AN INCREMENTAL ELICITATION APPROACH TO LIMITED-PRECISION AUCTIONS

by

Alexander Kress

A thesis submitted in conformity with the requirements
for the degree of Master of Science
Graduate Department of Computer Science
University of Toronto

Abstract

An Incremental Elicitation Approach to Limited-Precision Auctions

Alexander Kress

Master of Science

Graduate Department of Computer Science

University of Toronto

2004

Auction-based mechanisms are increasingly being used for automating resource allocation among large numbers of agents. To make these sort of mechanisms viable one needs to consider the issues of communication and computation expenditure required by these protocols as well as their stability. In this thesis we study limited-precision, iterative mechanisms with dominant strategy equilibria designed for allocation of a single good. Our goal is to limit the communication between the players and the mechanism, reduce the amount of information revealed by the players, as well as minimize the players' computational costs. We accomplish this by placing a number of operational constraints that permit the above objectives. We prove several necessary conditions that severely restrict the space of mechanisms satisfying our criteria. We develop a number of mechanisms and show that with a large and variable number of players, in the case of limited-precision, iterative mechanisms are superior to single-shot mechanisms.

Acknowledgements

A number of people have greatly helped with writing of this thesis. The person I would like to thank the most is of course my supervisor, Craig Boutilier. Without his support, advice, and encouragement this work would be impossible. During the last three years Craig has been not only a supervisor but also a mentor and a friend, his advice was always most valuable. It was a great pleasure to work with Craig and it would be great to continue working together in the future.

The second group of people that deserves much thanks are my friends. They have endured my often cryptic descriptions of yet another attempt at some proof or approximation. Throughout all this they have maintained faith that graduation will eventually come, and that faith had kept me going. Special thanks go to Anya for braving strange notation and reading through many of the proofs; and to Tolik for proofreading the entire thesis.

It is impossible to describe the amount of support and encouragement I have received from my parents during these years. They have always listened, believed and helped. My mom and dad are the people who deserve the most credit, thank you.

Contents

Chapter 1

Introduction

Fast evolution of chip-based technologies has greatly moved the computational boundary of problems considered feasible compared to even a few decades ago. Yet, with greater potential came a greater demand for computational resources. Recent applications increasingly have to deal with a significant number of users competing for the same resources. Multi-user operating systems, file and web servers are perhaps the most common examples, however this issue is by no means restricted to the universe of software and hardware developers. This problem is starting to become quite common in day-to-day business and personal life as individuals come to rely on the computer networks to provide communication and negotiation channels. Even with the current improvements in computational speeds, it becomes difficult to effectively satisfy the demands of everyone involved.

As the transactions become more complex the companies and individuals are starting to rely on sophisticated software agents that mediate these transactions. However, since the interests of the parties on whose behalf such agents act generally conflict, one can expect the agents to selfishly try to achieve their objective even if it comes at the expense of every other participant. From the perspective of an entity managing the resources, this is often an undesirable form of behavior.

Economics traditionally deals with the problem of deriving the most good from a limited set of resources. The situation described here is no different. So it comes as no surprise that economics provides us with tools to analyze the behavior of agents competing for limited resources. The theoretical treatment of the strategic interactions of rational agents, known as *Game Theory*, provides tools for predicting and guiding the agents' behavior. The assumption of agent rationality – while usually suspect in an interaction between individuals – is easily accepted when in comes to software agents. As such, recent research in computer science and economics has focused on design of economic agents and the mechanisms through which they interact.

1.1 Outline of the problem

While quite often one has to design or analyze an agent in a predefined environment, in other circumstances it is possible to design the environment itself. A subset of game theory concerned with the problem of designing the environment (i.e. the rules of the game), is called *Mechanism Design* [20]. Mechanism design has played a major role in much of the research at the intersection of computer science and economics. Key results in mechanism design, such as the revelation principle, have had a strong influence on the direction taken by research relying on both of these disciplines.

Recently, however, limitations of standard approaches to mechanism design have been identified, and are starting to be addressed. Chief among these is the computational complexity of the problems faced by software agents who must interact. For instance, mechanisms based on the revelation principle must accurately and fully reveal a significant amount of information about themselves (generally, their utility function). This presents a problem in circumstances where utility functions are large and difficult to communicate effectively and/or hard to compute accurately. Although initially not considered in the game theory literature, this problem is quite common in real applications. One can easily come up with examples where computation or communication constrains affect agents' behavior. For instance, consider the task of estimating the value of an undeveloped coal mine. The quality of the computed estimate often depends directly on the amount of resources spend on the initial survey. Computing the exact value of the mine can likely be prohibitively expensive for most agents. Thus the agents may have to base their actions on the partial knowledge of the true value.

As an example of an application where communication complexity plays a role, consider a protocol designed to allocate a large number of small resources, such as network packets or CPU time slots [1]. In this case, if an agent were to state her utility function she would need to fully describe her value for every resource. This means that if full valuation revelation is attempted the communication overhead might become more costly than the value of the resource itself.

In addition to communication and computational difficulties, another equally serious problem is the unwillingness of the participants to reveal certain information about themselves. The primary reason for that behavior is that the amount of revealed information may depend on certain strategic considerations. This problem usually arises in the context of repeated games. There, any information revealed by the player might place her at a disadvantage at future rounds. Even ignoring any other strategic considerations, the player would prefer to reveal information only if it positively affects her overall utility.

Combining all of these issues together one can see a need for a mechanism that would permit resource allocation without placing costly computational, communication and revelation requirements on the participants. In this document we will look at one approach to developing such mechanisms.

1.2 Related work

Recent research has begun to examine methods involving limited elicitation of player's preferences to avoid some of the difficulties mentioned above [1, 2, 32, 25, 27, 7], specifically in the context of (single-good or combinatorial) auctions.

Of particular note is the work by Blumrosen, Nisan and Segal [1, 2]. Their approach to limited communication auctions is to look at single-shot mechanisms, that is, mechanisms restricted to one iteration. They develop two mechanisms which can be shown to be optimal (social welfare, seller's revenue) among all two-player single-shot limited precision mechanisms and asymptotically optimal in the n players case. At various points in this document we will refer to this work. Most notably, however, in Chapter 5 we will show how our treatment of the problem can results in significant efficiency gains over the protocols proposed by Blumrosen et. al.

Another important body of work is the *iBundle* auction developed by D. Parkes [25]. This mechanism, although applicable to combinatorial auctions, something we do not consider here, supports our belief that multi-step mechanisms are superior for dealing with limited communication.

We describe these and other approaches to limited communication in the next chapter.

1.3 Approach

In this document we take the approach of using an auction mechanism as a tool for resource allocation when the players' preferences are unknown. We look at iterative mechanisms with dominant strategy equilibria, a combination we refer to as *incremental* mechanisms. More specifically, we conjecture that incremental mechanisms are in many ways superior to other recently developed limited communication mechanisms. This belief is based on the following reasons. The existence of dominant strategy implies that a player has a best plan of action independent of the way everyone else acts. This greatly simplifies life for both the players, who are always assured that they can't do any better, and for the agent operating the protocol, who can expect a more stable protocol performance. The other side of this design, the iterative feature of the mechanism, allows both the players and the mechanism to react to or to evolve with the game. In particular, it should be possible for the mechanism to determine early on (with small amount of communication) which players are more likely to win the good, and act on this information.

As was noted we call the combination of iterativity and dominant strategy an incremental auction. Generally, an incremental mechanism is the one that sequentially refines its estimate of each player's desire for the good based on their bids; as such it requires that the players' messages provide truthful information. Since any mechanism with dominant strategies can be thought of as truthful, this makes our desired mechanism incremental. These concepts will be introduced more formally and in more detail in the next chapter.

In this document we focus on single-good auctions, that is auctions designed to allocate one item among a group of players. To restrict the space of possible mechanisms we introduce a number of constraints, which are natural in the context of limited communication auctions. We consider deterministic[1] protocols which treat the players fairly, that is the good can only be allocated to a player if the mechanism knows with certainty that this player is uniquely superior to everyone else, according to some criterion. We also place a participation constraint that removes any bidder from the auction if this bidder provably can not be the winner. We assume that the players are *ex-post* rational, and require (informally) that the resulting auctions have an intuitive form. Under these constraints we prove a number of results, which are described in the next section.

1.4 Main contributions

We show that any mechanism satisfying the above constraints[2] and having a dominant strategy equilibrium, must belong to a class of *increasing price mechanisms*, a concept similar to *ascending price mechanisms*. Subsequently we prove a number of necessary conditions which must be satisfied by any mechanism with dominant strategy equilibrium respecting the above requirements. Using these necessary conditions we show that any such mechanism must have

[1] We partially relax this assumption in Chapter 5.

[2] This result also holds without the participation constraint.

a very specific form, where in most cases the space of bidders' actions can be safely restricted to just two actions. We arrive at this conclusion by showing that in most cases all players will act similarly (that is, follow largely equivalent strategies). This finding is important since it allows us to approximately specify the form of any potential mechanism, given our initial assumptions.

Next, to provide an example of how an increasing price mechanism, satisfying the above requirements and necessary conditions, would operate we study one such protocol in detail. Our *adaptive, symmetric, incremental auction (ASIA)* is parameterized by a (conditional) sequence of price levels, but has simple dominant strategies based on any such thresholds. *ASIA* restricts the players actions to just two legal choices, and at each iteration the player is required to state whether she wants to continue participating or to drop out.

An interesting feature of the protocols that we propose is that their threshold values can be optimized with respect to specific priors without losing the benefits of dominant strategies.[3] In this document we accomplish this using a straightforward Markov decision process formulation. In this respect, we combine the spirit of traditional mechanism design (where dominant strategies exist independent of the priors), optimal auction design [22] (where the mechanism varies in a parameterized way with the prior), and automated mechanism design [8] (where the mechanism is optimized using the specific priors). Importantly, we can optimize the mechanism to account for cost of communication (or computation) as well [27].

One unsatisfying feature of the protocols outlined above is that they are not fully efficient. That is, there exist some situations where these protocols would act suboptimally in terms of social welfare. As a possible solution to this problem we introduce an alternative protocol, which is derived by partially relaxing the original assumption of full determinism. As a result we develop a stochastic version of *ASIA* mechanism, which we call *ST-ASIA*. We introduce a fixed price update rule *Divide(k)* which allows us to show that under certain conditions *ST-ASIA* is superior to any threshold based single-shot auction. It is also more efficient than the original

[3]As long as they remain *increasing price* mechanisms.

(ASIA). In addition we demonstrate that *ST-ASIA* can also be easily optimized for most social choice functions using a similar MDP type optimization.

In short, the major contribution of this document is that it demonstrates that under a number of natural restrictions all mechanisms must have a certain, easily recognizable, form which, although restrictive, possesses some nice properties. Furthermore, a relaxation of just one of these conditions produces a mechanism that maintains all of these desirable properties but at the same time is superior to some of the known (nearly) optimal single-shot mechanisms.

1.5 Structure

This document is structured as follows. In Chapter 2 we review the necessary background on game theory and mechanism design and provide a summary of the current research that impacts the topic of this document. Chapter 3 introduces a simple single-shot auction. The protocol developed in this chapter often performs worse than some of the known single-shot auctions, however it serves as a good introduction and motivation to the main results located in the following chapters. Chapter 4 contains the theoretical contributions of this document. There we develop the necessary conditions possessed by all mechanisms proposed in this document. This chapter also shows how these conditions help define the form of the mechanism. In addition in this chapter we present two examples of auction mechanisms and analyze one of them. In Chapter 5 we develop a more efficient mechanism and demonstrate it's benefits. We show that this mechanism is superior under certain conditions to single-shot mechanisms, specifically the ones presented in [1, 2]. Finally, Chapter 6 summarizes the contributions of this document as well as provides directions for future research.

Chapter 2

Background and Related Work

This chapter provides an overview of the theoretical background underlying the results of this document as well as the summary of the current literature on limited precision auctions.

2.1 Notation

The following notation will be used throughout this document. The notation will also be introduced in the text as needed.

$C : \mathcal{S} \times A \rightarrow \Re$ - the cost function

$f : \vec{\Theta} \rightarrow O$ - the social choice function (SCF), $\vec{\Theta} = \Theta_1 \times ... \times \Theta_n$

$\Phi(\vec{\theta})$ - the probability distribution function over the type profiles

$g : \vec{\Sigma} \rightarrow O$ - the outcome rule, $\vec{\Sigma} = \Sigma_1 \times \Sigma_2 \times ... \times \Sigma_n$

h_i^t - the private history of player $i \in I$ up to iteration t

I - the set of players participating in a game.

$\mathcal{M}$ - a mechanism

M^t - the set of legal moves or messages for any player $i \in I$ at iteration t

n - the number of players participating in a game $n = |I|$

O - the set of possible outcomes

$R : \mathcal{S} \rightarrow \Re$ - the reward function

θ_i - the type of player $i \in I$, $\theta_i \in \Theta_i$

Θ_i - the space of possible types of player $i \in I$

$\vec{\theta}$ - the type profile, $\vec{\theta} \in \Theta_1 \times \Theta_2 \times ... \times \Theta_n$

$\vec{\theta}_{-i}$- the type profile of all players, except player $i \in I$

$\mathcal{S}$ - the state space

s_i - the sequence of moves of player $i \in I$, with $s_i[k]$ denoting a length *k* initial sequence, $s_i[k] \in M_i^1 \times M_i^2 \times ... \times M_i^k$

s_{-i} - the set of sequences of moves of all players except player $i \in I$

σ_i - the strategy of player $i \in I$

Σ_i - the space of possible strategies of player $i \in I$

$\vec{\sigma}$ - the strategy profile, $\vec{\sigma} \in \Sigma_1 \times \Sigma_2 \times ... \times \Sigma_n$

$\vec{\sigma}_{-i}$- the strategy profile of all players, except player $i \in I$

$\tau_i : \vec{\Sigma} \rightarrow \Re$ - the transfer function, for single-good auction mechanisms

$V_t^{\pi} : \mathcal{S} \rightarrow \Re$ - the value of being in a state after *t* stages under policy π

v_i - the valuation of player $i \in I$ for some good, (equivalent to type).

$u_i(o, \theta_i)$ - the utility of player $i \in I$ with type $\theta_i \in \Theta_i$, given outcome $o \in O$

$x_i : \vec{\Sigma} \rightarrow [0, 1]$ - the allocation function, for single-good auction mechanisms

2.2 Mechanism Design

Many real-life situations can be viewed as games where participants follow a certain set of rules to achieve their own individual objectives. For example, consider the task of commuting to and from work with the objective of minimizing the total traveling time. Each individual's actions are restricted by the rules of the game, in this case, the available means of transportation, their capacity, schedule, and cost. The actual traveling time is affected by both the individual's actions as well as the actions of others. Given a game modeled in this way, an external observer might try to design a set of rules in order to optimize the traveling time of some particular individual or group, or to achieve some other objective.

Mechanism design focuses on the task of constructing a *protocol* or *game rules* that guide the interaction of a number of self-interested agents in such a way as to optimize some global objective of the designer [20]. One of the underlying assumptions of mechanism design is that the individuals will act rationally and strategically in order to increase their own utility. This assumption allows the use of game-theoretic tools to model the interaction between the players.

