ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ
ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Randomized truthful mechanisms
(Πιθανοτικοί φιλαλήθεις μηχανισμοί)

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Ιωάννης Παναγέας

Επιβλέπων: Ευστάθιος Ζάχος
Καθηγητής

Αθήνα, Οκτώβριος 2010

ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ
ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
Τομέας Τεχνολογίας Πληροφορικής και Υπολογιστών

Randomized truthful mechanisms
(Πιθανοτικοί φιλαλήθεις μηχανισμοί)

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Ιωάννης Παναγέας

Επιβλέπων: Ευστάθιος Ζάχος
Καθηγητής

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 29η Οκτωρίου 2010

.....	.....	.....
Ευστάθιος Ζάχος	Δημήτρης Φωτάκης	Άρης Παγουρτζής
Καθηγητής	Λέκτορας	Επίκουρος Καθηγητής

Αθήνα, Οκτώβριος 2010

.....
Ιωάννης Λ. Παναγέας
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π

Περίληψη

Στη διπλωματική αυτή, αντιμετωπίζουμε προβλήματα του mechanism design, χρησιμοποιώντας πιθανοτικούς αλγορίθμους ως ιδέα επίλυσης. Εξετάζουμε πως οι πιθανοτικοί μηχανισμοί συμπεριφέρονται στις ψηφοφορίες, ένας σημαντικός τομέας της θεωρίας κοινωνικής επιλογής. Θεωρούμε διαφορετικά είδη ψηφοφοριών, και προσπαθούμε να προσεγγίσουμε το βέλτιστο score του υποψηφίου που κερδίζει. Οι πιθανοτικοί μηχανισμοί ουσιαστικά είναι το κλειδί για να πετύχουμε φιλαλήθεις μηχανισμούς. Επιπλέον κάνουμε μία φιλική εισαγωγή στο Differential privacy που είναι ουσιαστικά μία οικογένεια πιθανοτικών αλγορίθμων που έχει επίσης μια επιπρόσθετη παράμετρο ϵ που επηρεάζει το μέσο κόστος ενός αλγορίθμου και την πολυπλοκότητά του. Οι Differential privacy μηχανισμοί ικανοποιούν τον περιορισμό ότι λίγες αλλαγές στα (κρυφά) δεδομένα επίσης προκαλλούν μικρές αλλαγές στην πυκνότητα πιθανότητας των αποτελεσμάτων των αλγορίθμων. Αυτή η ιδέα είναι σημαντική για να πετύχουμε προσεγγιστική φιλαλήθεια επειδή το να παρεκκλίνει κάποιος προκαλλεί επίσης μικρές αλλαγές στη συνάρτηση χρησιμότητας των παικτών.

Λέξεις κλειδιά

Θεωρία Παιγνίων, Σχεδιαστικοί μηχανισμοί, Πιθανοτικοί Αλγόριθμοι, Differential Privacy

Abstract

In this thesis, we deal with problems of mechanism design, using randomization as a solution concept. We examine how randomization behaves in voting rules, a very important subfield of social choice theory. We consider different voting rules, and try to approximate the optimal score of the winner alternative. Randomization actually is the key to achieve strategy-proof mechanisms. Moreover we make a friendly introduction to Differential privacy which is actually a family of randomized algorithms that has additionally a parameter ϵ that affects the expected cost of the algorithm and its complexity. Differential privacy mechanisms satisfy the constraint that few changes in the (private) data also cause small changes at the probability distribution of the outcome of the algorithm. This idea is important to achieve approximate truthfulness because deviating causes small changes in the utility function of each agent too.

Keywords

Game Theory, Mechanism Design, Randomization, Differential Privacy

Acknowledgements

I would like to deeply thank my three teachers and members of the committee Prof. S.Zachos, Prof. D. Fotakis and Prof. A. Pagourtzis. I feel very lucky I had the chance to be taught by them; I feel they have influenced me in a very positive way not only academically but also on a personal level. Particularly I want to thank the two supervisors of this thesis, Prof. S. Zachos and Prof. D. Fotakis for their guidance and their persistence in helping me make this thesis better.
Additionally, I specifically want to thank my friend Andreas Galanis for being such an inspiration for me from the early ages, when we participated in Mathematical Olympiads. Many thanks to the members of the Corelab team for the interesting discussions we made. Finally, i want to thank my parents for their support since the first day of my birth.

25/10/2010, Ioannis Panageas

Contents

List of Algorithms

Chapter 1

Introduction

In this chapter, we make a brief introduction to Game Theory and Mechanism Design, knowledge that we will need to proceed with the rest of this diploma thesis. Additionally, we make a friendly introduction to Randomized Algorithms, because this is the point of view from which we try to examine different problems of Mechanism Design.

1.1 Introduction to Game Theory

1.1.1 Definitions

Game Theory is considered a subfield of applied mathematics and has a lot of applications to economics and computer science. It aims to mathematically capture situations, or *games*, in which players(agents) interact and their decisions affect each other's outcomes. Agents are modeled to be *rational* and *intelligent*, namely they are self-interested and are aware of all the existing knowledge of the game and capable of making all the logical inferences.

As far as the games are concerned, we will examine those that are known as *one-shot simultaneous*. Generally, a game consists of $N = \{1, 2, ..., n\}$ agents, each of whom choose a way of playing from a set of possible ways, namely a way of interacting to the game. This way of playing for agent i, is usually denoted by $s_i \in S_i$ and is called agent i's *strategy*, where S_i is a well defined set of available strategies of i. The question that arises is what are the criteria for which the agents choose their strategies. The answer is simple, each agent i has a *preference* profile and her preferences determine her strategy (choice).

We distinct the strategies of each agent to *pure* strategies and *mixed* strategies.

A pure strategy for agent i is just an element of S_i and a mixed strategy is a probability distribution over the elements of S_i.

We use the notation $s = \{s_1, ..., s_n\}$ for the strategy vector of the agents. Additionally, we also use the notation $s = (s_i, s_{-i})$, where $s_{-i} = \{s_1, ..., s_{i-1}, s_{i+1}, ..., s_n\}$, namely the strategy vector of the agents except of agent i.

As each agent selects her strategy, then the outcome of the game is determined. In order to specify the game and also the preference profile of each agent, we have to "measure" the outcome of the game for each agent. Thus, we define the function $u_i : S \to \mathcal{R}$ for agent i, where $S = S_1 \times ... \times S_n$. u_i is called *utility* function and shows the "happiness" of agent i for the selected vector s of strategies, thus the outcome of the game.

1.1.2 NE concept

There are lots of solution concepts to deal with a game, described above. The most intuitive and common solution concept is the *Nash Equilibrium* notion, where each agent tries to maximize his $u_i(s_i, s_{-i})$ given that the strategies of the other agents, namely s_{-i} is fixed. Formally, we have the following definition:

Definition 1: A strategy vector $s \in S$ is said to be a Nash equilibrium if for all players i and each alternate strategy $s'_i \in S_i$, we have that:

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

In words, there is no agent i such that, she can increase her payoff by changing her strategy. In pure strategies notion, whether there is a Nash equilibrium or not depends on the game. However, in mixed strategies, as John Nash proved, there is always a mixed strategy Nash equilibrium (we take expectation over utilities) and hence the model of mixed strategies equilibria is very useful and important, even though finding such an equilibrium is hard to compute ($PPAD$-complete problem).

A stronger notion than Nash equilibrium, is *Dominant strategy* equilibrium, where each agent chooses her strategy to maximize her payoff(utility function), independently of all other agents strategies. Formally we have the following definition.

Definition 2: A strategy $s_i \in S_i$ is a dominant strategy if for each alternate strategy $s'_i \in S_i$ and each $s_{-i} \in S_{-i} = \{S_1 \times ... \times S_{i-1} \times S_{i+1} \times ...S_n\}$, we have that:

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

If all agents have a dominant strategy, then the strategy vector which is formed, is a dominant strategy equilibrium.

It is rather straightforward to see that a dominant strategy equilibrium is also a Nash equilibrium. Additionally, a dominant strategy equilibrium and thus a Nash equilibrium doesn't yield the maximum payoff the agents can have, namely the maximum payoff that the agents can have is not usually a stable solution (there is at least one agent such that he can increase her utility by changing her strategy). For the rest of the diploma thesis, we search for dominant strategy equilibrium to the problems we deal with. Below we give some examples to make our claims more clear.

1.1.3 Examples

Prisoners Dilemma: (Dominant-strategy equilibrium) The following matrix expresses the years of jails for agents P_1, P_2 for all the possible selections of strategies for the two agents.

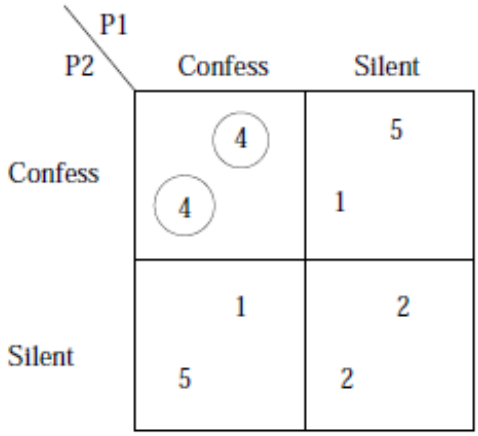

Cost matrix for prisoners dilemma

The game we describe is a symmetric game. Each player has available strategies $S_i = \{\text{Confess}, \text{Silent}\}$. In this game, each agent wants to minimize the years of jail. We assume that $u_i(x) \geq u_i(y)$ if $x \leq y$ (x, y are years of jail for agent i). Thus for $s_1 = \text{Confess}$ we have that $u_1(\text{Confess}, \text{Silent}) \geq u_1(\text{Silent}, \text{Silent})$ and $u_1(\text{Confess}, \text{Confess}) \geq u_1(\text{Silent}, \text{Confess})$, thus for agent 1, the strategy "confess" is a dominant strategy. Additionally, by symmetry this also holds for agent 2. Hence the vector strategy $s = \{\text{Confess}, \text{Confess}\}$ is a dominant strategy equilibrium and hence a Nash equilibrium. Notice that the strategy vector $s' = \{\text{Silent}, \text{Silent}\}$ gives a better payoff, namely 2 years of jail for each agent,

however it's not a stable solution, since for example agent 1 can change her strategy and have 1 year of jail instead of 2.

Firms price: (No dominant-strategy equilibrium) The following matrix expresses payoff of Firms 1,2 for all the possible strategies.

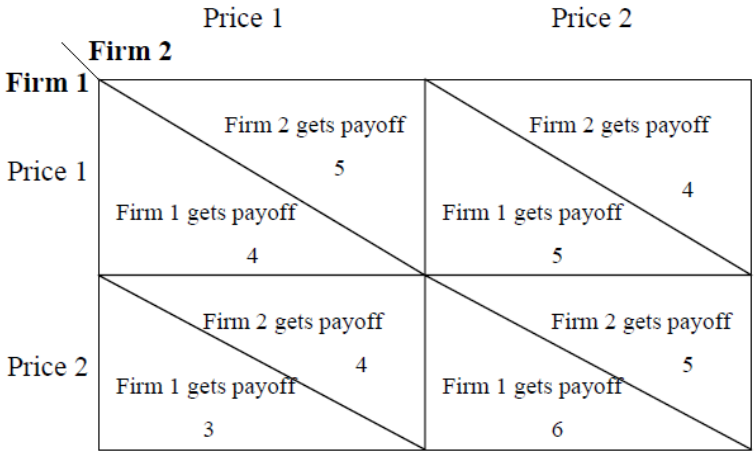

Cost matrix for Firms

This game has two Nash equilibria, namely $s = \{\text{Price 1}, \text{Price 1}\}$ and $s' = \{\text{Price 2}, \text{Price 2}\}$ (this can be easily proved if considering all the cases). However there is no dominant-strategy equilibrium. Consider the first Nash equilibrium s. Then strategy "Price 1" is not a dominant strategy for Firm 1 and hence s is not a dominant strategy equilibrium. Consider the second Nash equilibrium s'. Then strategy "Price 2" is not a dominant strategy for Firm 1 and hence s' is not a dominant strategy equilibrium.

Generally, dominant strategy equilibrium is more desired than Nash equilibrium because, each agent make no assumptions about the strategies of the other agents, she doesn't need any knowledge about the way the others play.

1.2 Introduction to Mechanism Design

Mechanism design constitutes the intersection of the field of economics and the field of game theory. This happens since we talk about self-interested and intelligent agents that want to maximize their payoff and we are searching for stable solutions (NE concept). Additionally, if we add to this concept algorithmic design (and analysis) and complexity we are talking about *Algorithmic* Mechanism Design. In essence, Mechanism design attempts implementing desired social choices in a

strategic setting. A social choice constitutes aggregations of the preferences of the different agents toward a single joint decision. It is remarkable that usually the preferences of the agents are private, that's why such strategic design is necessary.

1.2.1 Definitions

As we have already mentioned in previous section, each agent shows preferences over the different outcomes of the game (an outcome of the game is determined since every agent chooses a strategy). This showing preference is denoted by the *type* of agent i, $\theta_i \in \Theta_i$. For example assume two possible outcomes $a_j, a_k \in A$, where A is the set of all possible outcomes. Then given θ_i, we can conclude whether agent i prefers a_j to a_k or not. Thus θ_i is a transitive relation over the different outcomes and it can be seen as a partially ordering set over outcomes. Generally, we give a real number for each outcome for agent i, namely we "measure" that outcome with respect to agent i. This becomes attainable with the use of utility function.

Definition 1: Let $u_i : \theta_i \times A \to R$ be the utility function of i. We consider $u_i(\theta_i, a_j) \geq u_i(\theta_i, a_k)$ if i (weakly) prefers outcome a_j to a_k.

The goal of Mechanism design, is to find functions that given the types of the agents, selects the optimal outcome (in some criteria).

Definition 2: A function $f : \times_i \Theta_i \to A$ is called a social choice function.

A mechanism $\mathcal{M}$ given the game, defines the set of feasible strategies for the agents, namely defines $S_1, ..., S_n$ and additionally determines the outcome rule of the game, namely defines a function g, such that $g(s)$ is the outcome where $s = \{s_1, ..., s_n\}$ is the strategy vector played by the agents. In order to make a connection between a mechanism and the definitions described above, we have additionally to define the following:

Definition 3: A mechanism $\mathcal{M}$ *implements* social function f, if for every type vector of agents $\theta = (\theta_1, ..., \theta_n)$, we have that $f(\theta) = g(s_1(\theta_1), ..., s_n(\theta_n))$, where $s_i : \Theta_i \to S_i$ gives the strategy agent i plays, given her profile and also $(s_1(\theta_1), ..., s_n(\theta_n))$ is an equilibrium of the game.

Specifically, we want the equilibrium to be a dominant strategy equilibrium, because it makes the least assumptions about the agents. Designing mechanisms that implement social functions is a very difficult task, because each agent acts strategically and she may try to manipulate the mechanism so as to increase her

utility or expected utility. For example, if we assume that $S_i = \Theta_i$, namely the strategy of each agent is to report her type, an agent may misreport her type in order to increase her utility and then the outcome of the game may be far from the desired one. Below, we describe an example in order to make more clear the definitions mentioned above.

1.2.2 Examples

Let Alice and Bob be two energy consumers. An energy authority is charged with choosing the type of energy to be used by Alice and Bob. The different kind of energies are $\{\text{gas}, \text{oil}, \text{nuclear power}, \text{coal}\}$. Let us suppose that there are two possible states of the world. In state 1, the consumers place relatively little weight on the future, i.e., they have comparatively high temporal discount rates. In state 2, by contrast, they attach a great deal of importance to the future, meaning that their rates of discount are correspondingly low. Formally we have that $\Theta_A = \Theta_B = \{\text{state 1}, \text{state 2}\}$ and $A = \{\text{gas}, \text{oil}, \text{nuclear power}, \text{coal}\}$. The matrix below shows the consumers' energy rankings in the two states:

State 1		State 2	
Alice	**Bob**	**Alice**	**Bob**
gas	nuclear	nuclear	oil
oil	oil	gas	gas
coal	coal	coal	coal
nuclear	gas	oil	nuclear

The energy authority wants to select an energy source such that both Alice and Bob are reasonably happy with. Hence a social choice function f that makes both consumers reasonable happy is $f(\text{state 1}) = \text{oil}$ and $f(\text{state 2}) = \text{gas}$. However, the energy authority doesn't know which state holds and by asking the consumers which is the state of the world, both Alice and Bob have the incentive to lie. Indeed, since Alice prefers gas to oil, she will report that the state of the world is the second, without caring about Bob's answer. Additionally, since Bob prefers oil to gas, he will answer that the state of the world is the first. Hence the only reasonable thought is to flip a coin between the two states and then decide the alternative (type of energy). However, in that case, the probability of returning the optimal solution is 50%, namely the optimal type of energy conditionally on what state holds.

Let us suppose, therefore, that the authority has the consumers participate in the mechanism given by matrix below:

		Bob: Left	Bob: Right
Alice	Top	oil	coal
	Bottom	nuclear	gas

Each energy consumer must choose between two strategies. For Alice the two available strategies are $\{\text{Top}, \text{Bottom}\}$ and for Bob are $\{\text{Left}, \text{Right}\}$. We consider the following cases:

- The state of the world is the first. Then Bob will choose the strategy Left, no matter what Alice chooses, since he prefers oil and nuclear to coal and gas. Thus strategy Left is a dominant strategy for Bob. Additionally, Alice since Bob selects Left, she will select strategy Top, since she prefers oil to nuclear. Hence for state 1, the expected way of playing is $(s_A, s_B) = (\text{Top}, \text{Left})$ since each the only Nash equilibrium for that state.
- The state of the world is the second. Then clearly, Alice will choose the strategy Bottom, no matter what Bob chooses, since she prefers nuclear and gas to oil and coal. Thus strategy Bottom is a dominant strategy for Alice. Additionally, Bob since Alice plays Bottom, he will select the strategy Right, since he prefers gas to nuclear. Hence for state 2, the expected way of playing is $(s_A, s_B) = (\text{Bottom}, \text{Right})$, which is the only Nash equilibrium for state 2.

Thus in both states, the mechanism described above achieves the optimal outcome even though the energy authority doesn't know the actual state and the consumers are interested in their preferences. Finally, observe that the mechanism implements authority's social choice in Nash equilibrium since the Nash equilibrium outcomes of the mechanism coincide with the optimal outcomes in each state.

1.2.3 More Definitions

Definition 1: (*Pareto optimality or efficiency*). A function f is pareto-optimal if there is no alternative outcome (from the returned outcomes) that makes at least one agent better off without making any other agent worse off.

Intuitively, the social function f must not return an outcome a_i such that there exists an outcome a_j where all the agents have larger utilities than at a_i. Consider for example the Firms price problem. Nash equilibrium $s = \{\text{Price } 2, \text{Price } 2\}$ must be the only choice of f. This outcome is called pareto-optimal solution. It

is remarkable that a pareto-optimal solution may not be stable, namely an equilibrium. In the example of prisoners' dilemma, $\{\text{Silent}, \text{Silent}\}$ is a pareto-optimal solution but is not a Nash equilibrium.

Definition 2: (*direct-revelation mechanism*). A direct-revelation mechanism $\mathcal{M}$ restricts agent's i strategies $S_i = \Theta_i$ for all i and also defines the outcome $g(\hat{\theta}_1, ..., \hat{\theta}_n)$ of the game, where $(\hat{\theta}_1, ..., \hat{\theta}_n)$ is agents' reported vector-type.

In words, a direct-revelation mechanism is mechanism in which the only available strategy for the agents is to report their type, hence $s_i(\theta_i) = \hat{\theta}_i$ ($\hat{\theta}_i$ is the reported preference for agent i). If agent i *truthfully* reports her type, then $s_i(\theta_i) = \theta_i$.

Definition 3: (*Incentive-compatible Mechanism*). An incentive-compatible mechanism is a direct-revelation mechanism in which agents report truthful information about their preferences in equilibrium, namely $s_i(\theta_i) = \theta_i$ for every i.

Definition 4: (*Strategy-proof*). An incentive compatible mechanism is strategy proof, if for every agent, truthfully reporting her preference is a dominant-strategy.

Namely, a strategy-proof mechanism, is an incentive compatible mechanism in which agents report truthful information about their preferences in specifically dominant-strategy equilibrium.

There are games, such as auctions, where additionally there is a payment vector introduced to the game. For this reason we have to generalize our definition of utility functions and mechanisms.

Definition 5: A quasi-linear utility function for agent i is denoted by

$$u_i(\theta_i, a_j) = v_i(\theta_i, a_j) - p_i$$

where θ_i, v_i are agent's i type and valuation function respectively and a_j the outcome of the game.

In a similar way to the previous definitions, a quasi-linear mechanism $\mathcal{M}$ given the game, defines the set of feasible strategies for the agents, namely defines $S_1, ..., S_n$ it determines the outcome rule of the game, namely defines a function g, such that $g(s)$ is the outcome where $s = \{s_1, ..., s_n\}$ is the strategy vector played by the agents and also defines the payment $p_i = t_i(s)$ of agent i (payment vector $p = (p_1, ..., p_n)$). To make things more clear consider the following auction:

Example in auctions

Assume we have an auction of a single item, and a set N of n bidders. The set of outcomes is the winners of the auction (notice that we have a single winner). Formally $A = \{i\text{-wins} \mid i \in N\}$. Also, the valuation of each bidder is denoted by $v_i(i\text{-wins}) = w_i \geq 0$ (how much money she is willing to pay for the item, formally it shows the type of agent i) and $v_i(j - wins) = 0$, $j \neq i$ (intuitively, every bidder wants to win). Let g be a direct-revelation mechanism that selects the winner bidder. Thus g also determines the payment vector $p = (p_1, ..., p_n)$ for the bidders. Hence the utility function of bidder i is denoted by $u_i = v_i - p_i$. In order for the auction to be reasonable, we have to demand that $p_i \geq 0$. In case i wins then $u_i(i\text{-wins}) = v_i(i\text{-wins}) - p_i(i\text{-wins}) \geq 0$, or equivalently $w_i \geq p_i(i\text{-wins})$. Additionally, if $j \neq i$ wins then $u_i(j\text{-wins}) = 0 - p_i(j\text{-wins}) \geq 0$ and hence $p_i = 0$. In words, bidder i pays money iff i wins. Consider the following strategy-proof mechanism:

(*Vickrey's second price mechanism*)
Let the winner be the player i with the highest bid (w_i) and i pays p, which is the second highest declared bid, namely $p = \max_{j \neq i} w_j$. All the other bidders pay zero.

