 2021, arXiv:2106.03457.

[8] George Christodoulou, Aris Filos-Ratsikas, Soren Kristoffer Stiil Frederiksen, Paul W. Goldberg, Jie Zhang and Jinshan Zhang, *Social Welfare in One-Sided Matching Mechanisms*, 2016, arXiv:1502.03849.

[9] Sissi Jiang, Ndiame Ndiaye, Adrian Vetta and Eggie Wu, *The Price of Anarchy of the Asymmetric One-Sided Allocation Problem*, 2023, arXiv:2301.01367.

[10] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen and Jie Zhang, *Social Welfare in one-sided matchings: Random Priority and beyond*, 2014, arXiv:1403.1508.

[11] Atila Abdulkadiroglu and Tayfun Sönmez, *Matching Markets: Theory and Practice*, Advances in Economics and Econometrics: Tenth World Congress, Volume I: Economic Theory, 2011, Pages 1-47, ISBN 9781107016040, doi:10.1017/CBO9781139060011.002.

[12] Roth, A. E. and A. Postlewaite, *Weak versus Strong Domination in a Market with Indivisible Goods*, Journal of Mathematical Economics, 4, 131-137 (1977).

[13] Roth, A. E., *The Economics of Matching: Stability and Incentives*, Mathematics of Operations Research, 7, 617-628 (1982).

[14] Roth, A. E., *The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory*, Journal of Political Economy, 92, 991-1016 (1984).

[15] Shapley, L. and H. Scarf, *On Cores and Indivisibility*, Journal of Mathematical Economics, 1, 23-28 (1974).

[16] Hylland, A. and R. Zeckhauser, *The Efficient Allocation of Individuals to Positions*, Journal of Political Economy, 87, 293-314 (1979).

[17] Echenique, F., Immorlica, N., Vazirani, V. V. (eds.), *One-Sided Matching Markets*, In: Online and Matching-Based Market Design, Cambridge University Press, 2023, Pages 37-65.

[18] Gale, D. and L. Shapley, *College Admissions and the Stability of Marriage*, American Mathematical Monthly, 69, 9-15 (1962).

[19] Ma, J., *Strategy-Proofness and the Strict Core in a Market with Indivisibilities*, International Journal of Game Theory, 23, 75-83 (1994).

[20] Arnosti, N. *Ex ante vs ex post guarantees (+ open problem)*, https://nickarnosti.com/blog/exanteexpost/ (2023, January 28).

[21] Ruta Mehta, Debmalya Mandal, and Ruta Bhattacharya, *Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation*, arXiv preprint arXiv:2005.14122, 2020, arxiv:2005.14122.

[22] Yeon-Koo Che and Fuhito Kojima, *Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms*, Cowles Foundation Discussion Papers, 2008, elischolar:1991.

[23] Eshwar Ram Arunachaleswaran, Nina Balcan, and R. Ravi, *Fairness and Utilization in Allocating Resources with Uncertain Demand*, arXiv preprint arXiv:1906.09050, 2019, arxiv:1906.09050.

[24] John Nash, *Non-Cooperative Games*, Annals of Mathematics, 54, 286-295, 1951.

[25] Elias Koutsoupias and Christos Papadimitriou, *Worst-case equilibria*, In Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), 1999, cgi.di.uoa.gr/ elias/papers/paper-kp99.pdf.

[26] Procaccia, A. D., and M. Tennenholtz, *Approximate Mechanism Design without Money*, ACM Transactions on Economics and Computation, 1, 4, Article 18, 1-26 (2013).

[27] Bhalgat, A., Chakrabarty, D., and Khanna, S., *Social Welfare in One-Sided Matching Markets Without Money*, In APPROX-RANDOM, 87-98 (2011).

[28] Feige, U. and Tennenholtz, M., *Responsive lotteries*, In SAGT, pages 150--161 (2010).

[29] Peters, D., *Simple implementation of the Probabilistic Serial fair division mechanism*, `https://gist.github.com/DominikPeters/943e102c6d633e4adb024b086b5b5042`, Accessed on September 10, 2024.

[30] Haoqiang Huang, Zihe Wang, Zhide Wei, and Jie Zhang, *Bounded incentives in manipulating the probabilistic serial rule*, Journal of Computer and System Sciences, Vol. 140, 2024, doi:10.1016/j.jcss.2023.103491.