More formally, suppose we have a collection *I* of players or agents. Each agent $i \in I$ has a *type* $\theta_i \in \Theta_i$, usually known only to herself. In essence the agent's type defines her preferences over a set of possible world states. To formalize, suppose each agent has a utility function $u : O \times \Theta_i \rightarrow \Re$, where *O* is a set of possible outcomes (world states). One of the advantages of formulating the agent's preferences this way is that it provides an easy way to compare any two outcomes from the agent's point of view; that is, agent *i* prefers outcome o_1 to o_2 if $u(o_1, \theta_i) > u(o_2, \theta_i)$. (It is of course prudent to question the existence of such a function [28], however in classical game theory we simply assume that one exists [24]).

Next we define the concept of *strategy*, one of the most fundamental concepts in game theory. Most generally a strategy can be defined as follows:

Definition 2.1 A strategy σ_i of player $i \in I$ is a description of a complete action plan that defines which action the player will take given any possible (legal) type of player i and the current history of play available to i (the history of play is a current description of the world as known to player i).

The function σ_i can be thought of as a mapping from the player's type and all the additional historical information available to her into legal action(s). The information available to the player might be some acquired facts about the other players, current world state, the history of moves or any other relevant information. In one-stage games the strategy will determine a single action of the player, while in sequential games the strategy will determine a sequence of actions, taking into account any new information that becomes available to the player. We will use $\sigma_i(\theta_i, h)$ to denote player *i*'s strategy with type θ_i and history h, and $\sigma_i(\theta_i)$ for single-stage games.

In this document we will be mostly concerned with *pure* strategies, which deterministically select an action for every type and history. It is also possible to work with *mixed* strategies that select actions by defining a probability distribution over *pure* strategies.

We can now define the notion of a *mechanism*. Informally, a mechanism is simply a set of rules that dictate what each player is allowed to do for every game history and determine the consequence for each combined set of players' actions.

Definition 2.2 A *mechanism* $\mathcal{M} = (\vec{\Sigma}, g(\cdot))$ defines a set of possible *strategy profiles* $\vec{\Sigma} = \Sigma_1 \times \Sigma_2 \times ... \times \Sigma_n$ and an *outcome function* $g : \vec{\Sigma} \rightarrow O$.

A *strategy profile* $\vec{\sigma} \in \vec{\Sigma}$ is a vector of strategies, one for each player. This means that the space of strategy profiles is a set of all possible vectors of strategies over the actions allowed by the mechanism $\mathcal{M}$. Intuitively, the mechanism defines the set of legal moves (actions, bids, etc.) with some moves possible only in certain circumstances or implementable only by certain players. Then, for any player $i \in I$ one can construct a set Σ_i of all possible strategies based on legal moves.

An *outcome function* $g : \vec{\Sigma} \rightarrow O$ is a mapping from the players' strategy profile $\vec{\sigma}$ into an outcome space O. This function determines the outcome of the game based on the combined actions of all players. Note that the function is only concerned with the players' actual actions, not their full strategies (which are more expressive). The mapping from the strategy space is for notational convenience only.

2.2.1 Selecting social goals

The task of constructing a mechanism is based on some specific goal that the designer wants the mechanism to achieve or at least come close to. Some examples of such goals are revenue maximization in auctions [22, 1], congestion reduction in network routing [30, 19], and minimization of latency in a shared system [29]. To formalize these goals we will define a concept of *social choice function* and then describe a number of common *social choice functions*.

Definition 2.3 A *social choice function* $f : \vec{\Theta} \rightarrow O$ describes the desired outcome of the game given the players' type profile $\vec{\theta}$, where $\vec{\theta}$ is a vector of types of all players participating in the game.

Although any *social choice function* can be used as a goal in designing a mechanism, most literature [1, 2, 25, 22, 8] deals with one of the following four functions: *social welfare*, *Pareto efficiency*, *budget balance*, *discriminatory utility maximization*.

Definition 2.4 A social choice function leads to a *Pareto efficient* outcome if it is impossible to find any other outcome that makes some agent better off without making some other agent worse off. (Given a Pareto efficient $f(\vec{\theta}) = o$, if $\exists o' \in O$ and $i \in I$ s.t. $u_i(o', \theta_i) > u_i(o, \theta_i)$ then $\exists k \in I$ s.t. $u_k(o', \theta_k) < u_k(o, \theta_k)$)

Definition 2.5 A social choice function maximizes *social welfare* if it selects the outcome that maximizes the sum of players utilities, $f(\vec{\theta}) = argmax_o \sum_{i \in I} u_i(o, \theta_i)$. Note that a function that maximizes social welfare is also *Pareto efficient*.

Definition 2.6 A social choice function achieves a *discriminatory utility maximization* if it selects an outcome which results in a maximum possible utility for some group of players, the resulting utilities of all the other players are not considered. For some $S \subseteq I$, $f(\vec{\theta}) = argmax_o \sum_{i \in S} u_i(o, \theta_i)$.

In auction mechanism design, the most often used version of *discriminatory utility maximization* is *Seller's Revenue Maximization*. This social choice function always select the most profitable outcome for the seller. Also note that with $S = I$, Definition 2.6 reduces to Definition 2.5.

Definition 2.7 In games that have payments as a part of the outcome space, a social choice function achieves *budget balance* if there are no net flows in or out of the game.

Definition 2.8 A mechanism $\mathcal{M}$ implements a social choice function $f(\cdot)$ if there is an *equilibrium strategy profile* $\vec{\sigma}$ of the game induced by $\mathcal{M}$ such that $g(\vec{\sigma}(\vec{\theta})) = f(\vec{\theta}), \forall \vec{\theta} \in \vec{\Theta}$.

The notion of *equilibrium strategy* was purposefully left undefined for the time being; it will be presented in the next section.

2.2.2 Solution approaches

To evaluate the performance of a mechanism one needs to determine its stable point(s) (if any). These are the points where every player has converged on some strategy that she deems the best and has no wish to change it. There are a number of ways in which this can occur; we will outline the three most relevant – in the context of this document – forms of equilibria. To simplify the notation in the following definitions we will use $(\sigma_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i}))$ to denote the outcome $g(\sigma_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i}))$.

Definition 2.9 A *Nash equilibrium* is a strategy profile $\vec{\sigma} = (\sigma_i, \vec{\sigma}_{-i})$ such that for any player i,

$$u_i((\sigma_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i})), \theta_i) \geq u_i((\sigma'_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i})), \theta_i) \quad \forall \sigma'_i \tag{2.1}$$

where a type profile $\vec{\theta} = (\theta_i, \vec{\theta}_{-i})$ is public knowledge.

Nash Equilibrium requires perfect information, which means that all agents share the same knowledge about the world and each other [14, 24]. This assumption is often unrealistic since frequently an agent has only partial information about the other agents. A related equilibrium concept, *Bayes-Nash Equilibrium* relaxes this assumption. Under *Bayes-Nash Equilibrium* we assume that the agents know only their own type, but not the types of the other players. Instead, the players have a common probability distribution, denoted $\Phi(\vec{\theta})$, over the complete type profiles. Due to that, the equilibrium solution depends on the expected utility, with the expectation taken over the type profile of the other players; as outlined in the equation each player takes the expectation conditional on her own type.

Definition 2.10 A *Bayes-Nash equilibrium* is a strategy profile $\vec{\sigma} = (\sigma_i, \vec{\sigma}_{-i})$ such that for any player i with type θ_i

$$E_{\vec{\theta}_{-i}|\theta_i}\left(u_i((\sigma_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i})), \theta_i)\right) \geq E_{\vec{\theta}_{-i}|\theta_i}\left(u_i((\sigma'_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i})), \theta_i)\right) \quad \forall \sigma'_i \tag{2.2}$$

There also exist stronger notions of Nash Equilibrium and Bayes-Nash Equilibrium, called *Subgame-Perfect Nash Equilibrium* and *Perfect Bayes Equilibrium* respectively. These forms of equilibria are applicable in multi-stage games. They require that the strategies constitute an equilibrium (*Nash* or *Bayes-Nash*) in any subgame, where a subgame is any remaining portion of the game, such that the prior history of play is known to all the players. These forms of equilibria will not be relevant here. Our major findings rely on an even stronger version of stability called *Dominant strategy equilibrium*.

Definition 2.11 A *Dominant strategy equilibrium* is a strategy profile $\vec{\sigma} = (\sigma_i, \vec{\sigma}_{-i})$ such that for any player i with type θ_i

$$u_i((\sigma_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i})), \theta_i) \geq u_i((\sigma'_i(\theta_i), \vec{\sigma}_{-i}(\vec{\theta}_{-i})), \theta_i) \quad \forall \sigma'_i, \vec{\sigma}_{-i} \tag{2.3}$$

With any of the three described forms of equilibrium it is possible to have multiple stable points in a single mechanism. This can cause a serious problem for solutions that rely on *Nash* or *Bayes-Nash* type of equilibriums, because the players not only need to converge on an equilibrium set of strategies, but also they need to converge on the same one. This can be further complicated by the fact that players can have preferences over the stable points. With dominant strategies this is not an issue, since any dominant strategy is as good as any other and furthermore, the players' choices are not affected by the actions of other players.

We can now extend Definition 2.8 to include the concept of *Dominant strategy equilibrium*.

Definition 2.12 A mechanism $\mathcal{M}$ implements a social choice function $f(\cdot)$ in *dominant strategies* if there is a dominant strategy equilibrium profile $\vec{\sigma}$ of the game induced by $\mathcal{M}$ such that $g(\vec{\sigma}(\vec{\theta})) = f(\vec{\theta}), \forall \vec{\theta} \in \vec{\Theta}$.

It is often desirable to have a dominant strategy implementation as it offers a number of significant advantages over other equilibrium concepts [20]. A (weakly) dominant strategy for a player guarantees that the player will receive the greatest possible utility regardless of the actions of other players. Firstly, this means that unlike the requirements of *Bayes-Nash*, the player need not care about the types of other players. This is beneficial since a player often needs to spend a significant amount of resources in order to estimate these types and how they will effect the actions of her opponents [17]. Additionally, if a player happens to rely on an incorrect information regarding her opponents this will not affect her payoff. Secondly, a rational player would always be expected to follow (one) of her dominant strategies. This simplifies the task of evaluating a mechanism's performance since we can determine the outcome of the game unambiguously (as opposed to *Nash* family of equilibria).

2.2.3 Direct revelation and the revelation principle

The task of creating and evaluating any mechanism may seem daunting given the size of space of possible mechanisms. This is especially evident if as a part of the design process one is trying to decide which actions the mechanism should make available to the participants, the space of possibilities is almost infinite. Luckily, without losing any generality, the focus can be restricted to a smaller set of mechanisms. It has been shown that any mechanism inducing a *Dominant Strategy* or *Bayes-Nash* equilibrium can be transformed into a *direct revelation*, *incentive compatible* mechanism which will be outcome equivalent to the original. This result is known as the *revelation principle* [20]. We will now formally define the notions of *direct revelation* and *incentive compatibility* and then state the *revelation principle*.

Definition 2.13 A *direct revelation* mechanism $\mathcal{M} = (\vec{\Sigma}, g())$, is a mechanism in which the space of strategies is restricted to the space of types, $\vec{\Sigma} = \vec{\Theta}$, and the outcome rule is defined as $g : \vec{\Theta} \rightarrow O$.

Definition 2.14 An *incentive compatible* mechanism $\mathcal{M} = (\vec{\Sigma}, g())$, is a mechanism that induces an equilibrium in which the best strategy of every player is to announce her true type, $\sigma_i(\theta_i) = \theta_i$. These types of strategies are called *truth revealing*.

Combining the above definitions we obtain a mechanism in which the actions of every player are restricted to announcing some type and furthermore in equilibrium all players will announce their true types. Now, suppose there exists some mechanism $\mathcal{M}$ which is not *direct revelation* or *incentive compatible*, but does induce an equilibrium. Then one can construct a *direct revelation* mechanism $\mathcal{M}'$ which will convert the players' types into their equilibrium strategy profile for mechanism $\mathcal{M}$, and then run this profile on $\mathcal{M}$. Under such rules it is the best strategy for every player to announce her true type to $\mathcal{M}'$, which makes $\mathcal{M}'$ an *incentive compatible* mechanism. This is the basis for the *revelation principle*. In the following proposition we will state the *revelation principle* for *dominant strategies*, however this principle is much broader than that and can be proven for any equilibrium concept mentioned in this document.

Proposition 2.1 Suppose there exists a mechanism $\mathcal{M}$ that implements a social choice function $f(\cdot)$ in *dominant strategies* then there exists a *direct revelation*, *incentive compatible* mechanism $\mathcal{M}'$ which implements $f(\cdot)$ in dominant strategies.

From the results presented so far it may seem that the task of implementing a social choice function in dominant strategies is not very difficult. This however, is not the case, as is well illustrated by the impossibility result presented in the next section.

2.2.4 Important impossibility result

There are a number of impossibility results relating to the dominant strategy implementation [20]. Here we present the Gibbard-Satterthwaite Impossibility Theorem [11], which is arguably one of the most effective in showing that it is often impossible to find a mechanism that would implement an arbitrary *social choice function* in dominant strategies.

The Gibbard-Satterthwaite Impossibility Theorem shows that under a certain set of conditions the only *social choice functions* that can be implemented in dominant strategies are dictatorial, an often undesired property.

Definition 2.15 A *social choice function* $f(\cdot)$ is *dictatorial* if there exists a player *i* such that $f(\cdot)$ always selects an outcome which is optimal for *i*.

Proposition 2.2 (*The Gibbard-Satterthwaite Impossibility Theorem*) In settings where there are at least two agents ($n \geq 2$), at least three different outcomes ($|O| \geq 3$), and the agents' types belong to the space of all possible types (all possible strict orderings over *O*) it is the case that a *social choice function* $f(\cdot)$ is implementable in *dominant strategies* if and only if it is *dictatorial*.

The immediate conclusion of this theorem is that in the general case it is impossible to implement in dominant strategies many of the social choice functions that maximize some global version of optimum, such as social welfare.

Fortunately, there are a number of ways in which the negative conclusions of this theorem can be avoided. In particular, by placing a restriction on the players' type space it is possible to design a mechanism which implements a non-dictatorial social choice function in dominant strategies with 3 or more outcomes [20]. This is the approach we will be taking in this document.

2.3 Auction design

Single object auction design is the primary focus of this document. In this section we will discuss in some detail the classical theory underlining single object auction design, list some of the commonly known auction types, and provide a brief overview of multi-object auctions.

2.3.1 From mechanism design to auction design

The study of auction design is a subset of mechanism design, with all the concepts discussed in the previous section directly applicable in this context as well. There are, however, a number of basic properties common to all auctions. These properties make it necessary to impose a number of restrictions on some of the previously stated definitions and introduce new notation. We will explore these changes immediately after stating the definition of an auction.

Definition 2.16 An *auction* is a *mechanism* designed to allocate one or more goods (either divisible or indivisible) to some subset of players. An auction defines a set of legal actions consisting of possible bids and an outcome function which determines the object(s) allocation and the monetary transfers.

The first major change to our notation comes directly from Definition 2.16. Every outcome, in the range of the outcome rule $g(\vec{\sigma})$, will now have two components, the *allocation* and the *transfer*.

The *allocation* component specifies which player(s) receive the good(s), given their bids or more generally their strategies. For deterministic single good auctions we define the allocation function as $x_i : \vec{\Sigma} \rightarrow [0, 1]$, where $x_i(\vec{\sigma}) = 1$ if player *i* receives the good and 0 otherwise, with a restriction that $\sum_{i=1}^{n} x_i(\vec{\sigma}) = 1$. Similarly, the *transfer* component specifies for each player the payment amount she has to make to the mechanism (which can be negative). We define the transfer function as $\tau_i : \vec{\Sigma} \rightarrow \Re$. Combining these two components together we obtain the new definition for the outcome rule.