To prove that truthtelling is a dominant strategy for the mechanism above, let w_i be i's valuation and p^* the second largest valuation. Consider the following cases (w_i' is i's reported valuation):

- i wins the auction. If $w_i' > p^*$ then i would still win thus $u_i(w_i') = w_i - p^* = u_i(w_i)$. If $w_i \leq p^*$ then i would lose the auctions thus $0 = u_i(w_i') < w_i - p^* = u_i(w_i)$.
- i loses the auction, namely there exists j such that $w_i < w_j$ and j is the winner. If $w_i' < w_j$ then i still loses, thus $u_i(w_i) = 0 = u_i(w_i')$. If $w_i' > w_j$ then i would win the auction and $u_i(w_i') = w_i - w_j < 0 = u_i(w_i)$.

Hence truthtelling is a dominant strategy for every bidder i, thus truthtelling is a dominant strategy equilibrium (namely the mechanism is strategy-proof).

The mechanism above belongs to a family of mechanisms, called VCG-mechanisms. The VCG mechanisms are direct-revelation mechanisms with payments vectors and are defined as follows. The outcome $a(\hat{\theta})$ - where $\hat{\theta} = (\hat{\theta}_1, ..., \hat{\theta}_n)$ is the reported type vector- that is chosen is the one that maximizes the *social welfare* (sum of valuations) namely

$$a(\hat{\theta}) = \max_{a_j \in A} \sum_i v_i(\hat{\theta}_i, a_j)$$

Additionally, the payment is denoted by

$$p_i(\hat{\theta}) = h_i(\hat{\theta}_{-i}) - \sum_{j \neq i} v_i(\hat{\theta}_j, a(\hat{\theta}))$$

where $h_i : \Theta_{-i} \to \mathcal{R}$ is an arbitrary function. Notice that the amount agent i pays doesn't depend on her type (valuation) but on all other agents. Agent's i utility is equal to $u_i(\theta_i, a(\hat{\theta})) = v_i(\theta_i, a(\hat{\theta})) + \sum_{j \neq i} v_i(\hat{\theta}_j, a(\hat{\theta})) - h_i(\hat{\theta}_{-i})$. This is the key to achieve strategy-proofness since each agent has incentive to maximize the social welfare (since h_i doesn't depend on i) and this is feasible by telling the truth. Vickrey [29], Clarke [6], Groves [12] proved that the family of mechanisms described above are strategy-proof and that these mechanisms are the only strategy-proof mechanisms that maximize the sum of the utility functions.

Commonly, we define $h_i = \max a \in A \sum_{j \neq i} v_i(a, \theta_i)$ which is called *Clarke pivot rule*. Notice that in the example described above, we consider the Clarke pivot rule to solve the problem.

1.2.4 Important Theorems

Theorem 1: (Revelation Principle) If there exists an arbitrary mechanism $\mathcal{M}$ that implements social choice function f in dominant strategies, then there exists an incentive compatible mechanism that implements f in dominant strategies.

Proof: Since $\mathcal{M}$ implements f then $g(s_1(\theta_1), ..., s_n(\theta_n)) = f(\theta_1, ..., \theta_n)$, where $s = (s_1(\theta_1), ..., s_n(\theta_n))$ is a dominant-strategy equilibrium. We define a direct-revelation mechanism $\mathcal{M}'$ with $g'(\theta'_1, ..., \theta'_n) = f(\theta'_1, ..., \theta'_n)$, $\forall \theta' = (\theta'_1, ..., \theta'_n)$ and we will prove that $\mathcal{M}'$ is strategy-proof. Since $(s_1(\theta_1), ..., s_n(\theta_n))$ is a dominant strategy we have that

$$u_i(\theta_i, g(s_i(\theta_i), s'_{-i}(\theta_{-i}))) \geq u_i(\theta_i, g(s'_i(\theta_i), s'_{-i}(\theta_{-i}))) \; \forall s'_i \in S_i, s'_{-i} \in S_{-i}$$

We substitute $s'_{-i}(\theta_{-i})$ for $s_{-i}(\theta_{-i})$ and also $s'_i(\theta_i)$ for $s_i(\hat{\theta}_i)$ we have that the following holds:

$$u_i(\theta_i, g(s_i(\theta_i), s_{-i}(\theta_{-i}))) \geq u_i(\theta_i, g(s_i(\hat{\theta}_i), s_{-i}(\theta_{-i}))) \; \forall \hat{\theta}_i \in \Theta_i, \theta_{-i} \in \Theta_{-i}$$

or equivalently

$$u_i(\theta_i, g'(\theta_i, \theta_{-i})) \geq u_i(\theta_i, g'(\hat{\theta}_i, \theta_{-i})) \; \forall \hat{\theta}_i \in \Theta_i, \theta_{-i} \in \Theta_{-i}$$

Hence, for every agent i, reporting her true type is a dominant strategy, thus direct-revelation mechanism $\mathcal{M}'$ implements f in dominant strategies.

□

In words, Revelation Principle states that any mechanism $\mathcal{M}$ that implements social function f in dominant strategies can be transformed into a strategy-proof mechanism that implements the same social function. It is remarkable, that this theorem give us no knowledge of how to create that mechanism. However it is a very important theorem since it restricts our attention to direct-revelation and truthful mechanisms, since if a social function f can be implemented in dominant strategies by any mechanism, then it can be also implemented by an incentive compatible one. Another point that has to be stressed is that Revelation principle also stands for mechanisms that have also payment vectors, namely to mechanism design with money.

Definition 1: Social function f is a *dictatorship*, if there is an agent i such that for all possible outcomes $a_j \in A$, we have that $u_i(\theta_i, f(\theta_i, \theta_{-i})) \geq u_i(\theta_i, a_j)$.

Namely, a social function is a dictatorship for agent i, if the outcome of f is always the most preferable for agent i, no matter what the types of the other agents are.

Definition 2: Social function f is onto A if for every $a_j \in A$, there is a vector type θ such that $f(\theta) = a_j$.

Below we mention a very important theorem for social choice theory, which is an impossibility result.

Theorem 2: (Gibbard-Satterthwaite) There is no f that is incentive compatible social choice function onto A, where $|A| \geq 3$ and is not a dictatorship.

This theorem actually states that (because we are interested in dominant strategies equilibria) that there isn't a strategy-proof mechanism that is "reasonable", in a sense that every outcome can be returned and simultaneously not to be a dictatorship. There are many ways to circumvent Gibbard-Satterthwaite's [10],[26] theorem such as the restriction of the structure of agents preferences, the design of mechanisms where finding a manipulation is computationally hard or the use of randomization. In this diploma thesis, we will try to circumvent this impossibility result using randomization for different games.

1.3 Randomization

There exists lots of problems in mathematics and computer science, that are inefficiently solvable (NP-hard), or polynomially solvable, but with large complexity. Additionally, some algorithms are really difficult to be implemented. Thus in order to deal with this fact, we turn our attention to randomization. A randomized algorithm, is an algorithm that makes random choices, with respect to a *random number generator*. Generally, randomized algorithms are simply to implement and quicker than the deterministic ones (otherwise there wasn't a reason for examining them). However, their expected outcomes usually approximate the optimal outcome in optimization problems, and have a probability error in decision problems. In this section, we make a friendly introduction to probabilities and randomized algorithms. Subsection 1.3.1 exists for completeness reasons. The reader should see Motwani et al. [19], before continuing with the rest of this diploma thesis.

1.3.1 Introduction to Probabilities and Randomization

Definition 1: A real-valued random variable X on a sample space Ω is a function X mapping the elementary events of Ω to real numbers, namely $X : \Omega \to \mathcal{R}$.

Since an algorithm makes random choices, we don't have a unique outcome, thus we take the expectation of the possible outcomes (for optimization problems). Formally, expectation is the following:

Definition 2: Let X be a discrete random variable and let $f : \mathcal{R} \to \mathcal{R}$. Then the expectation is denoted by:

$$\mathbb{E}[f(X)] = \sum_x \Pr[X = x] f(x)$$

Additionally, we can define the conditional expectation, namely $\mathbb{E}[Y|Z = z] = \sum_y y \Pr[Y = y|Z = z]$.

Definition 3: Let X be a random variable with expectation μ_X. Then *variance* is denoted by $\sigma_X^2 = \mathbb{E}[(X - \mu_X)^2]$.

Definition 4: Let X, Y be random variables. X, Y are independent iff $\forall\ x, y$ we have that $\Pr[X = x \wedge Y = y] = \Pr[X = x] \cdot \Pr[Y = y]$.

Below, we mention some useful inequalities that are used in analyzing of the performance of the randomized algorithms.

Theorem 1: *Markov's* inequality: Let X a random variable, $t > 0$, the following holds:

$$\Pr[X \geq t] \leq \frac{\mathbb{E}[X]}{t}$$

Proof:

$$\begin{aligned}\mathbb{E}[X] &= \sum_x x \Pr[X = x] \\ &\geq \sum_{x \geq t} x \Pr[X = x] \\ &\geq \sum_{x \geq t} t \Pr[X = x] \\ &= t \Pr[X \geq t]\end{aligned}$$

□

It is remarkable that Markov inequality is tight. Consider the random variable X_t which is t if $X_t \geq t$ and 0 otherwise. Then $\mathbb{E}[X_t] = t \cdot \Pr[X_t \geq t]$.

Theorem 2: *Chebychef's* inequality: Let X a random variable with expectation μ_X and variance σ_X^2, the following holds:

$$\Pr[|X - \mu_X| \geq t \cdot \sigma_X] \leq \frac{1}{t^2}$$

Proof: Let $Y = (X - \mu)^2$. Then from Markov we have that $\Pr[Y \geq t^2 \cdot \sigma_X^2] \leq \frac{\mathbb{E}[Y]}{t^2 \sigma_X^2}$. Equivalently, $\Pr[|X - \mu_X| \geq t \cdot \sigma_X] \leq \frac{1}{t^2}$.

□

Theorem 3: *Chernoff Bounds* Let $X_1, ..., X_n$ be independent random variables with
$\Pr[X_i = 1] = p_i$ and $\Pr[X_i = 0] = 1 - p_i$. Then it holds that $\forall\ 0 \leq \epsilon \leq 1$:

$$\Pr[X > (1 + \epsilon)\mu_X] \leq e^{-\epsilon^2 \mu_X / 3} \text{ and also } \Pr[X < (1 - \epsilon)\mu_X] \leq e^{-\epsilon^2 \mu_X / 2}$$

where $X = \sum_{i=1}^n X_i$.

Proof: Let $M_X(t) = e^{tX}$, $t > 0$. Then from Markov's inequality we have that $\Pr[M_X(t) \geq e^{at}] \leq \frac{\mathbb{E}[M_X(t)]}{e^{at}}$. Equivalently, since $X_1, ..., X_n$ are independent, we have that $\mathbb{E}[M_X(t)] = \prod_{i=1}^n [(1 - p_i) + p_i e^t]$. Moreover, $1 - p_i + p_i e^t \leq e^{p_i (e^t - 1)}$ since $e^x \geq x + 1$. Hence, $\mathbb{E}[M_X(t)] \leq e^{(e^t - 1)(\sum_i p_i)} = e^{(e^t - 1)\mu_X}$. Thus for $a = (1 + \epsilon)\mu_X$

we have that $\Pr[X \geq (1+\epsilon)\mu_X] \leq \left(\frac{e^{(e^t-1)}}{e^{(1+\epsilon)t}}\right)^{\mu_X}$. Finally, for $t = \ln(1+\epsilon)$ is minimized, and since $(1+\epsilon)\ln(1+\epsilon) - \epsilon \geq \epsilon^2/3$ (it comes from simple calculus), we have that $\Pr[X > (1+\epsilon)\mu_X] \leq \Pr[X \geq (1+\epsilon)\mu_X] \leq \left(e^{\epsilon-(1+\epsilon)\ln(1+\epsilon)}\right)^{\mu_X} \leq e^{-\mu_X\epsilon^2/3}$.

For the second inequality, we define $M'_X(t) = e^{-tX}$, $t > 0$. Thus similarly, we have that $\Pr[X \leq a] = \Pr[M'_X(t) \geq e^{-at}] \leq \frac{e^{(e^{-t}-1)\mu_X}}{e^{-at}}$.we substitute a for $(1-\epsilon)\mu_X$, thus we have that $\Pr[X \leq (1-\epsilon)\mu_X] \leq \left(\frac{e^{(e^{-t}-1)}}{e^{-(1-\epsilon)t}}\right)^{\mu_X}$. Similarly, for $t = -\ln(1-\epsilon) > 0$ is minimized thus, since $-\epsilon-(1-\epsilon)\ln(1-\epsilon) < -\epsilon^2/2$ we have that $\Pr[X < (1-\epsilon)\mu_X] \leq \Pr[X \leq (1-\epsilon)\mu_X] \leq e^{-\mu_X\epsilon^2/2}$

□

The intuitive part of the theorem, is that the probability that X takes values far from the expectation is exponentially low, that is we have a strong concentration around the expectation.

Definition 5: Let A be a randomized algorithm and $Cost_A(I)$ be A's cost of output (for specific choices) for input I. The expected score of A is denoted $\mathbb{E}[Cost_A(I)]$. The approximation ratio of a randomized mechanism is defined as follows, namely we consider the *worst-case* scenario:

Definition 6: The approximation ratio of a randomized algorithm is denoted by for maximization problems

$$\min_I \mathbb{E}\left[\frac{Cost_A(I)}{OPT(I)}\right] \text{ or } \max_I \mathbb{E}\left[\frac{OPT(I)}{Cost_A(I)}\right]$$

and for minimization problems

$$\max_I \mathbb{E}\left[\frac{Cost_A(I)}{OPT(I)}\right] \text{ or } \min_I \mathbb{E}\left[\frac{OPT(I)}{Cost_A(I)}\right]$$

Examples for this scenario can be seen in the rest of the diploma thesis and also in subsection 1.3.2. There are also other scenarios, such as *average-case* where we examine the performance of the algorithm in average-case. One common example for this is Quicksort with random pivot. In the average case, the complexity of Quicksort is $O(n \log n)$.

In decision problems, we usually find the probability of the algorithm to be correct over the worst input, namely

$$\min_I \Pr[A(I) \text{ answers correctly}]$$

Usually, one of the answers "YES" or "NO" is correct with probability 1. Thus assume that if A answers "YES" then is correct and if A answers "NO" and the probability of error is $p_e < 1$. Applying A a lot of times (let n) (depends on p_e, the experiments must be independent) then we can make the probability of error very small, namely p_e^n (negligible). A common example is Primality test, or checking matrix multiplication. This family of randomized algorithms are called *Monte-Carlo*. Additionally, there are the *Las Vegas* algorithms, where the output is always correct, however the complexity of the algorithm may be unbounded (in that case we check the expected complexity).

In order to find lower bounds with respect to approximation ratio of randomized algorithms, a common way is to apply Yao's Minimax Principle[1]. In words, the cost of the best randomized algorithm over the worst deterministic input cannot be better than the cost of the best deterministic algorithm over the worst (randomized) input. If we substitute approximation ratio for the cost of the algorithm, it follows that the approximation ratio of the best randomized algorithm cannot be better than the approximation ratio of the best deterministic algorithm (over the worst randomized input).

1.3.2 Examples

Example 1 - Max cut (maximization). Let $G(V, E)$ be an undirected graph. Find a partition of $V = V_1 \cup V_2$ such that the edges between V_1 and V_2 is maximized. Consider the following algorithm:

Algorithm (E1)
1. **For** each vertex $v \in V$ 2. **flip** a fair coin

We will prove that $E1$ has approximation ratio of $\frac{1}{2}$. Let $I(V, E)$ be an arbitrary input. Firstly, we have that $\mathbb{E}\left[\frac{Cost_{E1}}{OPT}\right] \geq \frac{\mathbb{E}[Cost_{E1}]}{|E|}$. Let X_e be a random variable which is 1 if e is between V_1, V_2 and zero otherwise. Then $\mathbb{E}[Cost_{E1}] = \mathbb{E}[\sum_{e \in E} X_e] = \sum_{e \in E} \mathbb{E}[X_e]$. Additionally, $\mathbb{E}[X_e] = \Pr[e \text{ between } V_1, V_2] = \frac{1}{2}$. Hence $\mathbb{E}[Cost_{E1}] = \sum_{e \in E} \frac{1}{2} = \frac{|E|}{2}$, from which follows that $\mathbb{E}\left[\frac{Cost_{E1}}{OPT}\right] \geq \frac{1}{2}$.

Example 2 - checking matrix multiplication (decision). Let A, B, C be three $n \times n$ matrices. We want to check whether $AB = C$ holds or not. The easy

[1] See appendix for statement and proof

deterministic way, has complexity $O(n^3)$, however we can do better using randomization. Consider the following algorithm:

Algorithm (E2)

1. Pick a vector from $\mathbf{d} \in \{0, 1, ..., S\}^n$ uniformly at random
2. **if** $AB \cdot d = C \cdot d$ output "YES?"
3. **else** output "NO"

The complexity if the randomized is actually $O(n^2)$ since we use associative property to calculate $AB \cdot d = A(B \cdot d)$. Obviously, if $E2$ outputs "NO", then $E2$ is correct. We have to find the probability that $E2$ is wrong conditionally on outputting "YES". Let $S = AB - C, S \neq 0$ and assume $S_{ij} \neq 0$. Then if $E2$ outputs "YES" (thus $E2$ is wrong) then it holds that

$$d_j = -\frac{\sum_{t \neq j} d_t S_{it}}{S_{ij}} \text{ (i)}$$

So only if $E2$ selects d_j such that (i) holds, then it outputs wrongly. Hence $\Pr[\text{error}] \leq \frac{1}{S+1}$. Thus, if we run $E2$ m independent times (each time $E2$ outputs "YES"), then the probability of failure is at most $\left(\frac{1}{S+1}\right)^m$.

Example 3 - Fingerprinting (decision). Let $a = a_0 a_1 ... a_{n-1}$,$b = b_0 b_1 ... b_{n-1}$ be the numbers of Alice and Bob respectively. Alice and Bob want to check if $a = b$, without transmitting all n bits of the numbers to each other. Consider the following randomized algorithm.

Algorithm (E3)

1. Alice **picks** a prime p from $\{2, ..., T\}$ uniformly at random
2. She **sends** to Bob p, $t_a = a \mod p$
3. Bob calculates $t_b = b \mod p$
4. **if** $t_a = t_b$ output "YES?"
5. **else** output "NO"

Obviously, if $E3$ outputs "NO", then $E3$ outputs correctly. Our goal is to find the probability of error, if $E3$ answers "YES". From the analysis, we will find the size of the data that has to be transmitted ($= O(\log T)$), such that the probability of error is negligible. Let $E3$ outputs "YES". Then p divides $|a - b|$. $|a - b|$ is a n-bit number, thus it has at most n prime factors (each prime ≥ 2). Hence, the probability of error is at most $\frac{n}{\pi(T)}$, where $\pi(x)$ is the number of primes less than or equal to x. Hence $\Pr[\text{error}] \leq \frac{n}{\pi(T)}$. Since it holds that $\frac{x}{\ln x} \leq \pi(x)$, $\forall x \geq 17$, we have that $\Pr[\text{error}] \leq \frac{n}{\pi(T)} \leq \frac{n \ln T}{T}$. Thus for $T = cn \ln n$, we

have that Pr[error] $\leq \frac{1}{c} + o(1)$ and the number of bits transmitted is at most $\log T = O(\log n)$.

1.4 Purpose of thesis

The purpose of this thesis is to deal with problems of mechanism design, using randomization as a solution concept. In chapter 2, we examine how randomization behaves in voting rules, a very important subfield of social choice theory. We consider different voting rules, and try to approximate the optimal score of the winner alternative. Randomization actually is the key to achieve strategy-proof mechanisms. Moreover, in chapter 3, we make a friendly introduction to Differential privacy which is actually a family of randomized algorithms that has additionally a parameter ϵ that affects the expected cost of the algorithm and its complexity. Differential privacy mechanisms satisfy the constraint that few changes in the (private) data also cause small changes at the probability distribution of the outcome of the algorithm. This idea is important to achieve approximate truthfulness because deviating causes small changes in the utility function of each agent too. Finally, we examine applications of differential privacy to combinatorial optimization problems in order to understand in practice how actually works and the intuitive part of it.

Chapter 2

Randomization in voting

2.1 Introduction

The Gibbard-Satterthwaite [10],[26] theorem makes it impossible to have a "reasonable" and also strategy-proof voting. Bartholdi et al. [3], Conitzer & Sandholm [7], Hemaspaandra & Hemaspaandra [14] showed that manipulation is computationally hard for a variety of voting rules. However, this approach happens in the worst-case scenario, but most of the times, the prominent voting rules are usually easy to manipulate. Thus, in this chapter, we try to circumvent Gibbard-Satterthwaite theorem using randomization, with respect to score-based voting rules. Even though a strategy-proof randomized voting is a combination of non-"reasonable" deterministic strategy-proof mechanisms [11], the existence of randomization makes the situation more fair. More precisely, we examine randomized mechanisms that choose an alternative whose *expected* score is an approximation of the optimal score, namely the score of the alternative who should win [25]. Finally, we will examine another special case (problem) of approval voting and try to approximate the optimal solution using randomization.

