Εθνικο Μετσοβιο Πολυτεχνειο
Τμημα Ηλεκτρολογων Μηχανικων και Μηχανικων Υπολογιστων

Τομεας Τεχνολογιας Πληροφορικης και Υπολογιστων
Εργαστηριο Λογικης και Επιστημης Υπολογιστων

Σχεδιασμός Μηχανισμών υπό Περιορισμούς Ρευστότητας

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

ΔΗΜΗΤΡΙΟΥ Η. ΤΣΙΠΡΑ

Επιβλέπων: Δημήτρης Φωτάκης
Επίκουρος Καθηγητής Ε.Μ.Π.

Αθήνα, Οκτώβριος 2014

Εθνικο Μετσοβιο Πολυτεχνειο
ΤΜΗΜΑ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ
ΕΡΓΑΣΤΗΡΙΟ ΛΟΓΙΚΗΣ ΚΑΙ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΤΩΝ

Σχεδιασμός Μηχανισμών υπό Περιορισμούς Ρευστότητας

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

του

ΔΗΜΗΤΡΙΟΥ Η. ΤΣΙΠΡΑ

Επιβλέπων: Δημήτρης Φωτάκης
Επίκουρος Καθηγητής Ε.Μ.Π.

Εγκρίθηκε από την τριμελή εξεταστική επιτροπή την 22η Οκτωβρίου 2014.

....	....	....
Δημήτρης Φωτάκης	Ευστάθιος Ζάχος	Αριστείδης Παγουρτζής
Επίκουρος Καθηγητής Ε.Μ.Π.	Καθηγητής Ε.Μ.Π.	Επίκουρος Καθηγητής Ε.Μ.Π.

Αθήνα, Οκτώβριος 2014.

....
ΔΗΜΗΤΡΙΟΣ Η. ΤΣΙΠΡΑΣ
Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Ευχαριστίες

Θα ήθελα να ευχαριστήσω θερμά τους καθηγητές του Εργαστηρίου Λογικής και Επιστήμης Υπολογιστών, κ. Ζάχο, κ. Παγουρτζή και κ. Φωτάκη για την υποστήριξη και τη συνεργασία τους, τόσο κατά τις προπτυχιακές μου σπουδές όσο και κατά μείζονα λόγο κατά την εκπόνηση της διπλωματικής μου εργασίας.

Ιδιαίτερα θα ήθελα να ευχαριστήσω τον επιβλέποντα κ. Φωτάκη για την ανεκτίμητη καθοδήγηση και στήριξη που μου παρείχε κατά την προσπάθεια αυτή, τόσο σε προσωπικό όσο και σε ακαδημαϊκό επίπεδο. Μέσω της εργασίας αυτής με εισήγαγε με ουσιαστικό τρόπο στην ερευνητική διαδικασία, και ο χρόνος που διέθεσε σε αυτό όσο και η εμπιστοσύνη που έδειξε στις επιλογές μου, ήταν καθοριστικής σημασία για την πορεία μου.

Επίσης θα ήθελα να ευχαριστήσω όλα τα μέλη του εργαστηρίου Λογικής και Επιστήμης Υπολογιστών για την στήριξή τους και το εξαιρετικό περιβάλλον εργασίας και συνεργασίας που μου προσέφεραν. Ιδιαίτερα θα ήθελα να ευχαριστήσω τους Μανώλη Ζαμπετάκη και Χρήστο Τζάμο, για την τεράστιας σημασίας άμεση και έμμεση συνεργασία στο ερευνητικό τμήμα της παρούσας εργασίας. Οι τεχνικές λείανσης του χώρου βελτιστοποίησης που παρουσιάζονται οφείλονται σε δική τους προσπάθεια.

Τέλος, ένα μεγάλο ευχαριστώ οφείλω στην οικογένειά μου και τους ανθρώπους που ήταν κοντά μου καθ΄όλη τη διάρκεια της εργασίας αυτής και γενικότερα των προπτυχιακών μου σπουδών. Η στήριξη και συμπαράστασή τους ήταν ανεκτίμητη.

Περίληψη

Κατά το σχεδιασμό μηχανισμών, ο στόχος είναι η δημιουργία δημοπρασιών που επιτυγχάνουν (ή προσεγγίζουν) τη βέλτιστη κοινωνική ευημερία. Προκειμένου να γίνουν γνωστές οι ιδιωτικές πληροφορίες που παρέχονται από στρατηγικές οντότητες, επιτρέπεται η χρήση χρηματικών ανταλλαγμάτων. Στο μεγαλύτερο μέρος της βιβλιογραφίας, τα ανταλλάγματα αυτά αντιμετωπίζονται απλά ως μέσο απόκτησης των πληροφοριών.

Ωστόσο, σε πολλές περιπτώσεις, υπάρχουν όρια στο ύψος των χρημάτων που μπορούν να χρησιμοποιηθούν στα πλαίσια της δημοπρασίας, περιορίζοντας έτσι το πλήθος των δυνατών σχεδιαστικών επιλογών. Στην εργασία αυτή μελετάμε την επίδραση των περιορισμών αυτών και τρόπους παράκαμψης των δυσκολιών που δημιουργούν. Παρουσιάζουμε τα πρόσφατα αποτελέσματα σε δημοπρασίες με περιορισμένη ρευστότητα από την πλευρά των πλειοδοτών, και αντίστροφες δημοπρασίες με περιορισμένο προϋπολογισμό από τη μεριά του αγοραστή. Στη συνέχεια αναπτύσσουμε μία νέα προσέγγιση μείωσης των χρηματικών ανταλλαγμάτων που απαιτούνται και παρουσιάζουμε τις εφαρμογές της σε υπάρχοντα περιβάλλοντα.

Λέξεις Κλειδιά: Σχεδιασμός Μηχανισμών, Δημοπρασίες, Περιορισμοί Προϋπολογισμού, Περιορισμένη Ρευστότητα

Abstract

In Mechanism Design, the goal is to come up with mechanisms that achieve (or closely approximate) the optimal social welfare achievable. In order to truthfully extract private information from the strategic entities holding them, the mechanism is allowed to impose monetary transfers upon them. In the majority of research in algorithmic mechanism design, payments are treated solely as a means to extract this private information.

However, in many cases there are limits to the amount of monetary tranfers allowed, thus limiting the space of allowed mechanisms. In this thesis, we study the effect of imposing budget constraints upon the entities of the mechanism (the agents or the auctioneer) and the possible ways of circumventing the obstacles created by them. We survey recent results in auctions with budget constrained bidders and procurement auctions with a budget constrained center. We then develop a novel approach to reducing payments in order to respect payment bounds, and present applications to existing settings.

Keywords: Mechanism Design, Auctions, Procurement, Budget Constraints, Liquidity Constraints

Contents

Chapter 1

Introduction

The fact that we belong in a society, inevitably leads to the need for common decisions affecting the society as a whole. However, more often than not, the incentives of the society members are not aligned wth each other. The selfishness of the society members comes in the way of the decision making process affecting it in a variety of ways.

A great amount of research has been conducted to come up with decision making procedures called mechanisms that incentivize the participants to behave in a predictable and desirable way. In this thesis, we present some of the classic results in mechanism design, namely the seminal VCG mechanism along with its benefits and drawbacks. Along the way, we define most of the notation and terminology used in the mechanism design literature and provide insight to the already existant results.

We then procced to study the effect of liquidity constraints imposed when participants in the mechanism have an upper bound to the amount of monetary tranfers they can charged. We survey the results in classic auctions with budget constrained bidders, presenting the developed mechanisms as well as some intuiton on their inner workings and rationale. To justify the efficiency of these designs we measure their performance with respect to various proposed objectives such as Pareto Optimality and Liquid Welfare.

Additionally we survey the results in procurement auctions, where a center with a limited budget wishes to purchase goods from the participants. This is a challenging setting with applications to the modern electronic market where companies wish to purchase services offered by a large amount of individuals. By studying these results we will also meet a certain amount of auctions following the very successful paradigm of posted price auctions, that is auctions where agents are faced with simple take-it-or-leave-it decision, severely limiting the complexity of the agent-center interface.

Finally we propose a novel approach to handling liquidity constraints, through uniformly smoothing the space of mechanism outcomes. Through simple mathematical constructs we are able to create mechanisms that treat agents' values in a smooth way, thus removing a configurable amount of power from them. This way we can provide hard limits to the amount of payments charged. We present results of applying this approach to previously studied settings, as well as objectives rising from imposing different liquidity constraints on the auction.

1.1 Mechanism Design

Game Theory is the field concerned with the study of strategic agents' actions when participating in a common game. Each agent is selfish, in the sense that his sole goal is maximizing his own benefit. In the context of decision making, the selfish nature of the agents, leads them to attempt to manipulate the decision making process in order to achieve an outcome they desire. This manipulation however, often leads to outcomes that are *suboptimal* for the society as a whole.

A great deal of effort has been spent on quantifying the inefficiency of such behavior, that is how much worse the outcome gets when players act selfishly compared to acting in an optimal coordinated way. In many natural games, the outcome can be arbitrarily bad compared to the optimum, a result often seen in everyday life as the tragedy of the commons.

Mechanism Design can be thought as reverse game theory. It aims towards the design of decision making procedures, in which participants have a certain predictable and desirable behavior. Specifically, since agents are strategic and pursue their own goals, we aim to align their individual interests with the designer's interests.

1.2 Notation and Preliminaries

Formally, there is a set of outcomes O, a set of n agents N and each agent has a private valuation function $u_i : O \to \mathbb{R}$ (sometimes also referred to as type) from a set of possible valuations V_i. The designer or center of the mechanism wants to enforce the result of a function $f : V_1 \times \ldots \times V_n \to O$, called the allocation rule, from the agent utility functions to the set of outcomes. We will call a vector valuation function $u = (u_1, \ldots, u_n)$ a valuation profile and will denote by u_{-i} the same vector without the i-th coordinate. With a slight notation abuse we will denote u as (u_{-i}, u_i). Finally we denote as V the space of valuation profiles $V_1 \times \ldots \times V_n$ and as V_{-i} the set $V_1 \times \ldots \times V_{i-1} \times V_{i+1} \times \ldots \times V_n$.

However, the valuation of each agent is *private* and he will not report his true value if it is not profitable for them to do so. Since, f is designed to be computed on the valuation profile, allowing agents to falsely submit their types can lead to undesirable outcomes. In order for the computation to be performed on the actual problem input we will require that every agent is incentivized to report his true valuation.

Definition 1. *An allocation rule is* truthful *iff* $\forall i \in N,\ u_i \in V_i,\ u_{-i} \in V_{-i}, u'_i \in V_i$*:*

$$u_i(f(u_{-i}, u_i)) \geq u_i(f(u_{-i}, u'_i))$$

From another point of view, since the input of the allocation rule is provided from the agents in a totally uncontrolled fashion we need to make some behavioral assumptions on how the agents will act. The weakest possible assumption is that if every agent has a *dominant strategy* they will act according to it.

Definition 2. *The type* $t \in V_i$ *is a dominant strategy for agent* i *if* $\forall\, u_{-i} \in V_{-i}, t' \in V_i$*:*

$$u_i(f(u_{-i}, t)) \geq u_i(f(u_{-i}, t'))$$

Notice that the definition does not require $t = u_i$, and one might naturally wonder if designing rules that are not truthful but have dominant strategies for each agent allow for a richer design space. Moreover, could there be more complex interactions with the agents, rather than simply have them disclose their types, that enable different designs. Mechanisms where agents simply disclose their type are called *direct*. The answer to both questions is negative.