Definition 2.17 In the auctions context the *outcome rule* is defined as

$$g(\vec{\sigma}) = (x_1(\vec{\sigma}), \ldots, x_i(\vec{\sigma}), \ldots, x_n(\vec{\sigma}); \tau_1(\vec{\sigma}), \ldots, \tau_i(\vec{\sigma}), \ldots, \tau_n(\vec{\sigma})) \tag{2.4}$$

where $x_i(\cdot)$ and $\tau_i(\cdot)$ specify the allocation and transfer functions for player *i*.

A very common assumption in auction design literature [15, 25, 1, 18] is to restrict the set of utility functions to those *quasi-linear* in money. This simply means that the utility functions must be linear in money. In historical context, this assumption was instrumental in defining the important family of *Groves-Clarke* mechanisms [20] which, under certain restrictions, are the only mechanisms that provide the most efficient allocation as well as a dominant strategy equilibrium [12]. An example of a *Groves-Clarke* mechanism, the *Second Price Auction*, will be presented in the next section. We will now present and describe the formal definition of a *quasi-linear* utility function.

Definition 2.18 A *quasi-linear utility function* has the following form

$$u_i(x, \tau_i, \theta_i) = v_i(x) - \tau_i \tag{2.5}$$

Where x is some allocation, τ_i is a payment, and the player's type θ_i is fully described by the valuation function $v_i(\cdot)$.

As can be seen from the above definition we now distinguish between the value of possessing the object and the value of possessing some amount of money. For the former we define a concept of *valuation* denoted $v_i(x) \rightarrow \Re$, for player *i*'s valuation for the allocation $x = x_1(\vec{\sigma}), \ldots, x_i(\vec{\sigma}), \ldots, x_n(\vec{\sigma})$. The simplest interpretation of $v_i(x)$, is the maximum amount of money agent *i* is willing to pay to insure allocation x. If it is assumed (as we do in this document) that the agents are self interested (they are only concerned with the items that are allocated to them) $v_i(x_i(\vec{\sigma}))$ is sufficient to describe player's *i* valuation. Furthermore, since, by the above definition, the player's utility for money equals its monetary value, the valuation function becomes equivalent to the player's type. To simplify notation we will use v_i to denote player *i*'s valuation for being allocated the object.

Valuations, as types, are not always known to all players and therefore introduce a certain level of uncertainty into the game. In the auction context we distinguish between three types of uncertainty. The most familiar situation, in which each bidder fully knows her type but is not certain about the types of others is referred to as *private values*. It is also often the case that the bidder does not fully know the value of the object, in this situation any additional information, such as the information possessed by other bidders, may be used to refine the expected value of the object. Such setup is called *interdependent values*. A subset of *interdependent values* auctions called *common values* are auctions in which the values of all bidders for the item are identical although unknown at the time of the bidding. This can occur if the value of the auctioned object is dictated by the market price [15] for example.

We now turn to the second component of the utility function, the value of money. As was noted in Section 2.2.4, one of the ways to escape the conclusions of Gibbard-Satterthwaite Impossibility Theorem is to restrict the *type space*. The assumption of *quasi-linear* form of the utility functions allows us to do that.

We assume the quasi-linear form of utility functions throughout this document.

The last basic auction concept to be defined is the *reserve price*. There are a number of ways a mechanism designer can restrict the bidding behavior of the players. One often placed restriction is the *reserve price*.

Definition 2.19 The *reserve price* is the minimum amount below which no bids are allowed.

Next we will outline some of the most famous auction mechanisms.

2.3.2 Classic single-object auctions

The term auction usually brings up the image of a room with an auctioneer and a number of people calling increasingly higher prices for some item. This type of auction is referred to as *Open Ascending Price* or *English* auction. The rules of the *English* auction dictate that the players can openly announce prices, which must always be higher than the last announced price. The auction ends when no one is willing to announce a higher price. At this point the *highest bidder* (the last player to call out a price) receives the item and pays her bid.

Another open price auction, called *Open Descending Price* or *Dutch* auction, operates counter to the *English* auction. There, the auctioneer initially announces a high price, and then proceeds to lower it until one of the participants reports acceptance. At this point this individual receives the item and pays the price equal to the last announcement made by the auctioneer.

So far we have outlined the classical *open bid* auctions. Auction literature also discusses two classical *sealed bid* auctions. In a *sealed bid* auction the bidders must submit their bids in sealed envelopes, with the allocation determined based on all submitted bids. In practice, the most often encountered *sealed bid* auction is the *First Price Sealed-bid auction*. Under the rules of this auction the item is awarded to the highest bidder for the price of her bid. A no less theoretically famous, although not used as much in practice, is the *Second Price Sealed-bid auction* or *Vickrey Auction* [33]. The rules of the *Vickrey Auction* state that the item is also awarded to the highest bidder but for the price equal to the *second highest* bid.

Next we will look at the equilibriums of the above auctions. It turns out that the Dutch auction is strategically equivalent to the First Price Sealed-bid auction, that is for every strategy in one auction there exists an outcome equivalent strategy in the other auction. There also exists an equivalence between the English auction and the Second Price Sealed-bid auction, although a much weaker one. Under the assumption of *private values* the optimal strategies in two auctions are equivalent [15]. Due to these types of similarities we will only look at the equilibrium strategies of the sealed-bid auctions.

Informally, the equilibrium strategy of any player *i* with valuation v_i in the *First Price Sealed-bid auction* is to bid the highest expected valuation of her opponents, assuming that all their valuations are below v_i. This type of strategy constitutes a best response to other players actions if all players follow the same strategy, which implies that the equilibrium of the First Price Sealed-bid auction belongs to the *Bayes-Nash* family of equilibriums (Definition 2.10).

The *Second Price Sealed-bid auction* has a very intuitive *Dominant strategy* (Definition 2.11) equilibrium. It is weakly dominant for every player to bid her valuation. To see why this is the case note that any deviation from this strategy by any player *i* will result in an outcome with a lower or the same utility for *i* independent of the other players' actions. Also note that this auction is *direct* and *incentive compatible*.

The equilibrium strategy profiles of the above two auction mechanisms appear quite different. However, under certain assumptions, these two auctions, have the same expected revenue. Even more generally, it can be shown that a large class of auctions called *standard auctions* have equal expected revenue for the seller; this result is called *the revenue equivalence principle* [15].

Definition 2.20 An auction is *standard* if the auction's rules state that the person bidding the highest amount is awarded the object.

Proposition 2.3 (*The revenue equivalence principle*) Assume that the players' values are independently and identically distributed and all players are risk neutral. Under these assumptions any symmetric and increasing equilibrium of any *standard* auction, in which the expected payment of a player with valuation zero is zero, results in the same expected payment for the seller.

In the above definition a player is *risk neutral* if she is indifferent between a lottery with some expected utility and an action which brings the same utility but with certainty. A *symmetric and increasing equilibrium* simply means that in equilibrium all bidders with the same valuation and beliefs should bid the same amount, and a bidder with higher valuation should bid at least as much as a bidder with lower valuation.

2.3.3 Optimal single-object auctions

In section 2.2.1 we have reviewed a number of *social choice* functions often used as goals in designing mechanisms. In auction mechanism design the two most desired goals are *efficiency* and *seller's revenue maximization*. The goal of *efficiency* is described by Definition 2.5 and the goal of *seller's revenue maximization* or simply *revenue optimization* falls under Definition 2.6. Under the assumption of *private values* we know of two classic results showing the existence of an auction optimal in *social welfare* and an auction optimal in *revenue*.

The auction achieving *optimal social welfare* is the already familiar *Second Price Sealed-bid* or *Vickrey* auction [33]. As was described in the previous section, for each player the dominant strategy in this auction is to bid her true valuation. The auction is also *standard*, it always allocates the good to the highest bidder. This implies that the person receiving the good is always the person who values it the most, therefore the *social welfare* is always maximized.

The auction developed by R. Myerson [22], appropriately named the *Myerson auction*, attains the optimal seller's revenue. The following is a brief outline of this auction mechanism. The operation of this mechanism is based on the concept of *virtual valuation*, with virtual valuation of player *i* defined as

$$\omega_i(v_i) = v_i - \frac{1 - \Phi_i(v_i)}{\phi_i(v_i)}. \tag{2.6}$$

where Φ_i denotes the cumulative distribution function over player *i*'s valuation, $\phi_i = \Phi_i'$ is the density function.

Assume that the seller has a certain valuation v_s for the object being sold. Each buyer *i* submits her bid to the mechanism, by stating some valuation v_i. The mechanism then computes the maximum *virtual* valuation $\omega^* = \max_i \omega_i(v_i)$ over all bids and finds the buyer $i^* = \arg\max_i \omega_i(v_i)$ with this *virtual* valuation.

If $v_s > \omega^*$ the seller retains the good, otherwise player i^* receives the good and pays the minimum valuation τ^* that would still make her the winner.

$$\tau^* = inf\{v_i | \omega_i(v_i) \geq v_s \text{ and } \omega_i(v_i) \geq \omega_j(v_j), \forall j \neq i\} \tag{2.7}$$

The *Myerson auction* has a dominant strategy equilibrium, where each buyer truthfully reveals her valuation.

There is a strong connection between the *Myerson auction* and the *Vickrey auction*. Under the assumption that all bidders are symmetric (valuations are IID) and all $\omega_i(\cdot)$ are strictly increasing the *Myerson auction* reduces to the *Vickrey auction* where the seller also submits a bid equal to $\omega_i^{-1}(v_s)$ [22]. This type of *Vickrey auction* is equivalent to a *Vickrey auction* with the optimal *reserve price*. Note that under these conditions the *Vickrey auction* is no longer guaranteed to be efficient [22].

2.3.4 Multi-unit and combinatorial auctions

In this section we will briefly define multi-item and combinatorial auctions. Although these auctions are not the focus of this document, we do refer to literature [25] pertaining to these types of auctions.

In the previous sections we have explored the auctions designed to allocate a single item. A much more complicated problem arises when the task is the allocation of multiple items, especially if they have some correlations between them. For example the items could be identical, but the buyers' utility functions for these items might exhibit *diminishing marginal utility*, that is after some point each additional item will bring less utility than the previous one. Alternatively the auctioned items could be different but some of these items might be *substitutes* which implies that the possession of both item 1 and item 2 provides no more utility then the possession of a one of these items. They might also be *complements* which means that the utility of having two (or more) complementary items simultaneously exceeds the sum of having them individually. One can come up with a lot more possible correlations, stemming from anything such as economic incentives to personal preferences.

The difficulty in designing a multi-item auction arises from the fact that the buyers now have a much larger strategy space, which usually makes it much more difficult to design protocols with easy to find equilibrium points, especially protocols with dominant strategy equilibria.

A special case of a multi-item auction is a *combinatorial auction*. In a combinatorial auction a buyer is allowed to place a bid on a collection of goods, as opposed to being restricted to always placing one bid per good. In some implementations the buyer is allowed to place multiple bids, on various (possibly intersecting) sets of goods. The auction mechanism allocates the goods based on the buyers' bids, assigning each object to at most one bidder, so as to optimize some social choice function.

Implementing combinatorial auctions was shown to be a difficult task. The most prominent problem is the selection of the winning bids based on all players' submissions. This problem has been shown to be NP-hard [10]. The other, no less difficult problem, is transmitting all the players' bids to the auctioneer, considering that there may be exponentially many of them, if we allow multiple bids on intersecting sets of goods. A number of solutions to these problems have been proposed [4, 6, 23], these and related problems are currently a very active area of research.

2.4 Limited communication auctions

In classic auction design (Section 2.3) it is commonly assumed that each player is capable of fully expressing her type or valuation and if necessary transmitting it to the mechanism. This assumption is often unrealistic for various reasons, some of which include bandwidth bounds on communication, players' inability (due to computational constraints, for example) or unwillingness to fully describe their types, etc. The primary focus of this document is on dealing with a subset of mechanisms embracing this restriction. In this section we will explore related literature on the same subject.

2.4.1 Priority games

This section is dedicated to exploring two papers, *Auctions with severely bounded communication* by L. Blumrosen and N. Nisan [1] and its extension *Multi-Player and Multi-Round Auctions with severely bounded communication* by L. Blumrosen, N. Nisan and I. Segal [2]. These papers deal with the problem of designing efficient one-shot (single iteration) auctions that are restricted to using a very small amount of communication. The first paper tries to solve this problem in the context of 2 player games, the second provides an extension to n players. The topic explored in these papers is very closely related to the topic of this document, for this reason we will review the papers in significant detail, starting with [1].

Assumptions and Definitions. The players are assumed to be risk-neutral, with independent – but not necessarily identically distributed – private values, and quasi-linear utilities. The valuation of every player *i* is assumed to be distributed in the [0,1] interval according to a commonly known probability function ϕ_i. Furthermore, all the mechanisms described in the paper are *ex-post Individually Rational (IR)* (utility of at least 0 is guaranteed to each player). To describe her bid each player *i* is allowed to send $t_i = lg(k_i)$ bits, that is, she must chose one of k_i possible messages. The set of possible bids of player *i* is denoted $B_i = 0, 1, \ldots, k_{i-1}$, $b = b_1, \ldots, b_n$ is the vector of players' reported bids.

Definition 2.21 $\mathcal{M}_{n,k}$ denotes the set of all possible n-player mechanisms where each player must choose one among k ($k = |B_i| \, \forall i$) possible bids.

Definition 2.22 $\sigma_i : [0, 1] \to \{0, 1, \cdots, k_{i-1}\}$ is a strategy for player i in a game $m \in \mathcal{M}_{n,k}$.

Definition 2.23 Let a real vector $c = (c_1, c_2, \ldots, c_k)$ be a vector of *threshold-values* if $c_1 \leq c_2 \leq \ldots \leq c_k$. A strategy σ_i is a *threshold-strategy* based on vector of threshold values c if $c_0 = 0$, $c_k = 1$ and for every $c_j \leq v_i \leq c_{j+1}$, $\sigma_i(v_i) = j$.

Informally, Definition 2.23 states that a *threshold-strategy* is based on some predefined vector of *threshold-values* c. When a player becomes aware of her valuation v_i, she determines where v_i falls within c and reports this position.

The primary goal of the paper is to show how close, asymptotically, limited communication auctions can come to the auctions with unlimited communication. The auctions are compared to mechanisms optimal in social welfare and to mechanisms optimal in revenue. In order to do this, we need the concept of an optimal limited communication auction. Let $w^{opt}_{n,k}$ be the maximum possible expected welfare in any *n*-player game with any vector of strategies where each player has *k* possible bids. Let $r^{opt}_{n,k}$ be the maximum possible expected revenue in any *n*-player game with any vector of strategies where each player has *k* possible bids.

The results of the paper are based on showing that the two priority games defined below are in fact $w^{opt}_{2,k}$ and $r^{opt}_{2,k}$. We will first define the mechanisms and then proceed to the actual results.

Definition 2.24 A 2-player *priority-game* $PG_k(x, y) \in \mathcal{M}_{2,k}$ is a *standard* auction that allocates the item to the highest bidder, with ties (the bids of two players are equal) broken according to some predefined order. $x = (x_1, x_2, \ldots, x_k)$ and $y = (y_1, y_2, \ldots, y_k)$ are the vectors of threshold value assigned to two participating players *A* and *B*, respectively. Based on the pair of bids (i, j) from *A* and *B*, *A* pays x_{j+1} whenever she wins and *B* pays y_i whenever he wins.