2.2 Introduction to Voting Setting

In this chapter, we examine the traditional social choice point of view, namely voting. We assume that we have a set of agents(voters) that want to elect a winner from a set of alternatives(candidates). Our goal is to elect the most desirable alternative (based on some criteria). Formally, we have the following model:

Let N be the set of agents and A the set of alternatives where $|N| = n$ and $|A| = m$ respectively. Each agent has a preference $\prec_i$, which is a line order over the alternatives (let $\mathcal{L}$ the set of line orders over the A). A voting rule, is a function $f : \mathcal{L}^n \to A$, namely a social choice function.

There are various voting systems some of which are examined in sections below. There are voting systems such that the number of voters is far larger than the number of alternatives (political election case). Additionally, there are voting systems that the previous case is reversed (web search framework). Another case, is that there maybe weights on the voters, namely some agents' preferences affect the outcome more than others.

Some of the most common voting rules are *Plurality, Borda, Veto, Condorcet, Dodgson, Copeland.*

In Condorcet rule, the winner must satisfy the *condorcet criterion*, namely a winner of an election is the one who would beat all other alternatives in head-to-head elections (more than half of the agents must prefer the winner to every other alternative). However, consider the following instance, known as condorcet's paradox: Let v_1, v_2, v_3 be 3 agents and a, b, c be 3 alternatives. We have the following preference profiles: $v_1 : c \prec b \prec a$, $v_2 : a \prec c \prec b$ and $v_3 : b \prec a \prec c$. $b \prec a$ since v_1, v_3 prefer a to b. For similar reasons $c \prec b$ and $a \prec b$. However this is inconsistent. Notice that for every winner, at least half of the agents would prefer another alternative and they would want to change the outcome.

Another example is Dodgson's rule, the winner alternative pf which (let a) is the one with the minimum *score*, where a's equals to the smallest number of sequential exchanges in the voters' preference lists that suffices to make a, a Condorcet winner. Observe that the rules described above and also the ones that will be seen in next section, satisfy pareto optimality, namely if $\forall i \in N$, we have that $b \prec_i a$ then b can't be a winner.

Unfortunately, recalling Gibbard's-Satterthwaite theorem, we can't have a strategy-proof f such that f is not a dictatorship and also for every candidate i, there is a profile preference vector that f chooses i to be the winner, namely we can't have a strategy-proof and also reasonable voting rule. We will turn our attention to vector - positional scoring rules and also approval voting rules, because using randomization, we can choose an alternative that approximates the optimal score of the alternative that should be the winner.

2.3 Vector - Positional Scoring Rules

2.3.1 Definitions

For this chapter, we assume that that the voting rules are randomized, that is $f : \mathcal{L}^n \to \Delta(A)$, where $\Delta(A)$ is a probability distribution over the set of alternatives A. In words, f chooses alternative i with probability p_i, according to a probability distribution P.

Definition 1: *Vector - Positional scoring:* Given a m-dimensional vector $\mathbf{a} = [a_1, ..., a_m]$ of non-negative and non-decreasing coordinates and the preference profile of the agents (let $\prec$), the score of alternative x is denoted by

$$scr(x, \prec) = \sum_{j=1}^{n} a_{\prec_j(x)}$$

where $\prec_j (x) = k$ iff x is the k-th most preferable alternative for agent j.

Below we mention some well-known scoring rules:

- Plurality: is defined by the vector $\mathbf{a} = [1, 0, ..., 0]$.
- Borda: is defined by the vector $\mathbf{a} = [m-1, m-2, ..., 0]$.
- Veto: is defined by the vector $\mathbf{a} = [1, 1, ..., 1, 0]$.

For example, assume we have 4 agents and 3 alternatives, and let the preference profiles of the agents be $(\prec_1, \prec_2, \prec_3, \prec_4) = (\{1, 3, 2\}, \{1, 2, 3\}, \{2, 1, 3\}, \{3, 2, 1\})$. Thus for plurality rule, we have that $scr(1) = 2, scr(2) = 1, scr(3) = 1$, for Borda rule we have that $scr(1) = 5, scr(2) = 4, scr(3) = 3$ and for Veto rule $scr(1) = 3, scr(2) = 3, scr(3) = 2$.

Additionally, each agent i comes with a utility function $u_i : A \to \mathcal{R}$ that shows how preferable is every alternative for i. In order to make sense, we have to assume that function u_i "agrees" with $\prec_i$ that is $y \prec_i x \Rightarrow a_{\prec_i(y)} \geq a_{\prec_i(x)} \Rightarrow u_i(y) \geq u_i(x)$.

Our goal is to find strategy-proof mechanisms that choose an alternative, whose expected positional score is close to optimal, namely $\max_{x \in A} scr(x, \prec)$. Consider the following rather easy algorithm:

Choose alternative x uniformly at random

Hence, we choose x with probability $\frac{1}{m}$. For plurality, since $\sum_{i \in A} scr(i) = n$, then the $\mathbb{E}[scr(x)] = \frac{n}{m}$. However, $\max_{x \in A} scr(x) = n$ (every agent has x as her first preference), thus we have $\frac{1}{m}$ approximation ratio. It is remarkable that the mechanism above is trivially strategy-proof.

2.3.2 General strategy-proof Mechanism

In this section, we describe a general randomized strategy-proof mechanism and find the approximation ratio with respect to Plurality, Borda and Veto. Consider the following algorithm, as it can be seen in Procaccia [25]:

General Mechanism 1 (GM1)
1. Pick agent i uniformly at random 2. Choose alternative x with probability proportional to $a_{\prec_i(x)}$

The algorithm works as follows. At first step we choose an agent (let i) with probability $p = \frac{1}{n}$ and then we choose an alternative x with probability $p' = \frac{a_{\prec_i}(x)}{\sum_{y \in A} a_{\prec_i}(y)}$. It is rather straightforward to prove that $GM1$ is truthful.

Theorem 1: *GM1* is strategy-proof.

Proof: Let i be an agent that is not selected. Then agent i can't increase her utility by deviating, because she doesn't affect the outcome of the voting, thus she can't manipulate. Suppose i is the selected agent. Then her expected utility is equal to (conditionally on choosing i)

$$\mathbb{E}[u_i] = \sum_{x \in A} \frac{u_i(x) \cdot a_{\prec_i}(x)}{\sum_{y \in A} a_{\prec_i}(y)} = \frac{1}{\sum_{j=1}^{m} a_j} \sum_{x \in A} u_i(x) \cdot a_{\prec_i}(x)$$

Thus using Rearrangement inequality[2], we have that the expected utility is maximized. Thus, if i deviates, her expected utility will decrease or stay the same. Hence, $GM1$ is strategy-proof.

□

Observe that we have made no assumptions for the vector $\mathbf{a}$, that is $GM1$ is truthful for every positional scoring rule. Below we analyze the expected score of the outcome of $GM1$, without making any assumption for vector $\mathbf{a}$ too.

[2]See appendix for the statement and the proof of Rearrangement inequality

Theorem 2: $\mathbb{E}_{GM1}[scr(x,\prec)] \geq \Omega\left(\frac{1}{\sqrt{m}}\right) \cdot OPT$, where OPT is the largest score.

Proof: Thus we want to calculate $\mathbb{E}_{GM1}[scr(x,\prec)] = \sum_{y\in A} \Pr[GM1 \text{ chooses } y] \cdot scr(y,\prec)$ (definition). Additionally,

$$\begin{aligned}\Pr[GM1 \text{ chooses } y] &= \frac{1}{n}\sum_{j=1}^{n} \frac{a_{\prec_j(y)}}{\sum_{x\in A} a_{\prec_j(x)}} \\ &= \frac{1}{SUM}\sum_{j=1}^{n} a_{\prec_j(y)} \\ &= \frac{scr(y,\prec)}{\sum_{x\in A} scr(x,\prec)}\end{aligned}$$

where $SUM = n\sum_{i=1}^{m} a_i$. Hence $\mathbb{E}_{GM1}[scr(x,\prec)] = \dfrac{1}{SUM}\sum_{y\in A}\left(scr(y,\prec)\right)^2$. Equivalently, assuming a is the alternative that has the largest score then $\mathbb{E}_{GM1}[scr(x,\prec)] = \frac{OPT^2}{SUM} + \frac{1}{SUM}\sum_{y\in A\setminus\{a\}}\left(scr(y,\prec)\right)^2$. Using BCS inequality[3], it occurs that $\mathbb{E}_{GM1}[scr(x,\prec)] \geq \frac{OPT^2}{SUM} + \frac{1}{SUM}\frac{(SUM-OPT)^2}{m-1}$. Finally, taking the derivative of the function $f(x) = \frac{x}{SUM} + \frac{1}{SUM}\frac{(SUM-x)^2}{x(m-1)}$, it follows that it takes the minimum when $x = \frac{SUM}{\sqrt{m}}$, and we conclude that $\mathbb{E}_{GM1}[scr(x,\prec)] \geq OPT \cdot f(\frac{SUM}{\sqrt{m}}) = OPT \cdot \left(\frac{1}{\sqrt{m}} + \frac{\sqrt{m}(1-\frac{1}{\sqrt{m}})^2}{m-1}\right) = OPT \cdot \Omega\left(\frac{1}{\sqrt{m}}\right)$.

□

It is remarkable to mention that $OPT = \frac{SUM}{\sqrt{m}}$ may not be a feasible equation (it depends on the rule) as we will see in the next section. However, since f (of the proof) is convex, we manage to have approximation ratio $\left(\dfrac{\mathbb{E}[scr(x,\prec)]}{OPT}\right)$ of at least $\Omega\left(\dfrac{1}{\sqrt{m}}\right)$. In next sections, we consider each scoring rule separately and prove bounds for the approximation factor.

2.3.3 Upper and Lower Bounds

First of all we examine the upper bounds of the fraction $\dfrac{OPT}{\mathbb{E}_{GM1}[scr(x,\prec)]}$ (approximation ratio). We consider the following cases for vector **a**.

[3]See appendix for the statement and the proof of BCS inequality

1. Let $\mathbf{a} = [1, 0, ..., 0]$, namely we examine Plurality rule. Using Theorem 2 from section 2.3.2 and observing the fact that $OPT = \frac{SUM}{\sqrt{m}} = \frac{n}{\sqrt{m}}$ is a feasible equation we have that we can approximate the OPT using $GM1$ to a factor $c \geq \Omega(\frac{1}{\sqrt{m}})$.

2. Let $\mathbf{a} = [m-1, m-2, ..., 1, 0]$, namely we examine Borda rule. Since $SUM = \frac{nm(m-1)}{2}$ and $OPT \leq n(m-1)$ we have that $\frac{OPT}{SUM} \leq \frac{2}{m}$. Since function f from Theorem 2 of section 2.2.2 is convex and $\frac{2SUM}{m}$ is less than $\frac{SUM}{\sqrt{m}}$, we have that the minimum is taken when $OPT = \frac{2SUM}{m}$ (the equation is feasible from the previous claim) and hence we can approximate OPT to a factor $c \geq f\left(\frac{2SUM}{m}\right) \geq \frac{1}{2} + \frac{1}{2m} = \frac{1}{2} + \Omega(\frac{1}{m})$. We should mention that the approximation ratio being $\frac{1}{2}$ with respect to Borda is an easy fact, just choose an alternative at random. Then obviously the expected score is $\frac{n(m-1)}{2} \geq \frac{OPT}{2}$.

3. Let $\mathbf{a} = [1, 1, ..., 1, 0]$, namely we examine Veto rule. Similarly as in previous cases, since $SUM = n(m-1)$ and $OPT \leq n$, we have that $\frac{OPT}{SUM} \leq \frac{1}{m-1}$. Since f is convex and $\frac{SUM}{m-1}$ is less than $\frac{SUM}{\sqrt{m}}$ we conclude that the approximation factor $c \geq f\left(\frac{SUM}{m-1}\right) \geq 1 - \frac{1}{m-1} = 1 - O\left(\frac{1}{m}\right)$.

As you can imagine, proving lower bounds generally for the scoring rules above is far more difficult. To proceed with the proofs for lower bounds, we firstly have to mention some helpful definitions and a theorem proved by Gibbard [11] (the weak version of the theorem), we are going to use.

Definition 1: *Duple.* A voting rule f is duple if its range is at most two, that is there are $x, y \in A$ such that $\forall \prec \in L^n$, we have that $f(\prec) = \{x, y\}$.

Theorem 1: (Gibbard 1977) Every strategy-proof voting rule is a probability mixture of rules each of which is either dictatorship or duple.

In words, for every strategy-proof mechanism f, there are $f_1, ..., f_k$ rules, each of which is dictatorship or duple and $a_1, ..., a_k$ with $\sum_{i=1}^{k} a_i = 1$ such that for every preference profile $\prec$ it follows that $\Pr[f(\prec) = f_j(\prec)] = a_j$.

Using now Theorem 1, we can find a lower bound of strategy-proof mechanisms with respect to Plurality and Borda.

Theorem 2: There is no strategy-proof mechanism that achieves an approximation ratio of $\omega\left(\frac{1}{\sqrt{m}}\right)$ with respect to Plurality.

Proof: Assume $n = m = k^2$ and let f be a strategy-proof randomized mechanism. We will create a preference profile $\prec$ such that the approximation of f

is less than $O\left(\frac{1}{k}\right) = O\left(\frac{1}{\sqrt{m}}\right)$. Let p be the probability f is duple (thus $1-p$ is the probability f is a dictatorship) (Gibbard). Assume f is a duple and let $D_1, .., D_g$ be the sets of two alternatives that the duples $f_1, ..., f_g$ (the aggregate of them creates f when it is duple) return respectively. Then we define $q_x = \sum_{i=1}^{g} \Pr[x \in D_i | f \text{ is duple } f_i] \cdot \Pr[f \text{ is duple } f_i]$, that is q_x is the probability x is elected conditionally on f is duple. Apparently,

$$\begin{aligned}\sum_{y \in A} q_y &= \sum_{i=1}^{g} \sum_{y \in A} \Pr[y \in D_i | f \text{ is duple } f_i] \cdot \Pr[f \text{ is duple } f_i] \\ &= \sum_{i=1}^{g} 2 \Pr[f \text{ is duple } f_i] = 2\end{aligned}$$

Hence there is a set of $\frac{m}{k} = k$ alternatives, let A' such that $\sum_{y \in A'} q_y \leq \frac{2}{k}$. Thus the probability x is elected and $x \in A'$ by union bound is less than or equal to $\sum_{y \in A} q_y \leq \frac{2}{k}$ (i).

Assume now that f is a dictatorship. Then there is a set $N' \subset N$ with k agents such that the probability an agent $x \in N'$ is a dictator is less than or equal to $\frac{1}{n/k} = \frac{1}{k}$ (ii). We construct a preference $\prec$ profile as follows: Each alternative in $A \backslash A'$ is ranked first by exactly one agent of $N \backslash N'$ (the other alternatives are placed arbitrarily). Additionally, since $|A'| = k$, it follows that there is a $x^* \in A'$ such that

$$\Pr[f(\prec) = x^* | f \text{ dictatorship over } N \backslash N'] \leq \frac{1}{k} \text{ (iii)}$$

So, we complete $\prec$ in a way that every agent $\in N'$ ranks x^* first and the other alternatives arbitrarily. Hence from (ii) and (iii) it follows that

$$\begin{aligned}\Pr[f(\prec) = x^* | f \text{ dictatorship}] = {} & \Pr[f(\prec) = x^* | f \text{ dictatorship over } N'] \\ & \cdot \Pr[f \text{ dictatorship over } N' | f \text{ dictatorship}] \\ & + \Pr[f(\prec) = x^* | f \text{ dictatorship over } N \backslash N'] \\ & \cdot \Pr[f \text{ dictatorship over } N \backslash N' | f \text{ dictatorship}] \\ & \leq \frac{1}{k} + \frac{1}{k} = \frac{2}{k} \text{ (iv)}\end{aligned}$$

Finally from (i),(iv) it occurs that

$$\begin{aligned}\Pr[f(\prec) = x^*] = {} & \Pr[f(\prec) = x^* | f \text{ dictatorship}] \cdot \Pr[f \text{ dictatorship}] \\ & + \Pr[f(\prec) = x^* | f \text{ duple}] \cdot \Pr[f \text{ duple}] \\ & \leq (1-p)\frac{2}{k} + p\frac{2}{k} = \frac{2}{k}\end{aligned}$$

Thus $\mathbb{E}[scr(x, \prec)] \leq k\frac{2}{k} + \left(1 - \frac{2}{k}\right) \cdot 1 < 3$ and $OPT = k$. So we conclude that the approximation ratio is $\leq \frac{3}{k} = O\left(\frac{1}{\sqrt{m}}\right)$.

□

It is remarkable to mention that $GM1$ has a tight approximation ratio with respect to Plurality (is $\Theta\left(\frac{1}{\sqrt{m}}\right)$).

Theorem 3: There is no strategy-proof mechanism that achieves an approximation ratio of $\frac{1}{2} + \omega\left(\frac{1}{\sqrt{m}}\right)$ with respect to Borda.

Proof: Assume f an arbitrary strategy-proof voting rule. Then f is a probability mixture of rules each of which is either dictatorship or duple (Theorem 1), let $f_1, ..., f_k$. We will apply Yao's minimax principle [31][4], with $\mathcal{A} = \{f_1, ..., f_k\}$ and $\mathcal{I} = \mathcal{L}^n$ and $k = \dfrac{scr(f(\prec), \prec)}{\max_{x \in A} scr(x, \prec)}$ (maximization case), hence the expected score of f (in worst-case) is less or equal to the expected score of the best deterministic (dictatorship or duple) mechanism over a specific distribution of preference profiles. We will consider "difficult" inputs such that the expected score of every deterministic rule is upper bounded by $\frac{1}{2} + O\left(\frac{1}{\sqrt{m}}\right)$ and the theorem will hold. Consider the following distribution over the profiles:

Let $n = m - 1$, assume w.l.og that $\sqrt{m} \in \mathcal{N}$ and consider the following procedure which produces a probability distribution over $\mathcal{L}^n$:

1. Choose (fix) x^* uniformly at random (with probability $\frac{1}{m}$).
2. For each agent i, choose k_i uniformly at random from $\{1, 2, ..., \sqrt{m}\}$. Agent i ranks x^* in position k_i.
3. Choose a permutation of $\{1, ..., m-1\}$ uniformly at random, let π.
4. Let $A' = A \backslash \{x^*\} = \{x_1, ..., x_{m-1}\}$. The preference profile of agent 1, is the permutation π with inserting in index k_1, the alternative x^*. For each agent $j \geq 2$, x^* is inserted in k_j position and the other $m-1$ are inserted in a cyclic way (counter clock-wise with respect to agent $j-1$).

Firstly, observe that $(m-1)^2 = (m-1)n \geq scr(x^*, \prec) \geq (m - \sqrt{m})n = (m - \sqrt{m})(m-1)$ since x^* is ranked in the first $\sqrt{m}$ positions for every agent and also $scr(y, \prec) \leq \binom{m}{2}$ for every $y \in A'$ (y is ranked in position j at most once). We consider now two cases. The first case is that the best deterministic rule is

[4]See appendix for statement and proof of Yao's minimax principle

duple (let g). Then the probability x^* is chosen is at most $\frac{m-1}{\binom{m}{2}} = \frac{2}{m}$ from which follows that:

$$\begin{aligned}\mathbb{E}\left[\frac{scr(g(\prec),\prec)}{\max_{x\in A} scr(x,\prec)}\right] &\leq \mathbb{E}\left[\frac{scr(g(\prec),\prec)}{scr(x^*,\prec)}\right] \\ &\leq \frac{\frac{2}{m}(m-1)^2 + \left(1-\frac{2}{m}\right)\cdot\binom{m}{2}}{(m-1)(m-\sqrt{m})} \\ &= \frac{1}{2} + O\left(\frac{1}{m}\right)\end{aligned}$$

The second case is that the best deterministic rule is dictatorship (let g) and i be the dictator. Then we have that for every $y \in A$ in the first $\sqrt{m}$ positions, $\Pr[y = x^*] = \frac{1}{\sqrt{m}}$ (x^* can be in any position of the first $\sqrt{m}$ with the same probability), hence $\Pr[g(\prec) = x^*] = \frac{1}{\sqrt{m}}$, if $g(\prec)$ belongs to the first $\sqrt{m}$ positions. Additionally, if $g(\prec)$ doesn't belong to the first $\sqrt{m}$ positions, then $\Pr[g(\prec) = x^*] = 0$. Thus $\Pr[g(\prec) = x^*] \leq \frac{1}{\sqrt{m}}$. Furthermore,

$$\begin{aligned}\mathbb{E}\left[\frac{scr(g(\prec),\prec)}{\max_{x\in A} scr(x,\prec)}\right] &\leq \mathbb{E}\left[\frac{scr(g(\prec),\prec)}{scr(x^*,\prec)}\right] \\ &\leq \frac{\frac{1}{\sqrt{m}}(m-1)^2 + \left(1-\frac{1}{\sqrt{m}}\right)\cdot\binom{m}{2}}{(m-1)(m-\sqrt{m})} \\ &= \frac{1}{2} + O\left(\frac{1}{\sqrt{m}}\right)\end{aligned}$$

□

It is remarkable that in case of Borda, there is a gap since $GM1$ manages to have approximation ratio of at least $\frac{1}{2} + \Omega\left(\frac{1}{m}\right)$ and there is no strategy-proof mechanism that can have approximation ratio of $\frac{1}{2} + \omega\left(\frac{1}{\sqrt{m}}\right)$.

2.4 Other scored-based Voting Rules

In this section, we examine scored-based voting rules that the score of an alternative doesn't depend on a specific vector, but depends only on his relative position. We will examine Copeland, Maximin rules.

2.4.1 Copeland - Maximin

Let $P(x,y) = |i \in N : y \prec_i x|$, that is the number of agents such that x is preferred to y. *Copeland*$_a$ rule's score for alternative x is denoted by

$$scr(x, \prec) = |\{y \in A \backslash \{x\}\} : P(x,y) > n/2| + a \cdot |\{y \in A \backslash \{x\}\} : P(x,y) = n/2|$$

where $a \in [0,1]$.