Theorem 3 (Revelation Principle)**.** *Any mechanism that implements a certain allocation rule* f *in dominant strategies, can be transformed to a* direct truthful *mechanism implementing* f*.*

Proof. Suppose that each agent has a certain set of strategies S_i and a function s_i such that $s_i(u_i)$ is the dominant strategy the agent has. The allocation rule is a function $g : \prod_i S_i \to O$ operating on the agent strategies and achieves that $g(s_1(u_1), \ldots, s_n(u_n)) = f(u_1, \ldots, u_n)$ by hypothesis. Since the functions s_i are known we can create a *direct* mechanism h, defined as $h(u_1, \ldots, u_n) = g(s_1(u_1), \ldots, s_n(u_n))$. By definition, if the initial rule g implements f, then h also implements f. Moreover since h exactly simulates the behavior of players in the execution of g, players have no incentize to lie to h, since that would be equivalent to lying to themselves when faced with g. $\square$

Therefore we will limit our attention to direct truthful rules without loss of generality. A seminal result by Gibbard and Satterthwaite leaves no hope for the existence of truthful non-trivial rules in this general setting when there are more than 3 alternatives.

Theorem 4 ([14, 25])**.** *There are no allocation rules that simultaneously satisfy:*

- $|O| \geq 3$
- *the rule is* onto*, that is for any outcome* o *there is a valuation profile* $v \in V$ *such that* $f(v) = o$
- *the rule is* truthful

- *the rule is not* dictatorial, *that is there is no agent i with valuation u_i such that for any $u_{-i} \in V_{-i}$, $f(u_i, u_{-i}) = \text{argmax}_o u_i(o)$*

The requirements of the rule in the above theorem are quite weak. The second assumption simply rules out "dummy" outcomes. The third requirement rules out the rules that simply output the preferred outcome of an arbitrary agent.

To escape this impossibility, two main alternatives are widely researched. The first is the restriction of the utility functions to match some of the environment properties, and has produced a variety of important results in areas such as facility location. The second is allowing the designer to impose payments of the agents. In this thesis we will consider mainly the second.

Formally, a mechanism M is a pair of functions (f, p) where f is the allocation rule $f : V \to O$ as above and p is the payment rule $p : V \to \mathbb{R}^n$. Under this payment rule agent i has to pay an amount $p_i(v)$. Now the agents are not only interested in maximizing his utility $u_i(o)$ but also in minimizing his payments $p_i(v)$. We assume agents are *quasi-linear* and aim to maximize their *surplus*, that is the amount $u_i(f(u_i, u_{-i})) - p(u_i, u_{-i}))$. Naturally this changes the definition of truthfulness for mechanisms.

Definition 5. *A mechamism $M = (f, p)$ is defined as truthful iff $\forall i \in N,\ u_i \in V_i,\ u_{-i} \in V_{-i},\ u'_i \in V_i$:*

$$u_i(f(u_{-i}, u_i)) - p(u_{-i}, u_i) \geq u_i(f(u_{-i}, u'_i)) - p(u_{-i}, u'_i)$$

As we mentioned earlier, the imposition of paymetns allows us to construct truthful mechanism for a variety of not truthful allocation rules.

Definition 6. *An allocation rule f is* impementable *if there exists a payments scheme p such that the mechanism (f, p) is truthful.*

The mechanism is designed in order to output an outcome with certain designed properties. In order to be able to quantify how good the mechanism performs with respect to its purpose, we will assume that there is some objective function from outcomes to real numbers (depending on the agent true values) that the center wants to maximize. Fox example the center may wish to maximize the amount of money collected, also known as the *revenue* of the mechanism.

Definition 7. *For a mechanism $M = (f, p)$, we define as the* revenue *of the mechanism the quantity*

$$Rev(u) = \sum_i p_i(v)$$

In this work we concentrate on the most commonly used objective function, the social welfare. This captures in some sense the total amount of "happiness" the society enjoys for this particular outcome.

Definition 8. *For an outcome $o \in O$ we define as* social welfare *(SW) the quantity*

$$SW(o) = \sum_i u_i(o)$$

1.3 The VCG mechanism

1.3.1 Mechanism Description

An early result by Vickrey, Clarke and Groves, shows that maximizing the social welfare is always an implementable rule, independently of the agents' valuations. Definition 5 states that implementable mechanisms are exactly those which optimize for each agent over his set of valuations. Intuitively this means that in order to achieve truthfulness we need to align the incentives of each agent with the incentives of the mechanism designer.

Since we are interested in SW maximization it seems natural to award each agent for the utilities enjoyed by the rest of the agents. Defining the payment scheme as

$$p_i(u) = -\sum_{k \neq i} u_i(f(u))$$

makes each agent interested in the utilities of others exactly as much as in his own.

Definition 9 (Groves mechanisms)**.** *The mechanisms defined by*

$$f(u) = \text{argmax}_{o \in O} \sum_i u_i(o)$$

$$p_i(u) = h(u_{-i}) - \sum_{j \neq i} u_j(f(u))$$

in truthful and exactly optimizes SW.

Proof. The SW optimization follows trivially by the definition of the mechanism. The surplus of agent i for the outcome of the mechanism o

$$u_i(o) - p_i(o) = u_i(o) + \sum_{j \neq i} u_j(o) - h(u_{-i}) = \sum_i u_i(o) - h(u_{-i}) \geq \sum_i u_i(o') - h(u_{-i})$$

for any other outcome o' by the definition of the allocation rule. Therefore the chosen outcome results in maximal utility for each agent amongst outcomes, rendering the mechanism truthful. □

Notice how the term $h(u_{-i})$ is of no importance as far as truthfulness is concerned since agent i cannot affect it in any way. It is up to the designer to define a suitable function h, therefore we refer call the mechanisms of the definition a family of mechanisms.

We will now demand some additional properties from the mechanism. In most cases, users have the choice to not participate in a mechanism. In order to incentivize them to do so, we need to guarantee that they cannot lose by participation. This is a quite natural requirement, known as individual rationality.

Definition 10. *We call a mechanism* individually rational *(IR) if*

$$\forall i, u : u_i(f(u)) - p(u) \geq 0$$

Moreover, in many settings it not desirable for the mechanism to hand out money to the agents, and we will require that $p_i(u) \geq 0$.

A payment scheme that satisfies these requirements in many settings was proposed by Clarke and is known as the *Clarke tax* or the *pivot rule*. It is defined by setting

$$h_i(u_{-i}) = \max_o \sum_{j \neq i} u_j(o) = f(u_{-i}) \tag{1.1}$$

Intuitively, h outputs the outcome of the mechanism if it was applied to the set of agents, with i removed from it.

The Groves mechanism with h defined as in (1.1) is called the *VCG mechanism*, due to the contributions of Vickrey, Clarke and Groves.

Definition 11 (VCG mechanism [28, 7, 16])**.** *The VCG mechanism is defined by* (f, p) *such that*

$$f(u) = \text{argmax}_o \sum_i u_i(o)$$

$$p_i(u) = \sum_{j \neq i} u_j(f(u_{-i})) - \sum_{j \neq i} u_j(f(u))$$

The payment of each agent under the VCG mechanism is exactly equal to his *externalities*, that is how much he affects the utility of others by participating. In most natural settings, such as auctions and public project decisions, the VCG mechanism is IR, and this can be proven with certain mild requirements for the setting. Also, in many settings, the mechanism never pays money to the agents, however there is also a variety of settings, mostly procurement auctions, where the agents suffer costs, for being selected by the mechanism, and have to be paid to ensure IR.

1.3.2 Examples

Single Indivisible Good Auction

Suppose now, that we wish to sell a single item to a set of players. Each agent has a value v_i if he recieves the item and 0 otherwise. We wish to maximize the social welfare, so we will allocate the item to the highest bidder. Obviously, requiring the winner to pay his bid is not a truthful auction, since he has the incentive to understate his bid to decrease his payments, while continuing to recieve the item.

The mechanism is maximizing the SW, so we can calculate payments according to the VCG mechanism. The utilities of losing bidders are zero, while the mechanism when run without the highest bidder will allocate the item to the second highest bidder. Therefore the payments of the winner will be the second highest bid. Assuming w.l.o.g. that $v_1 \geq \ldots \geq v_n$, and with some arbitrary tie-breaking rule.

$$p_i(u) = \begin{cases} v_2 & \text{if } i = 1 \\ 0 & \text{otherwise} \end{cases}$$

One can verify the truthfulness of this payment scheme independently, as well as the fact that $\forall i,\ 0 \leq p_i(u) \leq u_i(f(u))$. This auction is known as the Vickrey auction or the second price auction and was discovered before the VCG framework.

Public Project

Suppose that the goverment is considering building a bridge that is estimated to cost a certain amount of money C. Each citizen i enjoys utility $v_i \in \mathbb{R}^+$ if the bridge is built and 0 otherwise. We can model the goverment as an extra agent with utility $-C$ if the bridge is built and 0 otherwise. The goverment is interested in maximizing the total utility so it builds the bridge iff $\sum_i v_i \geq C$. However in this game, every agent has a dominant stategy to report C instead of u_i so that the bridge is built anyway, so in order to ensure truthfulness the goverment has to enforce payments.

Since the objective is SW, the VCG payments will work and are calculated as follows

$$p_i(u) = \begin{cases} C - \sum_{j \neq i} v_i & \text{if } \sum_{j \neq i} v_j < C \leq \sum_j v_j \\ 0 & \text{otherwise} \end{cases} \tag{1.2}$$

These equations make the reason why this payment scheme is called pivot rule apparent. The only players the are assigned non-zero payments, are the ones that are *pivotal*, that is they can change the outcome of the mechanism by their presence. It easy to verify by the definition of payments that the mechanism is IR and never pays money to the agents. For every agent i with non-zero payments

$$\sum_{j \neq i} v_j < C \leq \sum_j v_j \Leftrightarrow 0 < C - \sum_{j \neq i} v_j \leq u_i \Leftrightarrow 0 < p_i(u) \leq u_i$$

1.3.3 Drawbacks of the VCG

Since VCG solves the problem of truthful SW maximization in the general setting, one might assume that this field is closed. However, VCG is not without is downsides, some of which we mention here.

VCG is not always tractable

One of the biggest issues arising in practice is the fact that VCG requires choosing the outcome maximizing social welfare. There are many setting, for example combinatorial auctions (selling m heterogenous items to n players), where the number of possible outcomes rises exponentially with the bidders or the items sold. When valuations are arbitrary even describing the valuation function to the mechanism is impossible in practice. Alternatively, the maximization of SW could require solving an NP-hard problem. Even though for many of them approximation algorithms exists, applying the VCG mechanisms with such rules is not truthful in general.

VCG is not group strategyproof

Although we proved that VCG cannot be manipulated by a single agent in a profitable way, agents can collide and coordinate to manipulate the mechanism. Consider the case of the public project, outlined above. If a pair of agents report that their values are both C, the goverment will build the bridge. Moreover, none of the agents are pivotal anymore, so all paymetns are 0. This outcome is at least a good as the one produced by truthful bidding for both the misreporting agents.

VCG is not frugal

Consider the case of the public project as described above and the instance with n agents each with value for the bridge

$$v_i = v = \frac{C}{n-1} - \frac{C}{n(n-1)}$$

. The value of $n-1$ agents for the bridge is equal to

$$(n-1)v = C - \frac{C}{n} < C$$

and the value of n agents

$$nv = n\frac{C}{n-1} - \frac{C}{n-1} = \frac{(n-1)C}{n-1} = C$$

Therefore every agents is pivotal and will pay according to Equation 1.2

$$p_i = C - (n-1)v = \frac{C}{n}$$

Therefore the total amount of payments collected is $\sum p_i = n\frac{C}{n} = C$, equal to the cost of the bridge! In the end the amount charged to the citizens to enforce truthfulness is enough to actually build the bridge, despite the fact that the funds were supposed to be provided by other sources.

Chapter 2

Budget Constraints

In chapter 1 we treated the payments imposed by the mechanism as a means of truthfully extracting private information from the agents. However, in real-life situations, bidders are not able to match any payment the mechanism requires and in fact may not be able to match their own value for the good in money. This leads to the introduction of budget constraints in auctions, these *hard* bounds to the amount an agent can be charged.

Depending on the specific setting, the nature of these liquidity constraints can vary. We may have public or private budgets, applying to each agent or the set of them as a whole, or even have a budget constrained designer wishing to purchase goods or services from strategic agents. In the extreme case, in certain scenarios it may be infeasible, immoral or even legally prohibited to impose any charge to the agents of the mechanism. Examples of such environment are political elections, kidney exchange procedures, assigning students to public school and many more natural cases.