Definition 2.25 A 2-player *modified priority-game* $MPG_k(x, y) \in \mathcal{M}_{2,k}$ has the same allocation as described in Definition 2.24 except no allocation is performed when all players bid 0. Using the same vector of threshold values x, y and a pair of bids (i, j), *A* pays x_{j+1} whenever she wins and whenever *B* wins he pays y_i when $i > 0$ and y_1 when $i = 0$.

Both in $PG_k(x, y)$ and in $MPG_k(x, y)$ player *B* is preferred to *A*, that is in case of a tie the object is allocated to player *B*.

Note that in both mechanisms the payment of the winning player is determined by the action of the opponent. Furthermore, the payment of player $C \in \{A, B\}$ can never be greater than c_i if she reports *i* and her threshold vector is c. This demonstrates that the presented *threshold games* have certain similarity to *Second Price Sealed-Bid Auction*.

Proposition 2.4 *(Proposition 4.1)* For every pair of threshold vectors x, y, the threshold strategies based on these threshold vectors are dominant in both $PG_k(x, y)$ and $MPG_k(x, y)$, and these mechanisms are ex-post individually rational.

We can now proceed to the main result. The mechanism $PG_k(x^w, y^w)$ is $w^{opt}_{2,k}$, that is it achieves optimal expected welfare among all 2 player *IR* limited communication mechanisms ($\mathcal{M}_{n,2}$), and incurs a loss of $O(\frac{1}{k^2})$ compared to the welfare optimal mechanism without communication bounds. Similarly, the mechanism $MPG_k(x^r, y^r)$ is $r^{opt}_{2,k}$, it achieves optimal expected revenue among all $\mathcal{M}_{n,2}$ mechanisms, and incurs a revenue loss of $O(\frac{1}{k^2})$ compared to the revenue optimal mechanism without communication bounds. To achieve optimality the threshold vectors x^w, y^w for players *A* and *B* must be *mutually-centered*, which is defined as $x^w = (x_1^w = 0, x_2^w, \ldots, x_{k-1}^w, x_k^w = 1)$, $y^w = (y_1^w = 0, y_2^w, \ldots, y_{k-1}^w, y_k^w = 1)$, $\forall 2 \leq i \leq k-1, x_i = E(v_B | y_{i-1} \leq v_B \leq y_i)$ and $\forall 2 \leq i \leq k-1, y_i = E(v_A | x_i \leq v_A \leq x_{i+1})$. For $MPG_k(x^r, y^r)$, x^r, y^r are defined similarly but based on virtual valuations (Equation 2.6).

The above shows that given a space of limited communication 2-player single shot auctions, where each bidder is restricted to *k* possible messages, it is possible to find an auction mechanism with dominant strategies that will be $O(\frac{1}{k^2})$ far in expected welfare from the optimal (with unlimited communication); and this is the best that can be done. Similarly for expected revenue. This is a very powerful result as it fully defines what can be possibly achieved with 2-player single shot mechanisms under limited communication.

As an aside note, that all the optimal mechanisms described above are asymmetric, since bidder *B* is preferred to bidder *A*. Blumrosen and Nisan prove that symmetric mechanisms incur greater loss in welfare and revenue than the asymmetric ones.

In their second paper [2] Blumrosen, Nisan and Segal extend the above results to multiple players. As in [1] they show that *priority games* and modified priority games are $O(\frac{1}{k^2})$ far from optimal with unrestricted communication in welfare and revenue (respectively). However, it remains unknown whether these mechanisms are optimal among all *n*-player limited communication mechanisms. Note that the bound $O(\frac{1}{k^2})$ holds only for a fixed *n*; with variable *n*, the upper bound grows at least *exponentially* with *n*.

The paper also presents a result concerning *multi-round* auctions. It shows that with a specific class of multi-round auction mechanisms it is possible to achieve better results than with single-shot, but the extra gain is limited. More specifically, they show that a multi-round mechanism (where the players send their bits one at a time, in alternating order, and have perfect knowledge about the game history) can achieve the same welfare/revenue as a single-shot priority game but with reduction is communication, bounded by a factor of 2.

2.4.2 Ascending price auctions and iBundle auction

This section presents *iBundle*, an efficient ascending price bundle auction developed by David Parkes [25]. *iBundle* is an iterative ascending combinatorial auction that guaranties an optimal bundle allocation with a best-response agent bidding strategy.

Definition 2.26 An auction is an *ascending price auction* if for any iterations t, k, s.t. $t < k$ and any item or bundle b, it is the case that if the asking price for b at iteration k is p then p is greater than or equal to the asking price at iteration t. Here the asking price is the price a winning bidder would pay for b.

iBundle operates roughly as following (for the complete description see [25]). The bidders are allowed to submit their bids in one of two forms ***OR*** or ***XOR***. An ***OR*** bid, $B_i^{OR} = \{(S_1, P_1), (S_2, P_2)\}$, implies that *i* is willing to buy the bundle S_1 for P_1, S_2 for P_2, or S_1 and S_2 for $P_1 + P_2$. An ***XOR*** bid, $B_i^{XOR} = \{(S_1, P_1), (S_2, P_2)\}$, implies that *i* is willing to buy only one of S_1 or S_2, for the given price. The two types of bids are used to express both additive utilities ($u(P_1) + u(P_2) = u(P_1, P_2)$) and non-additive utilities ($u(P_1) + u(P_2) \neq u(P_1, P_2)$).

At the beginning of each iteration the mechanism announces the *ask* prices for some subset of bundles. The *ask* prices serve as reserve prices, so that lower bids will not be accepted (with one exception, an agent is allowed to bid lower if it was assigned the bundle in the provisional allocation). Given all the players' bids, the mechanism generates a *provisional allocation*, by allocating the bundles so as to maximize revenue, at the same time making sure that the allocation is feasible and consistent with the agents' bids. The provisional allocation is publicly announced. At the next iteration the mechanism provides a new set of *ask* prices. The *ask* price for a bundle is at least the price at which this bundle is *provisionally allocated*. It might also be higher if there exists an agent that desires this bundle for a relatively high price, but did not receive it in the provisional allocation. The auction terminates if the bid of every agent is satisfied, or if all the agents submit the same bids in two successive rounds. At this point the provisional allocation becomes the final allocation.

Parkes shows that the best-response bidding strategy for an agent with non-additive utilities is $B_i^{XOR} = \{(S, P_i(S)) | v_i(S) - P_i(S) \geq 0, v_i(S) - P_i(S) + \epsilon \geq \max_{S'}\{v_i(S') - P_i(S')\}\}$ where ϵ is the minimum bid increment, similarly for an agent with additive utility the strategy is $B_i^{OR} = \{(S, P_i(S)) | v_i(S) - P_i(S) \geq 0\}$. Given that the players follow the above strategies the *iBundle* mechanism will generate an optimal allocation. The drawback of this result is that the strategies are myopic. That is, they do not take into account strategic moves of other players, and therefore do not constitute a game theoretic equilibrium. This sort of approach is different from the one we use in this document. Nevertheless, *iBundle* mechanism is important to the results presented in this document, as it demonstrates that an ascending iterative auction can be made efficient even though the amount of information revealed by the players is limited (players do not have to submit bids on every bundle).

In his other work [27, 26] Parkes provides further evidence to the benefit of using *ascending auctions* when it is difficult for an agent to fully determine her valuation or when the task of preference elicitation is costly. This goes in parallel with our results as we show that a generalization of the ascending price auctions, the *increasing price auctions* are beneficial under these and related settings.

2.4.3 Other relevant research

Here we will briefly outline a few other papers exploring similar topics.

Sandholm and Gilpin [32] develop a *TLA* (take-it-or-leave-it) mechanism with the goal of reducing valuation revelation in a one-good auction. *TLA* announces a sequence of offers to all sellers, where each offer specifies the amount and the player to whom the offer is made. The player receiving the offer has a choice of refusing, in which case the next offer is made, or accepting, by paying the requested amount. The paper shows that with multiple players there exists a *Perfect Bayes Nash* equilibrium for this game, it also provides an algorithm for calculating the equilibrium strategies. Furthermore, they show that with the optimized offers *TLA* comes close to the optimal revenue auction [22] and with two bidders the optimal *TLA* incurs the same loss as the optimal *priority auction* [1].

Grigorieva, Herings, Müller and Vermeulen present the *bisection* auction. The *bisection* auction is an iterative version of the *Vickrey* auction with a binary search type price update rule. The auction starts with an initial belief that all players' valuations are located in some interval $[a, b]$, and announces a price equal to the middle of that interval. The players have the option of saying *yes* or *no*. The auction adjusts it's belief interval according to the answers it receives, if all players state *no* the belief interval is changed to $[\frac{a+b}{2}, b]$. If there are at least 2 players stating *yes* it becomes $[a, \frac{a+b}{2}]$ with the players saying *no* removed from the auction. When there is exactly one player stating *yes* this player is declared winner (although this information is kept secret by the auctioneer), and the auction enters the *price-determination phase*. During this phase the auction determines, up to *maximum* allowed precision, the valuation of the *second highest* bidder for the good. Once that is accomplished the auction terminates and allocates the good to the highest bidder for the price equal to that valuation. They show that this mechanism has a truth revealing dominant strategy equilibrium where a player states *yes* if the current price is at least her valuation. They also demonstrate that the *bisection auction* is computationally more efficient than the *English* auction. However, it is doubtful this mechanism will ever be implemented given some of it's drawbacks. Firstly, there is a considerable amount of communication after the winner has already been determined. Secondly, also due to the *price-determination phase*, the losing players are required to reveal a significant amount of information to the mechanism, which is often undesirable from the players' point of view. Removing the *price-determination phase* will however result in the *bisection* auction losing it's dominant strategy equilibrium.

2.5 Markov Decision Processes

Although not the focus of this document *Markov Decision Processes* (MDPs) are used to model best response price policies in Chapter 4. Here we will provide a very short and limited review of MDPs, for a much more extensive review please refer to [5].

In the context of planning under uncertainty, *Markov Decision Processes* and their extensions have in recent years become a popular way to model the planning problems [5, 3, 13] (and many other examples). MDPs allow incorporating into the model different rewards for being in various world states, costly actions, uncertain action effects. Extensions of MDPs such as POMDPs (Partially Observable Markov Decision Processes) further generalize this model by allowing uncertainty in observations. Solution methods for MDPs and POMDPs can, in theory, find an optimal plan that maximizes the user's reward. This plan can be determined either for a finite number of steps also called *finite-horizon*, or for an *infinite horizon*. In this section we will explore an MDP formulation and a basic solution method for finding policies under *finite horizon* assumption.

2.5.1 Formulation

We formulate the MDP model as follows. The system is assumed to always be in one of the predefined world states. A state constitutes a complete description of the system at some point in time, where this description provides sufficient information for the decision maker to act on. We denote the set of system states by $\mathcal{S}$, assume that $\mathcal{S}$ is finite, and use the random variable S^t to denote the system's state at time *t*.

Since the goal of this model is to encapsulate a changing system, we view it as changing state at each time period. To describe this we need some way to predict where the system will move given the current state. We use a probability distribution $\Pr(S^t|S^{t-1}, \cdots, S^0)$ to formalize the chance of moving to state S^t given the prior history $S^{t-1}, \cdots, S^0$. To simplify this model, it is assumed that the current state always holds sufficient information to calculate the probability of the next state, this is the *Markov assumption*. Given this assumption the above equation reduces to $\Pr(S^t|S^{t-1})$. Most generally a system also requires a probability distribution over the initial states, for the purposes of this document the system will always have one starting state.

The above transition model by itself does not incorporate the fact that an agent is capable of performing an action that can alter the state of the system. To extend the model we allow each agent *i* to chose an action a_i at time *t* from a *feasible* set A_i^t. To capture the effect of this action we will use a *transition function* $T(s^{t-1}, a_i, s^t) = \Pr(S^t = s^t | S^{t-1} = s^{t-1}, A_i^t = a_i)$ to denote the probability that the system transitions from state s^{t-1} to state s^t when the action a_i is executed.

The *reward* and *cost* functions are used to specify the preferences and costs for an agent. The reward function $R : \mathcal{S} \rightarrow \Re$ specifies how much an agent gains from being in a state. Similarly, the cost function $C : \mathcal{S} \times A \rightarrow \Re$ specifies the cost of performing an action in some state.

2.5.2 Value of a policy

Given a specification for an MDP, an agent would like to determine her best action in each system state. The agent's strategy, a function mapping the current state and stage into an action is called a *policy*, denoted by π. The goal of the agent is to find the best policy that maximizes her utility as determined by the reward and cost functions. To compare the various policies one needs to compute the *value* of a policy. For *finite-horizon* problems an expected value of any state at stage (iteration) $t < T$ can be computed by

$$V_t^\pi(s) = R(s) + C(\pi(s,t)) + \sum_{s' \in S} \Pr(s' | \pi(s,t), s)\, V_{t+1}^\pi(s') \tag{2.8}$$

A policy π is optimal among all T-stage policies if $\forall s, \forall \pi',\ V_1^\pi(s) \geq V_1^{\pi'}(s)$. Next we discuss a dynamic programming algorithm that allows computing approximately optimal policies.

2.5.3 Computing a policy

The fundamental algorithm for computing the *finite horizon* policies, called *Value Iteration*, is based on iteratively calculating Equation 2.8. At stage T (the last stage of the system) let

$V_T^*(s) = R(s)$. Then for each stage $t \leq T$ (in sequence) compute $V_t^*(s)$ as follows

$$V_t^\pi(s) = R(s) + \max_a \{ C(a) + \sum_{s' \in S} \Pr(s'|a,s) \, V_{t+1}^\pi(s') \} \tag{2.9}$$

As the above values are computed it is easy to construct a corresponding policy.

Chapter 3

One-shot Limited Precision Mechanism

This chapter introduces the concept and provides motivation for studying the limited precision auctions. We then proceed to develop a simple one-shot mechanism that serves as an introduction to the main results of this document.

3.1 Limited precision

In the previous chapter we have reviewed a number of classic auctions, some of which are provably optimal. We have also outlined the concept of *revelation principle* which states that the attention of the mechanism designer can be restricted to *direct, incentive compatible* mechanisms (Section 2.2.3). This result is not completely general. The problem with the *revelation principle* is that it assumes that the players are capable and are willing to fully report their valuations to the mechanism.

One can imagine a number of social interactions where this assumption would not hold. For example, in an auction of a natural resource, a player may be unable to fully determine her valuation without a costly computation. Similarly, in a repeated type of interaction, a player may not be willing to fully reveal her valuation due to strategic reasons. These limitations imply that the auction mechanisms must be able to deal with reduced communication, or more generally with reduced *precision*. A *precision* of a message is said to be limited if the message alone does not fully describe a player's valuation or type. In the previous chapter (Section 2.4) we have presented a number of current results dealing with this problem. For a more detailed description of the limitations of revelations principle under the condition of limited precision we refer the reader to [9].

Next we will present an example of a simple single-shot limited precision mechanism.

3.2 Limited precision TIOLI mechanism

In this section we will define a one-shot (single iteration) limited precision take-it-or-leave-it (LP-TIOLI) mechanism. We would like to note outright that this mechanism is in many ways suboptimal to the already known mechanisms. Our goal here is to show that it is possible to build an intuitive single-shot mechanism with some desirable properties, this mechanism will also serve as a stepping stone to the next chapter.

3.2.1 Desired form

As was illustrated above we wish to design a mechanism that would be able to perform under the following three restrictions:

- Limited revelation
- Bounded computation
- Limited communication

The first point refers to willingness of the players to reveal the information about their type. Ideally, the mechanism should behave in such a way that the players are asked to reveal information only if absolutely necessary. *Bounded computation* refers to the difficulty players face when computing their own type; again the mechanism should only force the players to compute if absolutely necessary. The last point implies a direct limitation on the amount of communication that can take place between the mechanism and the players.