Additionally, *Maximin* rule's score for alternative x is denoted by

$$scr(x, \prec) = \min_{y \in A \backslash \{x\}} P(x,y)$$

For this special case of scored-based voting rules, we will describe another strategy-proof mechanism that achieves $\frac{1}{2} + \Omega\left(\frac{1}{m}\right)$ approximation ratio with respect to Copeland$_{1/2}$. Consider the following rather easy randomized mechanism.

General Mechanism 2 (GM2)
1. Pick two alternatives x, y uniformly at random 2. **if** $P(x,y) > n/2$ then the winner is x 3. **else if** $P(y,x) > n/2$ then the winner is y 4. **else** toss a fair coin.

Theorem 1: $GM2$ is strategy-proof w.r.t Copeland and Maximin.

Proof: Assume that x, y are the chosen alternatives. Then let N_x be the agents that prefer "more" x than y. Obviously, we are interested in the relative position of x, y and not the rank of them, thus w.l.o.g we restrict the preference profile to $A' = \{x, y\}$, namely we ignore all the other alternatives. If $i \in N_x$, that is she ranks x before y, if she changes her profile and ranks y before x, she will decrease or leave the same her utility function (as $u_i(x) \geq u_i(y)$), because then she increases the chances of y to be the winner. Hence, for every agent i, every pair of alternatives must remain their relative position. Formally, assuming i's true preference profile is $\prec_i = \{x_1, ...x_m\}$, then reporting $\prec'_i$, it must hold that $x_k \prec'_i x_l$ for every $k, l \in \{1, 2, ..., m\}$ with $l < k$, thus $\prec'_i = \prec_i$.

□

Theorem 2: $GM2$ gives $\frac{1}{2} + \Omega\left(\frac{1}{m}\right)$ approximation ratio with respect to Copeland$_{1/2}$

Proof: First of all, we will prove by induction on m that $\sum_{y \in A} scr(y, \prec) = \binom{m}{2}$. Since $P(x,y) + P(y,x) = n$, then the pair (x,y) gives one to the total sum (this

is rather straightforward if we consider cases if $P(x,y) = P(y,x)$ or not). Thus since we have $\binom{m}{2}$ pairs, then $SUM = \sum_{y \in A} scr(y, \prec) = \binom{m}{2}$. Hence, the probability of selecting alternative x is $\frac{1}{\binom{m}{2}} \left(\sum_{y \in A \setminus \{x\} : P(x,y) > n/2} + \frac{1}{2} \sum_{y \in A \setminus \{x\} : P(x,y) = n/2} \right) = \frac{scr(x, \prec)}{\sum_{y \in A} scr(y, \prec)}$. Thus, observing the fact that $OPT \leq m - 1$ and $SUM = \binom{m}{2}$ we have that $\frac{OPT}{SUM} \leq \frac{2}{m}$. Since function f from Theorem 2 of section 2.2.2 is convex and $OPT = \frac{2SUM}{m}$ is less than $\frac{SUM}{\sqrt{m}}$, the approximation ratio is at least $f(\frac{2SUM}{m}) \geq \frac{1}{2} + \frac{1}{2m} = \frac{1}{2} + \Omega\left(\frac{1}{m}\right)$.

□

Theorem 3: $GM2$ gives $\Omega\left(\frac{1}{nm^2}\right)$ approximation ratio with respect to Maximin

Proof: As previously, the probability alternative x is chosen is $\frac{1}{\binom{m}{2}} \left(\sum_{y \in A : P(x,y) > n/2} + \frac{1}{2} \sum_{y \in A : P(x,y) = n/2} \right)$, thus the expected score of $GM2$ is $\frac{1}{\binom{m}{2}} \sum_{x \in A} \left(\sum_{y \in A : P(x,y) > n/2} + \frac{1}{2} \sum_{y \in A : P(x,y) = n/2} \right) \cdot scr(x, \prec)$. First of all, we have that $OPT \leq n - 1$. We will prove that $\exists x \in A$ such that the probability x is chosen is larger or equal to $\frac{1/2}{\binom{m}{2}}$ and also $scr(x, \prec) \geq 1$. Observe that we have $\binom{m}{2} n$ pairwise "battles" between alternatives. Thus there is an alternative y such that he wins at least $\frac{\binom{m}{2} n}{m} = n(m-1)/2$ times. Thus there is an alternative z such that $P(y,z) \geq \frac{n(m-1)/2}{m-1} = n/2$. If $scr(y, \prec) \geq 1$, then we are done. Else there is an alternative y' such that $P(y,y') = 0$, or equivalently $P(y',y) = n$. If $scr(y', \prec) \geq 1$ we are done (since y' is in front of y). Else we do the same procedure until we find an alternative w such that $scr(w, \prec) \geq 1$ (obviously we can't have zero score for every alternative).

Thus the expected score is larger or equal to $\frac{1}{2\binom{m}{2}}$ and hence the approximation ratio is at least $\Omega\left(\frac{1}{nm^2}\right)$.

□

Observe that this bound depends also on the number of agents n, thus there is no constant $c(m)$ such that the approximation ratio is larger than $c(m)$. We believe that we can achieve better bound for $GM2$ because the analysis is not tight.

2.4.2 Lower Bounds for Copelan-Maximin

In this section we examine how tight is the approximation ratio we found for Copeland$_{1/2}$ and Maximin using $GM2$ mechanism. Namely, we describe theorems about the upper bound of any strategy-proof randomized mechanism with respect to Copeland$_{1/2}$ and Maximin. Both proofs depend on the weak version of theorem of Gibbard [11].

Theorem 1: There is no strategy-proof randomized voting rule that can approximate Copeland$_{1/2}$ to a factor of $\frac{1}{2} + \omega\left(\frac{1}{m}\right)$

Proof: Assume $n = m! + m - 1$ and let f be a strategy-proof randomized mechanism. We will create a preference profile $\prec$ such that the approximation of f is less than $1/2 + O\left(\frac{1}{m}\right)$. Let p be the probability f is duple (thus $1 - p$ is the probability f is a dictatorship). Assume f is a duple and let $D_1, .., D_g$ be the sets of two alternatives that the duples $f_1, ..., f_g$ (the aggregate of them creates f when it is duple) return respectively. Then we define $q_x = \sum_{i=1}^{g} \Pr[x \in D_i | f \text{ is duple } f_i] \cdot \Pr[f \text{ is duple } f_i]$, that is q_x is the probability x is elected conditionally on f is duple. Apparently,

$$\begin{aligned}\sum_{y \in A} q_y &= \sum_{i=1}^{g} \sum_{y \in A} \Pr[y \in D_i | f \text{ is duple } f_i] \cdot \Pr[f \text{ is duple } f_i] \\ &= \sum_{i=1}^{g} 2 \Pr[f \text{ is duple } f_i] = 2\end{aligned}$$

Hence there is a set of $m/2$ alternatives, let A' such that each alternative $y \in A'$ has the property that $q_y \leq \frac{4}{m}$ (i).

Assume now that f is a dictatorship. Then there is a set N' with $m - 1$ agents such that the probability an agent $x \in N'$ is a dictator is less than or equal to $\frac{1}{n/(m-1)} = \frac{m-1}{m!+m-1}$ (ii). We construct a preference $\prec$ profile as follows: Each agent's preference profile in $N \backslash N'$ is a permutation of $\{1, 2, ..., m\}$ (notice that $N \backslash N' = m!$) Additionally, since $|A'| = m/2$, it follows that there is a $x^* \in A'$ such that

$$\Pr[f(\prec) = x^* | f \text{ dictatorship over } N \backslash N'] \leq \frac{2}{m} \text{ (iii)}$$

So, we complete $\prec$ in a way that every agent $\in N'$ ranks x^* first and the other alternatives cyclically. Hence from (ii) and (iii) it follows that:

$$\begin{aligned}\Pr[f(\prec)=x^*|f \text{ dictatorship}] &= \Pr[f(\prec)=x^*|f \text{ dictatorship over } N'] \\ &\cdot \Pr[f \text{ dictatorship over } N'|f \text{ dictatorship}] \\ &+ \Pr[f(\prec)=x^*|f \text{ dictatorship over } N\backslash N'] \\ &\cdot \Pr[f \text{ dictatorship over } N\backslash N'|f \text{ dictatorship}] \\ &\leq \frac{m-1}{m!+m-1}+\frac{2}{m} \leq \frac{4}{m} \text{ (iv)}\end{aligned}$$

Finally from (i),(iv) it occurs that

$$\begin{aligned}\Pr[f(\prec)=x^*] &= \Pr[f(\prec)=x^*|f \text{ dictatorship}]\cdot\Pr[f \text{ dictatorship}] \\ &+ \Pr[f(\prec)=x^*|f \text{ duple}]\cdot\Pr[f \text{ duple}] \\ &\leq (1-p)\frac{4}{m}+p\frac{4}{m}=\frac{4}{m}\end{aligned}$$

Now observe that the score of x^* is $m-1$ since it $P(x^*,y)=\frac{m!}{2}+m-1>n/2$ and all the other alternatives are tied, thus $scr(y,\prec)=\frac{1}{2}\cdot(m-2)$, $\forall y\in A\backslash\{x^*\}$. Thus $\mathbb{E}[scr(x,\prec)]\leq(m-1)\frac{4}{m}+\frac{1}{2}\cdot(m-2)\left(1-\frac{4}{m}\right)\cdot<\frac{1}{2}m+4$ and $OPT=m-1$. So we conclude that the approximation ratio is $\leq\frac{\frac{1}{2}m+4}{m-1}=\frac{1}{2}+O\left(\frac{1}{m}\right)$.

□

It is remarkable to mention that $GM2$ has a tight approximation ratio with respect to $\text{Copelan}_{1/2}$ (is $\frac{1}{2}+\Theta\left(\frac{1}{m}\right)$). Additionally, observe that the proof is similar to the one of Plurality.

Theorem 2: There is no strategy-proof randomized voting rule that can approximate Maximin to a factor of $\omega\left(\frac{1}{m}\right)$

Proof: Assume f an arbitrary strategy-proof voting rule. Then f is a probability mixture of rules each of which is either dictatorship or duple (Gibbard), let $f_1,...,f_k$. We will apply Yao's minimax principle[5], with $\mathcal{A}=\{f_1,...,f_k\}$ and $\mathcal{I}=\mathcal{L}^n$ and $k=\dfrac{scr(f(\prec),\prec)}{\max_{x\in A}scr(x,\prec)}$ (maximization case), hence the expected score of f (in worst-case) is less or equal to the expected score of the best deterministic (dictatorship or duple) mechanism over a specific distribution of preference profiles. We will consider "difficult" inputs such that the expected score of every deterministic rule is upper bounded by $O\left(\frac{1}{m}\right)$ and the theorem will hold. Consider the following distribution over the profiles:

[5]See appendix for statement and proof of Yao's minimax principle

Let $n = m - 1$, assume w.l.og that m is even and consider the following procedure which produces a probability distribution over $\mathcal{L}^n$:

1. Choose (fix) x^* uniformly at random (with probability $\frac{1}{m}$).
2. For each agent i, choose k_i uniformly at random from $\{1, 2, ..., m/2\}$. Agent i ranks x^* in position k_i.
3. Choose a permutation of $\{1, ..., m-1\}$ uniformly at random, let π.
4. Let $A' = A \backslash \{x^*\} = \{x_1, ..., x_{m-1}\}$. The preference profile of agent 1, is the permutation π with inserting in index k_1, the alternative x^*. For each agent $j \geq 2$, x^* is inserted in k_j position and the other $m-1$ are inserted in a cyclic way (counter clock-wise with respect to agent $j-1$).

Firstly, observe that $scr(x^*, \prec) \geq m/2$ since each alternative $y \neq x^*$ is ranked below $m/2$ position for $m/2$ times. Additionally, $scr(y, \prec) = 1$ since for every $y \neq x^*$, there is $y' \neq x^*$ who is ranked before y for $n-1$ times. We consider now two cases. The first case is that the best deterministic rule is duple (let g). Then the probability x^* is chosen is at most $\frac{m-1}{\binom{m}{2}} = \frac{2}{m}$ from which follows that

$$\begin{aligned}
\mathbb{E}\left[\frac{scr(g(\prec), \prec)}{\max_{x \in A} scr(x, \prec)}\right] &\leq \mathbb{E}\left[\frac{scr(g(\prec), \prec)}{scr(x^*, \prec)}\right] \\
&\leq \frac{\frac{2}{m}(m-1) + 1 \cdot \left(1 - \frac{2}{m}\right)}{m/2} < \frac{6}{m} \\
&= O\left(\frac{1}{m}\right)
\end{aligned}$$

The second case is that the best deterministic rule is dictatorship (let g) and i be the dictator. Then we have that for every $y \in A$ in the first $m/2$ positions, $\Pr[y = x^*] = \frac{1}{m/2}$ (x^* can be in any position of the first $m/2$ with the same probability), hence $\Pr[g(\prec) = x^*] = \frac{1}{m/2}$, if $g(\prec)$ belongs to the first $m/2$ positions. Additionally, if $g(\prec)$ doesn't belong to the first $m/2$ positions, then $\Pr[g(\prec) = x^*] = 0$. Thus $\Pr[g(\prec) = x^*] \leq \frac{1}{m/2}$. Furthermore,

$$\begin{aligned}
\mathbb{E}\left[\frac{scr(g(\prec), \prec)}{\max_{x \in A} scr(x, \prec)}\right] &\leq \mathbb{E}\left[\frac{scr(g(\prec), \prec)}{scr(x^*, \prec)}\right] \\
&\leq \frac{\frac{1}{m/2}(m-1) + 1 \cdot \left(1 - \frac{1}{m/2}\right)}{m/2} < \frac{6}{m} = O\left(\frac{1}{m}\right)
\end{aligned}$$

□

Thus, there is a gap between the approximation ratio of $GM2$ and Theorem 2, with respect to Maximin.

2.5 Special case for Approval Voting

Approval voting [4],[16], is a subfield of social choice theory, where each agent's i preference profile, is a subset of alternatives that i approves and that set hasn't a specific size (consider a generalization of Veto rule, with variant number of ones). In this section, we consider that set N and A, namely agents(voters) and alternatives respectively, coincide and that we can have multiple winners (alternatives that are selected). This problem is rather common, and has its applications in social networks. The goal of this section is to describe strategy-proof randomized mechanism that selects $k < n$ agents who are the winners, work that can be seen in Alon et al [1].

2.5.1 Model

Let $N = \{1, ..., n\}$ be the set of agents, which coincides with set A, the set of alternatives, that is $N = A$. Each agent i comes with an approval profile $a_i \subset N$, that is a set of agents that i approves. Assume d_i is the number of agents that "like" agent i. Our goal is for a given $k < n$ (if $k = n$ we select them all), to select a set $S \subset N$, with $|S| = k$ such that $\sum_{i \in S} d_i$ is maximized. It is remarkable that the "approval" relation is not symmetric. Additionally, each agent wants to be selected, thus agent i has the property that $u_i = 1$ if $i \in S$ and $u_i = 0$ otherwise. From the definition of the model, we can represent an instance of the problem, with a directed graph $G(V, E)$ where $V = N$ and $(i, j) \in E$ iff $j \in a_i$, namely i "likes" j. The rather strange fact is that even for $k = 1$, there is no deterministic strategy-proof mechanism f that can approximate OPT ($\max_{S \subset N, |S|=k} \sum_{i \in S} d_i$) to a factor $c > 0$ (we assume $c = \frac{Cost_f}{OPT}$). See section 2.4.2 for a proof.

2.5.2 Deterministic approach doesn't work

As we mentioned above, we have made restrictions to the preference profile of the agents and also we have assumed that each agent is interested in being selected, namely only for himself. Thus, it is an obvious idea to think that we can circumvent the impossibility result of Gibbard-Satterthwaite [10],[26]. However the theorem, below shows that we can't have deterministic strategy-proof mechanism

that has a good approximation with respect to the selection problem we discuss in this section. Formally, there is no (deterministic) strategy-proof function f that has a non-trivial, that is > 0, approximation ratio with respect to the optimal outcome. The fraction we deal with is $\frac{Cost_f}{OPT}$.

Theorem 1: Let $N = \{1, ..., n\}$, $n \geq 2$, and $k \in \{1, ..., n-1\}$. Then there is no deterministic strategy-proof k-selection mechanism that gives non-zero approximation ratio.

Proof: Assume the contrary and let f be a function that gives > 0 approximation ratio. Let G^* be the empty graph of N vertices. Obviously, since $k < n$, there is an agent i such that $i \notin f(G^*)$ (w.l.o.g assume $i = n$). We restrict function f to star graphs, where the only vertex with incoming edges is n (the vertex that doesn't belong to $f(G^*)$). See the figure below.

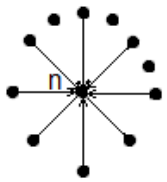

Star graph, with n the centered vertex

This kind of graphs, can be represented with a $(n-1)$-vector $\mathbf{x} = (x_1, ..., x_{n-1})$, where $x_i \in \{0, 1\}$ ($x_i = 1$ iff i approves n). Assume $\mathbf{x} \neq \mathbf{0}$. Since f has approximation ratio > 0 and the only vertex with incoming edges is n, then $n \in f(\mathbf{x})$ (i) (we restricted f's domain). Additionally, since f is a strategy-proof mechanism, then $i \in f(\mathbf{x})$ iff $i \in f(\mathbf{x}')$ where $\mathbf{x} = (x_1, ..., x_{n-1})$ and $\mathbf{x}' = (x_1, ..., x_{i-1}, (x_i + 1) \mod 2, x_{i+1}, ..., x_{n-1})$ (because agent i then could increase her utility (being selected) by deviating, namely changing her opinion about n). Hence, by partitioning the vectors $\mathbf{x}$ into two sets, depending on the $x_i \mod$, we conclude that $|\mathbf{x} : i \in f(\mathbf{x})|$ is an even number $\forall i \neq n$ (ii).

By Fubini's principle, we count with two ways the aggregate number of selections of agents for every graph of the form of the star, and must be equal. Obviously, the number of graphs is 2^{n-1} and since we select k agents for each graph, we have that the number of selections of agents is $k2^{n-1}$, which is even. Moreover, from (ii) the number of selections of each agent $i \neq n$ is even and the number of selections of n is $2^{n-1} - 1$ which is odd. Thus the aggregate number of selections of agents is odd. Contradiction.

□

Thus, we must turn our attention to randomization, in order to approximate the selection problem.

2.5.3 Randomized approach

Alon et al [1] described an algorithm, based on partitioning, called m-RP (Random Partition) with parameter m, which gives $1/4$ approximation in case $m = 2$ and $1/\lceil k^{1/3} \rceil$ for $m > 2$ and can be seen below.

m-RP Algorithm

1. Assign each agent independently and uniformly at random to one set of $S_1, ..., S_m$
2. Let $T \subset \{1, ..., m\}$ be a random subset of size $k - m \cdot \lfloor k/m \rfloor$.
3. If $t \in T$, select the $\lceil k/m \rceil$ agents from S_t with highest indegrees based only on edges from $N \backslash S_t$. If $t \notin T$, select the $\lfloor k/m \rfloor$ agents from S_t with highest indegrees based only on edges from $N \backslash S_t$. Break ties lexicographically in both cases. If one of the subsets S_t is smaller than the number of agents to be selected from this subset, select the entire subset.
4. If only $k' < k$ agents were selected in Step 3, select $k - k'$ additional agents uniformly from the set of agents that were not previously selected.

The algorithm works as follows. It separates the set of agents into m sets uniformly at random and then chooses the k/m with the highest indegree agents from each set S_i (without calculating the edges with endpoints in S_i, which is the main idea of giving truthfulness)(ties are broken lexicographically). Since ($k \mod m$) may not be zero, we have also step 2, to insure that in the end we select exactly k agents. Additionally, in case step 3 selects less than k agents (let k'), we fill with other $k - k'$ uniformly. We claim that m-RP is strategy-proof and has non-trivial approximation ratio.

Theorem 1: m-RP is strategy-proof.

Proof: Assume an arbitrary agent i, that is assigned in set S_j. Then, the outcoming edges of i to the agents of the set $N \backslash S_j$ don't affect the selection of i, thus i doesn't increase the probability of being selected by misreporting these edges. Additionally, the selection of i depends on the outcoming edges of the agents of the set $N \backslash S_j$ to set S_j, thus i doesn't increase the probability of being selected by misreporting the edges that have the other endpoint also in S_j.

□

The Theorem below shows that 2-RP gives an expected output of $OPT/4$. The main idea of the proof lies in a well-known theorem which says that for a given graph $G(V, E)$ and a random bipartition of G, the expected number of edges from the one partitioned set to the other is $\frac{|E|}{2}$.

Theorem 2: 2-RP gives 1/4 approximation ratio.