In this chapter we will mostly consider agents with individual budget constraints and survey results in this direction.

2.1 The Infinitely Divisible Good Scenario with Budgets

The main setting we will consider is the auction of a single infinitely divisible item to budget constrained bidders. Although the development of these auctions initiated from the multi-unit indivisible case, we find it mathematically convenient to focus on the divisible setting. We first define the setting formally.

Each agent reports a value $v_i \in \mathbb{R}$ to the mechanism which represent his utility for acquiring the whole item. We assume these utilities are *additive*, and therefore the utility of an agent for receiving a fraction x_i of the item is $x_i \cdot v_i$.

Moreover, each agent has a budget B_i which is the maximum amount he can be charged by the mechanism. We will not change the utiliity function to accomodate for this modification and will simply treat outcomes where payments exceed budgets as infeasible. We will differantiate the case where budgets are public knowledge and the case where they are private information and as such have to be reported to the mechanism in a strategic way.

We will require out mechanisms to be:

- truthful, agents maximize their surplus $x_i v_i - p_i$ by bidding truthfully
- individually rational, that is agents never receive negative surplus
- budget feasible, $p_i \leq B_i$

Since the value of each agent is $x_i v_i$ this is a single parameter environment and Myerson's Lemma applies. We fix the values of all agents except i and denote the allocation and payment of agent i as $x(v_i), p(v_i)$.

Theorem 12 (Myerson's Lemma [21]). *For any single parameter environment the following are equivalent:*

- *the allocation rule is* monotone*, that is*

$$v_i \leq v_i' \Rightarrow x(v_i) \leq x_(v_i')$$

- *the unique payment scheme implementing the allocation is*

$$p(v_i) = x(v_i) \cdot v_i - \int_0^{v_i} x(z) \mathrm{d}z$$

Our environment is monotone and as such the design of truthful auction reduces to the design of monotone allocation rules.

We will first describe two auctions for the public budget case and the intuition behind them. Then, we will show why the classic social welfare objective cannot be used to measure performance guarantees and study two different proposed objectives. Finally we will measure the auctions' performance with respect to these objectives. For the rest of the chapter we will assume the values of the agents are pairwise distinct. In the opposite case we can infitinesimally perturb the input or choose some tie-breaking rule. We will ignore such technical details and focus on the case of pairwise distrinct input.

2.2 The Market Clearing Price Auction

First we will describe an auction inspired by market equilibrium theory. The description and properties of the auction follow [10]. The market clearing price is the price where the supply and the demand of the good are equal. In our setting we naturally define the supply as

$$S = 1$$

Moreover, if we sell the good at price p per unit and $p \geq v_i$ then player i will not desire to purchase any amount of the good. If $p < v_i$ then the player wants to buy as much fraction as he can. It is therefore natural to define the demand at price p as

$$D_i(p) = \begin{cases} \frac{B_i}{p} & \text{if } p < v_i \\ 0 & \text{otherwise} \end{cases}$$

The idea here is to calculate the maximum price p^* where $\sum_i D_i(p^*) \leq S$. Notice that we cannot necessarily ensure that demand equals supply, since the demand drops abruplty when the price reaches some agent's value.

Assume that agents are sorted in decreasing value order, that is $v_1 > v_2 > \ldots > v_n$. We wish to find the maximum price p such that

$$\sum_i D_i(p) \leq S \Leftrightarrow \sum_{i:v_i>p} \frac{B_i}{p} \leq 1 \Leftrightarrow \sum_{i:v_i>p} B_i \leq p$$

The demand of an agent drops in two ways

- By increasing the price and staying below the value ($D_i(p) = B_i/p$)
- By exceeding the value ($D_i(p) = 0$ if $p \geq v_i$)

This directly leads to 2 cases

1. If the demand drops because of the first reason then $\sum D_i(p) = S$ and $p = \sum_{i=1} kB_i$ where k is the largest integer such that $\sum_{i=1}^{k} B_i \leq v_k$.
2. If the demand drops due to an agent falling out of the auction the price will be that agents value. Let k be the largest integer such that $\sum_{i=1}^{k} B_i \leq v_k$. Then the agent falling out will be agent $k+1$ and therefore agent $1, \ldots, k$ will get allocated a share at price v_{k+1}. In this case the total demand of the first k players at price v_{k+1} will be below 1 and we will allocate the remaining good to agent $k+1$.

By the description above it becomes apparent that, if we set k as the maximum integer such that $\sum_{i=1}^{k} B_i \leq u_k$ the case we are in depends on whether or not the agent $k+1$ wishes to clinching good or not, which in turns depens on whether $\sum_{i=1}^{k} B_i \leq u_{k+1}$ or not.

Surprisingly, charging players the clearing price is not truthful, since in some cases players have incetives to understate their values to acquire smaller fractions of the good but at much lower prices. We will therefore resort to Myerson's Lemma for the calculation of the payments. We first state the complete description of the auction following the observations above.

Definition 13 (Market Clearing Price (MCP) Auction)**.** *Assume agents are sorted in decreasing value order ($v_1 > \ldots > v_n$). Let k by the maximum integer such that $\sum_{i=1}^{k} B_i \leq u_k$*

1. *if $\sum_{i=1}^{k} B_i > v_{k+1}$, then $x_i = \frac{B_i}{\sum_{j=1}^{k} B_j}$ for $i \leq k$ and $x_i = 0$ otherwise*

2. *if $\sum_{i=1}^{k} B_i \leq v_{k+1}$, then $x_i = \frac{B_i}{v_{k+1}}$ for $i \leq k$, $x_i = 0$ for $i > k+1$ and $x_{k+1} = 1 - \sum_{i=1}^{k} x_i$*

Payments are calculated through Myerson's Lemma (Theorem 12)

We now show that this auction has the desired properties

Lemma 14. *The MCP Auction rule is monotone.*

The proof is derived directly through case analysis. As a result of charging Myerson payments we get the following.

Corollary 15. *The MCP Auction is truthful in the public budget environment.*

Lemma 16. *The payments charged by the MCP Auction do not exceed the agent budgets.*

For the sake of concreteness we show an example of the MCP Auction.
Example 1. Suppose we have three agents with values $v_1 = 2$, $v_2 = 5$, $v_3 = 7$ and budgets $B_1 = 1$, $B_2 = 2$, $B_3 = 2$. In Figure 2.1 one can see the demand curve a the pricve clock ascends. At price $p = 4$, $D(4) = 1$ so we allocate according to that, agents 2 and 3 get half the good, $x_1 = 0$, $x_2 = 0.5$, $x_3 = 0.5$. Obviously agent 1 pays nothing, $p_1 = 0$. We now analyse the payments for agent 2. His allocation as a function of v_2 is as follows:

$$x_2(v_2) = \begin{cases} 0 & \text{if } v_2 \leq 2 \\ 1 - \frac{2}{v_2} & \text{if } 2 < v_2 \leq 4 \\ 0.5 & \text{if } v_2 > 4 \end{cases}$$

This allows us to calculate the payments though Myerson's lemma:

$$p_2 = 0.5 \cdot 5 - \int_2^4 (1 - \frac{2}{x})\mathrm{d}x - \int_4^5 0.5\mathrm{d}x =$$

$$= 2.5 - 2 + 2[\ln x]_2^4 - 0.5 = 2\ln 2$$

Symmetrically we also get that $p_3 = 2\ln 2$. Notice that both players pay below the clearing price per unit.

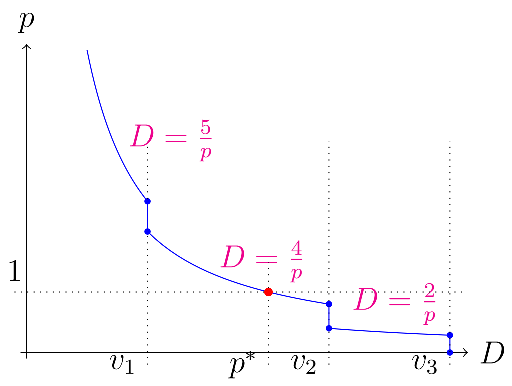

Figure 2.1: Running the MCP auction for the instance of example 1.

2.3 The Adaptive Clinching Auction

2.3.1 Description

The Clinching auction was developed in [1] and was adapted to the budgeted case in [9]. The auction was described for the multi-unit indivisible item case and although the case of divisible goods was described intuitively, the study and full description of the auction was done in [4].

The intuition behind this auction is that instead of calculating the clearing price and allocating goods at this price, it allocates the good *as the price clock rises*. That is players "clinch" different amounts of the goods at different prices. We will slowly raise the price clock and will allocate goods to an agent i whenever increasing the price will result in $\sum_{j \neq i} D_i < S$. Before we formally describe the auctio, we will show the thought process one can go through to derive the auction.

The auction will maintain and gradually update a set of variables as the price clock rises.

- the remaining supply of good $S(p)$, initialized at $S(0) = 1$
- the amount clinched by each agent i $x_i(p)$ initialized at $x_i(0) = 0$
- the amount charged to each agent i $p_i(p)$ initialized at $p_i(0) = 0$
- the remaining budget of each agent $b_i(p)$ initialized at $b_i(0) = B_i$

To ensure consistency of of the action we need to enforce an amount of invariants. We denote at $f'(p)$ for a function f of p the derivative with respect to p. Since we do not throw away any good we require that

$$S(p) = 1 - \sum_i x_i(p) \tag{2.1}$$

Moreover players do not burn their money and either keep them or give them to the auction.

$$b_i(p) = B_i - p_i(p) \tag{2.2}$$

We require that agents purchase the good at the current price per unit and therefore

$$p_i'(p) = p \cdot x_i'(p) \tag{2.3}$$

During the auction we will place agents in sets according to their role in the auction. We define the sets of active bidders at price p, $A(p)$ and the set of clinching bidders at price p, $C(p)$ as follows. We use the same definition for demand as in the previous section and denote as $D_{-i} = \sum_{j \neq i} D_i$.

$$D_i(p) = \begin{cases} \frac{b_i(p)}{p} & \text{if } p < v_i \\ 0 & \text{otherwise} \end{cases} \tag{2.4}$$

$$A(p) = \{i \mid v_i > p\} \tag{2.5}$$

$$C(p) = \{i \mid i \in A(p), D_{-i}(p) = S(p)\} \tag{2.6}$$

We will gradually raise a price clock starting from zero and we will allocate items in order to ensure $S(p) \leq D_{-i}(p)$. Notice that agents in the clinching set are the ones tightly satisfying this constraint. In therefore makes sense to allocate fractions of the good to them in order to ensure the preservation of the constraint.

$$i \in C(p) \Rightarrow S(p) = D_{-i}(p) \tag{2.7}$$

In the stages of the auction where demands are continuous and differentiable (that is if $p \notin v_1, \ldots, v_n$), deriving Equation 2.7 suffices to ensure this, however we need to take care of the function discontinuities. For a function f and price p we define $f(p^-) = \lim_{x \to p^-} f(x)$ and $f(p^+) = \lim_{x \to p^+} f(x)$. Moreover we define $a^+ = \max\{a, 0\}$.

Notice that if some player has a value of v_i and we simply drop the bidder at price v_i the discontinuity for some agent with value v_j is

$$\delta_j^i = \left[S(v_i^-) - \sum_{k \in A(v_i)} D_k(v_i^-) \right]^+ \tag{2.8}$$

We will therefore sell the excess good produced by the sudden demand drop to the remaining active players. Notice that no matter which player we choose to sell first, the supply and the demand are reduced by the same amount and therefore the rest of the agents are sold their respective quantity δ_j^i.