The last objective can easily be achieved by limiting the *action space*, the space of legal messages or bids. Fortunately, the same limitation can also help satisfy the other two objectives, as the players will now only have to compute or reveal up to one of the actions or messages of the action space.

We limit the space of mechanisms of interest to those that are intuitive and have a dominant strategy equilibrium. Many of the currently known mechanisms rely on a set of complicated or unintuitive rules, the *Myerson* auction [22] being the most prominent example, but arguably so are the *priority games*. The one widely known optimal auction which does have an intuitive set of rules and an easily provable dominant strategy equilibrium is the *Vickrey* auction [33]. The *Vickrey* auction is known to be optimal in welfare, but under a certain of conditions the *Vickrey* auction with a reserve price is provably optimal in revenue [15].

We aim to design an *incentive compatible* and *symmetric* single object one-shot mechanism with an intuitive set of rules and an easily provable dominant strategy equilibrium. Under these requirements it is a given that this mechanism can never be optimal in revenue or social welfare, due to its *symmetric* structure [1]. However, we would like to be able to optimize it for any objective including Social welfare and Seller's revenue.

As was stated in the previous chapter we will assume that the utility functions are quasi-linear in money, the players are individually rational and we have some prior distribution $\Phi(x)$ over the players' valuations.

3.2.2 LP-TIOLI

The following is a description of the proposed mechanism that satisfies the above requirements.

Limit the *action space* by fixing a set of k price thresholds, $0 \leq p_1 < p_2 < \ldots < p_{k-1} < p_k = 1$. Each bidder announces one of these prices to the mechanism. If only one bidder announces a unique highest price, the good is allocated to this bidder at the second highest bid (from among the k possible bids). If the highest bid is offered by two or more bidders (a condition called a *tie*), a random highest bidder is selected and made a take-it-or-leave-it

(TIOLI) offer for the good at the highest bid price. If the offer is rejected, the same offer is made to another, randomly chosen, highest bidder. If the offer is rejected by all, the good is not allocated (retained by the seller).

Given this structure, a player in *LP-TIOLI* auction with k price levels has to send $\log(k)$ bits in order to describe her bid; an amount of communication similar to the *priority games* (Section 2.4.1).

LP-TIOLI is a limited precision variant of the Vickrey auction, with the distinction that under some conditions the good is not allocated. Not surprisingly, the dominant strategy equilibrium of *LP-TIOLI* is also reminiscent of Vickrey auction. In addition the dominant strategy holds for any set of threshold prices, a feature we will take advantage of in the next section.

Proposition 3.1 In *LP-TIOLI* it is a weakly dominant strategy of any player *i* with valuation v_i to bid the *least* price $p \in \{p_1, \ldots, p_{k-1}, p_k = 1\}$ which is at least as great as v_i. Formally, $\sigma_i(v_i) = \arg\min_p \{p \in \{p_1, \ldots, p_{k-1}, p_k = 1\} | p \geq v_i\}$. In case of a tie it is dominant to always refuse the offer.

Proof: See Proposition A.1 in Appendix A.

3.2.3 Setting the price thresholds

In the previous section it was shown that *LP-TIOLI* has a dominant strategy equilibrium. An examination of the proof of that result reveals that the equilibrium holds for any set of threshold prices. Combining this fact with the prior knowledge of the probability distribution over players' valuations we get an opportunity for optimization. That is, it is possible to optimize the price thresholds based on the prior information to attain any social objective, without losing the benefit of dominant strategy. This approach to optimization is closely related to automated mechanism design (AMD) [8] in the sense that we wish to optimize the mechanism using specific distributional information; however, we restrict our attention to a class of mechanisms with specific parameterized dominant strategies, and simply optimize the parameters, rather than leaving the whole mechanism "up for grabs." The optimization task need not be limited to the traditional objectives, but can also incorporate "mechanism properties" such as expected amount of communication. In the next chapter on incremental mechanisms, we will study optimization with respect to welfare that accounts for the cost of communication.

To illustrate the potential for optimization we will provide two examples. In the following we will optimize the price thresholds with the objective of maximizing the seller's revenue, and in the second example, minimizing the probability of not allocating the good.

To simplify the analysis we assume that all bidders' valuations are independently and identically distributed (IID). Let $\Phi(x)$ be the prior (cumulative) probability distribution of a single bidder's valuation, let $\phi(x) = \Phi'(x)$ be the density function. The goal is to find a set of k price thresholds that optimize the objective (supposing that the communication is limited to $\log(k)$ bits per player).

Expected Seller's Revenue

The expected revenue of *LP-TIOLI* auction can be written as follows; to simplify presentation it is introduced in three steps.

The first sub-formula is the probability that given any $n-1$ bidders, all of them have valuation lower than or equal to p_i and at least one bidder has valuation between p_i and p_{i-1}

$$Prob(p_{i-1}, p_i) = \sum_{z=1}^{n-1} \binom{n-1}{z} * (\Phi(p_i) - \Phi(p_{i-1}))^z * \Phi(p_{i-1})^{n-1-z} \tag{3.1}$$

The expected revenue from a special case when the valuations of any $n-1$ players are below p_1.

$$ER^{SC}(p_1) = (1 - \Phi(p_1)) * \Phi(p_1)^{n-1} * p_1 * n \tag{3.2}$$

The complete expression for the expected revenue of a *LP-TIOLI* auctions with n bidders.

$$ER_n^L(p_1, ..., p_k) = ER^{SC}(p_1) + \sum_{i=2}^{k} Prob(p_{i-1}, p_i) * (1 - \Phi(p_i)) * n * p_i \tag{3.3}$$

There are a number of ways in which one can optimize the prices, given the above equations. We have employed a gradient ascent approach; to apply it, it was necessary to differentiate Equation 3.3 with respect to $p_1, p_2, \cdots, p_{n-1}$, a routine exercise, we omit the derivation.

The results demonstrating the expected revenue of *LP-TIOLI* with optimized prices are presented in section 3.2.5.

Probability of not allocating

Using a similar approach we have derived an expression for the probability of not allocating the good.

$$
\begin{aligned}
&probNotAllocated(p_1, ..., p_k) = \\
&\left(\sum_{i=1}^{k-1} \sum_{z=2}^{n} \binom{n}{z} * (\Phi(p_i) - \Phi(p_{i-1}))^z * \Phi(p_{i+1})^{n-z} \right) + \Phi(p_1)^n
\end{aligned}
\tag{3.4}
$$

There are two main reasons why this particular social choice function was used. Firstly, it shows that it is possible to optimize *LP-TIOLI* for objectives other than the traditional *social welfare* and *seller's revenue*. Secondly, the probability of not allocating constitutes an upper bound for the *expected loss* in social welfare. To see this, note that whenever *LP-TIOLI* allocates the good, it allocates it to the player who wants it the most, thus maximizing social welfare. In the cases that it does not allocate the good, the society incurs a loss in social welfare equal to the highest valuation among all participating players, this value is always less than 1.

Similarly to the previous example, we have used gradient descent to find the price thresholds minimizing Equation 3.4. Please refer to section 3.2.5 for the results.

3.2.4 2-bidder case

Under the restriction of two symmetric bidders LP-TIOLI displays some interesting properties.

If the prices are set with the goal of maximizing the seller's revenue, then *LP-TIOLI* auction becomes a discrete approximation of the *Vickrey auction with an optimal reserve price*, which is known to be optimal in revenue. More specifically, it must be the case that the lowest price threshold of *LP-TIOLI* is greater than or equal to the optimal reserve price of the *Vickrey* auction. In the limit, as the number of price thresholds increase, the expected revenue of *LP-TIOLI* approaches that of the *Vickrey* with an optimal reserve price. The following is the formal treatment of these claims.

The expected revenue of a standard 2-bidder *Vickrey* auction is

$$ER_2^V = \int_0^1 2 * x * (1 - \Phi(x)) * \phi(x) dx \tag{3.5}$$

The expected revenue of 2-bidder *LP-TIOLI* with price thresholds $0 \leq p_1 < p_2 < \ldots < p_{k-1} < p_k = 1$ is

$$\begin{gathered} ER_2^L(p_1, ..., p_k) = \sum_{i=1}^{k} 2 * p_i * (1 - \Phi(p_i)) * (\Phi(p_i) - \Phi(p_{i-1})) \\ \text{letting } p_0 = 0 \text{ and } \Phi(p_0) = 0 \end{gathered} \tag{3.6}$$

Suppose that the threshold prices are positioned uniformly on the [0,1] interval, then equation 3.6 is a discrete approximation of equation 3.5 and $\lim_{k\to\infty} ER_2^L(p_1, ..., p_k) = ER_2^V$.

Take the common part of the above two equations.

$$G(x) = x * (1 - \Phi(x)) \tag{3.7}$$

Since $\Phi(x)$ is continuous, $\forall x \in [0, 1], 0 \leq \Phi(x) \leq 1$ and $G(0) = 0$, $G(1) = 0$ it is the case that there exists some $p_r \in (0, 1)$ such that p_r is the *smallest* number in $(0, 1)$ such that $G(p_r)$ is a global maximum of $G(x)$ on $[0, 1]$.

Lemma 3.1 For any distribution $\Phi(x)$ it is always possible to find the set of threshold prices for a 2-bidder *LP-TIOLI* such that $ER_2^L(p_1, p_2, ..., p_k) > ER_2^V$ where $p_1 = p_r$ and p_r is as defined above.

Proof: See Lemma A.1 in Appendix A.

The price p_r is actually nothing more than the optimal reserve price of the *Vickrey* auction. From the auction theory [15] it is known that the optimal reserve price r^* must satisfy the following necessary condition

$$r^* * \frac{\phi(r^*)}{(1 - \Phi(r^*))} = 1 \tag{3.8}$$

Taking equation 3.7 and differentiating it with respect to x, we obtain the following

$$(1 - \Phi(x)) - x * \phi(x) = 0 \tag{3.9}$$

which is equivalent to equation 3.8. By construction p_r maximizes $G(x)$. Therefore, if $\frac{\phi(x)}{1-\Phi(x)}$ is increasing on $[0, 1]$ (which is the case for the uniform distribution, for example) the condition 3.8 is also sufficient, which makes p_r the optimal reserve price of the *Vickrey* auction.

Lemma 3.2 To optimize the seller's revenue with 2-bidders, the lowest price threshold p_1 of the *LP-TIOLI* mechanism should be set to at least p_r (set p_1 such that $p_1 \geq p_r$).

Proof: See Lemma A.2 in Appendix A.

With the first price threshold set to p_r and the rest distributed uniformly over the remaining interval, if p_r satisfies the above sufficient condition then *LP-TIOLI* provides a discrete approximation of the *Vickrey* auction with an optimal reserve price and as the number of threshold levels increases, approaches it in expected revenue.

3.2.5 Empirical observations

In this section we present several graphs that compare the performance of *LP-TIOLI* to that of *Vickrey* auction with unrestricted communication. These results also demonstrate the increase in the performance of *LP-TIOLI* due to the optimization of the price thresholds.

As was stated in section 3.2.3 we have used gradient ascent and gradient descent to optimize prices for the objective of maximizing the seller's revenue and minimizing the probability of not allocating the good, respectively.

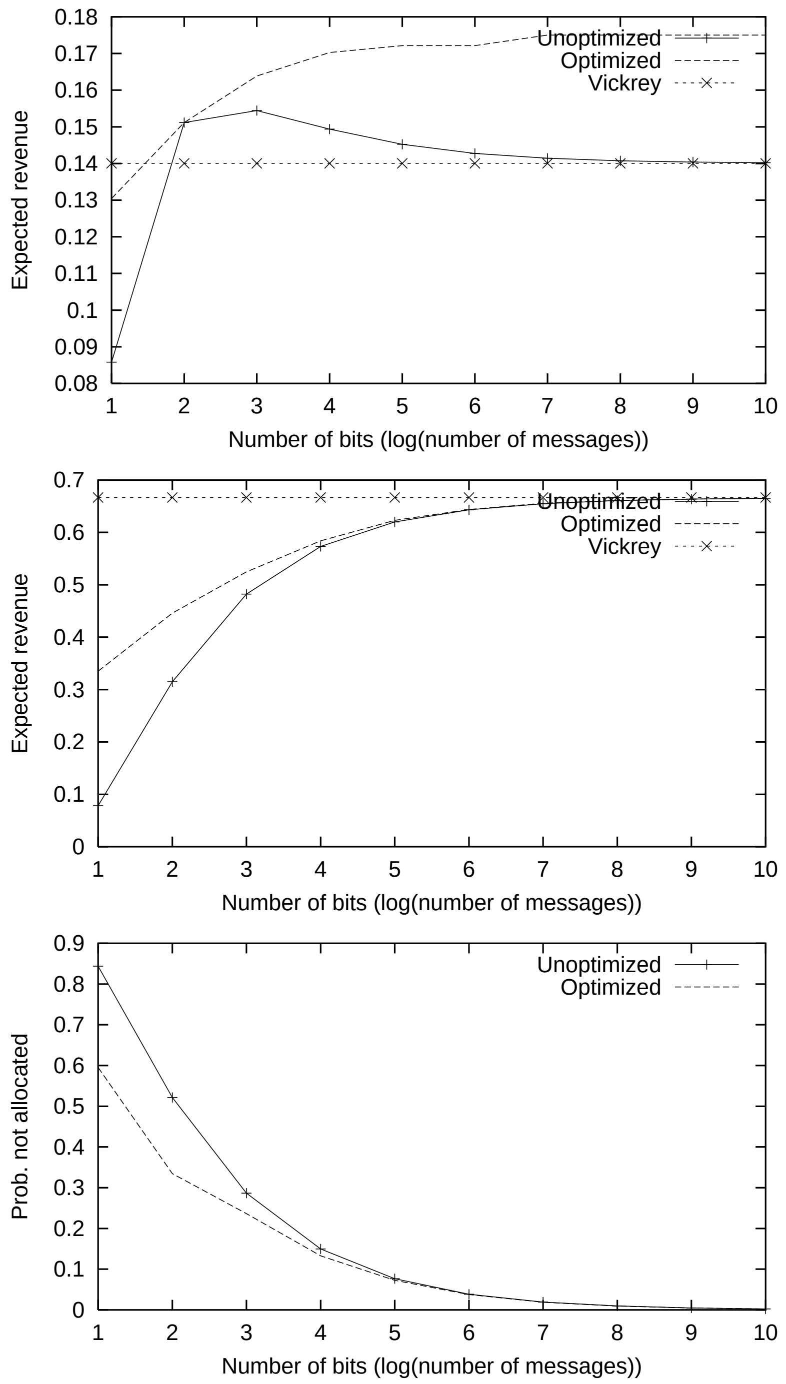

Figure 3.1: (a) Expected revenue (truncated half-Gaussian, $\mu = 0, \sigma = 0.3$); (b) Expected revenue (uniform distribution); (c) Probability of good unallocated (uniform distribution).

Figure 3.1(a) shows the expected revenue as a function of the number of bits sent by each player, with valuations drawn from a (normalized truncated) half-Gaussian density function. This graph was chosen based on two factors. First, it is the sort of distribution that can occur in a real example, since it is possible that most people have a low valuation for a good, except for a few. Second, this graph provides the connection to the results derived in section 3.2.4. The graph was plotted for the case of two bidders. The unoptimized prices (uniformly set at $1/k, 2/k, \ldots$) fare much worse than the optimized prices. As was mentioned above, this graph confirms the results of section 3.2.4. It shows that the expected revenue of *LP-TIOLI* with optimized prices strictly exceeds that of the standard *Vickrey* auction. One the other hand the expected revenue of *LP-TIOLI* with uniformly set price thresholds eventually approaches that of *Vickrey*.