Proof: Let $G(N, E)$ be the graph that represents the relation between the agents, K^* be an optimal set of agents, that is $OPT = \sum_{i \in K^*} d_i$ (maximum) and K_1^*,K_2^* is the partition (let π) of K^* such that $K_1^* \subset S_1$,$K_2^* \subset S_2$. Assume w.l.o.g $|K_1^*| \geq |K_2^*|$ and let $d_1(\pi) = |(i,j) \in E : i \in S_2, j \in K_1^*|$, $d_2(\pi) = |(i,j) \in E : i \in S_1, j \in K_2^*|$, namely the edges from S_2 to K_1^* and from S_1 to K_2^* respectively. Since $|K_2^*| \leq \lfloor k/2 \rfloor$ (assumption), the selection of the mechanism ensures that at the selected agents from S_2 have at least $d_2(\pi)$ incoming edges. Additionally, assuming k is odd (for the case k is even $T = \emptyset$ and the inequality below still holds) we have that the selection of the mechanism ensures that the selected agents from S_1 have at least incoming edges $\frac{\lceil k/2 \rceil}{k} \cdot d_1(\pi)$ if $T = \{1\}$ and $\frac{\lfloor k/2 \rfloor}{k} \cdot d_1(\pi)$ otherwise. Hence

$$\begin{aligned}
\mathbb{E}[\text{2-RP}|\pi] &\geq \frac{1}{2}\left(\frac{\lceil k/2 \rceil}{k} \cdot d_1(\pi) + d_2(\pi)\right) + \frac{1}{2}\left(\frac{\lfloor k/2 \rfloor}{k} \cdot d_1(\pi) + d_2(\pi)\right) \\
&\geq \frac{d_1(\pi) + d_2(\pi)}{2}
\end{aligned}$$

Since we count the incoming edges, we will find the expected number of incoming edges of set K^* (the number of incoming edges of K^* is OPT) after a given partition. Let $X_{(u,v)}$ be random variables such that $X_{(u,v)} = 1$ iff agents u, v are not in the same set after the partition. Then

$$\begin{aligned}
\sum_{\pi}[d_1(\pi) + d_2(\pi)] \cdot \Pr[\pi] = E\left[\sum_{(u,v):v \in K^*} X_{(u,v)}\right] &= \sum_{(u,v):v \in K^*} E[X_{(u,v)}] \\
&= \sum_{(u,v):v \in K^*} \Pr[u, v \text{ not in same set}] \\
&= \frac{1}{2} \sum_{(u,v):v \in K^*} 1 \\
&= \frac{OPT}{2}
\end{aligned}$$

Hence $\mathbb{E}[\text{2-RP}] = \sum_{\pi} \mathbb{E}[\text{2-RP}|\pi] \cdot \Pr[\pi] \geq \sum_{\pi} \frac{d_1(\pi)+d_2(\pi)}{2} \cdot \Pr[\pi] = \frac{OPT}{4}$.

□

Observe that the analysis made above is rather tight. To see this assume $k = 1$ and consider the graph $G(N, \{(1, n)\})$. In order to choose n, n must belong to

the partitioned set from which 2-RP chooses the one agent. This happens with probability 1/2. Since n belongs to that set, then in order to be selected, agent 1 must belong to the remaining set which happens with probability 1/2 (in all the other cases, since ties are broken lexicographically, n will not be selected). Hence the total probability of selecting n is $\frac{1}{4}$. Thus $\mathbb{E}[2\text{-}RP(G)] = \frac{1}{4}$ and $OPT = 1$.

Theorem 3: $\lceil k^{1/3} \rceil$-RP gives $1 - O(1/k^{1/3})$ approximation ratio.

Proof: Assume K^* be one optimal set (w.l.o.g let $K^* = \{1, 2, ..., k\}$). Let Z_i be a random variable and is equal to the number of agents (except i) that belong to both K^* and also to the same partition set as i. Additionally, as in previous proof, assuming that $d'_i = \{(j, i) \in E : j \in \text{partition set different from } i\text{'s}\}$, then if $Z_i = s_i$, agent i contributes $d'_i \cdot \frac{k^{2/3}}{s_i+1}$ incoming edges if $s_i + 1 > k^{2/3}$ and d'_i otherwise (we denote the coefficient by $\sigma_{s_i} = \min(1, \frac{k^{2/3}}{s_i+1})$). Hence

$$\begin{aligned}
\mathbb{E}[\lceil k^{1/3} \rceil\text{-RP}] &= \sum_{s_1,...,s_k} \mathbb{E}[\lceil k^{1/3} \rceil\text{-RP} | Z_1 = s_1 \wedge ... \wedge Z_k = s_k] \cdot \Pr[Z_1 = s_1 \wedge ... \wedge Z_k = s_k] \\
&\geq \sum_{s_1,...,s_k} \mathbb{E}[\sum_{i \in K^*} \sigma_{s_i} d'_i | Z_1 = s_1 \wedge ... \wedge Z_k = s_k] \cdot \Pr[Z_1 = s_1 \wedge ... \wedge Z_k = s_k] \\
&= \sum_{s_1,...,s_k} \sum_{i \in K^*} \mathbb{E}[\sigma_{s_i} d'_i | Z_1 = s_1 \wedge ... \wedge Z_k = s_k] \cdot \Pr[Z_1 = s_1 \wedge ... \wedge Z_k = s_k] (i)
\end{aligned}$$

Let $d(i, S)$ be the set of i's incoming edges from set S. The probability of $j \neq i$ and $j \in K^*$ be in the same set as i is $\frac{s_i}{k-1}$ and thus the probability of not being is $1 - \frac{s_i}{k-1}$. Additionally for $j \in N \backslash K^*$, the probability of not being in the same set is $\frac{k^{1/3}-1}{k^{1/3}}$. Hence $\mathbb{E}[\sigma_{s_i} d'_i | Z_i = s_i] = \sigma_{s_i} \left(d(i, K^*) \frac{k - 1 - s_i}{k - 1} + d(i, N \backslash K^*) \frac{k^{1/3} - 1}{k^{1/3}} \right)$

(ii). Since $\dfrac{\mathbb{E}[\sigma_{s_i} d'_i | Z_i = s_i]}{\Pr[Z_i = s_i]} = \displaystyle\sum_{s_1,...,s_{i-1},s_{i+1},...,s_k} \mathbb{E}[\sigma_{s_i} d'_i | Z_1 = s_1 \wedge ... \wedge Z_k = s_k]$
$\cdot Pr[Z_1 = s_1 \wedge ... \wedge Z_k = s_k]$, from (i),(ii) it occurs that

$$\begin{aligned}
\mathbb{E}[\lceil k^{1/3} \rceil\text{-RP}] &\geq \sum_{i \in K^*} \sum_{s_i} \Pr[Z_i = s_i] \cdot \sigma_{s_i} \left(d(i, K^*) \frac{k - 1 - s_i}{k - 1} + d(i, N \backslash K^*) \frac{k^{1/3} - 1}{k^{1/3}} \right) \\
&= \sum_{i \in K^*} d(i, K^*) \sum_{s_i} \Pr[Z_i = s_i] \cdot \sigma_{s_i} \frac{k - 1 - s_i}{k - 1} \\
&\quad + \sum_{i \in K^*} d(i, N \backslash K^*) \sum_{s_i} \Pr[Z_i = s_i] \cdot \sigma_{s_i} \frac{k^{1/3} - 1}{k^{1/3}}
\end{aligned}$$

Since $OPT = \sum_{i \in K^*} d(i, K^*) + d(i, N \backslash K^*)$, we have to show that
$\sum_{s_i} \Pr[Z_i = s_i] \cdot (1 - \sigma_{s_i} \frac{k-1-s_i}{k-1})$ and $\sum_{s_i} \Pr[Z_i = s_i] \cdot (1 - \sigma_{s_i} \frac{k^{1/3}-1}{k^{1/3}})$ are $O(1/k^{1/3})$.

For $s_i \leq k^{2/3}$ we have that $\sigma_{s_i} = 1$ and thus both sums till the additive term $s_i = k^{2/3}$ are $O(1/k^{1/3})$. Thus we have to show that $\sum_{s_i > k^{2/3}} \Pr[Z_i = s_i] \cdot (1 - \frac{k^{2/3}}{s_i+1} \frac{k-1-s_i}{k-1})$ and $\sum_{s_i > k^{2/3}} \Pr[Z_i = s_i] \cdot (1 - \frac{k^{2/3}}{s_i+1} \frac{k^{1/3}-1}{k^{1/3}})$ are $O(1/k^{1/3})$.

Equivalently, for the second term, we have to show that $\sum_{s_i > k^{2/3}} \Pr[Z_i = s_i] \cdot \frac{s_i+1-k^{2/3}}{s_i+1}$ is $O(1/k^{1/3})$. To prove this, we have to observe that $Z_i = \sum_{j \in N, j \neq i} X_j$ where X_j is 1 if j is in the same set with i, 0 otherwise. Hence $\Pr[X_j = 1] = \frac{1}{k^{1/3}}$. Now we can apply Chernoff bounds (see section 1.3), with $\epsilon \leftarrow \epsilon/\mu_{Z_i}$, namely $\Pr[Z_i - \mu_{Z_i} > \epsilon] \leq e^{-\epsilon^2/(3\mu_{Z_i})}$. We have that $\mu_{Z_i} = \sum_{j \in K*} \mathbb{E}[X_j] = k \cdot \frac{1}{k^{1/3}} = k^{2/3}$. Using telescopic idea we have that

$$\sum_{s_i > k^{2/3}} \Pr[Z_i = s_i] \cdot \frac{s_i + 1 - k^{2/3}}{s_i + 1} \leq \sum_{x=1}^{2\sqrt{\ln k}} \Pr[Z_i \geq k^{2/3} + (x-1)k^{1/3}] \cdot \frac{xk^{1/3} + 1}{k^{2/3} + xk^{1/3} + 1} + \Pr[Z_i \geq k^{2/3} + 2\sqrt{\ln k} \cdot k^{1/3}]$$

Hence for $\epsilon = (x-1)k^{1/3}$ and for $\epsilon = 2\sqrt{\ln k} \cdot k^{1/3}$ we conclude that $\Pr[Z_i \geq k^{2/3} + (x-1)k^{1/3}] \leq e^{-\frac{(x-1)^2}{3}}$ and $\Pr[Z_i \geq k^{2/3} + 2\sqrt{\ln k} \cdot k^{1/3}] \leq e^{-\frac{4 \ln k}{3}} \leq \frac{1}{k^{4/3}}$. Finally, $\sum_{x=1}^{2\sqrt{\ln k}} e^{-\frac{(x-1)^2}{3}} \cdot \frac{xk^{1/3}+1}{k^{2/3}+xk^{1/3}+1} + \frac{1}{k^{4/3}} \leq \frac{1}{k^{1/3}} \sum_{x=1}^{2\sqrt{\ln k}} e^{-\frac{(x-1)^2}{3}} 2x + \frac{1}{k^{4/3}} = O(1/k^{1/3})$.

With similar arguments we can prove the same for the first term. Thus, it follows that $\mathbb{E}[\lceil k^{1/3} \rceil\text{-RP}] \geq \left(1 - O\left(\frac{1}{k^{1/3}}\right)\right) \cdot OPT$.

□

A lower bound for any randomized strategy-proof mechanism is $1 - O(\frac{1}{k^2})$. To notice this, consider a randomized strategy-proof mechanism f and the graph $G(N, E)$ where for $1 \leq i \leq k$, $(i, i+1) \in E$, $(k+1, 1) \in E$ and E doesn't contain any other edge. Hence there is an agent (w.l.o.g assume 1) that the probability of being selected from f is at most $\frac{k}{k+1}$. We omit agent's 1 outgoing edge and form graph $G'(N, E \backslash \{(1,2)\}$. Then the probability of 1 being selected continues to be at most $\frac{k}{k+1}$ since f is strategy-proof. Thus the expected score of f is at most $k\frac{k}{k+1} + (k-1)\frac{1}{k+1}$. However $OPT = k$ and hence $\mathbb{E}[f] \leq (1 - \frac{1}{k^2+k}) \cdot OPT = (1 - O(\frac{1}{k^2})) \cdot OPT$.

2.5.4 GSP consideration

Sometimes, we are interested in whether there exists a group of people rather than a person, that can manipulate the voting, that is, some agents make a coalition

and gain from jointly misreporting their preference profiles (their outgoing edges). In this section we try to find mechanisms that are group strategy-proof (GSP), that is there is no such group of agents that can make a coalition and benefit from that coalition. Consider the following rather easy algorithm where S is a set of agents with size k:

Choose set S uniformly at random

This algorithm (ALG) is trivially GSP. Additionally, it is straightforward that the approximation ratio is $\frac{k}{n}$. The technique to prove this is classical. Let $X_{(i,j)}$ be a random variable where $X_{(i,j)} = 1$ if agent $j \in S$ and $X_{(i,j)} = 0$ otherwise. The cost the algorithm returns from selection is $\mathbb{E}[\sum_{(i,j)\in E} X_{i,j}]$. Additionally, $\Pr[j \in S] = \frac{\binom{n-1}{k-1}}{\binom{n}{k}} = \frac{k}{n}$. Hence

$$\begin{aligned}\mathbb{E}[\sum_{(i,j)\in E} X_{i,j}] &= \sum_{(i,j)\in E} \mathbb{E}[X_{i,j}] \\ &= \sum_{(i,j)\in E} \Pr[j \in S] \\ &= \frac{k}{n}|E| \\ &\geq \frac{k}{n} OPT\end{aligned}$$

We claim that $\frac{k}{n}$ is optimal (asymptotically), namely we claim that there is no GSP randomized that can have approximation ratio larger than $k \cdot \Theta(\frac{1}{n})$.

Theorem 1: There is no GSP randomized mechanism that gives $k \cdot \omega\left(\frac{1}{n}\right)$ approximation ratio.

Proof: Let f be GSP randomized mechanism. Then for the empty graph, namely $G^*(N, \emptyset)$ we have that $\exists i, j \in N$ such that $\Pr[i \in f(G^*)] \leq \frac{k}{n-1}$ and $\Pr[j \in f(G^*)] \leq \frac{k}{n-1}$. To prove this, we use the same technique as we did in previous proofs. Let X_i be a random variable which is 1 if i is selected else is 0. Thus we have that

$$\begin{aligned}\sum_{i\in N} \Pr[i \in f(G^*)] &= \sum_{i\in N} \mathbb{E}[X_i] \\ &= \mathbb{E}[\sum_{i\in N} X_i] = k\end{aligned}$$

Hence if at least $n-1$ agents had probability of being selected more than $\frac{k}{n-1}$ this would lead to contradiction. Assume now G'(N,{(i,j),(j,i)}). Thus since f is GSP,

then $\Pr[i \in f(G')] \leq \frac{k}{n-1}$ or $\Pr[j \in f(G')] \leq \frac{k}{n-1}$ else i, j would make a coalition (assume it holds for i w.l.o.g). Finally, we consider the graph $G''(N, \{(j, i)\})$. The probability of selecting i (f is GSP thus strategy-proof), is $\Pr[i \in f(G'')] \leq \frac{k}{n-1}$. Hence $\mathbb{E}[ALG] \leq 1 \cdot \frac{k}{n-1}$ and $OPT = 1$.

□

Chapter 3

Differential Privacy

3.1 Introduction

Search engines, hospitals etc possess huge amounts of personal sensitive information that have to be private. In our case, suppose we have some private data and we want to find an algorithm that returns a "good" solution as far as these private data are concerned. Differential privacy means the constraint that few changes in the private data also cause small changes at the output of the algorithm. Actually, we want the outputs of the algorithm not to be strongly correlated to each particular element of the input. We say outputs, because for a single input we have as an output a probability distribution over a range $\mathcal{R}$. That constraint of differential privacy is the key to achieve approximate truthfulness in mechanism design or even strong truthfulness. The idea is simple, if agent i misreports her utility function, then the "output" barely changes, so she gains little or nothing by deviating. Below, we give a formal definition for differential privacy.

3.2 Definitions

3.2.1 Definition

Let an abstract domain $\mathcal{X}$, n individuals that want their privacy to be preserved and $\mathcal{D} \in \mathcal{X}^n$ the private data - input. We assume for this chapter that an *algorithm* or a *mechanism* is a function $\mathcal{M} : \mathcal{X}^n \to \Delta(\mathcal{R})$, where $\Delta(\mathcal{R})$ is a probability

distribution over range $\mathcal{R}$ (randomized algorithm).

Definition 1. Let $\mathcal{D}_1$,$\mathcal{D}_2$ be two data sets. We say that $\mathcal{D}_1$,$\mathcal{D}_2$ are *neighboring* if and only if their symmetric difference equals to 1, that is $|\mathcal{D}_1 \triangle \mathcal{D}_2| = 1$ (they differ only on a single user's data).

Definition 2. A mechanism $\mathcal{M}$ preserves ϵ-*differential privacy* if for every pair of neighboring sets $\mathcal{D}_1$,$\mathcal{D}_2$ and for every set $S \subseteq \mathcal{R}$ we have the following inequality:

$$\Pr[\mathcal{M}(\mathcal{D}_1) \in S] \leq e^{\epsilon} \Pr[\mathcal{M}(\mathcal{D}_2) \in S]$$

Intuitively, in differential privacy every event changes in probability by a small factor for every element that is changed in $\mathcal{D}$. Generally, we think ϵ as a small constant. It is actually meaningless to think of $\epsilon = o(\frac{1}{n})$ because it follows that for every $\mathcal{D}$ we have approximately (w.h.p) the same output distribution. This is true because for $\mathcal{D}_1$,$\mathcal{D}_2$ such that $|\mathcal{D}_1 \triangle \mathcal{D}_2| \leq t$, it apparently follows that $\Pr[\mathcal{M}(\mathcal{D}_1) \in S] \leq e^{\epsilon t} \Pr[\mathcal{M}(\mathcal{D}_2) \in S]$ (the property is called *collusion resistance*). Thus substituting t for n and ϵ for $o(\frac{1}{n})$ it occurs that for every $\mathcal{D}_1$,$\mathcal{D}_2$, $\Pr[\mathcal{M}(\mathcal{D}_2) \in S] \leq \Pr[\mathcal{M}(\mathcal{D}_1) \in S] \leq \Pr[\mathcal{M}(\mathcal{D}_2) \in S]$ (when $n \to \infty$ then $n \cdot o(\frac{1}{n}) \approx 0$) and hence $\Pr[\mathcal{M}(\mathcal{D}_1) \in S] = \Pr[\mathcal{M}(\mathcal{D}_2) \in S]$. Another obvious observation that holds and has many applications on auctions and on databases (queries) is the following:

Composability: The sequential application of $\mathcal{M}_1$,$\mathcal{M}_2$ with ϵ_1,ϵ_2-differential privacy respectively gives $(\epsilon_1 + \epsilon_2)$-differential privacy.

McSherry and Talwar in [18] define the *exponential mechanism*, a mechanism that outputs privately objects from $\mathcal{X}$ and preserves differential privacy. In order to proceed to the algorithm, we first have to define what a quality function is. A quality function q is a function $q : \mathcal{X}^n \times \mathcal{R} \to \mathbb{R}$ that maps private data/output pairs to quality scores (for our case each user-agent wants as high a quality score as possible).

3.2.2 Exponential Mechanism

Exponential mechanism: For any quality function q and base measure μ over $\mathcal{R}$, we define

$$\mathcal{M}_\epsilon(q, \mathcal{D}) := \text{Choose } r \text{ with probability proportional to } exp(\epsilon q(\mathcal{D}, r)) \times \mu(r)$$

We will typically take $\mu(r)$ to be uniform, so we ignore it most of the time.

Theorem 1: *Exponential Mechanism* $\mathcal{M}_\epsilon(q, \mathcal{D})$ preserves $2\epsilon\Delta q$ differential privacy (we define Δq as the maximum possible *difference* in q for all pairs of neighboring sets $\mathcal{D}_1, \mathcal{D}_2$).

Proof: Assume $\mathcal{R}$ is discrete (in other case use integrals). Let $\mathcal{D}_1$,$\mathcal{D}_2$ two neighboring data sets. Then

$$\begin{aligned}
\frac{\Pr[\mathcal{M}_\epsilon(q, \mathcal{D}_1) = r]}{\Pr[\mathcal{M}_\epsilon(q, \mathcal{D}_2) = r]} &= \frac{\frac{\exp(\epsilon q(\mathcal{D}_1, r)) \times \mu(r)}{\sum_{r' \in \mathcal{R}} \exp(\epsilon q(\mathcal{D}_1, r')) \times \mu(r')}}{\frac{\exp(\epsilon q(\mathcal{D}_2, r)) \times \mu(r)}{\sum_{r' \in \mathcal{R}} \exp(\epsilon q(\mathcal{D}_2, r')) \times \mu(r')}} \\
&= \exp(\epsilon(q(\mathcal{D}_1, r) - q(\mathcal{D}_2, r))) \cdot \frac{\sum_{r' \in \mathcal{R}} \exp(\epsilon q(\mathcal{D}_2, r')) \times \mu(r')}{\sum_{r' \in \mathcal{R}} \exp(\epsilon q(\mathcal{D}_1, r')) \times \mu(r')} \\
&\leq \exp(\epsilon \Delta q) \cdot \frac{\exp(\epsilon \Delta q) \sum_{r' \in \mathcal{R}} \exp(\epsilon q(\mathcal{D}_1, r')) \times \mu(r')}{\sum_{r' \in \mathcal{R}} \exp(\epsilon q(\mathcal{D}_1, r')) \times \mu(r')} \\
&= \exp(2\epsilon \Delta q).
\end{aligned}$$

Thus $\dfrac{\Pr[\mathcal{M}_\epsilon(q, \mathcal{D}_1) \in S]}{\Pr[\mathcal{M}_\epsilon(q, \mathcal{D}_2) \in S]} \leq \exp(2\epsilon\Delta q)$.

□

In many cases, we consider function q so as to $\Delta q = 1$ and thus from the previous theorem it follows that we can achieve 2ϵ-differential privacy when using exponential mechanism. As we have already said, we are interested in maximizing q. Defining R_{OPT} to be the subset of range $\mathcal{R}$ such that $q(\mathcal{D}, r) = \max_r q(\mathcal{D}, r)$, the theorem below shows that the probability of a highly suboptimal output is low (exponentially low). This also can be seen in Anupam Gupta et al [13] and a similar version in McSherry et al. [18]. It is remarkable that we consider $\mathcal{R}$ normalized, namely $|\mathcal{R}| = 1$.

Theorem 2: Using *Exponential Mechanism* $\mathcal{M}_\epsilon(q, \mathcal{D})$ and having as an output r, the following inequality holds:

$$\Pr[q(\mathcal{D}, r) < \max_v q(\mathcal{D}, v) + \ln(|\mathcal{R}_{OPT}|)/\epsilon - t/\epsilon] \leq \exp(-t)$$

where $|\mathcal{R}_{OPT}|$ is the measure of $\mathcal{R}_{OPT}$.