Before stating the auction formally, we will use Equations 2.1 - 2.7 to derive the auction rules. We mentioned that during the phases, where the price is different from agents' values, the allocation is continuous and the only agents allocating good are those in the clinching set. So for $i \in C(p)$,

$$S(p) = D_{-i}(p) \Rightarrow S'(p) = D'_{-i}(p) = \sum_{j \in A(p) \setminus \{i\}} \left(\frac{b_j(p)}{p}\right)' = \sum_{j \in A(p) \setminus \{i\}} \left(-\frac{b_j(p)}{p^2} + \frac{b'_j(p)}{p}\right)$$

Through equation 2.2 we demand that

$$b'_i(p) = -p'_i(p)$$

And therefore through equation 2.3 we require

$$\frac{b'_i(p)}{p} = -x'_i(p)$$

By combining these together with equation 2.7 we get that

$$S'(p) = -\frac{S(p)}{p} - \sum_{j \in A(p) \setminus \{i\}} x'_j(p) \tag{2.9}$$

At the same time, from equation 2.1 we get

$$S'(p) = -\sum_i x'_i(p) = -\sum_{i \in A(p)} x'_i(p) \tag{2.10}$$

We subtract equations 2.9 and 2.10 to derive

$$x'_i(p) = \frac{S(p)}{p} \tag{2.11}$$

Additionally from equation 2.3 we get that

$$p'_i(p) = S(p) \tag{2.12}$$

This thought process leads to defining the auction formally as

Definition 17 (Adaptive Clinching Auction)**.** *The auction is defined by the set of differential equations presented below*

- *if* $p \notin \{v_1, \ldots, v_n\}$

$$x'_i(p) = \begin{cases} \frac{S(p)}{p} & \text{if } i \in C(p) \\ 0 & \text{if } i \notin C(p) \end{cases} \tag{2.13}$$

$$p'_i(p) = \begin{cases} S(p) & \text{if } i \in C(p) \\ 0 & \text{if } i \notin C(p) \end{cases} \tag{2.14}$$

- *if $p = v_j$ for some j*

$$x_i(p^+) = x_i(p) = \begin{cases} x_i(p^-) + \delta_i^j & \text{if } i \in A(p) \\ x_i(p^-) & \text{if } i \notin A(p) \end{cases} \tag{2.15}$$

$$p_i(p^+) = p_i(p) = \begin{cases} p_i(p^-) + \delta_i^j & \text{if } i \in A(p) \\ p_i(p^-) & \text{if } i \notin A(p) \end{cases} \tag{2.16}$$

The above equations fully describes the allocation and payment rule and hence the auction as a whole. The final allocation described by the limits $\lim_{p\to\infty} x_i(p)$ and $\lim_{p\to\infty} p_i(p)$. It should be clear by the auction design that the requirements of equations 2.1 - 2.7 are satisfied at all stages of the auction.

2.3.2 Properties of the auction

We will not give technical proofs for the following properties when the results are intuitive and easily understood by simple combinatorial arguments. For detailed proofs and analysis the reader can consult [4, 15, 9, 10].

Lemma 18. *The adaptive clinching auction is truthful for the public budget case.*

Proof. The equations deriving the auctions, depend mainly on the budget of the agents which is public. The only point values come into play is the price at which an agent drops out of the auction, in the sense that he decreases his demand to zero.

The agent is allocating good at price p per unit at various points of the auction. The derivative of the surplus is positive when

$$v_i \cdot x_i'(p) - p_i'(p) \geq 0 \Leftrightarrow (v_i - p) \cdot x_i'(p) \geq 0 \Leftrightarrow v_i \geq p$$

The surplus of the agent is therefore (weakly) increasing while $p \leq v_i$ and (weakly) decreasing while $p > v_i$. As a result the optimal point to quit the auction is exactly v_i which is the default behavior of the auction. Dropping out earlier by reporting a lower value can only decrease surplus, by missing out on good at a low price, and dropping out later can only decrease surplus by buying good at a high price. □

Lemma 19. *The adaptive clinching auction respects budgets and never gives out money*

Proof. Obvious by the definition of the auction. □

Notice that by equation 2.7 and the fact that the demand is always below the supply, only agents with maximal $b_i(p)$ can belong in the clinching set

Proposition 20. *We denote as $b_{max}(p) = \max\{b_i(p) \mid i \in A(p)\}$. Then*

$$C(p) \neq \emptyset \Rightarrow C(p) = \{i \in A(p) \mid b_i(p) = b_{max}(p)\}$$

The auction has a simple structure with respect to budgets. Initially the price rises until the clinching set becomes non-empty, and therefore contains the set of agents that are active and have maximal budgets. As the clinching goes on these agents' (common) budget get reduces, until it reaches the budget of another agent. Then this agent gets in the clinching set and the process goes on.

At the points when an active agent drops out of the active set due to the price getting above his value, agents in the clinching set (or about to get in it) receive some one-shot allocations. When an agent drops out of the clinching set, the rest of the demand matches the supply and all agents in $C(p)$ are allocated their demands at that price. By the above description the following proposition is easily derived.

Proposition 21. *If the budget of the highest bidder is above the second highest value, the whole good is allocated to the highest bidder. Else there is some bidder k such that the final price of the auction is $p_f = u_k$. Then:*

- $\forall i, x_i > 0 \Rightarrow v_i \geq v_k$
- $\forall i, v_i \geq v_k \Rightarrow p_i = B_i$

2.4 Measuring auction performance

So far, we have mentioned nothing about the performance of either the MCP auction or the clinching auction. The fact that they are truthful and budget respecting does not mean much as is, since the auction where we simply give the good to a random player and charge nothing also has these properties.

We would like to prove some performance guarantee with respect to the optimal social welfare. However it is easy to see that when $B_i \to 0$ we cannot do anything better that give the item randomly to some bidder, since we cannot the reported values at all. The only hope would be to prove that the auction performs as well as any truthful auction respecting the budgets, but such a task seems quite daunting. Instead we will study alternative notions of optimality and compare to them.

2.4.1 Pareto Optimality

We will first study the notion of Pareto Efficience, a concept emerging from the economic literature on the subject.

Definition 22. *An outcome (o, p) is Pareto Optimal or Pareto Efficient when there exists no other outcome o' and payment vector p' such that $\forall i$*

$$u_i(o) - p_i \geq u_i(o') - p'_i$$

$$\sum p'_i \geq \sum p_i$$

with at least one equality strict.

Notice that in the setting without budgets the only pareto optimal outcome is the optimal, or else the high bidder could "buy" good from another agent at a high enough rate such that both agents are satisfied, thus rendering this outcome not Pareto Optimal. In the budgeted setting the same intuition leads to a similar characterization .

Lemma 23. *An outcome is Pareto Optimal in the setting with budget constraints if* $\forall i, j$

$$x_i > 0, v_j > v_i \Rightarrow p_j = B_j$$

This lemma exactly captures the notion that in order for an outcome to be Pareto Optimal for any agent that receives share of the good, all agents with higher values have exhausted their budget, and therefore no trade amongst the agents after the auction or desirable for them.

From Proposition 21 it is easy to see that the Clinching Auction produces Pareto Optimal outcomes. In [9], the following is proved.

Lemma 24. *The Clinching Auction is the unique truthful auction for the setting of public budgets that produces Pareto Optimal outcomes. For the setting of private budgets there is no such auction.*

From this results it is apperent that the MCP auction does not satisfy Pareto Optimality and the reason lies in the fact that the prices are computed through Myerson's Lemma and are therefore below the budgets for many agents.

2.4.2 Liquid Welfare

Apart from Pareto optimality, the objective of Liquid Welfare maximization was first studied in this context in [10].

Definition 25. *For an outcome o the liquid welfare is defined by*

$$\bar{W}(o) = \sum_i \min\{u_i(o), B_i\}$$

and we denote as the optimal liquid welfare for an instance as $\bar{W}^*$*.*

The intuition behind this measure is that the quantity $\bar{W}^*$ is the maximum revenue and seller with infinite information can extract from this market. This is exactly the optimal welfare in the unbudgeted case. Following this remark we can easily calculate the LW optimizing allocation. In order to extract as much revenue as possible with full information we sell the item to the agent with the highest value with price per unit his value, and sell exactly the quantity the exhausts his budget. We continue until we exhaust the whole good.

Lemma 26. *If the players are sorted in non-decreasing value order $v_1 \geq \ldots \geq v_n$ the allocation maximizing LW is*

$$x_i = \min\{\frac{B_i}{v_i}, (1 - \sum_{j<i} x_j)^+\}$$

Lemma 27. *The optimal LW allocation can not be implemented truthfully.*

This is apparent from the fact that we wish to reduce the agent allocation as he increases his value, by raising his price. However both of the presented auction closely approximate the optimal LW achievable.

Proposition 28. *The MCP auction produces and outcome o such that*

$$LW(o) \geq \frac{1}{2}\bar{W}^*$$

and is therefore a 2-approximation to the optimal LW.

Proposition 29. *The Clinching auction produces and outcome o such that*

$$LW(o) \geq \frac{1}{2}\bar{W}^*$$

and is therefore a 2-approximation to the optimal LW.

Chapter 3

Smoothening VCG payments

In the previous chapter we defined and discussed the properties of the VCG mechanism, the mechanism choosing the best outcome for the agents and charges each one his externalities. We concluded the chapter noticing that this mechanism can charge the set of agents considerably high prices. In this chapter, we will present a novel (to the best of our knowledge) way to reduce the total amount of payments by sacrifing a certain amount of social welfare. To achieve this we consider a broader class of mechanisms, by allowing the mechanism to use randomization.

3.1 Randomized Mechanisms

In Chapter 1 we defined the allocation rule of the mechanism as a function $f : V \to O$ outputting an outcome. Now we will consider function $f : V \to \Delta(O)$, where $\Delta(O)$ is the simplex of probability distribution over outcomes, that is we allow the mechanism to produce a probability distribution over outcomes and then we sample this distribution to select one of them. In other words, the mechanism assigns a probability f_i to each outcome o such that $\sum_i f_i(u) = 1$. and then we select outcome o with probability $f_i(u)$. In addition to the allocation, the mechanism outputs a price vector p, as in the deterministic case, which may or may not depend on the realization of the random flips. In the case where payments depend on the outcome, the payment of each agent is also a random variable.

The addition of randomization introduces uncetrainty about the payoff of each agent, since the utility from the outcome of the mechanism is a random variable. Different agents may have different values when faced with the same random variable based on what we call *risk attitude*. To make this notion concrete consider the scenario where we run a lottery with a prize of 1000 \$, each ticket has a 1% chance of winning and it costs c dollars. The expected value for participating in the lottery is obviously $1000 \cdot 0.01 = 10\$$. What should the price c be in order for the agent to consider lottery a good opportunity? The answer is, it depends.

Conditioning on the risk attitude of each agent we will examine the cases separately:

- The player is *risk averse.* He prefers lotteries with lower prizes but higher chances of success, that is he recieves negative utility for the random variable's standard deviation. This player prefers to keep the 10$ than give them for a ticket and would only buy with a lower value, for example 8$.
- The player is *risk neutral.* He cares only about his expected utility and as a result would buy a ticket if $c \leq 10\$$.
- The player is *risk seeking.* He is willing to participate in lotteries with a chance of high reward, even if his expected surplus is negative. Therefore he is willing to buy a ticket even when the price is above 10$ (for example 20$) hoping to win the grand prize.

For the rest of this work we will concentrate exclusively on risk neutral agents. Besides mathematical convenience, one of the main reasons for this, is the complexity of accuretely modeling the utilities of risk averse agents, and therefore the robustness of the mechanism when faced with different utilities. Moreover, risk neutrality is a reasonable assumption when the amount of money considered is modest for the agents. We will therefore modify the definition of truthfulness to take into account randomization, assuming agents are risk neutral.

Definition 30. *A randomized mechanism $M = (f, p)$ is* truthful in expectation *if* $\forall i \in N,\ u_i,\ u_i' \in V_i,\ u_{-i} \in V_{-i}$*:*

$$\mathbf{E}_{o \sim f(u)}[u_i(o) - p_i(u, o)] \geq \mathbf{E}_{o \sim f(u_{-i}, u_i')}[u_i(o) - p_i((u_{-i}, u_i), o)]$$

A stronger definition of truthful is the notion of universally truthful mechanisms

Definition 31. *A randomized mechanism is* universally truthful *if it is randomization over deterministic truthful mechanisms.*

This means that a mechanism is truthful even if the random bits are known to the agent beforehand. Notice that universal truthfulness is independent of the agents' risk attitudes and may seem more attractive in that sense. However, some easy manipulations of the payment scheme (assuming the center is risk neutral) can alleviate this and truthful in expectation mechanism can be transformed to remain truthful even in the presence of risk averse agents with only mild assumptions. For a more detailed discussion we refer the reader to [12].