Figure 3.1(b) shows the results from running *LP-TIOLI* with 5 bidders with uniformly distributed valuations. It also demonstrates that the optimized price thresholds allow greater expected revenue. However in the limit both the optimized and the unoptimized versions of *LP-TIOLI* approach the expected revenue of the *Vickrey* auction, this can be explained by the fact that with more players the probability of a tie (not allocating the good) outweighs any gains from discretizing the price space.

Figure 3.1(c) demonstrates the results from optimizing *LP-TIOLI* for the second social objective. Again, the optimized version performs better than the unoptimized one. The graph does not include an entry for *Vickrey* auction, since it always allocates the good.

Chapter 4

Incremental Elicitation Limited Precision Mechanisms

In chapter 3 we have stated a number of requirements and restrictions, and developed a simple single-shot auction that satisfies them. In this chapter we will place similar restrictions on the desired auction form; in particular we will search for a mechanism with dominant strategy equilibria, but in the setting of *incremental elicitation limited precision* auctions. We will show a number of necessary conditions that must hold in order for an auction to have these properties, and demonstrate how these conditions dictate the auction's structure. Next a number of examples of possible mechanisms will be provided as well as a way to optimize them for various social choice functions.

4.1 Incremental mechanisms

This section defines the incremental auctions and provides evidence of potential advantages these auctions have over the single-shot mechanisms.

Definition 4.1 An auction belongs to a class of *iterative mechanisms* if it may take the mechanism more than one *operational step* to obtain all the necessary information from the players to make a decision. An *operational step* is a time interval in which the players' and the mechanism's knowledge about the world remains constant.

Iterative or *multi-step* mechanisms have been shown to be significantly better than single-shot mechanisms in certain cases. For example, it can be shown that in some settings one-shot protocols can require exponentially more communication than iterative protocols [16]. We conjecture that iterative mechanisms will also be superior in the limited communication/revelation setting.

For the purposes of auctions, which allocate the good based on the utility information elicited from the participants, we propose a refinement of iterative mechanisms – the *incremental mechanisms*.

Definition 4.2 An *incremental* mechanism is an *iterative* mechanism which interprets each player's message as a (possibly partial) revelation of her type. An *incremental* mechanism restricts the players actions in such a way that in equilibrium at every iteration for each player the space of possible valuations is refined.

One of the main themes of this document is the search for mechanisms with dominant strategy equilibria. It is interesting to note that when a player acts according to her dominant strategy she has no choice but to reveal certain information about herself. Even if a dominant strategy direct a player to ”lie”, this would be known and expected by the mechanism. Therefore, the combination of an iterative mechanism and a dominant strategy equilibrium would fall under the above definition of an *incremental mechanism*.

Restricting the focus to incremental mechanisms allows the use of such concepts as *limited participation* (to be defined later) which naturally results in less revelation and therefore greater privacy and smaller computational costs for the players. In the following section we will review all the assumptions placed on the desired auction type.

4.2 Assumptions

We consider incremental mechanisms for the allocation of a single good. As with single-shot auctions (Section 3.2.1), these mechanisms have to operate under the restrictions of *Limited revelation*, *Bounded computation*, and *Limited communication*. We restrict the class of mechanisms by imposing a number of requirements that seem natural in the space of single-good auctions, as will be demonstrated.

4.2.1 Valuations and Utility functions

Assume that each agent *i* has valuation $v_i \in [0, 1]$. Further, assume that all agents have *quasi-linear* utility functions (see Definition 2.18). These are standard assumptions in auction literature.

4.2.2 Actions

At every iteration *t* the mechanism needs to specify a *space* of legal actions or *messages* M^t. The players are only allowed to reveal one of the messages from the *message space*.

Finite sequences of messages are assumed to be comparable. In other words, there exists a total order $\leq$ such that either $s_i[t] \leq s_j[t']$ or $s_j[t'] \leq s_i[t]$ for any length t message sequence $s_i[t]$, any length t' message sequence $s_j[t']$ and any players i, j. As a consequence, for any t, there is a minimum and maximum sequence of length t, and a minimum and maximum "extension" of any such length t sequence to length $t + k$. This allows message sequences to be interpreted as bids, and the mechanism to be interpreted as limited precision[1].

A strategy σ_i for agent i in such an incremental mechanism requires a choice of message $m_i^t \in M^t$ at each round t as a function of its type v_i and its history h_i^{t-1} up to that point.

[1]This is not to say that this is the only way of realizing incremental elicitation. Arbitrary query languages (e.g., asking agents to communicate upper and lower bounds on valuations) are certainly possible [27]. We consider only mechanisms that can be viewed as allowing "limited-precision bids."

4.2.3 Termination

We restrict our attention to protocols with the following termination condition. The auction terminates at iteration t whenever t is such that some bidder has specified a unique greatest message sequence. The good is allocated to that bidder. At any iteration where a unique highest bidder is not discovered the auction is allowed to terminate, but no allocation can be made.

Note that this assumption implies that we restrict the space of mechanisms to fully deterministic mechanisms that only allocate to the unique highest bidder.

This approach has a number of advantages. First it results in a fair protocol where each player is assured that she will be treated equally (with respect to allocation) and no allocation decisions will be left to chance. Secondly, by terminating immediately after a unique highest bidder has been discovered, it ensures that no resources are wasted. Finally, mechanisms with this restriction often have a unique (in form) dominant strategy equilibrium, as will be shown in the following sections.

4.2.4 Rationality

We assume that all bidders are *ex-post* rational. This means that in every game a player picks her actions in such a way that she never regrets the outcome of the game (that is, she would never want to pick a different set of actions even knowing the moves of her opponent). First of all, this implies that, if participation is optional or there is some way to avoid winning, only the winner will make a payment. Secondly this forms a foundation for dominant strategy equilibrium, as all dominant strategies are *ex-post* rational.

4.2.5 Intuitive form

The last restriction placed on the mechanism is that it is *intuitive*, that is easy to understand and implement. This is, of course, a very subjective measure and as such will not be used in any of the formal proofs, however we claim that the example mechanisms (presented in Section 4.4) do have an intuitive form. This is partly an attribute of these mechanisms having a dominant strategy equilibrium (which is often easier to understand and follow than other forms of equilibriums, such as Bayes-Nash equilibrium), but dominant strategy alone in not sufficient to make a mechanism intuitive as is well illustrated by the *Myerson* auction [22].

4.3 Necessary conditions

One of the goals of this chapter is to specify an approximate form of the auction mechanism that satisfies the above constraints and has a dominant strategy equilibrium. To accomplish this we will introduce and prove a number of necessary conditions that every such mechanism must satisfy. We will show that these conditions taken together define the approximate form of the kind of protocols we are interested in. In the course of developing the necessary conditions we will also introduce two new concepts: limited participation[2] and the last profitable iteration.

The first necessary condition demonstrates that all the auctions of interest must belong to a restricted set, referred to as *increasing price mechanisms*.

Definition 4.3 An iterative mechanism is an *increasing price mechanism* if the following condition holds for any player i. Let v_i be i's valuation and let s_{-i} be any sequence of moves of all other players (with $s_{-i}[k]$ denoting the length k initial segment). Let t be the least iteration at which i can win the good playing against sequence s_{-i} and pay a price $p < v_i^{\max}$ ($v_i^{\max}$ is the maximum valuation player *i* can possibly have), and s_i a sequence of moves that achieves this. Then for any other sequences $\hat{s}_{-i}$, $\hat{s}_i$ where $\hat{s}_{-i}[t] = s_{-i}[t]$, if i wins playing $\hat{s}_i$ against $\hat{s}_{-i}$ at some stage $k > t$ for price $\hat{p}$, then $\hat{p} \geq p$.

Intuitively, an increasing price mechanism has the following property for any player i: if we fix the moves of the opponents, and let t be the earliest round at which i could win, then

[2]The general concept of limited participation is well know, and not a part of the new contributions developed in this document.

the price paid by i at round t (against these fixed opponents), should she choose a strategy that wins at t, must be no more than the price paid if i wins at any later round against the same opponent moves.

Proposition 4.1 All single-good, iterative auctions with a *dominant strategy equilibrium* satisfying the conditions listed in section 4.2 are *increasing price mechanisms*.

Proof: See Proposition B.1 in Appendix B.

The *increasing price mechanisms* are similar to the *ascending price mechanisms* (Definition 2.26), although more general. In ascending price mechanisms given any two iterations t, k at which some player i can potentially win such that $k > t$, if the price player i pays for the good at iteration t is p_t and p_k at iteration k, then it must be the case that $p_k \geq p_t$. In increasing price mechanisms this condition holds only for the *first* iteration at which the player can potentially win and any later iteration.

4.3.1 Limited participation

We now introduce another important restriction on the class of incremental mechanisms, this restriction is made possible by the iterative nature of the elicitation process. We say that an incremental mechanism *limits participation* iff no player is allowed to participate once her utility function has been refined to the extent necessary to permit optimization of the mechanism's objective. In our limited-precision auction setting, this corresponds to the following activity rule: a player remains *active* as long as her message sequence is at least as great as that of any other player. This implies that if the auction has not yet terminated, then i is active iff her bids at all prior rounds have been tied with other highest bids. Limiting participation is one of the main factors that allows us to achieve a reduction in both communication and revelation. Intuitively, it ensures that we deal only with players whose valuations are high enough to potentially win and remove everyone else from the auction.

Relaxing this assumption can potentially remove all the benefit of using incremental auctions over the single-shot ones. Blumrosen, Nisan and Segal [2] show that for any 2-player sequential mechanism with communication complexity *m*, where the players send their bits one at a time in alternating order and are aware of the all the previously sent bits, there exists a one-shot mechanism that achieves at least the same expected welfare with communication complexity $2m - 1$. Although this result applies to a very specific class of mechanisms in both the form and the number of bidders, it shows that allowing all bidders to remain active for the duration of the auction is not the best strategy if we are concerned with the amount of communication.

From here on all the protocols of interest will be restricted to those implementing the *limited participation* constraint.

Proposition 4.2 Let player i have dominant strategy σ_i and valuation v_i. Suppose there exists an opponents' strategies profile σ_{-i} and valuations v_{-i} such that if $\sigma_i(v_i)$ is played against $\sigma_{-i}(v_{-i})$, the mechanism terminates at iteration t with i winning and paying $p_i < v_i$. Then for any $v_i' \geq v_i$ and any dominant strategy σ_i', if i plays $\sigma_i'(v_i')$ against $\sigma_{-i}(v_{-i})$, we must have:

(a) i wins and pays p_i (as with $\sigma_i(v_i)$)

(b) the mechanism terminates at iteration t (as with $\sigma_i(v_i)$)

Proof: For the proof of part (a) see Proposition B.2 in Appendix B, for the proof of part (b) see Proposition B.3 in Appendix B.

Intuitively, Proposition 4.2 shows that with limited participation, if i has valuation v_i and wins at iteration t with price p using a dominant strategy, then given any greater valuation it must win at the same iteration with the same payment using any dominant strategy, against fixed opponents.

Next we introduce a concept of *Last Profitable Iteration* or LPI which lets us define another necessary condition that further restricts the form of possible mechanisms.

Definition 4.4 The *Last Profitable Iteration* for player i, given history h_i^{t-1} and valuation v_i is defined as follows: we say $LPI(v_i, h_i^{t-1}) = now$ if there exist moves of other players such that i can profitably (with positive utility) win at round t, but cannot profitably win at any future round; $LPI(v_i, h_i^{t-1}) = \mathit{future}$ if there exist moves of other players such that i can profitably win at some round later than t; $LPI(v_i, h_i^{t-1}) = past$ otherwise. h_i^{t-1} is the history of play known to player *i* up to and including iteration $t-1$.

Proposition 4.3 Let player *i* have dominant strategy σ_i and valuation v_i. Suppose there exists some σ_{-i} and v_{-i} such that, if *i* plays $\sigma_i(v_i)$ against $\sigma_{-i}(v_{-i})$ (inducing history h_i^{t-1}), then $LPI(v_i, h_i^{t-1}) = \mathit{future}$ and player *i* is active at iteration *t*. Then for any other $v_i' \geq v_i$ and any dominant strategy σ_i' it must be the case that if *i* plays $\sigma_i'(v_i')$ against $\sigma_{-i}(v_{-i})$ (inducing history $\hat{h}_i^{t-1}$), we must have $h_i^r = \hat{h}_i^r$ and $\sigma_i(v_i, h_i^r) = \sigma_i'(v_i', \hat{h}_i^r)$, for all $r \leq t$.

Proof: See Proposition B.4 in Appendix B.

Intuitively, Proposition 4.3 has the following implications. If player *i*, with valuation v_i, is facing a history h_i^{t-1} where $LPI(v_i, h_i^{t-1}) = \mathit{future}$, then any dominant strategy must choose the same actions (against the same opponent moves) at all stages from 0 to t for *any* valuation $v_i' \geq v_i$. Thus, if $LPI(v_i, h_i^t) = \mathit{future}$, i is not required to "bid" so much as she must simply signal her willingness to participate. The only time an "interesting" message is proposed is when $LPI(v_i, h_i^t) = now$. This severely restricts the space of limited-precision mechanisms that admit dominant strategies.

Proposition 4.3 restricts the form of potential mechanisms to those that in most situations simply elicit a *yes* or *no* response. It does not, however, exclude the mechanisms that allow participation past a player's point of *LPI*. Allowing this kind of participation is wasteful, since no rational player will want to purchase the good passed her point of *LPI* and therefore all the extra incurred communication could be avoided. In the next proposition we show that one can safely restrict the space of mechanisms to those that do not allow participation after the point of *LPI*.

Proposition 4.4 Suppose i has a dominant strategy. Suppose that for some history h_i^{t-1} and valuation v_i, $LPI(v_i, h_i^{t-1}) = now$. Then there is a dominant strategy σ_i in which $\sigma_i(v_i, h_i^k) = minbid$ for any $k > t+1$ and history h_i^k s.t. $h_i^k[t-1] = h_i^{t-1}$ ($minbid = \min_b b \in M^k$).

Proof: See Proposition B.5 in Appendix B.

Consider any auction mechanism that does allow participation passed a player's *LPI*. If all players follow the dominant strategy described in Proposition 4.4 then the outcome of the game would be the same, independent of whether the auction permits such participation or not (keeping everything else equal). Furthermore, since the described strategy is dominant, there is no reason to believe that the players will not use it; therefore any mechanism that allows participation passed the point of *LPI* wastes resources by forcing extra communication.

In the next section we will show that the results presented so far can easily be extended to multiple players given a symmetry assumption.

4.3.2 Symmetry

In Section 4.2.3 we introduced the *termination* assumption. This assumption implies that the players are treated equally by the allocation function, so that the good can only be allocated to the player with the highest valuation and is never allocated randomly. In this section we extend this assumption to the payments made by the players. We say that the mechanism is *fully symmetric* if it is unable to discriminate based on the identity of the player. This implies that given some set of players, if the outcome of the game is *o*, then if we replace any player *i* from this set by another player *j* such that *j* will fully imitate the actions of *i* then the outcome of the game will be exactly the same (same payments and allocation).

Proposition 4.5 Given a *fully symmetric* auction mechanism with dominant strategy equilibrium, if σ_i is a dominant strategy of player *i*, then it is also a dominant strategy of any other player $j \neq i$.