Proof: Let S be the set of outputs r such that $q(\mathcal{D}, r) < \max_v q(\mathcal{D}, v) + \ln(|\mathcal{R}_{OPT}|)/\epsilon - t/\epsilon$. The probability of choosing element $r \in S$ is obviously $\dfrac{e^{\epsilon q(\mathcal{D}, r)}}{\sum_{v \in \mathcal{R}} e^{\epsilon q(\mathcal{D}, v)}}$ which is less than $\dfrac{e^{\epsilon(\max_v q(\mathcal{D}, v) + \ln(|\mathcal{R}_{OPT}|)/\epsilon - t/\epsilon)}}{\sum_{v \in \mathcal{R}} e^{\epsilon q(\mathcal{D}, v)}}$. Moreover, $|S| \leq |\mathcal{R}| = 1$ and also

$\sum_{v\in\mathcal{R}} e^{\epsilon q(\mathcal{D},v)} \geq |\mathcal{R}_{OPT}| \cdot e^{\epsilon \max_v q(\mathcal{D},v)}$. Thus the probability of choosing any element of S by union bound is $\leq |S| \cdot \frac{e^{\epsilon \max_v q(\mathcal{D},v)+\ln(|\mathcal{R}_{OPT}|)-t}}{|\mathcal{R}_{OPT}| \cdot e^{\epsilon \max_v q(\mathcal{D},v)}} = |S| \cdot e^{-t} \leq e^{-t}$.

□

When set $\mathcal{R}$ is infinite, we use integrals and the proof is similar. However, the measure of $\mathcal{R}$ must be finite ($< \infty$).

Dwork et al [9] introduced the *Laplace mechanism* which is based on *Laplace Distribution*. *Laplace Distribution* is defined by the following probability density function

$$Lap(x|b) = \frac{1}{2b} e^{\frac{-|x|}{b}}$$

Definition 3. Given function $f : \mathcal{X}^n \to \mathbb{R}^k$ the *Laplace mechanism* and a vector of independent random variables $(Y_1, ..., Y_k)$ which every Y_i follows $Lap(x|\Delta f/\epsilon)$ is defined as

$$\mathcal{M}_L(\mathcal{D}, f, \epsilon) = f(\mathcal{D}) + (Y_1, ..., Y_k)$$

That is $\Pr[\mathcal{M}_L(\mathcal{D}, f, \epsilon) = (x_1, ..., x_k)] = \prod_{i=1}^{k} \frac{\epsilon}{2\Delta f} e^{\frac{-\epsilon|f(\mathcal{D})_i - x_i|}{\Delta f}}$. It is rather straightforward (similar proof to *Theorem 1*, that's why we omit it) to prove that the *Laplace mechanism* gives ϵ-differential privacy. Here, Δf is the maximum possible *distance* in f for all pairs of neighboring sets $\mathcal{D}_1, \mathcal{D}_2$. It is remarkable that the *Laplace mechanism* can be implemented by *Exponential Mechanism* by taking $q(\mathcal{D}, x) = -||f(\mathcal{D}) - x||$. More generally, let $\mathcal{N}$ be a differential private mechanism and let $p(x) = \Pr[\mathcal{N}(\mathcal{D}) = x]$. By taking $q(\mathcal{D}, x) = \ln(p(x))$ we can implement $\mathcal{N}$ using the *Exponential Mechanism*. Thus *Exponential Mechanism* captures the full class of differential privacy mechanisms.

Application to pricing

Assume we organize an auction, where we want to sell at most one item out of k at a *fixed* price (each bidder will buy the specific item). We have n bidders and each bidder i has a non-increasing bid function $b_i^j : [0, 1] \to [0, a]$ referred to item j. We define revenue $Rev = p \sum_i b_i^j(p)$ where j is the item we sell and p its price. The goal is to choose the item that maximizes the *revenue*. Considering $\mathcal{D} = (b_1, ..., b_n)$, $\mathcal{R} = \{1, ..., k\} \times [0, a]$, $q(\mathcal{D}, (j, p)) = p \sum_i b_i^j(p)$ and assuming $p \cdot b_i^j(p) \leq 1 \ \forall i, j, p$, we use the *exponential mechanism* to solve the problem in a private manner.
It is straightforward that *exponential mechanism* gives 2ϵ-differential privacy for

the problem. Suppose two neighboring sets $\mathcal{D}_1, \mathcal{D}_2$. These sets differ on a single bidder i. By changing the bid of bidder i, we have that the revenue changes by $p \cdot (b_i^j - \beta_i^j) \leq p \cdot \max(b_i^j, \beta_i^j) \leq 1$. Thus using *Theorem 1* we finish the proof. In order to find the expected *revenue* we pick an appropriate t and make use of *theorem 8* as it can be seen in Mc-Sherry et al. [18]. Selecting $t = \ln(e + \epsilon^2 \mathcal{OPT} km)$, where m is the number of sold items and using the fact that b_i^j is non-increasing for every i, we conclude that $\mathbb{E}[p \cdot \sum_i b_i^j(p)] \geq \mathcal{OPT} - 3\ln(e + \epsilon^2 \mathcal{OPT} km)$.

3.3 Applications to combinatorial optimization problems

There are a lot of well-known combinatorial optimization problems, some of them are polynomially (efficiently) solvable and other are NP-hard. The *minimum cut*, *vertex cover*, *k-median* and *Combinatorial Public Project* (CPP)[6] are some of them. In this scenario, the input consists of sensitive information about individuals, who don't want to reveal them. So privacy, is an important goal in some cases and algorithms that preserve differential privacy and give suboptimal solutions are preferred to non-private optimal algorithms. The purpose of this section is to examine in practice (giving examples) how we face problems with private data, an interesting work that can be viewed in the paper of Anupam Gupta et al [13].

3.3.1 Unweighted Vertex Cover

Let $G = (V, E)$ be an undirected graph. The problem is to pick a set $S \subset V$ such that for every $(u, v) \in E$ we have $u \in S$ or $v \in S$. The goal is to minimize the size of S. The set S is called *vertex cover* of G. It's remarkable to mention that the *vertex cover* problem is NP-complete. In our case, we want using randomization, to find a solution close to the optimal in a private manner, where private data are the edges (absence or presence of each edge). The difficult part of this problem is that we demand a set S to be the output and not only the size $|S|$. Suppose we have as output the vertex cover S and $u, v \notin S$, then we can conclude that $(u, v) \notin E$. Hence, how can we solve *Vertex Cover Problem* privately? The randomized algorithm that we mention below, surpass this problem, as it actually outputs a permutation of all vertices. In order to determine the *vertex cover* for an input G, for each edge (u, v) we put u at S if u comes before v in the output permutation, otherwise we put v at S (in the end, both u and v may belong to

[6] we examine CPP in a next section

S). Additionally, the algorithm has a good approximation with respect to the *minimum vertex cover* and also gives ϵ-differential privacy. It is remarkable that UVC can be implemented with the exponential mechanism (see section 3.2.2 where we give a general method of implementing every ϵ-differential private mechanism with the exponential mechanism).

Algorithm Unweighted Vertex Cover (UVC)

1. $n \leftarrow |V|$, $V_1 \leftarrow V$, $E_1 \leftarrow E$
2. **for** $i = 1, 2, ..., n$ **do**
3. $\quad w_i \leftarrow (4/\epsilon) \cdot \sqrt{n/(n-i+1)}$
4. $\quad$ **pick** a vertex $v \in V_i$ with probability proportional to $d_{E_i}(v) + w_i$
5. $\quad$ **output** v. $V_{i+1} \leftarrow V_i \backslash \{v\}$,$E_{i+1} \leftarrow E_i \backslash \{v\} \times V_i$
6. **end for**.

The algorithm works as follows. As initialization, we have $G_1(V_1, E_1) = G(V, E)$. At step i, let $G_i(V_i, E_i)$ be the graph with $n - i + 1$ vertices remaining. We omit vertex u with probability proportional to $d_{G_i}(u) + w_i$ and obviously all its incident edges and we move to the next step. The following theorems show that *UVC* algorithm is $(2 + 16/\epsilon)$-approximation and also preserves ϵ-differential privacy.

Theorem 1: *UVC* gives ϵ-differential privacy.

Proof: In such cases, we try to bound the fraction $\dfrac{\Pr[UVC_\epsilon(E) = r]}{\Pr[UVC_\epsilon(E') = r]}$ with e^ϵ (see *Theorem 1* from section *3.2.2*, similar technique), where r is a permutation of the vertices. Let d_i,m_i, $d_i(E_j)$ be the degree of i-th vertex in permutation r (in graph G_i), the edges of G_i and the degree of i-th vertex in r in graph G_j respectively. Then $\Pr[UVC_\epsilon(E) = r] = \prod_{j=1}^{n} \dfrac{w_j + d_j}{\sum_{k=j}^{n}(w_j + d_k(E_j))} = \prod_{j=1}^{n} \dfrac{w_j + d_j}{(n-j+1)w_j + 2m_i}$. Thus $\dfrac{\Pr[UVC_\epsilon(E) = r]}{\Pr[UVC_\epsilon(E') = r]} = \prod_{j=1}^{n} \dfrac{w_j + d_j}{w_j + d'_j} \cdot \dfrac{(n-j+1)w_j + 2m'_j}{(n-j+1)w_j + 2m_j}$. As $|E \triangle E'| = 1$, let the t-th vertex in r be the one endpoint of edge e that belongs to precisely one of E, E'. Then for every $j \leq t$ we have that $m_j = m'_j + 1$ or $m'_j = m_j + 1$ and for $j > t$ we have that $m_j = m'_j$. Additionally $d_j = d'_j$ for every $j \neq t$ and $|d_t - d'_t| = 1$. We consider two cases:

- $d_t = d'_t + 1 \Rightarrow m_j = m'_j + 1$ for $j \leq t$. Hence

$$\begin{aligned}\frac{\Pr[UVC_\epsilon(E) = r]}{\Pr[UVC_\epsilon(E') = r]} &= \frac{w_t + d'_t + 1}{w_t + d'_t} \prod_{j=1}^{t} \frac{(n-j+1)w_j + 2m'_j}{(n-j+1)w_j + 2m'_j + 2} \\ &< \frac{w_t + d'_t + 1}{w_t + d'_t} \leq \frac{w_t + 1}{w_t} \\ &\leq \exp(1/w_t)\end{aligned}$$

(the first inequality comes from the fact that the second factor is less than 1 and the second from $e^x \geq x + 1$ for every $x \in \mathbb{R}$). Using the fact that $1/w_t \leq \epsilon/4$ it follows that $\dfrac{\Pr[UVC_\epsilon(E) = r]}{\Pr[UVC_\epsilon(E') = r]} \leq e^{\epsilon/4} \leq e^\epsilon$.

- $d'_t = d_t + 1 \Rightarrow m'_j = m_j + 1$ for $j \leq t$. Thus

$$\begin{aligned}\frac{\Pr[UVC_\epsilon(E) = r]}{\Pr[UVC_\epsilon(E') = r]} &= \frac{w_t + d_t}{w_t + d_t + 1} \prod_{j=1}^{t} \frac{(n-j+1)w_j + 2m_j + 2}{(n-j+1)w_j + 2m_j} \\ &< \prod_{j=1}^{t} \frac{(n-j+1)w_j + 2m_j + 2}{(n-j+1)w_j + 2m_j} \\ &< \prod_{j=1}^{t} \frac{(n-j+1)w_j + 2}{(n-j+1)w_j} \\ &< \prod_{j=1}^{t} \exp\left(\frac{2}{(n-j+1)w_j}\right)\end{aligned}$$

(similar analysis). Using the fact that $\sum_{j=1}^{t} \frac{1}{\sqrt{n(n-j+1)}} \leq \sum_{j=1}^{n} \frac{1}{\sqrt{n(n-j+1)}}$ $\leq \frac{1}{\sqrt{n}} + \int_1^n \frac{1}{\sqrt{n(n-x+1)}} dx = 2 - \frac{1}{\sqrt{n}} < 2$, occurs that $\dfrac{\Pr[UVC_\epsilon(E) = r]}{\Pr[UVC_\epsilon(E') = r]} \leq e^\epsilon$.

□

Theorem 2: *UVC* is $(2 + 2avg_{i \leq n} w_i) \leq (2 + 16/\epsilon)$-approximation.

Proof: We will use induction on the number of vertices n. During the proof, we make no assumptions for the values of the sequence w_i. For $n = 2$ trivially holds. Assume that it holds for every graph with $n = k$ vertices. We will prove that it also holds for $n = k + 1$. Suppose that $OPT(G) \leq \frac{k+1}{2}$ (the other case trivially holds). Let v be an arbitrary vertex that was picked in step 1. Then the expected cost of *UVC* is the following:

$$\mathbb{E}[UVC(G)] = 1 \cdot \Pr[v \text{ is incident to an edge at } G] + \mathbb{E}_v[\mathbb{E}[UVC(G \backslash v)]]$$

The first term is equal to $\frac{\sum_{t\in V,d_t(G)>0}(w_1+d_t(G))}{\sum_{j=1}^{k+1}(w_1+d_j(G))}$. Assuming $|E|=m$, as each edge has two endpoints, we conclude that $|t : t \in V,\ d_t > 0| \leq 2m$. Thus $\frac{\sum_{t\in V,d_t(G)>0}(w_1+d_t(G))}{\sum_{j=1}^{k+1}(w_1+d_j(G))} \leq \frac{2mw_1+2m}{(k+1)w_1+2m}$. Additionally, the probability $\Pr[v \in OPT(G)]$, where $OPT(G)$ is a minimum vertex cover of G, is equal to

$$\begin{aligned}\Pr[v \in OPT(G)] &= \frac{\sum_{t\in OPT(G)}(w_1+d_t(G))}{\sum_{j=1}^{k+1}(w_1+d_j(G))} \\ &\geq \frac{|OPT(G)|w_1+m}{(k+1)w_1+2m} \\ &\geq \frac{m}{(k+1)w_1+2m} \text{ (i)}\end{aligned}$$

(the second from the end inequality comes from the fact that $OPT(G)$ "covers" the m edges, thus $\sum_{t\in OPT(G)} d_t(G) \geq m$). Hence $\Pr[v$ is incident to an edge at $G] \leq \frac{(2w_1+2)m}{(k+1)w_1+2m} \leq (2w_1+2)\Pr[v \in OPT(G)]$ (ii).

Furthermore, it is straightforward to prove that $\Pr[v \in OPT(G)] \leq \mathbb{E}_v[|OPT(G)| - |OPT(G\backslash v)|] = |OPT(G)| - \mathbb{E}_v[|OPT(G\backslash v)|]$ (iii). To do this, consider that if $|OPT(G)| = |OPT(G\backslash v)|$ then obviously $v \notin OPT(G)$. Using (ii),(iii) we conclude that $\mathbb{E}[UVC(G)] \leq (2w_1+2)(|OPT(G)| - \mathbb{E}_v[|OPT(G\backslash v)|]) + \mathbb{E}_v[\mathbb{E}[UVC(G\backslash v)]$. Now, using inductive hypothesis on graph $G\backslash v$ which has k vertices, it occurs that $\mathbb{E}[UVC(G)] \leq (2w_1+2)(|OPT(G)|) + (2avg_{1<i\leq k+1}w_i - 2w_1)\mathbb{E}_v[|OPT(G\backslash v)|]$.Also from inequalities (iii),(i) obviously

$$\begin{aligned}\mathbb{E}_v[|OPT(G\backslash v)|] &\leq |OPT(G)| - \Pr[v \in OPT(G)] \\ &\leq |OPT(G)| - \frac{|OPT(G)|w_1+m}{(k+1)w_1+2m} \leq (1-\frac{1}{k+1})|OPT(G)|\end{aligned}$$

since $\frac{k+1}{2} \geq |OPT(G)|$. Finally, from the last inequality we have that $\mathbb{E}[UVC(G)] \leq |OPT(G)|(2+2avg_{1\leq i\leq k+1}w_i)$. In order to prove that UVC is $(2+16/\epsilon)$-approximation, we use the following: $2avg_{1\leq j\leq n}w_j = \sum_{j=1}^{n}\frac{2w_j}{n} = \sum_{j=1}^{n}\frac{8}{\epsilon\sqrt{n(n-j+1)}} \leq$ $\frac{8}{\epsilon}\left(\frac{1}{\sqrt{n}} + \int_1^n \frac{1}{\sqrt{n(n-x+1)}}dx\right) = \frac{8}{\epsilon}(2-\frac{1}{\sqrt{n}}) < 16/\epsilon.$

□

3.3.2 Min-Cut

Let $G = (V, E)$ be an undirected graph. The problem is to partition the set V into S, S^c such that set $C = \{(u, v) : (u, v) \in E, u \in S, v \in S^c\}$ is minimized. Minimum cut problem has been solved efficiently (polynomial complexity). A rather common way to solve the problem is to make a reduction to the *max-flow* problem, that is create a graph $G' = (V, E')$ where $E' = \{(u, v), (v, u) : (u, v) \in E\}$ and for every pair of vertices x, y, find the maximum flow of G' with $x \rightarrow$ sink, $y \rightarrow$ source and output the smallest (which will be the size of min-cut, namely |C|). From the residual graph comes the set C. Below we present an algorithm that gives 4ϵ-differential privacy and expected cost $OPT(G) + O(\log n/\epsilon)$ of the cut, where $OPT(G)$ is the size of the minimum cut.

Algorithm Min-Cut (PMC)

1. **Let** $H_0 \subset H_1 ... \subset H_{\binom{n}{2}}$ be arbitrary strictly increasing sets of edges on V .
2. **Choose** index i with probability proportional to $\exp(-\epsilon|OPT(G \cup H_i) - 8 \ln n/\epsilon|)$.
3. **Choose** cut C with probability proportional to $\exp(-\epsilon Cost(G \cup H_i, C))$.
4. **Output** C.

From the technical analysis of Anupam Gupta et al [13], the writer wanted to ensure that OPT is at least $4 \ln n/\epsilon$ (w.h.p) in order to bound the expected cost of PMC with the use of the theorem of Karger, see [15][7]. For this purpose, he adds some more edges in graph G at the first part of the algorithm. He uses the exponential mechanism to achieve that. In the second part, he also uses the exponential mechanism to output the cut, where quality function q is the cost of the cut.

Theorem 1: *PMC* gives 4ϵ-differential privacy.

Proof: Let G,G' be two graphs which differ only in a single edge. Then the probability for a certain cut C to be an output is the following:

$$\begin{aligned}
\Pr[PMC_\epsilon(G) = C] &= \sum_i \frac{e^{-\epsilon|OPT(G \cup H_i) - 8 \ln n/\epsilon|}}{\sum_j e^{-\epsilon|OPT(G \cup H_j) - 8 \ln n/\epsilon|}} \times \frac{e^{-\epsilon Cost(G \cup H_i, C)}}{\sum_{C'} e^{-\epsilon Cost(G \cup H_i, C')}} \\
&\leq e^{4\epsilon} \times \sum_i \frac{e^{-\epsilon|OPT(G' \cup H_i) - 8 \ln n/\epsilon|}}{\sum_j e^{-\epsilon|OPT(G' \cup H_j) - 8 \ln n/\epsilon|}} \times \frac{e^{-\epsilon Cost(G' \cup H_i, C)}}{\sum_{C'} e^{-\epsilon Cost(G' \cup H_i, C')}} \\
&= e^{4\epsilon} \Pr[PMC_\epsilon(G') = C]
\end{aligned}$$

[7] Karger's theorem: For any graph G with min cut C, there are at most n^{2a} cuts of size at most aC

This is true because $OPT(G \cup H_i) - 1 \leq OPT(G' \cup H_i) \leq OPT(G \cup H_i) + 1$ and also $Cost(G \cup H_i, C) - 1 \leq Cost(G' \cup H_i, C) \leq Cost(G \cup H_i, C) + 1$. Thus for any set S of cuts we have that $\dfrac{\Pr[PMC_\epsilon(G) \in S]}{\Pr[PMC_\epsilon(G') \in S]} \leq e^{4\epsilon}$.

□

Theorem 2: $\mathbb{E}[PMC_\epsilon] \leq OPT + O(\ln n/\epsilon)$.

Proof: First of all, we will prove that $(OPT \cup H_i) > 4 \ln n/\epsilon$ with probability at least $1 - \frac{1}{n^2}$ (i is the selected index). From Theorem 2 of section 3.2.2, for $t = 2 \ln n$, $|R_{OPT}| = \frac{2}{n(n-1)}$ and $max \approx 0$ we have that $|OPT(G \cup H_i) - 8 \ln n/\epsilon| \geq 4 \ln n/\epsilon$ with probability at most $\frac{1}{n^2}$. Thus $\Pr[OPT(G \cup H_i) \leq 4 \ln n/\epsilon] \leq \frac{1}{n^2}$, from which follows that $\Pr[OPT(G \cup H_i) > 4 \ln n/\epsilon] \geq 1 - \frac{1}{n^2}$. Let C be a cut with $|C| = OPT(G \cup H_i) + t$. The probability of choosing C is $\dfrac{e^{-\epsilon(OPT(G \cup H_i)+t)}}{\sum_{C'} e^{-\epsilon(Cost(G \cup H_i), C')}} \leq \dfrac{e^{-\epsilon(OPT(G \cup H_i)+t)}}{e^{-\epsilon OPT(G \cup H_i)}} = e^{-\epsilon t}$ (i). Additionally, from Kerger's theorem, there are at most $n^{2\left(1+\frac{t}{OPT(G \cup H_i)}\right)}$ cuts [8] with size at most $\left(1 + \dfrac{t}{OPT(G \cup H_i)}\right) OPT(G \cup H_i) = OPT(G \cup H_i) + t$ (ii). Hence for some b, from (i),(ii) we have that from union bound (assuming c_t the number of cuts of size at most $OPT(G \cup H_i) + t$ and $c_t - c_{t-1}$ is the number of cuts with size $OPT(G \cup H_i) + t$)

$$
\begin{aligned}
\Pr[Cost(G \cup H_i, C) > OPT + b] &\leq \sum_{t>b} e^{-\epsilon t}(c_t - c_{t-1}) \\
&\leq \sum_{t>b} e^{-\epsilon t} c_t - e^{-\epsilon} \sum_{t>b} e^{-\epsilon t} c_t \\
&= (1 - e^{-\epsilon}) \sum_{t>b} e^{-\epsilon t} c_t \\
&\leq (1 - e^{-\epsilon}) \sum_{t>b} e^{-\epsilon t} n^2 e^{\frac{2t \ln n}{OPT(G \cup H_i)}} \\
&\leq (1 - e^{-\epsilon}) \sum_{t>b} e^{-\epsilon t/2} n^2 \\
&\leq n^2 \frac{(1 - e^{-\epsilon})}{e^{\epsilon/2} - 1} e^{-\frac{\epsilon b}{2}} \leq n^2 (e^{\epsilon/2} + 1) e^{-\frac{\epsilon b}{2}}
\end{aligned}
$$

Thus, for $b = 8 \ln n/\epsilon$ we have that $\Pr[Cost(G \cup H_i, C) > OPT + 8 \ln n/\epsilon] \leq \frac{(e^{\epsilon/2}+1)}{n^2}$. Finally, the expected cost of PMC is at most $(1 - \frac{1}{n^2}) \times (1 - \frac{(e^{\epsilon/2}+1)}{n^2}) \times (OPT +$

[8] Substitute a for $1 + \dfrac{t}{OPT(G \cup H_i)}$ in Kerger's theorem

$8\ln n/\epsilon) + \frac{1+(e^{\epsilon/2}+1)-\frac{e^{\epsilon/2}+1}{n^2}}{n^2} \times \binom{n}{2} \leq OPT + 8\ln n/\epsilon + 3 = OPT + O(8\ln n/\epsilon)$.