Moreover, in a similar notion we modify the definition of individual rationality to match the randomized fashion of the mechanism.

Definition 32. *A randomized mechanism is* individually rational in expectation *if* $\forall i \in N,\ u \in V$*:*

$$\mathbf{E}_{o \sim f(u)}[u_i(o) - p_i(u, o)] \geq 0$$

Definition 33. *A randomized mechanism is* ex post individually rational *if* $\forall i \in N,\ u \in V,\ o \in O$*:*

$$u_i(o) - p_i(u, o) \geq 0$$

Notice that when a mechanism is IR in expectation, some "unlucky" coin flips can cause agent to receive negative utility.

3.2 The setting

We will now restrict the setting we are considering by adding some requirements that are quite reasonable in a variety of natural environments. While defining our setting we have in mind mainly public projects and auctions so we make the assumptions that fit to these environments. Suppose that there are m alternatives to choose from and we restrict agents to having positive valuations over them. Therefore the valuation of each agent is a vector $y \in \mathbb{R}_+^m$ such that y_i represent his utility for outcome i. Moreover we treat the utilities of each agent in a particularly fair way, by considering only strongly anonymous mechanisms.

Definition 34. *A (randomized) mechanism with allocation rule* $g : (\mathbb{R}_+^m)^n \to \Delta(O)$ *is* strongly anonymous *if there exists a function* $f : \mathbb{R}_+^m \to \Delta(O)$ *such that:*

$$g(x_1, \ldots, x_n) = f(\sum_i x_i)$$

This definition captures the notion, that we don't care from which agent the utility for each outcome comes.

It follows that the expected utility of an agents is equal to $x_i \cdot f(\sum_i x_i)$. Since the utilities of the agents act as a unit of measurement we will require our mechanisms to be scale invariant.

Definition 35. *A mechanism with allocation rule f is* scale invariant *if* $\forall x \in \mathbb{R}_+^m$

$$f(ax) = f(x),\ \forall a \in (0, +\infty) \tag{3.1}$$

Finally we demand that the mechanisms are truthful in expectation, ex post individually rational and that the payments are always positive (i.e. the mechanism never pays the agents).

In this setting the definitions of truthfulness will be slightly modified and are as follows:

Definition 36. *A strongly anonymous mechanism is* truthful *if* $\forall x, y, y' \in \mathbb{R}_+^m$*:*

$$y \cdot f(x + y) - p(x, y) \geq y \cdot f(x + y') - p(x, y') \tag{3.2}$$

Definition 37. *A strongly anonymous mechanism is* individually rational in expectation *if* $\forall x, y \in \mathbb{R}_+^m$*:*

$$y \cdot f(x + y) - p(x, y) \geq 0 \tag{3.3}$$

3.3 The space of impementable mechanisms

We will now characterize the mechanisms that are implementable in this setting and the possible payment schemes they admit.

3.3.1 Implementable allocation rules

Proposition 38. *Every strongly anonymous and scale invariant mechanism is implementable iff there is a set $S \subset \mathbb{R}_+^m$ such that*

$$f(x) \in argmax_{s \in S} x \cdot s \tag{3.4}$$

Proof. We require that mechanisms are truthful in expectation and therefore

$$y \cdot f(x+y) - p(x,y) \geq y \cdot f(x+y') - p(x,y'),\ \forall x, y, y' \in \mathbb{R}_+^m \tag{3.5}$$

symmetrically we get that

$$y' \cdot f(x+y') - p(x,y') \geq y' \cdot f(x+y) - p(x,y),\ \forall x, y, y' \in \mathbb{R}_+^m$$

and by summing up the inequalities we get that all impementable mechanisms satisfy *weak monotonicity* (WMON).

$$(y' - y) \cdot (f(x+y') - f(x+y)) \geq 0 \tag{3.6}$$

Setting $y' = 0$, $x = ax$ and $y = (1-a)y$ for some $a \in (0,1)$ we get

$$(1-a)y \cdot f(ax + (1-a)y) \geq (1-a)y \cdot f(ax)$$

Since f is scale invariant and $a \neq 1$,

$$y \cdot f(ax + (1-a)y) \geq y \cdot f(x)$$

Taking the limit $a \to 1$,

$$y \cdot f(y) \geq y \cdot f(x)$$

Therefore if we define $S = \{s \mid \exists x f(x) = s\}$ we get that any truthfully implementable function satisfies $f(x) \in \text{argmax}_{s \in S} x \cdot s$.

In the opposite direction any allocation rule satisfying (3.4) is implementable if we define payments as $p(x,y) = -x \cdot f(x+y)$.

$$y \cdot f(x+y) - p(x,y) = y \cdot f(x+y) + x \cdot f(x+y) =$$

$$= (x+y) \cdot f(x+y) \geq^{(3.4)} (x+y) \cdot f(x+y') = y \cdot f(x+y') - p(x,y')$$

□

3.3.2 Possible Payment Schemes

We showed that any function that is a social welfare maximizer can be implemented with payments $p(x, y) = -x \cdot f(x+y)$. We will now modify these payments to achieve the required payment properties. This way we end up with the VCG mechanism.

Proposition 39. *For every strongly anonymous and scale invariart implementable allocation rule, the unique payments scheme satisfying individual rationality, truthfulness and positive payments is*

$$p(x, y) = x \cdot f(x) - x \cdot f(x + y) \tag{3.7}$$

Proof. We fix a single agent and the utilities of other agents, and let x be the summation of the utilities of other agents. We first prove two useful lemmas.

Lemma 40. *The payment rule imposing thuthfulness can only depend on the output of the mechanism, that is $p(y) = g(f(x + y))$.*

Proof of Lemma. Suppose there exist $y_1, y_2 \in \mathbb{R}^m_+$ such that $f(x+y_1) = f(x+y_2) = s$ and $p(y_1) \neq p(y_2)$ (assume $p(y_1) > p(y_2)$ w.l.o.g.). It follows that

$$y_2 \cdot f(x + y_2) - p(y_2) = y_2 \cdot f(x + y_1) - p(y_2) > y_2 \cdot f(x + y_1) - p(y_1)$$

violating truthfulnes. Therefore $p(y_1) = p(y_2)$ for every y_1, y_2 such that $f(x+y_1) = f(x + y_2)$. □

Lemma 41. *For any payments rules p, p' that implement f, there is a function h such that $p'(x, y) = p(x, y) + h(x)$.*

Proof of Lemma. We denote p^a the payments of the fixed agent when the outcome is a. There exists a function h such that $p'(x, y) = p(x, y) + h(x, y)$, since h is an arbitrary function and can therefore cancel out the term $p(x + y)$.

It suffices to prove that for any $y_1, y_2 \in \mathbb{R}^m_+$ such that $f(x + y_1) = a$ and $f(x + y_2) = b$, $h(x, y_1) = h(x, y_2) \Rightarrow p'^b - p'^a = p^b - p^a$.

We consider outcomes a, b as *close*, if for any ϵ there utilities y_a, y_b such that $f(x + y_a) = a$, $f(x + y_b) = b$ and $\forall c \in \mathbb{R}^m_+,\ |y_a \cdot c - y_b \cdot c| \leq \epsilon$.

Let a, b be arbitrary, close outcomes and y_a, y_b as in the definition of closeness for a fixed ϵ. By the definition of truthfulness we get

$$y_a \cdot b - y_a \cdot a \leq p^b - p^a \leq y_b \cdot b - y_b \cdot a$$

$$y_a \cdot b - y_a \cdot a \leq p'^b - p'^a \leq y_b \cdot b - y_b \cdot a$$

By the definition of closeness we get that, $|y_a \cdot a - y_b \cdot a| \leq \epsilon$ and $|y_a \cdot b - y_b \cdot b| \leq \epsilon$, therefore $|y_b \cdot b - y_b \cdot a + y_a \cdot a - y_a \cdot b| \leq 2\epsilon$. Since the bounds of $p^b - p^a$ and $p'^b - p'^a$ coincide and this holds for any ϵ we conclude that for close outcomes $p^b - p^a = p'^b - p'^a$.

Since the outcomes of the mechanism are probility distributions we get that for every outcome c, $|c| \leq 1$. Let an arbitrary valuation y_a such that $f(x+y_a) = a$ and any $\epsilon > 0$. Then for any valuation y_b s.t. $f(x+y_b) = b$ and $|y_a - y_b| \leq \epsilon$, a and b are close since $|y_a \cdot c - y_b \cdot c| \leq |y_a - y_b| \leq \epsilon$ for any outcome c.

This implies that since the valuation space is connected, for any valuations y_a, y_b with $f(x+y_a) = a$, $f(x+y_b)$, there is a sequence of valuations $y_a = y_1, \ldots, y_k = y_b$ such that $|y_i - y_{i+1}| \leq \epsilon$ for any $1 \leq i \leq k-1$Therefore. that outcomes $f(x+y_i)$, $f(x+y_{i+1})$ are close and as a result $p^{y_{i+1}} - p^{y_i} = p'^{y_{i+1}} - p'^{y_i}$. By adding the k equalities we get that $p'^b - p'^a = p^b - p^a$. □

We have showed that the payment scheme $p(x,y) = -x \cdot f(x+y)$ renders the mechanism truthful, and by the above lemma we know that the only possible payment scheme is $p(x,y) = h(x) - x \cdot f(x+y)$ for some h. By individual rationality and the positive payments requirement we get $0 \leq p(x,y) \leq y \cdot f(x+y)$. Since this holds for any y, setting $y = 0$ implies $p(x,0) = 0 \Rightarrow h(x) = x \cdot f(x)$.

It is easy to verify that for any mechanism that satisfies (3.4) this payment scheme gives positive payments. Individual rationality in expectation holds since

$$y \cdot f(x+y) - p(x,y) \overset{(3.2)}{\geq} y \cdot f(x) - p(x,0) =$$
$$= y \cdot f(x) - x \cdot f(x) + x \cdot f(x) = y \cdot f(x) \geq 0$$

□

3.3.3 Derandomizing rationality

We have shown that the mechanism is IR in expectation, however there are examples where players net negative utility for certain random outcomes. However the mechanism can be modified to be universally IR. Although we state the following theorem in terms of this setting, mainly for mathematical convenience, it does not depend on the scale invariance and strong anonymity assumption, and therefore holds in a variety of settings.

Theorem 42. *Any randomized mechanism with finite outcomes that is individually rational in expectation, can have its payment scheme modified so that it is universally invididually rational.*

Proof. Fix an agent with utility vector y. Let $P = x \cdot f(x) - x \cdot f(x+y)$ denote the expected payments that induce truthfulness. We need to compute a vector $p \in \mathbb{R}^m_+$, where p_i denotes the payment for realized outcome i. In order to achieve universal IR we need to satisfy the constrains:

$$y_i \geq p_i \tag{3.8}$$

Furthermore we require that in expectation the payment scheme remains the same , that is:

$$p \cdot f(x+y) = P \tag{3.9}$$

It now suffices to prove the constraints (3.8) and (3.9) are always satisfiable. Constraints (3.8) define an m-dimensional hypercube and constrain (3.9) defines an m-dimensional hyperplane. The cases $P = 0$ and $P = y \cdot f(x + y)$ are trivial since we can use $p = 0$ and $p = y$ respectively. We can therefore assume w.l.o.g. that $0 < P < y \cdot f(x + y)$, since payments are positive and the mechanism IR in expectation.

This implies that the points $(0, \dots 0)$ and y lie on different sides of the plane defined in Equation 3.9. Notice that these points are vertices of the cube defined in Equation 3.8. Therefore the hyperplane intersects the hypercube and the feasible region is non empty. □

A simple algorithm to achieve this, is traversing the edges of the hypercude until the plane is encountered. By using a variable P initialized to the expected payments $x \cdot f(x) - x \cdot f(x + y)$. Consider each outcome i sequentially and if $\frac{P}{f_i(x+y)} \leq y_i$ then set the price of outcome i to $\frac{P}{f_i(x+y)}$ and the rest of the prices to 0, else set the price to y_i and update P to $P - y_i \cdot f_i(x + y)$. The algorithm exhausts all of P before the outcomes are exhausted since $y \cdot f(x + y) \geq P$.