Proof: See Proposition B.6 in Appendix B.

Note that all results presented in Propositions 4.2, 4.3, 4.4 apply to a single player with various dominant strategies and valuations. Proposition 4.5 shows that under the full symmetry assumption these results can be extended to multiple players. Most notably, the extension of Proposition 4.3 implies that if some player *i* with dominant strategy σ_i and valuation v_i reports a message m at iteration t, then, given the same opponents, any other player *j* with valuation $v_j \geq v_i$ will report exactly the same message.

Based on this and other results we will define a number of mechanisms that satisfy all the requirements presented in this chapter. These mechanisms are introduced in the next section.

4.4 Mechanisms

We will now define a number of symmetric mechanisms that follow the requirements of Section 4.2 and satisfy the necessary conditions of the previous section.

All the presented mechanisms are structured in a such a way that it is possible to optimize them with respect to almost any social choice function. We present an *MDP* based optimization, which uses prior knowledge about the players, more specifically, the probability density over valuations, to guide the policy generation. This approach is similar to the *automated mechanism design* [8, 31], where the mechanism is optimized based on specific priors. Furthermore, in the spirit of this document, we optimize the mechanisms to account for the cost of communication (it is also possible to explicitly incorporate the cost of computation).

4.4.1 Adaptive Symmetric Incremental Auction

First we present a simple *Adaptive Symmetric Incremental Auction* or *ASIA*. The structure of this auction is similar to what is sometimes referred to as *Japanese* auction [21], which is a version of the *English* auction, Section 2.3.2.

ASIA operates as follows.

- Initially, all the players are active.

- At iteration *t*, the mechanism announces price p^t to all active players, with $p^t \geq p^{t-1}$.
- Each active player reveals either 1, indicating a willingness to purchase the good for p_t and to remain active (intention to participate), or 0, forcing the player to become inactive (desire to become inactive)
- The mechanism terminates when:
  - only one player bids 1, in which case that player receives the good and pays the last announced price
  - all active players bid 0, in which case the good is not assigned and no payments are made.

As required, *ASIA* has an *intuitive* dominant strategy equilibrium.

Proposition 4.6 Under the rules of ASIA, it is a weakly dominant strategy of any player *i* with valuation v_i to bid 1 at iteration *t* as long as $v_i > p^t$, and to bid 0 otherwise.

Proof: See Proposition B.7 in Appendix B.

This dominant strategy equilibrium is intuitive in a sense that given only two possible actions with one of them always resulting in utility of 0, it is easy for a player to see that she can't do any better than saying *yes* when her valuation is above the potential purchase price. Interestingly enough, neither this intuition nor the existence of dominant strategy is affected by the actual numerical values of the announced prices. The only condition that must be maintained is that the auction remains ascending. We will take advantage of this fact in deriving the best prices for optimizing the auction for various social choice functions (see Section 4.4.3).

We can now check that *ASIA* conforms with all requirements placed on the desired type of mechanism. In *ASIA*, the announced price is not lowered at every successive iteration, therefore it clearly belongs to the space of *increasing price mechanisms*. The presented dominant strategy equilibrium corresponds to the results derived in the previous section. Furthermore, the style of the mechanism follows the conclusion of Proposition 4.3. That is, the players are only allowed to report whether they wish to keep participating or to drop out. This is a somewhat extreme approach, since the players are not allowed any other options at the point of their *LPI*. In the next section we will present a mechanism that does explicitly take advantage of the knowledge that the player is at her *LPI*.

4.4.2 Second-price Symmetric Incremental Auction

The *Second-price Symmetric Incremental Auction* or *SPSIA* is similar to *ASIA* in design; it, however, incorporates the notion of *LPI* more directly. In essence, *SPSIA* is a combination of *ASIA* and *LP-TIOLI* (Section 3.2).

SPSIA operates as follows.

- Initially, all the players are active.
- At iteration *t*, the mechanism announces a price *interval* $[a_t, b_t]$ to all active players, the interval $[a_t, b_t]$ is such that $a_t < b_t$ and $a_t \geq b_{t-1}$.
- All active players reveal either 1, indicating an intention to participate, or 0, indicating a desire to become inactive.
- The mechanism terminates when only one player (the winner) bids 1 or all active players bid 0. In both cases the active players list remains unchanged from the previous iteration. Assuming this occurred at iteration *k* the mechanism then
  - announces a set of prices η^k where $a_k, b_k \in \eta^k$ and all prices are between a_k and b_k.
  - restrict η^k to be the same if either zero or one players reported 1 at iteration k (given the same history up to iteration $k-1$ for both situations).
  - as in LP-TIOLI, each active player is required to select one of these prices, with an exception; if a winner has already been determined, the mechanism enters a bid for her equal to b_k.

- if there is a unique highest bidder, she receives the good for the price of the second highest bid.
- otherwise the good is not assigned and no payments are made.

As can be seen from the above, *SPSIA* incorporates the idea of *LPI* by giving the player, who is at the point of her *LPI*, a last chance to refine her valuation. Since the communication complexity is important this option is provided only if the mechanism is unable to find the highest bidder or has just determined her identity. Note, that the later condition brings some inefficiency into the protocol since the mechanism forces a transmission from a number of players, while already aware of the identity of the winner. This communication is required in order to maintain the dominant equilibrium of the mechanism; in any case given that the mechanism should exclude most players by the time it reaches this point and it will elicit at most one extra message from each player the overhead should not be large.

Proposition 4.7 It is a weakly dominant strategy of any player *i* in SPSIA to bid 1 at iteration *t* as long as $v_i > b_t$ (equivalent to $LPI(h_i^t, v_i) = \mathit{future}$). Otherwise, the player should bid 0. If at any point the player is asked to choose from some set of prices she should pick the (announced) price which is just above her valuation.

Proof: See Proposition B.8 in Appendix B.

Like the protocol itself, the dominant strategy of a player in *SPSIA* is also a combination of dominant strategies of *ASIA* and *LP-TIOLI*. Each player follows the dominant strategy for *ASIA* until the point of her *LPI* and then switches to *LP-TIOLI*'s dominant strategy (if asked for more information).

Again similar to *ASIA* and *LP-TIOLI*, the dominant strategy of *SPSIA* does not depend on the exact numeric values of the prices. The optimization potential of this mechanism is greater however. This is so since the optimization can be separated into two components, the optimization of the *ASIA* type prices and the optimization of the *LP-TIOLI* type prices. While optimizing the second component we can take advantage of the knowledge that this is the last possible iteration, furthermore we can adjust the number of price thresholds based on the remaining communication capacity.

In the next section we present a procedure for adapting the above two mechanisms for various social choice functions.

4.4.3 Optimization of prices

In the previous section we have presented two mechanisms *ASIA* and *SPSIA* both of which have a dominant strategy equilibrium. As was noted, the dominant strategy equilibria of both of these mechanisms do not depend on the specific numeric values of the prices announced to the participants. Therefore, as with *LP-TIOLI*, this independence makes it possible to apply a straightforward optimization.

Contrary to *LP-TIOLI*, however, these mechanisms go through multiple iterations before terminating, with each iteration bringing more information to the mechanism. To take advantage of this fact, the price determination protocol should in fact be a policy that determines the best price based on the history of play from the previous iterations. We will solve a Markov Decision Process (MDP) to find the policies, some of which we will then evaluate empirically. We will begin with welfare optimization policy for *ASIA*.

Assume a finite horizon of T rounds, that is after the round T the auction will terminate (even if no unique highest bidder has been determined). Note, that this does not affect the dominant strategy equilibrium, since if the auction is forced to terminate at round T it will do so without an allocation. Assume a common prior density Φ over the bidders' valuations (this is for simplicity only). We set prices to optimize social welfare by formulating the MDP as follows: states are pairs $\langle p, m\rangle$, where p is the prior price threshold and m is the number of active bidders. At any stage $k < T$, and at any state $\langle p^{k-1}, m\rangle$ where $m > 1$, we can set any price $p^k > p^{k-1}$. If $m = 0$ or $m = 1$, the auction terminates. We define the optimal k-stage value function (reflecting expected welfare) EW^k and Q-function as follows:

$$EW^k(p^{k-1}, 0) = 0 \tag{4.1}$$

$$EW^k(p^{k-1}, 1) = E_\phi(v|v > p^{k-1}) \tag{4.2}$$

$$EW^k(p^{k-1}, m > 1) = \max_{p^k > p^{k-1}} Q^k(p^{k-1}, m, p^k) - c(m) \tag{4.3}$$

$$Q^k(p^{k-1}, m, p^k) = \Pr(1 \text{ active}|p^{k-1}, m, p^k) EW^{k+1}(p^k, 1) \tag{4.4}$$

$$+ \sum_{2 \le n \le m} \Pr(n \text{ active}|p^{k-1}, m, p^k) EW^{k+1}(p^k, n) \tag{4.5}$$

Where equation 4.1 is the expected welfare when zero players are active at the current iteration and the last announced price was p^{k-1}. Equation 4.2 is the expected welfare when exactly one player is active the last announced price was p^{k-1}. Equation 4.3 is the expected welfare (as well as mechanism's valuation) for selecting the best possible price when more than one player is active and the last announced price was p^{k-1}. Equation 4.4 calculates the mechanism's valuation for announcing a price p^k at an iteration where m players are active and the last announced price was p^{k-1}.

In equation 4.4

$$\Pr(n \text{ active}|p^{k-1}, m, p^k) = \binom{m}{n} * \frac{\left(1 - \Phi(p^k)\right)^n * \left(\Phi(p^k) - \Phi(p^{k-1})\right)^{m-n}}{(1 - \Phi(p^{k-1}))^m}$$

is the probability of exactly n bidders, among m bidders with valuations greater than p^{k-1}, having valuations greater than p^k.

To factor the communication cost into these equations we introduce a function $c(m)$, which calculates the cost based on the number of active players. The assumption in this case is that each *active* player transmits exactly one bit per iteration; this is due to the design of *ASIA*. As can be seen from the above equations, we don't need to restrict the form of the cost function as it does not affect the optimization; it does not even have to be the same from iteration to iteration. When deriving results, for the purposes of simplification, we have used the same linear cost function at all iterations.

The equations used to derive the revenue maximization policy are largely similar. We use ER^k to denote the k-stage value function.

$$\begin{aligned} ER^k(p^{k-1},0) &= 0 \\ ER^k(p^{k-1},1) &= p^{k-1} \\ ER^k(p^{k-1},m>1) &= \max_{p^k>p^{k-1}} Q^k(p^{k-1},m,p^k) - c(m) \\ Q^k(p^{k-1},m,p^k) = \sum_{1\le n\le m} &\Pr(n \text{ active}|p^{k-1},m,p^k) ER^{k+1}(p^k,n) \end{aligned} \tag{4.6}$$

By the prior assumptions, the players are ex-post rational, therefore they can not be forced to pay anything unless they are allocated the object. For that reason, in equation 4.6, we have decided to always extract the payment for communication from the seller. Alternatively, it should be possible to extract that payment from the winning bidder. This however, would complicate the protocol and possibly destroy the dominant strategy equilibrium, since the buyer's value for participating will now depend on the number of iterations, the number of participants, and their expected valuations.

The MDP for *SPSIA* is largely similar to *ASIA* save for the few variations due to the way *SPSIA* announces prices. As an exercise we derive the *MDP* for social welfare. We make the same assumptions about the finite horizon and the prior distribution. With *SPSIA* the state becomes a triple $\langle a, b, m\rangle$ where a is the lower bound of the previous interval, b is the upper bound, and m is the number of active bidders. The policy needs to select the price interval $[a^t, b^t]$ at every state $\langle a^{t-1}, b^{t-1}, m\rangle$ and any stage $k < T$. The auction terminates whenever $m = 0$ or $m = 1$.

Note that although we need to pick both a and b at each interval, the players' dominant strategies depend mostly on b. Furthermore, all probability calculations depend only on b. We can use these facts to simplify the value function by separating it into two components. We define the optimal k-stage value function (reflecting expected welfare) using two components EW^k and EW^k_L as follows:

$$
\begin{aligned}
EW_L^k(b^{k-2}, a^{k-1}, b^{k-1}, n, 0) &= EW_L(b^{k-2}, \eta^k(b^{k-2}, a^{k-1}, b^{k-1}, n), n, 0)\\
EW_L^k(b^{k-2}, a^{k-1}, b^{k-1}, n, 1) &= EW_L(b^{k-2}, \eta^k(b^{k-2}, a^{k-1}, b^{k-1}, n), n, 1)\\
EW^k(b^{k-1}, m > 1) &= \max_{b^k > a^k \geq b^{k-1}} Q^k(b^{k-1}, m, a^k, b^k) - c(m)\\
Q^k(b^{k-1}, m, a^k, b^k) &= \Pr(0 \text{ active}|b^{k-1}, m, b^k) EW_L^{k+1}(b^{k-1}, a^k, b^k, m, 0)\\
&\quad + \Pr(1 \text{ active}|b^{k-1}, m, b^k) EW_L^{k+1}(b^{k-1}, a^k, b^k, m, 1)\\
&\quad + \sum_{2 \leq n \leq m} \Pr(n \text{ active}|b^{k-1}, m, b^k) EW^{k+1}(b^k, n)
\end{aligned}
\tag{4.7}
$$

$EW_L(b^{k-2}, \eta^k, n, 0)$ and $EW_L(b^{k-2}, \eta^k, n, 1)$ denote the social welfare after the mechanism announces some set of prices η^k. The first equation refers to the situation when all the players have valuations below the highest price and the second when exactly one player has valuation above the highest price and everyone else below. In both cases the expected welfare is calculated under the assumption that each player's valuation is above b^{k-2}.

The formula for calculating $EW_L(b, \{p_1, p_2, \cdots, p_r\}, n, 0)$, $p_1 < p_2 < \cdots < p_r$ is similar to the expected revenue equation 3.3 from chapter 3. Let $Norm = \Phi(p_r) - \Phi(b)$. Therefore,

$$
\begin{aligned}
&EW_L(b, \{p_1, p_2, \cdots, p_r\}, n, 0) =\\
&\quad \left(\frac{1}{Norm}\right)^n * \sum_{i=1}^{r-1} \begin{array}{l}(\Phi(p_i) - \Phi(b))^{n-1} * (\Phi(p_{i+1}) - \Phi(p_i)) *\\ n * E_\phi(v|p_{i+1} < v < p_i)\end{array}\\
&\quad - c(n) * \log(|\{p_1, p_2, \cdots, p_r\}|)
\end{aligned}
\tag{4.8}
$$

The equation for $EW_L(b, \{p_1, p_2, \cdots, p_r\}, n, 1)$ can be written as follows

$$
\begin{aligned}
&EW_L(b, \{p_1, p_2, \cdots, p_r\}, n, 1) =\\
&\quad \left(\frac{1}{Norm}\right)^{n-1} * (\Phi(p_{r-1}) - \Phi(b))^{n-1} * E_\phi(v|v > p_r)\\
&\quad - c(n) * \log(|\{p_1, p_2, \cdots, p_r\}|)
\end{aligned}
\tag{4.9}
$$

By construction, the mechanism requires that the set of prices announced to the players are the same for both situations of one or zero players reporting 1, given the same history.