□

It is rather easy to prove, if we consider the instance of a $d = \Theta(\ln n/\epsilon)$-regular graph G that has $n = |V|$ min-cuts of size $d-1$ of the form $(V\backslash\{v\}, \{v\})$ (there are such graphs according to Anupam Gupta et al. [13]), that for any ϵ-differential private mechanism, the expected cost of that mechanism for the graph $G'(V, E\backslash A)$ (A is the set of neighbors of u, the selection of u is such that the probability of outputting the cut $(V\backslash\{u\}, \{u\})$ with input G is less than or equal to $\frac{1}{n}$)[9] is at least $(1 - \frac{1}{n^{4/3}})(d-1) = \Omega(\ln n/\epsilon)$[10]. On the other hand, the optimal cost is 0. Hence for every ϵ-differential private mechanism $\mathcal{M}$, $\mathbb{E}[\mathcal{M}] \geq OPT + \Omega(\ln n/\epsilon)$.

3.3.3 k-Median

Let V be a set of points in $\mathcal{R}^2$ ($|V| = n$), $d : V \times V \to \mathcal{R}$ be a metric function and $D \subseteq V$ be a set of (private) demand points. The goal is to find a set of points F, with $|F| = k$ such that $Cost(F) = \sum_{v \in D} d(v, F)$ is minimized. With $d(v, S)$ we denote the minimum distance of v and a point from S (even if it is wrong, we use the same notation for distance between points and distance between a point and a set of points). An easy way to deal with this problem, is to use "directly" the exponential mechanism, that is choose output $R = (p_1, ..., p_k)$ with probability proportional to $e^{-\epsilon\frac{Cost(R)}{2\Delta}}$ where Δ is the diameter of set V. It is rather easy to observe that this mechanism is ϵ-differential private[11]. Additionally, from Theorem 2 of section 3.2.2, for $t = k\ln n$ and $R_{OPT} = \frac{1}{\binom{n}{k}}$, it follows that $\Pr[Cost(R)/2\Delta > OPT/2\Delta + 2k\ln n/\epsilon] \leq \frac{1}{n^k}$. Thus $\mathbb{E}[Cost(R)] \leq OPT + 4\Delta k\ln n/\epsilon + \frac{1}{n^k} \times |D|\Delta \leq OPT + 4\Delta k\ln n/\epsilon + \frac{\Delta}{n^{k-1}} = OPT + O(\Delta k\ln n/\epsilon)$.

However the algorithm above is not efficient, as its complexity is $\Omega\left(\binom{n}{k}\right)$. Thus, we examine another another ϵ-differential private mechanism that even though its expected output is worse, it is efficient (polynomial running time), and it can be seen in Anupam Gupta et al [13].

[9] There is such u, it comes from Pigeonhole Principle
[10] To prove this, just use collusion resistance, the input G' differs at $d-1$ edges from G
[11] See theorem 1 from section 3.2.2, substitute ϵ for $\epsilon'/2\Delta$

Algorithm Private k-median (PKM)

1. **Let** $F_1 \subset V$ arbitrarily with $|F_1| = k$ and $\epsilon' \leftarrow \epsilon/(2\Delta(T+1))$.
2. **for** $i = 1$ to T **do**
3. Select $(x, y) \in F_i \times (V \backslash F_i)$ with probability proportional to $\exp(-\epsilon' \times Cost(F_i - \{x\} + \{y\}))$.
4. **Let** $F_{i+1} \leftarrow F_i - \{x\} + \{y\}$.
5. **end for**
6. Select $j \in \{1, 2, ..., T+1\}$ with probability proportional to $\exp(-\epsilon' \times cost(F_j))$.
7. **output** F_j.

The algorithm works as follows. We define $T+1$ $(F_1, ..., F_{T+1})$ possible set outcomes by using exponential mechanism. The first set F_1 is arbitrary and at step i, F_i comes from F_{i-1} by making a swap of a point $x \in F_{i-1}$ and a point $y \in V \backslash F_{i-1}$. Finally, we also use exponential mechanism to choose one of the $T+1$ possible outcomes. PKM gives ϵ-differential privacy and the expected cost of PKM is $6OPT + O(\Delta k^2 \log^2 n/\epsilon)$. This derives from a lemma of Arya et al [2]:

Lemma: For any set $F \subseteq V$ with $|F| = k$, there exists a set of k swaps $(x_1, y_1), ..., (x_k, y_k)$, (that is we swap x_i with y_i), such that

$$\sum_{i=1}^{k} (Cost(F) - Cost(F - \{x_i\} + \{y_i\})) \geq Cost(F) - 5OPT$$

Theorem 1: PKM gives ϵ-differential privacy.

Proof: Let $\mathcal{D}, \mathcal{D}'$ be two neighboring sets of points. Additionally, observe that $|Cost_D(F_i) - Cost_{D'}(F_i)| \leq \Delta$. Then (assuming F_1 is common) we have that for step 3 the probability that the outcome is $(F_1, ..., F_{T+1})$ gives $2\epsilon'\Delta T$-differential privacy (it comes from theorem 1 from section 3.2.2 and composability). Finally, step 6 of PKM gives another $2\epsilon'\Delta$-differential privacy since $|Cost_D(F_i) - Cost_{D'}(F_i)| \leq \Delta$ (also comes from use of theorem 1 from section 3.2.2). Hence, PKM gives $2\epsilon'\Delta(T+1)$-differential privacy (composability), namely ϵ-differential privacy.

□

Theorem 2: $\mathbb{E}[PKM] \leq 6OPT + O(\Delta k^2 \ln^2 n/\epsilon)$, for $T = 12k \ln n$.

Proof: We consider the case that $Cost(F_i) \geq 6OPT$. Then it is easy to see from Lemma above that there is a swap (x, y) such that

$$Cost(F_i) - Cost(F_{i+1}) \geq \frac{Cost(F_i) - 5OPT}{k}$$
$$\geq \frac{Cost(F_i) - (5/6)Cost(F_i)}{k} = \frac{Cost(F_i)}{6k} \Rightarrow$$
$$Cost(F_i)\left(1 - \frac{1}{6k}\right) \geq Cost(F_{i+1}) \text{ (i)}$$

By applying Theorem 2 from section 3.2.2, for $t = 2\ln n$ and $R_{OPT} = \frac{1}{n^2}$ (as there are n^2 swaps at each step) we have that $\Pr[Cost(F_{i+1}) > OPT_{Cost(F_{i+1})} + 4\ln n/\epsilon'] \leq \frac{1}{n^2}$, conditionally on choosing $(F_1, ..., F_i)$. Thus from (i) we have that $Cost(F_{i+1}) \leq OPT_{Cost(F_{i+1})} + 4\ln n/\epsilon' \leq (1 - \frac{1}{6k})Cost(F_i) + 4\ln n/\epsilon'$ with probability at least $(1 - \frac{1}{n^2})$. Assume now that $Cost(F_i) \geq 6OPT + 48k\ln n/\epsilon'$, then $Cost(F_{i+1}) \leq (1 - \frac{1}{6k} + \frac{1}{12k})Cost(F_i)$ with probability at least $1 - \frac{1}{n^2}$ (ii). Applying (ii) to $F_2, ..., F_{T+1}$ we have that $Cost(F_{i+1}) \leq (1 - \frac{1}{12k})^i \times Cost(F_1)$ with probability at least $(1 - \frac{1}{n^2})^i \geq 1 - \frac{i}{n^2}$ (iii). Hence since $Cost(F_1) \leq n\Delta$ and for $T = 12k\ln n$ we have that $Cost(F_{T+1}) \leq (1 - \frac{1}{12k})^T n\Delta \leq \frac{1}{e}^{\ln n} n\Delta = \Delta \leq 6OPT + 48k\ln n/\epsilon'$. Thus either $Cost(F_{T+1}) \leq 6OPT + 48k\ln n/\epsilon'$ or there exists a smaller index $j > 1$ such that $Cost(F_j) \leq 6OPT + 48k\ln n/\epsilon'$ with probability at least $1 - \frac{T}{n^2}$ (applying (iii) for $i = T$). Finally for step 6 of the algorithm, assuming the optimal cost is less than or equal to $6OPT + 48k\ln n/\epsilon'$ and using Theorem 2 from section 3.2.2 for $t = 2\ln n$ and $R_{OPT} = \frac{1}{T+1}$ we have that $PKM_{cost} \leq 6OPT + 48k\ln n/\epsilon' + 2\ln n/\epsilon' + \ln(T+1)/\epsilon'$ with probability at least $1 - \frac{1}{n^2}$. So, $\mathbb{E}[PKM] \leq 6OPT + O(k\ln n/\epsilon') + \frac{T}{n^2} \times n\Delta = 6OPT + O(\Delta k^2 \ln^2 n/\epsilon)$.

□

3.4 Differential Privacy and Truthfulness

In *Mechanism Design*, most of the time we are interested in finding mechanisms-algorithms that agents being truthful is a dominant strategy (incentive compatible mechanisms). In this subsection, we prove that *exponential mechanism* is approximate truthful and then we introduce the *gap* mechanism that achieves strong truthfulness.

3.4.1 Approximate Truthfulness

Shummer in [27] defines and studies ϵ-dominance. Using that definition, we prove that mechanisms that satisfy ϵ-differential privacy, make also truthfulness a $(e^{\epsilon}-1)$-

dominant strategy. Informally, ϵ-dominance means that there is no agent that can increase her utility function by more than an ϵ additive, when misreporting.

Definition 1. *ϵ-improvement*: Giving a utility function u, we say that agent i has an *ϵ-improvement* if there exists profile θ_i' such that $u_i(\theta_i, \theta_{-i}) + \epsilon \leq u_i(\theta_i', \theta_{-i})$ where θ_i is i's true profile-preference

Definition 2. *ϵ-dominance*: We say that function u provides *ϵ-dominance* if there is no agent i that has *ϵ-improvement*.

Defining *ϵ-dominance* and assuming $u_i \in [0, 1]$, we prove now that a mechanism that satisfies ϵ-differential privacy, makes truthtelling $(e^\epsilon - 1)$-dominant strategy. Because mechanisms that preserve differential privacy are randomized, we need to show that the *expected value* of the utility function of agent i changes at most by a factor e^ϵ if she deviates, from which follows that the *expected value* of the utility function increases at most by $e^\epsilon - 1$ (this is true because $x + e^\epsilon - 1 \geq e^\epsilon \cdot x$ for every $x \in [0, 1]$, $\epsilon > 0$).

Theorem 1: For any mechanism $\mathcal{M}$ giving ϵ-differential privacy, any non-negative function u and any neighboring sets $\mathcal{D}_1$,$\mathcal{D}_2$ the following holds:

$$\mathbb{E}[u(\mathcal{M}(\mathcal{D}_1))] \leq e^\epsilon \cdot \mathbb{E}[u(\mathcal{M}(\mathcal{D}_2))]$$

Proof: From the definition of $\mathbb{E}[u(x)] = \sum_{\mathcal{R}} \Pr[\textit{choose } x] \cdot u(x)$ occurs that $\mathbb{E}[u(\mathcal{M}(\mathcal{D}_1))] = \sum_{\mathcal{R}} \Pr_{\mathcal{D}_1}[\textit{choose } x] \cdot u(x) \leq \sum_{\mathcal{R}} e^\epsilon \cdot \Pr_{\mathcal{D}_2}[\textit{choose } x] \cdot u(x) = e^\epsilon \cdot \mathbb{E}[u(\mathcal{M}(\mathcal{D}_2))]$. The first inequality comes from the fact that $\mathcal{M}$ gives ϵ-differential privacy (definition). In case range $\mathcal{R}$ is infinite, we use integrals instead of sums.

□

Thus if agent i deviates, which means that neighboring sets $\mathcal{D}_1$,$\mathcal{D}_2$ differ at index i, she gains no more than a factor of e^ϵ (for the expected *utility* score) from which follows that she gains no more than an $(e^\epsilon - 1)$ additive term. A special case is when $\epsilon = 0$, where we achieve *strategy-proofness* (it makes sense here as opposed to ϵ in differential privacy).

Example

Suppose we have an auction where we want to sell an item and the mechanism we use gives 0.001-differential privacy. Bidder i can change the *expected* sell price of the item so that the *expected* sell price if the bidder was truthful was at most

$e^{0.001}$ times the *expected* sell price if the bidder was untruthful. Assuming that the bid function $b_i \in [0,1]$ and also $u_i = b_i - p$ (if the item is sold to bidder i) or else $u_i = 0$ then $u_i' \leq u_i \times e^{0.001}$ from which follows that $u_i' \leq u_i + (e^{0.001} - 1)$.

3.4.2 Combinatorial Public Projects

In CPP problem, we have n agents and m resources publicly known and a parameter k. Each agent i submits a (private) submodular[12] function f_i over the subsets of resources. The goal is to find a subset S of resources with $|S| = k$ so as to maximize $F(S) = \sum_{i=1}^{n} f_i(S)$. We assume w.l.o.g that the image of f_i is $[0,1]$. Papadimitriou et al. [23] introduced CPP problem and prove that there is no efficient truthful algorithm that guarantees approximation ratio better than $m^{\frac{1}{2}-\epsilon}$ unless $NP \subseteq BPP$. Here, we demonstrate an inefficient and approximate truthful mechanism that achieves a $(1 - 1/e)$ factor of the optimal minus $O(k \log m/\epsilon)$. The algorithm (called $ACPP$) uses exponential mechanism k times and has the same idea as k-median problem:

CPP Algorithm (ACPP)

1. $S_1 = \emptyset$, $\epsilon' \leftarrow \frac{\epsilon}{2k}$
2. **for** $i = 1$ to k **do**
3. Select resource $r \notin S_i$ with probability proportional to $\exp(\epsilon' \times (F(S_i \cup \{r\}) - F(S_i)))$.
4. $S_{i+1} \leftarrow S_i \cup \{r\}$
5. **output** S_{k+1}.

Theorem 1: ACPP gives ϵ-differential privacy

Proof: Let $\mathcal{D}, \mathcal{D}'$ be two sets that differ on a single agent i (the one set contains him and the other doesn't) and let $S_{k+1} = \{r_1, ..., r_k\}$ be an output of resources. Then for the one iteration (let i) conditionally the previous iterations ACCP chose the same elements, we have that $\Pr_D[\text{choose } r_i] \leq e^{2\epsilon'} \Pr_{D'}[\text{choose } r_i]$ (Theorem 1 from section 2, $\forall j. f_j \in [0,1]$). Thus $\Pr[ACCP_D = (r_1, ..., r_k)] \leq e^{2k\epsilon'} \Pr[ACCP_{D'} = (r_1, ..., r_k)]$. However, we have considered an ordered output, although the output is a set, so we use union bound. That is

[12] f is called submodular if for every $A \subset B \subset S$ and $x \in S$, we have that $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$

$$\begin{aligned}\Pr[ACPP_D = S_{k+1}] &= \sum_{\text{permutations}} \Pr[ACCP_D = (r'_1, ..., r'_k)] \\ &\leq e^{2k\epsilon'} \sum_{\text{permutations}} \Pr[ACCP_{D'} = (r'_1, ..., r'_k)] \\ &= e^{2k\epsilon'} \Pr[ACPP_{D'} = S_{k+1}]\end{aligned}$$

Hence, ACCP gives $2k\epsilon'$-differential privacy, namely ϵ-differential privacy.

□

Theorem 2: $\mathbb{E}[ACPP] \geq (1 - \frac{1}{e})OPT + O(k^2 \ln mn/\epsilon)$.

Proof: The proof consists of two parts. The first thing is to notice that F is submodular, thus at iteration i we have that there is a resource r such that $F(S_i \cup \{r\}) - F(S_i) \geq \frac{F(S_{OPT}) - F(S_i)}{k}$. To prove this assume $S_{OPT} \backslash S_i = \{j_1, ..., j_g\}$ and $S_i \backslash S_{OPT} = \{k_1, ..., k_v\}$. Then

$$\begin{aligned}F(S_{OPT} \cup S_i) - F(S_i) &= \sum_{t=1}^{g} [F(S_i \cup \{j_1, ..., j_t\}) - F(S_i \cup \{j_1, ..., j_{t-1}\})] \\ &\leq \sum_{t=1}^{g} [F(S_i \cup \{j_t\}) - F(S_i)]\end{aligned}$$

Additionally

$$\begin{aligned}F(S_{OPT} \cup S_i) - F(S_{OPT}) &= \sum_{t=1}^{v} [F(S_{OPT} \cup \{k_1, ..., k_t\}) - F(S_{OPT} \cup \{k_1, ..., k_{t-1}\})] \\ &\geq \sum_{t=1}^{v} [F(S_{OPT} \cup S_i) - F((S_{OPT} \cup S_i) \backslash \{k_t\})]\end{aligned}$$

Thus, by subtracting it occurs that $F(S_{OPT}) - F(S_i) \leq \sum_{t=1}^{g} [F(S_i \cup \{j_t\}) - F(S_i)] = \sum_{j \in S_{OPT} \backslash S_i} [F(S_i \cup \{j\}) - F(S_i)]$. Hence there is a $r \in S_{OPT} \backslash S_i$ such that $[F(S_i \cup \{r\}) - F(S_i)] \geq \frac{F(S_{OPT}) - F(S_i)}{g} \geq \frac{F(S_{OPT}) - F(S_i)}{k}$. So, there is resource r such that $F(S_{OPT}) - F(S_i \cup \{r\}) \leq (1 - 1/k)(F(S_{OPT}) - F(S_i))$.

The second thing is to use Theorem 2 from section 3.2.2 for $t = 3 \ln n$ and $R_{OPT} = \frac{1}{m}$. Let r be the chosen element at iteration i. We have that $F(S_i \cup \{r\}) \geq F(S_i \cup \{r_{opt}\}) - (\ln m + 3 \ln n)/\epsilon'$ with probability at least $1 - 1/n^3$. Hence $OPT - F(S_{i+1}) \leq (1 - 1/k)(F(S_{OPT}) - F(S_i)) + \ln n^3 m/\epsilon'$ with probability at least $1 - \frac{1}{n^3}$. So with probability at least $(1 - \frac{1}{n^3})^k \geq 1 - \frac{k}{n^3}$ (applying Bernoulli's inequality) we have that $OPT - F(S_{k+1}) \leq (1 - 1/k)^k (OPT - F(S_1)) + O(k \ln nm/\epsilon') \leq$

$\frac{OPT}{e} + O(k \ln nm/\epsilon')$. Thus $\mathbb{E}[ACCP] \geq (1 - 1/e)OPT + O(k^2 \ln nm/\epsilon)$ since $OPT \in o(n^3)$.

□

In the statement of the problem, we mentioned that there is no *truthful* mechanism that achieves good approximation ratio. Namely, we gave a game-theoretic notion to the problem. From the view of game theory and mechanism design, in CPP, the agents want to maximize their utility function, that is to maximize f_i (for agent i). In order to accomplish this, they may misreport their true submodular function f_i in order to have a better outcome. The ACCP algorithm, approximately prevent this. This fact occurs from Theorem 1 of section 3.4.1. That is $E[ACCP[f_i]] \leq e^{2\epsilon} E[ACCP[f_i']] \leq E[ACCP[f_i']] + e^{2\epsilon} - 1$ (the two input sets have symmetric difference 2, that's why it is $2\epsilon'$). So we have approximate truthfulness. Notice that the bound we proved above is assuming we have the true f_i functions.

3.4.3 Gap Mechanism

Even though with *exponential mechanism* we achieve approximate truthfulness ($(e^{\epsilon} - 1)$-dominance), there are problems where misreporting is a dominant strategy. Nissim et al. [22] give such an example, a simplified version of an unlimited supply auction where *exponential mechanism's* expected revenue is far from optimal. McSherry et al. [18] although, prove that the output of *exponential mechanism's* differs to a term at most $O(\log n)/\epsilon$ with respect to the optimal revenue. The problem is that McSherry et al. assumed that the bidders are truthful (however, the bidders are approximately truthful). The example is the following:

Assume an auction (digital goods) with n agents, each of whom wants to buy a single unit. Let $T = \{0.5, 1\}$ be the set of types and $S = \{0.5, 1\}$ be the set of alternatives (prices). Thus $OPT = \max_{s \in S}(s \cdot |i : t_i \leq s|)$. Introducing $q(\mathbf{b}, s) = s \cdot (s \cdot |i : b_i \leq s|)$ ($\mathbf{b}$ is the vector of announced bids), we have that according to McSherry and Talwar the expected revenue of exponential mechanism is $OPT - O(\log n/\epsilon)$ (if agents are truthful) and that it gives 2ϵ-differential privacy. However, a dominant strategy is everybody to bid 0.5. Thus for the worst case that everybody has valuation 1, the exponential mechanism will return as a price value 0.5 with probability $\approx 1 - e^{-\epsilon \frac{n}{2}}$, thus the expected revenue is at most $n(0.5 + e^{-\epsilon \frac{n}{2}})$, although the $OPT = n$.