Alternatively one can define a "discount rate" as follows for each player,

$$d = \frac{x \cdot f(x) - x \cdot f(x + y)}{y \cdot f(x + y)}$$

This can be computed by only knowing the expected payments and expected utility of the agent. Then the payments can be computed simply as

$$p_i = dy_i$$

for the specific realized outcome. By IR in expectation, follows that $d \in [0, 1]$ and therefore the new payments scheme is IR universally. It is trivial to verify that the expected payments remaining unaltered.

$$\mathbf{E}[p] = p \cdot f(x + y) = \left(\frac{x \cdot f(x) - x \cdot f(x + y)}{y \cdot f(x + y)} y \right) \cdot f(x + y) =$$
$$= \frac{x \cdot f(x) - x \cdot f(x + y)}{y \cdot f(x + y)} (y \cdot f(x + y)) = P$$

3.3.4 Payments are unique assuming nothing but rationality

Lets assume that the only requirent for the payment scheme, is $p(x, 0) = 0$. We will augment each coordinate of the y vector to transition between the bid 0 and y. We abuse notation and define v_{-i} for a vector v to be the same vector with the i-th coordinate zero.

Consider some implementable allocation f with the corresponding payment scheme p, and fix a vector x of other bidders. For arbitrary a, i, we define $h(a) =$

$p(x, y_{-i}+ay_i)$, $g(a) = f(x+y_{-i}+ay_i)$ and $y(a) = y_{-i}+ay_i$. Then by the definition of truthfulness:

$$y(a) \cdot g(a) - h(a) \geq y(a) \cdot g(a+e) - h(a+e)$$

$$y(a+e) \cdot g(a+e) - h(a+e) \geq y(a+e) \cdot g(a) - h(a)$$

By combining the inequalities we get:

$$y(a) \cdot (g(a+e) - g(a)) \leq h(a+e) - h(a) \leq y(a+e) \cdot (g(a+e) - g(a))$$

We divide by e and take the limit as e approaches 0. The two sides sandwich together and we get an equality. We denote by h' and g' the derivatives of the functions with respect to a.

$$h'(a) = y(a) \cdot g'(a)$$

Note that even if the functions are not differentiable, since they have a bounded range and we immediately integrate over a bounded domain, we are still valid. We integrate for a from 0 to 1 (integration by parts holds for vector valued functions and dot product).

$$\int_0^1 h'(a)da = \int_0^1 y(a) \cdot g'(a)da \Leftrightarrow$$

$$\Leftrightarrow h(1) - h(0) = y(1) \cdot g(1) - y(0) \cdot g(0) - \int_0^1 y'(a) \cdot g(a)da \Leftrightarrow$$

$$\Leftrightarrow p(x, y) - p(x, y_{-i}) = y \cdot f(x+y) - y_{-i} \cdot f(x+y_{-i}) - \int_0^{y_i} f_i(x+y_{-i}+z)dz$$

It is easy to see that writing down the equality for each i from 1 to m and summing them up, causes most of the terms to telescope, leading to the following relation

$$p(x, y) = y \cdot f(x+y) - \sum_{i=1}^{m} \int_0^{y_i} f_i(x+y_{-1\ldots i}+z)dz$$

This concludes the proof that any implementable rule in this setting has uniquely defined prices assuming normalization.

3.4 Achieving a payment-approximation trade-off

We have shown that payments are uniquely determined by the allocation rules. The amount each agent pays is exactly the amount he modifies the utilities of others by participating. Therefore optimizing SW as an allocation rule leaves no option to reduce the total payments, which we have already shown to be unpleasingly large. We will therefore try to create mechanisms that do not change their allocation too much when players participate.

3.4.1 The intuition

The characterization of Proposition 38 reduces the design of a truthful, strongly anonymous, scale invariant mechanism to the choice of a set of feasible solution to optimize on. Suppose we choose as solution space the simplex of probabilities over outcomes. Then the optimization problem for a given vector x becomes

$$\begin{aligned} \text{maximize} \quad & f \cdot x = \sum_i f_i x_i \\ \text{s.t.} \quad & |f| = \sum_i f_i = 1 \\ & f_i \geq 0 \end{aligned}$$

The optimization space is the hyperplane $\sum f_i = 1$ and the solution is the extremal point in the direction of x. Therefore assuming that the coordinates of x are pairwise distinct, the solution is $f_k = 1$ where $k = \text{argmax}_i x_i$ and $f_i = 0$ for $i \neq k$. This shows that the optimal solution is choosing the outcome with higher preference deterministically, which is exactly the classic VCG mechanism. In Figure 3.1 this result is obvious and moreover one can notice that the reason payments are large is the fact that a infinitesimal change in the direction of x can cause a different outcome to be chosen.

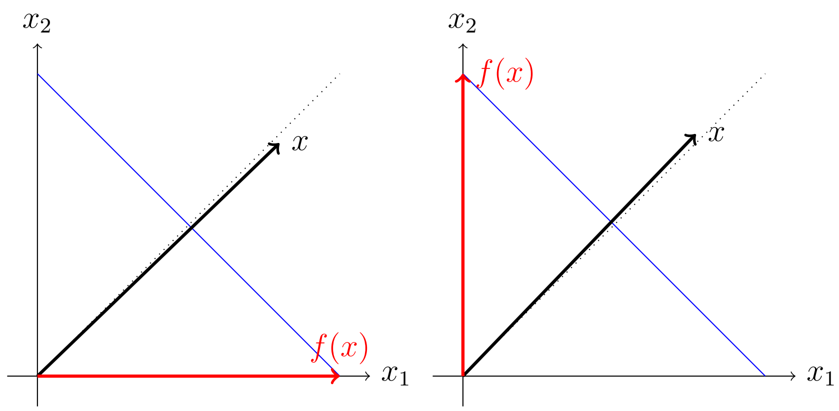

Figure 3.1: Optimizing on the plane, for $|O| = 2$

This issue is caused by the fact the pure outcomes are strictly better than mixed in this setting. We will therefore allow our mechanism to output probabilities f_i that do not sum up to 1. This assumption is reasonable in the sense that we can add one extra outcome, the *null outcome* that has zero utility for every agent and receives the probability unallocated in f. The existence of such an outcome is obvious in the case of auctions, where we simple sell nothing, and in public project and facility location problems, where we simply construct nothing. A similar approach, mostly in the setting without money is taken in [8] and [17].

By allowing the mechanism to output incomplete outcomes we can mix pure outcomes with the null outcome lowering their value in order to make the optimization produce mixed outcomes.

3.4.2 The mechanism

We define a family of mechanisms f_k which respectively maximize on the surfaces

$$S_k = \left\{ s \in \mathbb{R}_+^m \;\middle|\; ||s||_k \leq \frac{1}{m^{1-1/k}} \right\} \quad \text{where } ||\cdot||_k \text{ is the } L_k \text{ norm} \tag{3.10}$$

In Figure 3.2 we notice that the optimization on these surfaces produces outcomes the tend to be a more "equal mixture" of pure outcomes. The mechanism exhibits some kind of bias towards outcomes that distribute probabilities equally and therefore changes in x have a smoother change in $f(x)$, as compared to Figure 3.1

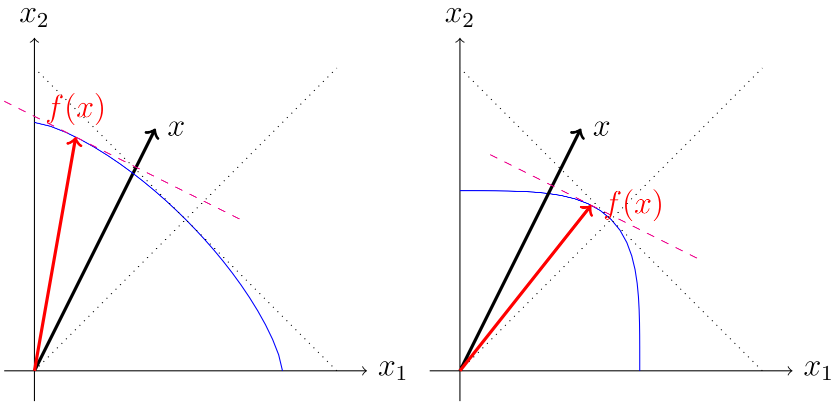

Figure 3.2: Optimizing on the smooth surfaces
($k = 1.4$ and $k = 4$ respectively)

The extreme case is observed when k tends to infinity. The surface we optimize then is a cube and the optimal solution is always the corner, that is the outcome $(\frac{1}{m}, \ldots, \frac{1}{m})$, as shown in Figure 3.3. In this case it is obvious that the value of each player is of no importance to the mechanism and as a result the payments are zero.

Notice that when $k = 1$ the resulting surface is the plane and we have the VCG mechanism which *exactly* optimizes the SW, but charges a large sum of payments. On the other extreme, when $k \to +\infty$, we only achieve an expected welfare that is m times below the optimal but charges no payments. The transition between the two extremes is continuous thus achieving a tradeoff between SW and payments.

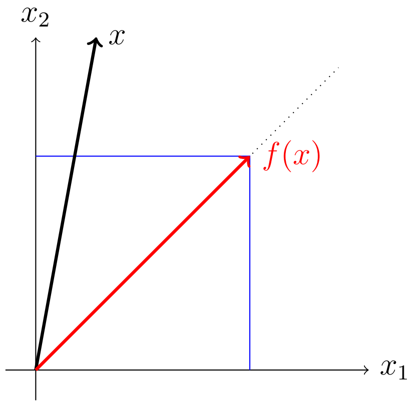

Figure 3.3: Optimizing on the square
The optimal solution is independent of x

We will know prove certain properties of our mechanism and for this we will need the mechanism in closed form. We denote as x^d for any vector $x \in \mathbb{R}^n_+$ the vector whose i-th coordinate is x_i^d, and as $|x|$ the sum of the vector's coordinates.

Proposition 43. *The closed form of the mechanism f_k is*

$$f_k(x) = \frac{1}{m^{1-1/k}} \frac{x^{\frac{1}{k-1}}}{||x^{\frac{1}{k-1}}||_k}$$

Proof. The outcome of the mechanism is the vector s the optimizes $x \cdot s$ subject to $||s||_k \leq m^{-\frac{k-1}{k}}$. The surface of Equation 3.10 defines a strictly convex space (convexity follows immediately from the triangle inequality for normed spaces, and strictness from the the fact that triangle inequality is strict for colinear vectors only).

Therefore the optimal point will be on the boundary of the space, at the extremal point in the direction of x. The boundary is defined by

$$||s||_k = \frac{1}{m^{1-\frac{1}{k}}} \Leftrightarrow ||s||_k^k = \frac{1}{m^{k-1}}$$

and since we seek the extremal point in the direction of x, x must be perpendicular to the boundary at the optimal point (else we would be able to follow the boundary in a direction the would increase the objective). Therefore at the optimal point s_* the gradient of the surface is in the direction of x, that is there is some c such the

$$\nabla(||s_*||_k^k) = cx \Leftrightarrow k{s_*}^{k-1} = cx \Leftrightarrow s_* = \left(\frac{cx}{k}\right)^{1/(k-1)} = \left(\frac{c}{k}\right)^{\frac{1}{k-1}} x^{\frac{1}{k-1}}$$

Moreover s_* need to be to be on the surface, and therefore

$$||s_*||_k^k = \frac{1}{m^{k-1}} \Leftrightarrow \left| \left(\frac{cx}{k}\right)^{\frac{k}{k-1}} \right| = \frac{1}{m^{k-1}} \Leftrightarrow \left(\frac{c}{k}\right)^{\frac{1}{k-1}} = \frac{1}{m^{\frac{k-1}{k}} ||x||_{\frac{k}{k-1}}^{\frac{1}{k-1}}}$$

Replacing in the above equation we get

$$s_* = \frac{x^{\frac{1}{k-1}}}{m^{\frac{k-1}{k}} ||x||_{\frac{k}{k-1}}^{\frac{1}{k-1}}} = \frac{1}{m^{1-1/k}} \frac{x^{\frac{1}{k-1}}}{||x^{\frac{1}{k-1}}||_k}$$

□

We are interested in mechanisms with S_k close to S_1, so we set $k = l/(l-1)$ for some $l \in \mathbb{R}$. The resulting mechanism is

$$f(x) = \frac{1}{m^{1/l}} \frac{x^{l-1}}{||x^{l-1}||_{\frac{l}{l-1}}} \tag{3.11}$$

We can verify the by setting $l = 1$ we get the uniform mechanism and by setting $l \to +\infty$ we get the Vickrey auction.

3.4.3 Social Welfare Approximation Ratio

We need a measure to evaluate how good our allocation rule is. We will compare the SW produced by our mechanism to the SW produced by the optimal allocation.

Definition 44. *A mechanism $M = (f, p)$ has an* approximation ratio *of a or is an a-approximation to social welfare for some $a > 1$ if $\forall x \in \mathbb{R}_+^m$*

$$x \cdot f(x) \geq \frac{1}{a} OPT(x)$$

where $OPT(x)$ is the SW of the best allocation rule.

For example, when a mechanism 2-approximates SW, the mechanism achieves at least half of the welfare achieved by the best (implementable or not) allocation rule. The VCG mechanism we analyzed in Chapter 1 is a 1-approximation while the uniform mechanism that randomly chooses an outcome is a m-approximation (consider the case where all agent unanimously prefer one specific outcome).

Lemma 45. *For any vector a holds that*

$$\frac{||a^l||_1}{||a^{l-1}||_{\frac{l}{l-1}}} = ||a||_l$$

Proof. We denote as a_i the i-th coordinate of the vector a

$$\frac{||a^l||_1}{||a^{l-1}||_{\frac{l}{l-1}}} = \frac{\sum a_i^l}{(\sum a_i^l)^{\frac{l-1}{l}}} = (\sum a_i^l)^{1/l} = ||a||_l$$

□

Proposition 46. *The mechanism of equation (3.11) has an approximation ratio of at most $m^{1/l}$.*

Proof. The optimal mechanism outputs with probability 1 the maximum coordinate of x, and as a result the optimal SW is $OPT = ||x||_\infty$.

$$\frac{||x||_\infty}{x \cdot f(x)} = m^{1/l} \frac{||x^{l-1}||_{l/(l-1)} ||x||_\infty}{x \cdot x^{l-1}} \overset{\text{Lemma 45}}{=} m^{1/l} \frac{||x||_\infty}{||x||_l} \leq m^{1/l}$$

□

Since $m^{1/l}$ is a decreasing function of l, as l ranges from 1 to $+\infty$, the approximation ratio ranges from m to 1.

3.4.4 Payment Bounds

We will now study the amount of payments charged by the mechanism. The payments of player i are computed as follows

$$\begin{aligned} p(x_{-i}, x_i) &= x_{-i} \cdot f(x_{-i}) - x_{-i} \cdot f(x) \\ &= x_{-i} \cdot f(x_{-i}) - x \cdot f(x) + x_i \cdot f(x) \\ &= \frac{1}{m^{1/l}} \left(||x_{-i}||_l - ||x||_l + \frac{x_i \cdot x}{||x||_l^{l-1}} \right) \end{aligned} \tag{3.12}$$

Therefore we can now bind the total amount of payments

Proposition 47. *The mechanism of Equation 3.11 charges the set of agents at most $\frac{|x|}{m^{1/l}}$.*

Proof. With use of Lemma 45 and basic properties of normed spaces we get

$$\begin{aligned}
\sum_{i=1}^{n} p(x_{-i}, x_i) &= \frac{1}{m^{1/l}} \sum_i \left(||x_{-i}||_l - ||x||_l + \frac{x_i \cdot x}{||x||_l^{l-1}} \right) \\
&= \frac{1}{m^{1/l}} \left(\sum_i ||x_{-i}||_l - n||x||_l + \frac{(\sum_i x_i) \cdot x}{||x||_l^{l-1}} \right) \\
&= \frac{1}{m^{1/l}} \left(\sum_i ||x_{-i}||_l - n||x||_l + ||x||_l \right) \\
&= \frac{1}{m^{1/l}} \left(||x||_l - \sum_i (||x||_l - ||x_{-i}||_l) \right) \qquad (3.13) \\
&\leq \frac{||x||_l}{m^{1/l}} \leq \frac{|x|}{m^{1/l}} \qquad (3.14)
\end{aligned}$$

□

The inequalities used in the last equation are not tight, so payments are strictly below the presented bound. However this bound expresses clearly the relation between payments and approximation guarantees. To sum up, we have constracted a mechanism that achieves at least a $m^{1/l}$ fraction of the SW and charges at most $\frac{|x|}{m^{1/l}}$. If we define $a = m^{1/l}$ we denote this family of mechanisms with approximation ratio a and payments bound $|x|/a$ as

$$\left[a, \frac{|x|}{a} \right]$$

Theorem 48. *For every $a \in [1, m]$ there exists a mechanism that achieves social welfare at least $1/a$ times the optimal and charges at most $|x|/a$.*

3.5 Towards Budget Feasibility

Suppose now there is a budget B denoting the maximum amount that the players can be charged. We call a mechanism budget feasible if the total amount of payments does not exceed the collective agents' budget B. Since our family of mechanisms exhibits a tradeoff between approximation ratio and payments, both increasing with l we wish to choose the maximum value of l such that the mechanism remains budget feasible. However, calculating such l from Equation 3.13 is challenging, so we will use the bound of Equation 3.14

Claim 49. $\frac{|x|}{m^{1/l}}$ *is a strictly increasing function of l.*

We can exploit this fact so that given any given B, we compute the maximum value of l so that the payment bound of equation 3.14 is below B.

Formally, we define a function g that computes such an l and we apply the mechanism f_l to the input. In other words we define a master mechanism

$$f(x,B) = f_{g(x,B)}(x) \quad \text{, where } g(x,B) = \frac{\ln m}{\ln(|x|) - \ln(B)}$$

The master mechanism derived by this procedure achieves

$$\left[\frac{|x|}{B}, B\right]$$

and is therefore budget feasible.

However this approach introduces an additional dependence of the mechanism to the agents' valuation which was not included in the above analysis of fixed l . The outcome of the master mechanism is a complex function of x that is no more truthful. To overcome this issue, we will employ a sampling technique that is heaviy used in the field. We will split the agents randomly into two equal size sets S_1, S_2, define x_1, x_2 as $x_1 = \sum_{y \in S_1} y$, $x_2 = \sum_{y \in S_2} y$. We will use x_1 to compute the optimal value of l, allocating nothing to the agents of S_1 and apply f_l to set S_2.

Definition 50. *We define the complete mechanism as*

$$F(x_1, \ldots, x_n, B) = f_{g(\sum_{i=1}^{\lfloor n/2 \rfloor} x_i, B)} \left(\sum_{i=\lceil n/2 \rceil + 1}^{n} x_i \right)$$

and apply in to a random permutation of the agents. The function g is defined as $g(x,B) = \sup\{l \mid l \in (1, +\infty), |x|/m^{1/l} \leq B\}$.

A subtle complication in the sampling phase is introduced if agents in the sample set actually receive utility from the outcome of the mechanism. In the case of public projects agents may misreport their true values hoping to end up in the sampling set and change the mechanism applied dramatically. This is of no concern in auctions since we can explicitly exclude agents in the sampling set from any allocation.

Proposition 51. *The mechanism F is truthful in expectation.*

Proof. Each agents is randomly placed in the sampling set or the allocation set. If the agents is in the sampling set, any valuation nets zero utility, leaving no incentive for misreport. Agents in the allocation set participate in the mechanism f_l for some l they cannot control, which is truthful in expectation. The mechanism F is a randomization of over truthful in expectation mechanisms and is therefore truthful in expectation. □

Notice that any function computed on the agents of the sampling set, has the same expected value if computed on the allocation set, due to the symmetry of the sampling. However, this is not enough to guarantee budget feasibility even in expectation.

Example. Suppose we have 2 outcomes and 4 agents with values $x_1 = (1000, 0)$, $x_2 = (10, 0)$, $x_3 = (0, 8)$, $x_4 = (0, 999)$ and $B = 9$. Sampling agents x_2, x_3 will chose $l = \infty$ since we can implement the VCG auction. Running this auction on agent x_1, x_4 will charge agent x_1 the second price 999 greatly overshooting the budget.

In order to overcome this issue we need to consider the properties of the environment and make certain assumptions about the valuations of agents. For the sake of concreteness we will show the case where the valuations of the agents are smooth, that is they are a constant factor from each other.

Definition 52. *A valuation profile $(x_1, \ldots, x_n) \in (\mathbb{R}^m_+)^n$ is* c-smooth *if for every pair of agents i, j:*

$$|x_i| \leq c|x_j|$$

Notice that for any c-smooth profile, the sum of any two equal partitions can be at most a factor of c apart.

Proposition 53. *For every c-smooth valuation profile $(x_1, \ldots, x_n) \in (\mathbb{R}^m_+)^n$, let $k = \lfloor \frac{n}{2} \rfloor$. Then*

$$|\sum_{k=1}^{k} x_i| \leq c|\sum_{i=k+1}^{n} x_i|$$

Proof. For all $i \leq k$, by the definition of c-smoothness

$$|x_i| \leq |x_{k+i}|$$

If n is even, we sum up these inequalities for $i \in [1, k]$ and we get

$$\sum_{i=1}^{k} |x_i| \leq c \sum_{i=k+1}^{n} |x_i|$$

In the case of even n we use the fact the $|x_n| \geq 0$ to derive the exact same relation. Changing the order of summation, by associativity we get

$$|\sum_{k=1}^{k} x_i| \leq c|\sum_{i=k+1}^{n} x_i|$$

□

Applying our mechanism F to smooth profiles, implies that during the sampling phase the mechanism learns the sum of agents' vectors in the allocation set at least within a factor of c. By discounting the budget by a factor of c we can guarantee budget feasibility. Notice that we will charge each agents the expected amount independently of the realized outcome, to ensure ex post budget feasibility.

Proposition 54. *Let F be the mechanism of Definition 50.Then for any c-smooth valuation profile $(x_1, \ldots, x_n) \in (\mathbb{R}_+^m)^n$ (randomly permuted), the mechanism*

$$F(x_1, \ldots, x_n, B/c)$$

achieves expected payments and approximation ratio bound by

$$\left[\frac{c|x|}{2B}, B\right]$$

Proof. Let x_1, x_2 be the valuations of the sampling and allocation set respectively. The mechanism selects l such the

$$\frac{|x_1|}{m^{1/l}} = \frac{B}{c} \Leftrightarrow m^{1/l} = \frac{c|x_1|}{B}$$

The expected approximation ratio of the mechanism is the expectation of $m^{1/l}$ taken over the randomization of the sampling, that is

$$\mathbf{E}[m^{1/l}] = \mathbf{E}\left[\frac{|x_1|c}{B}\right] = \frac{c\mathbf{E}[|x_1|]}{B}$$

The expectation of $|x_1|$ is the average sum of subsets of agents' vectors with cardinality $n/2$. In this summation the vector of each agent appears in exactly half the subsets so

$$\mathbf{E}[|x_1|] = \frac{|x|}{2}$$

Substituting in the formula above we get the the expected SW approximation ratio is at most

$$\frac{c|x|}{2B}$$

The payments are bound by

$$\frac{|x_2|}{m^{1/l}} = \frac{|x_2|B}{|x_1|c} \leq \frac{c|x_1|B}{|x_1|c} = B$$

where the inequality holds, due to Proposition 53. □

As a result of the smoothness assumption we were able to guarantee budget feasibility independently of the random outcome of sampling. In order to achieve this we lost a factor of c in social welfare approximation ratio. Notice that assuming some smooth income is not the only assumption leading to budget feasible mechanisms in the ex post sense with this technique. Any sampling procedure (together with some assumptions on the input values) that can *guarantee* an upper bound on the sum of the agent values (or some other payment bound) in the *ex post* sense can produce similar results.