So Nissim et al. [22] introduce a combined mechanism of *exponential mechanism* and another probabilistic mechanism that guarantees truthfulness and approximates well the optimal *social welfare*. Intuitively, the latter mechanism with probability q punishes the agents that deviate from reporting their true types. This is

feasible because in this model, we introduce for each agent i a set of post-actions A_i after the alternative has been chosen.

Definition 1. *Environment:* Suppose n agents want to choose among a set S of alternatives. Each agent i has a type T_i that is her private information. Choosing a social alternative s, agent i has a finite set of available post-actions A_i that can take advantage of picking s. Finally, each agent i has a utility function $u_i : T_i \times S \times A_i \to [0, 1]$. *Environment* is defined as the tuple (T, S, A, u).

It is remarkable that a *mechanism design* except of the social alternative, it will also choose a subset of A_i, $\forall$ agent i (an element of $2^{A_i} \backslash \{\emptyset\}$). Additionally, we must mention that each agent i announces a type b_i. If $b_i = t_i$ then agent i tells the truth (truthful) otherwise she misreports her true type. For example, suppose a social planner wants to make two choices (create) among $S = \{$train station , port , airport , road$\}$ in order to help the citizens to commute. Each citizen i has a type T_i that "shows" what they actually prefer. After the fact that the citizens announce their types and the social planner picks the "ways" of travelling, each citizen must choose which way to commute in order to maximize his utility function. The goal of social planner is to maximize the social welfare.

Suppose for the rest of this section that we deal with a non-trivial *environment*, that is for any i and $\forall t_i \neq t_i'$, there is some alternative s such that $a_i(t_i, s) \cap a_i(t_i', s) = \emptyset$ where $a_i(t_i, s)$ denotes the set of optimal post-actions of agent i with type t_i when alternative s has already been chosen.

Definition 2. *gap:* Let P be a probability distribution over the set of alternatives, then

$$GAP_P(T, S, A, u) = \min_{i, t_i \neq b_i} \mathbb{E}[u_i(t_i, s, a_i(t_i, s)) - u_i(t_i, s, a_i(b_i, s))]$$

If we take the maximum of $GAP_P(T, S, A, u)$ over all probability distributions P of S we define the *gap* of *environment* (T, S, A, u).

Intuitively, *gap* is the maximal loss (over all probability distributions) if an agent's strategy (i for example) is to report a type $b_i \neq t_i$ (t_i is i's true type). We are interested in non-trivial *environments* in order the *gap* to be non-zero.

Gap mechanism: Let P be the probability distribution that maximizes $GAP_P(T, S, A, u)$. $\mathcal{M}_{gap}$ mechanism is defined as follows :

$$\mathcal{M}_{gap}(T, S, A, u) = \text{choose } s \text{ according to } P$$

The post-action of agent i is then an element of $a_i(b_i, s)$ where s is the chosen alternative and b_i is i's announced type.

It is rather straightforward to prove that the mechanism above is truthful (truthtelling is a dominant strategy).

Lemma 1: For any agent i with type t_i and announced type $b_i \neq t_i$ and every announced type vector b_{-i} the following inequality holds:

$$\mathbb{E}_{\mathcal{M}_{gap}(t_i,b_{-i})}[u_i(t_i, s, a_i(t_i, s))] > \mathbb{E}_{\mathcal{M}_{gap}(b_i,b_{-i})}[u_i(t_i, s, a_i(b_i, s))] + GAP(T, S, A, u)$$

Proof: Let i be an arbitrary agent and P the probability distribution that maximizes $GAP(T, S, A, u)$. Then $\mathbb{E}_{\mathcal{M}_{gap}(t_i,b_{-i})}[u_i(t_i, s, a_i(t_i, s)) - u_i(t_i, s, a_i(b_i, s))] > GAP_P(T, S, A, u)$ (definition). Equivalently, $\mathbb{E}_{\mathcal{M}_{gap}(t_i,b_{-i})}[u_i(t_i, s, a_i(t_i, s))] > \mathbb{E}_{\mathcal{M}_{gap}(t_i,b_{-i})}[u_i(t_i, s, a_i(b_i, s))] + GAP_P(T, S, A, u)$. Finally by definition follows that $GAP_P(T, S, A, u) = GAP(T, S, A, u)$ and also $\mathbb{E}_{\mathcal{M}_{gap}(t_i,b_{-i})}[u_i(t_i, s, a_i(b_i, s))] = \mathbb{E}_{\mathcal{M}_{gap}(b_i,b_{-i})}[u_i(t_i, s, a_i(b_i, s))]$ because $\mathcal{M}_{gap}$ *mechanism* is independent of the announced types (it depends only on T, S, A, u)

□

We are ready now to introduce the combined mechanism that is strategy-proof. For the rest of this section we consider that *social welfare function* $F \in [0, 1]$ and F is 1-*sensitive*, that is for every pair of types t, t' differing on a single user $|F(t, s)-F(t', s)| \leq \frac{1}{n}$ so as to $n\Delta F \leq 1$. This is a necessary hypothesis in order the exponential mechanism $\mathcal{M}_{\epsilon/2}(nF, t)$ to preserve ϵ-differential privacy (this follows from *Theorem 2* in section *3.2.2*). It is remarkable that the *social welfare function* F also depends on the types of the agents. Additionally, we should mention that F is not strongly correlated to each particular agent (if agent i changes her type, F changes by a $O(\frac{1}{n})$). A common example is $F = \frac{\sum_i u_i}{n}$ where $u_i \in [0, 1]$ for every agent i.

Combined mechanism: Let q be a real number such that $0 \leq q \leq 1$. The *combined mechanism* is denoted by

$$\mathcal{M}_{comb}(F, t, \epsilon) = (1 - q)\mathcal{M}_{\epsilon/2}(nF, t) + q\mathcal{M}_{gap}(t)$$

where $\mathcal{M}_{\epsilon/2}(nF, t)$ is the *exponential mechanism* as defined in section *3.2.2*.

Theorem 1: Let a non-trivial *environment* (T, S, A, u) with gap γ, a *social welfare function* F and $q\gamma \geq 2\epsilon$. The *combined mechanism* is truthful.

Proof: Let i be an arbitrary agent with true type t_i and announced type $b_i \neq t_i$. Then $\mathbb{E}[u_i(t_i)] = (1 - q)\mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,(t_i,b_{-i}))}[u_i(t_i)] + q\mathbb{E}_{\mathcal{M}_{gap}(t_i,b_{-i})}[u_i(t_i)]$. Additionally from *lemma 1* and *Theorem 1* of section *3.3.1* follows that $\mathbb{E}[u_i(t_i)] > (1 - q)e^{-\epsilon}\mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,(b_i,b_{-i}))}[u_i(b_i)] + q(\mathbb{E}_{\mathcal{M}_{gap}(b_i,b_{-i})}[u_i(b_i)] + \gamma)$. By observing the fact

that $e^{-\epsilon}x \geq x + e^{-\epsilon} - 1 \geq x - 2\epsilon$ ($x \in [0,1]$, using derivatives) we have that $\mathbb{E}[u_i(t_i)] > (1-q)\mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,(b_i,b_{-i}))}[u_i(b_i)] + q(\mathbb{E}_{\mathcal{M}_{gap}(b_i,b_{-i})}[u_i(b_i)] - 2\epsilon(1-q) + q\gamma \geq \mathbb{E}[u_i(b_i)] + 2\epsilon q$. Thus $\mathbb{E}[u_i(t_i)] > \mathbb{E}[u_i(b_i)]$

□

Except of the truthfulness, we are interested in exploring how close to the optimal ($OPT = \max_{s'} F(t,s')$) is the outcome of the *combined mechanism*. The following theorem answers this question.

Theorem 2: Using the *combined mechanism* $\mathcal{M}_{comb}$ with $\epsilon = \sqrt{\frac{\gamma}{n}} \cdot \sqrt{ln(\frac{n\gamma e}{2}|S|)}$ and $q = 2\epsilon/\gamma$ the following inequality holds:

$$\mathbb{E}[F(s,t)] \geq \max_{s'} F(s',t) - 4\sqrt{\frac{1}{\gamma n}} \cdot \sqrt{ln(\frac{n\gamma e}{2}|S|)}$$

where $|S|$ is the number of the alternatives.

Proof: Firstly, using the fact that $F \geq 0$ we conclude that $\mathbb{E}[F(s,t)] \geq (1-q) \cdot \mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,t)}[F(s,t)]$. Let B be the set of alternatives such that $F(s,t) < \max_{s'} F(s',t) - \delta$. From *Theorem 2* in section *3.2.2*, assuming that $\mathcal{R}_{OPT}$ is normalized and choosing $2t/n\epsilon = \delta$ or equivalently $t = \frac{n\epsilon}{2}\delta$ it occurs that $\Pr[s \in B] \leq |B| \cdot e^{-\frac{n\epsilon}{2}\delta} \leq |S| \cdot e^{-\frac{n\epsilon}{2}\delta}$. Thus $\mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,t)}[F(s,t)] \geq (1 - |S|e^{-\frac{n\epsilon}{2}\delta}) \cdot (OPT - \delta)$. Since $F(s,t) \leq 1$ so is $OPT - \delta < 1$ and the inequality becomes $\mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,t)}[F(s,t)] \geq OPT - \delta - |S|e^{-\frac{n\epsilon}{2}\delta}$. By choosing $\delta = \frac{2}{n\epsilon}\ln(\frac{n\epsilon|S|}{2})$ it follows that

$$\mathbb{E}_{\mathcal{M}_{\epsilon/2}(nF,t)}[F(s,t)] \geq OPT - \frac{2}{n\epsilon}\ln(\frac{n\epsilon|S|e}{2})$$

So $\mathbb{E}[F(s,t)] \geq (1 - \frac{2\epsilon}{\gamma})(OPT - \frac{2}{n\epsilon}ln(\frac{n\epsilon|S|e}{2})) \geq OPT - \frac{2}{n\epsilon}\ln(\frac{n\epsilon|S|e}{2}) - \frac{2\epsilon}{\gamma}$. Thus substituting ϵ and using the fact that for large n, $\epsilon \leq \gamma$ we have that $\mathbb{E}[F(s,t)] \geq OPT - 4\sqrt{\frac{1}{\gamma n}}\sqrt{\ln(\frac{\gamma ne}{2}|S|)}$.

□

From our analysis, it is obvious that $\gamma < 1$. For larger γ from theorem above we have better approximation with respect to the optimal (maximum) *social welfare*. So, when we deal with problems, if we can find a lower bound for γ, assume l, then it follows that $\mathbb{E}[F(s,t)] \geq OPT - 4\sqrt{\frac{1}{ln}} \cdot \sqrt{ln(\frac{nle}{2}|S|)}$ (we substitute γ for l). Additionally, it is sometimes inefficient to find the probability distribution P on S which realizes the gap γ. Instead we can find another probability distribution P' (see problem below) such that $GAP_{P'}(T,S,A,u) = \gamma' \leq \gamma$ and then use $\mathcal{M}_{gap}$ and $\mathcal{M}_{comb}$ according to P' (then obviously we substitute γ for γ' to the Theorem 2).

Application to auctions for digital goods

Assume a seller faces n bidders each of whom wants to buy a unit of a digital good. The seller must decide a fixed price $p \in \{0, \frac{1}{m}, ..., \frac{m-1}{m}, 1\} = S$ such that the bidders will pay to acquire the item (if they are willing to). Each bidder has a type $T_i \in \{0, \frac{1}{m}, ..., \frac{m-1}{m}, 1\}$ that shows how willing he is to buy the good. The set of post actions for each bidder i is $A_i = \{"buy", "not\ buy"\}$, his post-action function and his utility function are denoted respectively by

$$a_i(b_i, p) = \begin{cases} "buy" & \text{if } b_i \geq p \\ "not\ buy" & \text{if } b_i < p \end{cases}$$

$$u_i(t_i, p, a_i) = \begin{cases} t_i - p & \text{if } a_i = "buy" \\ 0 & \text{if } a_i = "not\ buy" \end{cases}$$

Furthermore, assume that each bidder i can't announce $b_i > t_i$ (only $b_i \leq t_i$ is allowed) and let $F(p, t) = \frac{p}{n} \cdot |i : t_i \geq p|$ which is obviously 1-sensitive. Consider P distribution to be uniform over S. We use the *combined mechanism* $\mathcal{M}_{comb}$ where $\mathcal{M}_{gap}$ chooses $p \in S$ according to P. We will prove that $GAP_P = \gamma \geq \frac{1}{m(m+1)}$. Let an arbitrary bidder i and $t_i \neq b_i$. Equivalently $t_i \geq b_i + \frac{1}{m}$, thus for $b_i \leq p < t_i$ we have that $u_i(t_i, p, a_i(t_i, p)) - u_i(t_i, p, a_i(b_i, p)) = t_i - p \geq \frac{1}{m}$. The probability of choosing $p \in [b_i, t_i)$ is at least $\frac{1}{m+1}$ (in case $t_i = b_i + \frac{1}{m}$, else it is larger), thus $\gamma \geq \frac{1}{m(m+1)}$. So from Theorem 2 it follows that $\mathbb{E}[F(p, t)] \geq OPT - 4\sqrt{\frac{m(m+1)}{n}} \cdot \sqrt{\ln(\frac{ne}{2m(m+1)}|S|)}$.

3.5 Conclusion

In this chapter, we examined the idea of differential privacy. First of all we discussed about exponential mechanism, a mechanism that is "universal", a nice technique two solve problems in a privacy manner. Furthermore, we saw differential privacy's applications to combinatorial optimization problems and also to game-theoretic ones. Finally, even though exponential mechanism can achieve approximate truthfulness as far as mechanism design is concerned, we used gap mechanism in order to create a combined mechanism that has good approximation and is "exact" truthful.

References

[1] Alon Noga, Fischer Felix,Procaccia Ariel D., and Tennenholtz Moshe. Sum of Us: Strategyproof Selection from the Selectors. Working paper

[2] Arya V., Garg N., Khandekar R., Meyerson A., Munagala K., and Pandit V.. Local search heuristics for k- median and facility location problems. *SIAM J. Comput.*, 33(3):544-562, 2004.

[3] Bartholdi J.; Tovey, C. A.; and Trick, M. A. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6:227-241.

[4] Brams S. J. and Fishburn P. C. Approval Voting. Springer, 2nd edition, 2007.

[5] Caragiannis I. Kaklamanis C., Karanikolas N. and Procaccia A. D.. Socially Desirable Approximations for Dodgson's Voting Rule.

[6] Clarke E. H. Multipart pricing of public goods. *Public Choice*, 11(1), September 1971.

[7] Conitzer V., and Sandholm, T. 2003. Universal voting protocol tweaks to make manipulation hard. In *Proc. of 18th IJCAI*, 781- 788.

[8] Conitzer V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of 21st AAAI*, 627-634.

[9] Dwork C., McSherry F., Nissim K., and Smith A. Calibrating noise to sensitivity in private data analysis. In S. Halevi and T. Rabin, editors, TCC, volume 3876 of *Lecture Notes in Computer Science*, pages 265-284. Springer, 2006.

[10] Gibbard A. 1973. Manipulation of voting schemes: A general result. *Econometrica*, 41:587-602, July 1973.

[11] Gibbard A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45:665-681.

[12] Groves T. Incentives in teams. *Econometrica*, 41:617-631, 1973.

[13] Gupta A., Ligett K., McSherry F., Roth A., Talwar K.: Differentially Private Combinatorial Optimization. *SODA 2010*:1106-1125

[14] Hemaspaandra, E., and Hemaspaandra, L. A. 2007. Dichotomy for voting systems. *Journal of Computer and System Sciences* 73(1):73-83.

[15] Karger D. R.. Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (Austin, TX, 1993)*, pages 21-30, New York, 1993. ACM.

[16] Laslier Jean-Francois, Sanver M. Remzi. Handbook on Approval Voting

[17] Maskin Eric. Mechanism Design: How to Implement Social Goals. February 2008

[18] McSherry F. and Talwar K. Mechanism Design via Differential Privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94-103, 2007.

[19] Motwani R., Raghavan P.: Randomized Algorithms

[20] Nemhauser G. L. ,Wolsey L. A., and Fisher M. L. An analysis of approximations for maximizing submodular set functions ii. *Math. Programming Study 8*, pages 73-87, 1978.

[21] Nisan Noam, Roughgarde Tim, Tardos Eva, and V. Vazirani Vijay. Algorithmic Game Theory

[22] Nissim Kobbi, Smorodinsky Rann and Tennenholtz Moshe. Approximately Optimal Mechanism Design via Differential Privacy. April 2010. Draft version

[23] Papadimitriou C., Schapira M., and Singer Y. On the Hardness of Being Truthful. In *Foundations of Computer Science, 2008. FOCS'08. 49th Annual IEEE Symposium on*, 2008.

[24] Procaccia A. D., and Tennenholtz, M. 2009. Approximate mechanism design without money. In *Proc. of 10th EC*, 177-186.

[25] Procaccia A.D. Can Approximation Circumvent Gibbard-Satterthwaite? In *Proc. 24th AAAI Conference on Artificial Intelligence*, pp. 836-841, Jul 2010.

[26] Satterthwaite M. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187-217.

[27] Schummer J. Almost-dominant strategy implementation. *Games and Economic Behavior*, 48:154-170, 2004.

[28] Vazirani V. Approximation algorithms. Springer-Verlag, 2001.

[29] Vickrey W. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8-37, 1961.

[30] Vodak Jan. A note on concentration of submodular functions (2010)

[31] Yao A. C. 1977. Probabilistic computations: Towards a unified measure of complexity. In *Proc. of 17th FOCS*, 222-227.

[32] Notes on Randomized Algorithms, Berkeley.

[33] Αγγελής Γεώργιος. Computational Considerations of Voting Rules. Διπλωματική Εργασία, 2010.

Appendix

Missing Proofs

Theorem 1: *Rearrangement inequality:*

Let $a_1, ..., a_n,\ b_1, ..., b_n$ be two non-decreasing sequences of positive real numbers and $b'_1, ..., b'_n$ be a permutation of sequence b_n. Then the following inequality holds

$$\sum_{i=1}^{n} a_i b_i \geq a_i b'_i$$

Proof: Suppose the contrary. Let j be the smallest index such that $b'_j > b'_{j+1}$ and $a_j < a_{j+1}$ (if there is no such j then the inequality holds). Then the right part of the inequality is equal to $\sum_{i=1}^{j-1} a_i b'_i + a_j b'_j + a_{j+1} b'_{j+1} + \sum_{i=j+2}^{n} a_i b'_i$. However $a_j b'_j + a_{j+1} b'_{j+1} < a_j b'_{j+1} + a_{j+1} b'_j$ because $(a_{j+1} - a_j)(b'_j - b'_{j+1}) > 0$. This leads to contradiction because we found a larger sum by swapping b'_j, b'_{j+1}.

□

Theorem 2: *BCS inequality:*

Let $a_1, ..., a_n,\ b_1, ..., b_n$ be two sequences of positive real numbers. The following inequality holds

$$\left(\sum_{i=1}^{n} a_i^2\right) \cdot \left(\sum_{i=1}^{n} b_i^2\right) \geq \left(\sum_{i=1}^{n} a_i b_i\right)^2$$

Proof: After operations we have that $\left(\sum_{i=1}^{n} a_i^2\right) \cdot \left(\sum_{i=1}^{n} b_i^2\right) - \left(\sum_{i=1}^{n} a_i b_i\right)^2$ becomes equivalently $\frac{1}{2} \sum_{1 \leq i \leq n} \sum_{1 \leq j \leq n, j \neq i} (a_i b_j - a_j b_i)^2 \geq 0$

□

Theorem 3: (*Yao's Minimax Principle*) Let $\mathcal{A}$ (set of deterministic algorithms) and $\mathcal{I}$ (set of inputs) be nonempty sets and $k : \mathcal{A} \times \mathcal{I} \to \mathcal{R}$ be a cost function, and let σ and τ be probability distributions on $\mathcal{A}$ and $\mathcal{I}$. Let A_σ be a random variable with values in $\mathcal{A}$ and distribution σ and let I_τ be a random variable with values in $\mathcal{I}$ and distribution τ. Then we have

$$\min_{A \in \mathcal{A}} \mathbb{E}[k(A, I_\tau)] \leq \max_{I \in \mathcal{I}} \mathbb{E}[k(A_\sigma, I)]$$

Proof: Let $\mathbf{k} = \sum_{(A,I) \in \mathcal{A} \times \mathcal{I}} \Pr_\sigma[A] \cdot \Pr_\tau[I] \cdot k(A, I)$. We have the following:

$$\begin{aligned}
\min_{A \in \mathcal{A}} \mathbb{E}[k(A, I_\tau)] &= \min_{A \in \mathcal{A}} \sum_{I \in \mathcal{I}} \Pr{}_\tau[I] \cdot k(A, I) \\
&\leq \sum_{A \in \mathcal{A}} \Pr{}_\sigma[A] \sum_{I \in \mathcal{I}} \Pr{}_\tau[I] \cdot k(A, I) \\
&= \mathbf{k} \\
&= \sum_{I \in \mathcal{I}} \Pr{}_\tau[I] \sum_{A \in \mathcal{A}} \Pr{}_\sigma[A] \cdot k(A, I) \\
&\leq \max_{I \in \mathcal{I}} \sum_{A \in \mathcal{A}} \Pr{}_\sigma[A] \cdot k(A, I) \\
&= \max_{I \in \mathcal{I}} \mathbb{E}[k(A_\sigma, I)]
\end{aligned}$$

It is remarkable, that this inequality stands for minimization problems. For maximization ones, the inequality becomes

$$\max_{A \in \mathcal{A}} \mathbb{E}[k(A, I_\tau)] \geq \min_{I \in \mathcal{I}} \mathbb{E}[k(A_\sigma, I)]$$

and the proof is the same.

□