It follows from our sampling procedure that the complete mechanism is not scale invariant. Our mechanism solves the optimization problem on a space defined by sampling half the agents. Formally, if x_1, x_2 are the sampling and allocation sets

$$f(x) = \text{argmax}_{s \in S(x_1)} x_2 \cdot s$$

since the feasible region depends on the input of the mechanism.

3.6 An alternative approach - The exponential mechanism

The reason we are able to achieve the welfare-payment tradeoff is the fact that through the explicit choice of optimization space we limit ourselves to smooth distributions, by discriminating against pure outcomes. This way agents cannot change the welfare by too much, hence reducing the VCG payments. An alternative approach to achieving the same result is instead of changing the optimization space, to modify the optimization objective to take into account the "sharpness" of the outputted distribution, and penalize outcomes with non-smooth distributions. Such a measure of smoothness in probability distributions in the Shannon entropy

Definition 55. *For a probability distribution $p \in [0,1]^n$, the Shannon Entropy is defined as*

$$S(p) = \sum_i -p_i \log(p_i)$$

The definition captures the notion of predictability of random events following the distribution. The amount of "information" conveyed by a random event is minus the logarithm of the propability of it. Thus, the Shannon entropy is the *expected* information gain by conducting the random experiment. Distribution with high entropy are the ones where there is high uncertainty about the outcome of the experiment, and are therefore those which have probabilities of events close to each other. This captured exactly the notion of smoothness we utilized to achieve our tradeoff in the previous part of the chapter.

We will therefore want to bias the mechanism in a configurable way towards mixed outcomes. We define the allocation rule as

$$f(x) = \text{argmax}_{s \in \Delta(O)}(x \cdot s + aS(s))$$

In [18], it was proved that the result of this optimization objective is exactly the exponential mechanism

$$f(x) = \frac{e^{x/a}}{|e^{x/a}|}$$

where for a vector $x \in \mathbb{R}_+^n$ we denote as $e^x, \ln(x)$ the vectors with i-th coordinate $e^{x_i}, \ln(x_i)$ respectively.

It is easy to verify that as $a \leftarrow +\infty$, f becomes the uniform allocation, while as $a \leftarrow 0$ f becomes the optimal allocation.

The result of [18] was developed in order to prove that the exponential mechanism of [20] is indeed exaclty truthful. The exponential mechanism was developed in the radically different framework of differential privacy. In this framework, the desired goal is developing mechanism that are resilient to single agent deviations, thus ensuring the privacy of the participants.

Definition 56. *An allocation rule f is ϵ-differentially private if* $\forall i \in N, u_{-i} \in V_{-i}, u_i, u_i' \in V_i, o \in \{0,1\}^{|O|}$,

$$o \cdot f(u_{-i}, u_i) \leq e^{\epsilon} o \cdot f(u_{-i}, u_i')$$

The definition captures the amount a single agent can change the probability of a set of events, and consequently the information the outcome reveals about the input of individual agents. A survey of results on the topic can be found in [13]. The initial result of [20] proved that the exponential mechanism is $\frac{2}{a}$-differentially private, and by proving the mechanism is a maximizer, it allowed the derivation of a truthful payment scheme for it.

Lemma 57. *The payments that render the mechanism with allocation f truthful and IR are*

$$p(x,y) = -\frac{y \cdot e^{(x+y)/a}}{|e^{(x+y)/a}|} - a \ln\left(\frac{|e^{x/a}|}{|e^{(x+y)/a}|}\right)$$

Proof. Following the analysis of section 3 we get that the required payments are

$$p(x,y) = x \cdot f(x+y) + aS(f(x+y)) - x \cdot f(x) - aS(f(x))$$

We denote as $\vec{1}$ the all-one vector of dimension n. For any input x

$$S(f(x)) = -f(x) \cdot ln(f(x)) = -\frac{e^{x/a}}{|e^{x/a}|} \cdot \ln\left(\frac{e^{x/a}}{|e^{x/a}|}\right) =$$

$$= -\frac{e^{x/a}}{|e^{x/a}|} \cdot \left(\frac{x}{a} - \ln(|e^{x/a}|)\vec{1}\right) = -\frac{x \cdot e^{x/a}}{a|e^{x/a}|} + \ln(|e^{x/a}|) = -\frac{x \cdot f(x)}{a} + \ln(|e^{x/a}|)$$

Substituting in $p(x,y)$ we get

$$p(x,y) = x \cdot f(x+y) - (x+y) \cdot f(x+y) + a\ln(|e^{(x+y)/a}|) - x \cdot f(x) + x \cdot f(x) - a\ln(|e^{x/a}|) =$$

$$= -y \cdot f(x+y) - a \ln \left(\frac{|e^{x/a}|}{|e^{(x+y)/a|}} \right) = -\frac{y \cdot e^{(x+y)/a}}{|e^{(x+y)/a}|} - a \ln \left(\frac{|e^{x/a}|}{|e^{(x+y)/a|}} \right)$$

Notice the relation between a in the exponential mechanism and l in the poly-proportional mechanism of the previous sections. This relation reveals a deeper connection between payment amount and allocation smoothness, and is definitely worth examining deeper.

□

3.7 Applying smooth mechanism to the per-player budget case

We will now consider the case where each agent is constrained by a budget B as in the setting of Chapter 3. We will be restricting the setting to that of a single divisible good auction with n agents where agent i has a utility vector of $x_i = v_i \cdot e_i$ where e_i is the unit vector in the direction of axis i and $v_i \in [0, +\infty)$. This can be also interpreted as the model for randomized allocation rules with single minded bidders.

In this setting, for the mechanism of (3.11), it is obvious that the highest bidder has the highest allocation share due to the proportional nature of the mechanism. Moreover we can prove the following.

Lemma 58. *The mechanism with allocation rule $f(x)$ assigns the payments such that*

$$v_i \geq v_j \Rightarrow p(x_{-i}, x_i) \geq p(x_{-j}, x_j)$$

Proof. Assume $v_i \geq v_j$. The payments for player i are computed as

$$p(x_{-i}, x_i) = x_{-i} \cdot f(x_{-i}) - x \cdot f(x) + x_i \cdot f(x) =$$

$$= \frac{1}{m^{1/l}} \left(||x_{-i}||_l - ||x||_l + \frac{v_i^l}{||x^{l-1}||_{\frac{l}{l-1}}} \right)$$

We can assume w.l.o.g. that $|x^l| = 1$, since the mechanism is scale invariant and if the theorem holds for this case, then it holds for every multiple of x. Therefore $||x^{l-1}||_{l/(l-1)} = 1$ and

$$p(x_{-i}, x_i) - p(x_{-j}, x_j) \geq 0 \Leftrightarrow ||x_{-i}||_l - ||x_{-j}||_l + v_i^l - v_j^l \geq 0 \Leftrightarrow$$

$$\Leftrightarrow (1 - v_i^l)^{1/l} + v_i^l \geq (1 - v_j^l)^{1/l} + v_j^l$$

Let $|x^l - x_i^l - x_j^l| = s \Leftrightarrow v_i^l + v_j^l + s = 1$. Therefore

$$(1 - v_i^l)^{1/l} + v_i^l \geq (1 - v_j^l)^{1/l} + v_j^l \Leftrightarrow (s + v_j^l)^{1/l} + v_i^l \geq (s + v_i^l)^{1/l} + v_j^l \Leftrightarrow$$

$$\Leftrightarrow v_i^l - v_j^l \geq (v_i^l + s)^{1/l} - (v_j^l + s)^{1/l}$$

For the case of $s = 0$, this is equivalent to

$$v_i^l - v_j^l \geq v_i - v_j$$

which holds for $l \in [1, +\infty)$ since it holds for $l = 1$ and $v_i^l - v_j^l$ is increasing with respect to l.

Moreover, $(v_i^l + s)^{1/l} - (v_j^l + s)^{1/l}$ is decreasing with s since,

$$\frac{d}{ds}((v_i^l + s)^{1/l} - (v_j^l + s)^{1/l}) = \frac{1}{l}((v_i^l + s)^{1/l-1} - (v_j^l + s)^{1/l-1})$$

and

$$v_i \geq v_j \Leftrightarrow v_i^l + s \geq v_j^l + s \Leftrightarrow (v_i^l + s)^{1/l-1} \leq (v_j^l + s)^{1/l-1}$$

Therefore the inequality holds for every s and the proof is complete. □

This lemma allows us to consider only the payments of the highest bidder and setting them equal to B if $v_2 \geq B$. However we face the same problem we faced in Section 3.5, with the agents being able to affect the choice of l. We will need to develop some similar sampling techniques to derive a truthful complete mechanism. We strongly believe such a mechanism is a constant approximation to the LW.

Bibliography

[1] L. M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, pages 1452–1475, 2004.

[2] A. Badanidiyuru, R. Kleinberg, and Y. Singer. Learning on a budget: posted price mechanisms for online procurement. In B. Faltings, K. Leyton-Brown, and P. Ipeirotis, editors, *ACM Conference on Electronic Commerce*, pages 128–145. ACM, 2012.

[3] X. Bei, N. Chen, N. Gravin, and P. Lu. Budget Feasible Mechanism Design: From Prior-Free to Bayesian. *CoRR*, abs/1203.4455, 2012.

[4] S. Bhattacharya, V. Conitzer, K. Munagala, and L. Xia. Incentive Compatible Budget Elicitation in Multi-unit Auctions, Apr. 2009.

[5] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget Constrained Auctions with Heterogeneous Items, Mar. 2010.

[6] N. Chen, N. Gravin, and P. Lu. On the Approximability of Budget Feasible Mechanisms. *CoRR*, abs/1007.3801, 2010.

[7] E. Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

[8] R. Cole, V. Gkatzelis, and G. Goel. Mechanism Design for Fair Division. *CoRR*, abs/1212.1522, 2012.

[9] S. Dobzinski, R. Lavi, and N. Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

[10] S. Dobzinski and R. P. Leme. Efficiency Guarantees in Auctions with Budgets. *CoRR*, abs/1304.7048, 2013.

[11] S. Dobzinski, C. H. Papadimitriou, and Y. Singer. Mechanisms for complement-free procurement. In Y. Shoham, Y. Chen, and T. Roughgarden, editors, *ACM Conference on Electronic Commerce*, pages 273–282. ACM, 2011.

[12] S. Dughmi and Y. Peres. Mechanisms for Risk Averse Agents, Without Loss, June 2012.

[13] C. Dwork. Differential Privacy: A Survey of Results. *Theory and Applications of Models of Computation*, pages 1–19, 2008.

[14] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.

[15] G. Goel, V. S. Mirrokni, and R. P. Leme. Clinching Auction with Online Supply. In S. Khanna, editor, *SODA*, pages 605–619. SIAM, 2013.

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

[17] J. D. Hartline and T. Roughgarden. Optimal Mechansim Design and Money Burning. *CoRR*, abs/0804.2097, 2008.

[18] Z. Huang and S. Kannan. The Exponential Mechanism for Social Welfare: Private, Truthful, and Nearly Optimal. *CoRR*, abs/1204.1255, 2012.

[19] P. Lu and T. Xiao. Improved Efficiency Guarantees in Auctions with Budgets, Sept. 2014.

[20] F. McSherry and K. Talwar. Mechanism Design via Differential Privacy. In *FOCS*, pages 94–103. IEEE Computer Society, 2007.

[21] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[22] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. 2007.

[23] A. D. Procaccia. Can Approximation Circumvent Gibbard-Satterthwaite? In M. Fox and D. Poole, editors, *AAAI*. AAAI Press, 2010.

[24] A. D. Procaccia and M. Tennenholtz. Approximate Mechanism Design without Money. *ACM Trans. Economics and Comput.*, 1(4):18, 2013.

[25] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2):187–217, 1975.

[26] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.

[27] Y. Singer. Budget Feasible Mechanisms. In *FOCS*, pages 765–774. IEEE Computer Society, 2010.

[28] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–27, 1961.