Therefore we will determine the price set as follows

$$\begin{aligned}\eta^k(b^{-1}, a, b, n) = \{a, b\} \cup \\ \max_{r,\{p_2,\cdots,p_{r-1}\}} \Big(&\Pr(0 \text{ active}|b^{-1}, m, b) * EW_L(b^{-1}, \{a, p_2, \cdots, p_{r-1}, b\}, n, 0) + \\ &\Pr(1 \text{ active}|b^{-1}, m, b) * EW_L(b^{-1}, \{a, p_2, \cdots, p_{r-1}, b\}, n, 1) \Big)\end{aligned} \tag{4.10}$$

Note that although η^k has an iteration index, this is solely for notational purposes, the prices do not depend on the iteration number, but only on the parameters.

4.4.4 Empirical results

In this section we will review a number of performance results produced by running simulations of *ASIA*. We will also compare the performance of *ASIA* with optimized price policy versus an unoptimized one.

The price policies underlying these results are based on the *MDP* formulations presented in the previous section. The policies were computed for a varying number of (initially active) players, 2 to 50, and various communication cost values. We have made an assumption that all valuations are distributed uniformly. This is for simplicity only. It is possible to compute these policies under any other distribution over valuations, in this case however, these results are sufficient for corroborating our conclusions (to be presented). Finally, to keep the action space finite, we have subdivided the price interval into 1000 possible prices. This potentially makes the policies inefficient, however experiments with other (larger) values have shown that the difference is minor. Nevertheless, it is possible to easily discard this constraint by adopting methods for dealing with continuous state and action spaces. The performance of ASIA was measured by simulating 100,000 runs under each policy (drawing bidder valuations randomly).

To demonstrate the benefit of computing a policy versus using a fixed price selection rule, we have simulated the performance of *ASIA* under the $Divide(k)$ price update rule.

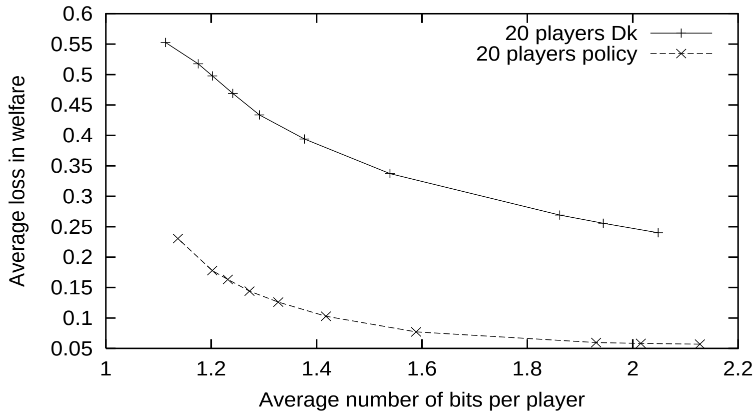

Figure 4.1: Comparison of loss in welfare in ASIA under optimized policy vs. ASIA under *Divide(k)* price update rule

Definition 4.5 Define the price update rule *Divide(k)* as follows. Given any fixed *k*, $k > 1$, an iterative mechanism operating under this rule would announce the price $p_0 = \frac{1}{k}$ at iteration 0. At any other iteration $t > 0$ the mechanism would announce $p_t = p_{t-1} + \frac{1-p_{t-1}}{k}$.

Although the choice of *Divide(k)* appears somewhat arbitrary, it is in fact similar to using uniformly placed prices in a one-shot setting. Furthermore, as will be shown in the next chapter, *Divide(k)* exhibits some beneficial properties which make it a feasible alternative to an optimized policy.

For the purposes of simulation, the k parameter was set in such a way as to force *ASIA* with *Divide(k)* to use the same average number of bits per player as *ASIA* with optimized policy. We compare the resulting loss in welfare. Figure 4.1 demonstrates that under the optimized policy *ASIA* consistently achieves a much smaller loss in welfare, the value of 20 bidders was chosen arbitrarily.

The next few graphs demonstrate the performance of *ASIA* under different initial numbers of *active* players. Figure 4.2(a) shows the relation between the cost of communication (per bit) and the expected amount of communication per player until mechanism's termination. As expected, the amount of communication decreases as cost increases. Note that given the

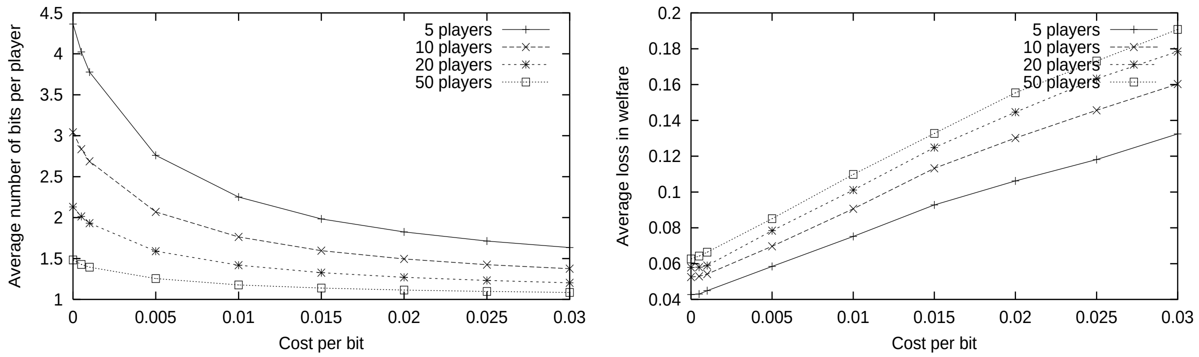

Figure 4.2: (a)Required bits per player as costs vary; (b)Average welfare loss as costs vary.

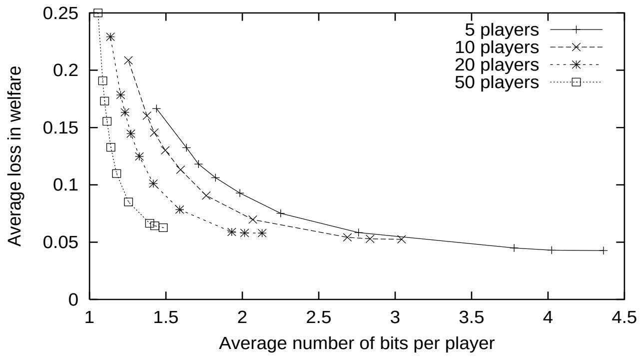

Figure 4.3: Average welfare vs. bits per player.

same cost the loss in welfare is smaller with more players. The next two graphs deal directly with the social welfare; we use the measure of *loss in welfare*, which equals to the difference between the obtained social welfare and the social welfare of the optimal mechanism. Figure 4.2(b) demonstrates the relation between the communication cost and the loss in social welfare. While this suggests that ASIA fares worse with increasing numbers of players, we can recast our results by considering the loss in welfare as a function of the amount of per-player communication. Figure 4.3 demonstrates that, in fact, with more players ASIA requires *less* per-player communication to achieve the same loss in welfare. This gain is due to the fact that with a greater number of players, the policy can initially be more "aggressive," eliminating more players right away. Once only a few active players remain, the policy can become more conservative, which results in a higher chance of allocating the good.

The presented results demonstrate that the performance of the mechanism increases with the number of players. This is a significant advantage over at least some single-shot mechanisms, in particular *LP-TIOLI* presented in the previous chapter. As can easily be noted from the definition of *LP-TIOLI* as the number of participants increase the probability of not allocating the good will also increase and therefore its performance (in both social welfare and revenue) will fall drastically. This drawback of the single-shot mechanisms is of course not limited to *LP-TIOLI*, we however, postpone this discussion until the next chapter, where we will present a more efficient version of an incremental mechanism.

Chapter 5

Efficiency gains through stochasticity

In this chapter we relax one of the major assumptions of Chapter 4, the restriction of the allocation to the unique highest bidder. We introduce a modified version of *ASIA* without this restriction and show that it is more efficient than the original. We also show that the modified version is provably more efficient under certain conditions than any limited precision threshold-based single-shot auction.

5.1 Stochastic ASIA

Blumrosen and Nisan [1] showed that in the context of single-shot auctions, the symmetric mechanisms, that is the mechanisms that treat the players equally, are suboptimal to the asymmetric ones. Their *priority games* (see Section 2.4.1) incorporate tie breaking rules which are based on a fixed ordering of the players. Essentially this implies that the mechanism has a preference ordering over the players. In the event that multiple players exhibit the same desire for the good, the mechanism always allocates the good to the most preferred player. In addition, priority game discriminate between the players in terms of prices. That is, the players might be required to pay different amounts for the same good, with the amount being decided based on the player's identity. This makes the mechanism unfair in a certain sense.

So far in this document we have constrained our attention to the mechanisms that are non-discriminatory, with this notion taken to the extreme. This means that whenever the mechanism is forced to end the auction but is unable to decide between the two (or more) highest bidders the good in not allocated. This is true for both *LP-TIOLI*[1] and incremental auctions. For incremental auctions this sort of allocation choice arises due to the *determinism* assumption, made in the previous chapter. Since, given the restrictions of Chapter 4, the good can only be allocated to a player with a unique highest valuation, any sort of randomness has to be excluded from this decision. Although this restriction allowed us to develop the proofs in the previous chapter, thus leading to the *ASIA* type of mechanisms, it certainly limits the space of possible mechanisms to those which are clearly not optimal.

To alleviate this problem we need to consider mechanisms which always allocate the good. Since with limited precision it might not always be possible to find the player who desires the good the most, it is necessary to relax the assumption of allocating only to the *unique* highest bidder. There are a number of ways of doing this. One way is to follow the style of *priority games* and introduce discrimination between players, which, for the purposes of this document, is not desirable. An alternative approach is to relax the determinism component but force the mechanism to be fair in expectation. This is the approach that we will follow.

We propose the following variant of *ASIA*, which we call *Stochastic ASIA* or *ST-ASIA*. *ST-ASIA* has a dominant strategy equilibrium, but, as opposed to *ASIA*, it always allocates the good. This mechanism operates as follows.

- Initially all the players are active. Denote the set of active players by A.
- At each iteration t, the mechanism randomly (and uniformly) *holds out* one of the active players, say, player i. It then announces price p^t to all active players, with $p^t \geq p^{t-1}$.
- All active players *except* the holdout player i are asked to reveal either 1, indicating a willingness to purchase the good for p_t (intention to participate), or 0, indicating a desire

[1] Although the rules of *LP-TIOLI* state that in case of a tie the good is offered randomly, it is the best strategy for any rational player to reject this offer, therefore it would be outcome equivalent to forgo making any offers and not allocate the good.

to become inactive.

- If all players in $A - \{i\}$ reveal 0, the good is allocated to *i* for price p^{t-1}.
- If at least one of the players in $A - \{i\}$ reveals 1, then *i* is also asked to reveal either 0 or 1 (with the same interpretation):
  - If only one active player (including the holdout) reveals 1, then this player receives the good and pays p^t.
  - If more than one *active* player (including the holdout) reveals 1, then the game moves into the next iteration.

Proposition 5.1 In ST-ASIA mechanism, it is a weakly dominant strategy of any player *i* with valuation v_i to bid 1 at iteration *t* as long as $v_i > p^t$, and to bid 0 otherwise.

Proof: See Proposition C.1 in Appendix C.

The conditions developed in the previous chapter are not known to be necessary for the existence of a dominant strategy equilibrium when the determinism assumption is relaxed. However, as can be easily verified, *ST-ASIA* belongs to the space of *increasing price mechanisms* and since its dominant strategy is exactly the same as *ASIA*'s it also satisfies all other necessary conditions. Furthermore, as with *ASIA* its dominant strategy is independent of the numerical values of the price thresholds, which allows us to use a similar optimization approach. Before proceeding to the performance results of *ST-ASIA* with optimized price policies, we wish to provide a general comparison of *ST-ASIA* to one-shot auctions.

5.2 Stochastic ASIA vs. one-shot auctions

In this section we analyze and compare the performance of two auction types under the assumption that the players' valuations are independently distributed according to the uniform distribution. The purpose of this section is to show that with a large number of players incremental auctions, specifically *ST-ASIA*, can be significantly better than limited precision threshold based one-shot auctions. In particular we compare *ST-ASIA* to the sort of mechanisms introduced by Blumrosen and Nissan [1]. Even though the analysis is performed under an unrealistic assumption of IID uniform distribution, it is sufficient for demonstrating that incremental auctions can be better than even the optimal limited precision one-shot threshold-based auctions.

Consider the behavior of a *fixed structure* (the auction structure is known before the number of participants is announced) one-shot limited-precision threshold-style auction as the number of participants is increased. As more players enter the auction, the expected highest valuation of this group rises. However, the expected welfare of a fixed structure one-shot auction will eventually start to fall or will stabilize at some fixed level (below the expected welfare of the optimal auction with unrestricted communication). This is very easy to see in the case of *LP-TIOLI* (Section 3.2.2). As the number of players increases so does the probability of having a tie, which implies that eventually the probability of not allocating the good will become large enough to outweigh any increase in the expected highest valuation. This problem is not as pronounced in the case of *priority games* (Section 2.4.1). However, even a priority game has the highest fixed price threshold, therefore as the number of players increases and the highest valuation of the group reaches this price threshold, the expected loss in welfare will be upper bounded by the distance between this price threshold and its closest neighbor, a fixed positive constant. The actual expected loss would be smaller than this upper bound, in particular under the uniform distribution, it will be non decreasing with more players and equal to about half of this interval.

Suppose we were to implement *ST-ASIA* with *Divide(k)* price update rule. *Divide(k)* was defined in Chapter 4, we will restate the definition.

Definition 5.1 Define the price update rule *Divide(k)* as follows. Given any fixed k, $k > 1$, an iterative mechanism operating under this rule would announce the price $p_0 = \frac{1}{k}$ at iteration 0.

At any other iteration $t > 0$ the mechanism would announce $p_t = p_{t-1} + \frac{1-p_{t-1}}{k}$.

The two immediately visible features of this price update rule is that it does not depend on the number of participating players (fixed mechanism structure) and the distance between the two consecutive prices decreases with each iteration. This price update rule also has another interesting property [2].

Proposition 5.2 Suppose that the players' valuations are IID according to uniform distribution. The expected total number of bits sent by each player participating in *ST-ASIA* with unrestricted number of iterations and price update rule *Divide(k)* is upper bounded by k.

Proof: See Proposition C.2 in Appendix C.

Intuitively, the above means that if we take any fixed k and implement *ST-ASIA* with *Divide(k)* such that we always allow it to run until termination (by finding the unique highest bidder or announcing a price above all valuations) on average each player would submit k or fewer bids. This result is independent of the number of participating players.

With this price update rule, if the auction terminates at iteration t then the loss in welfare would be upper bounded by $p_t - p_{t-1}$. From the definition of *Divide(k)* we know that, $p_t - p_{t-1} < p_{t-1} - p_{t-2}$ and $\lim_{t\to\infty}(p_t - p_{t-1}) = 0$. This means that, as the number of players increases the loss in welfare will fall. This is so, since with more players the expected highest valuation rises and therefore the ties, if any, would have to occur at a higher iteration. Note, that all this occurs while the number of bits submitted by each player on average is fixed.

As was noted above, the loss in social welfare of a fixed-structure threshold-based limited precision single-shot auction has a fixed non-zero lower bound for increasing number of players. Therefore, as the number of players is increased, the expected loss of welfare of a fixed-structure single-shot threshold-based mechanism will be greater then the expected loss of welfare of *ST-ASIA* (which can be arbitrarily close to 0), while the total expected com-

[2] While *Divide(k)* is best suited for the uniform distribution, is can be modified to work with any other form of distribution. The easiest modification would be to select the price so that the distance between the new price and the previous price contains $\frac{1}{k}$ of the remaining probability density.

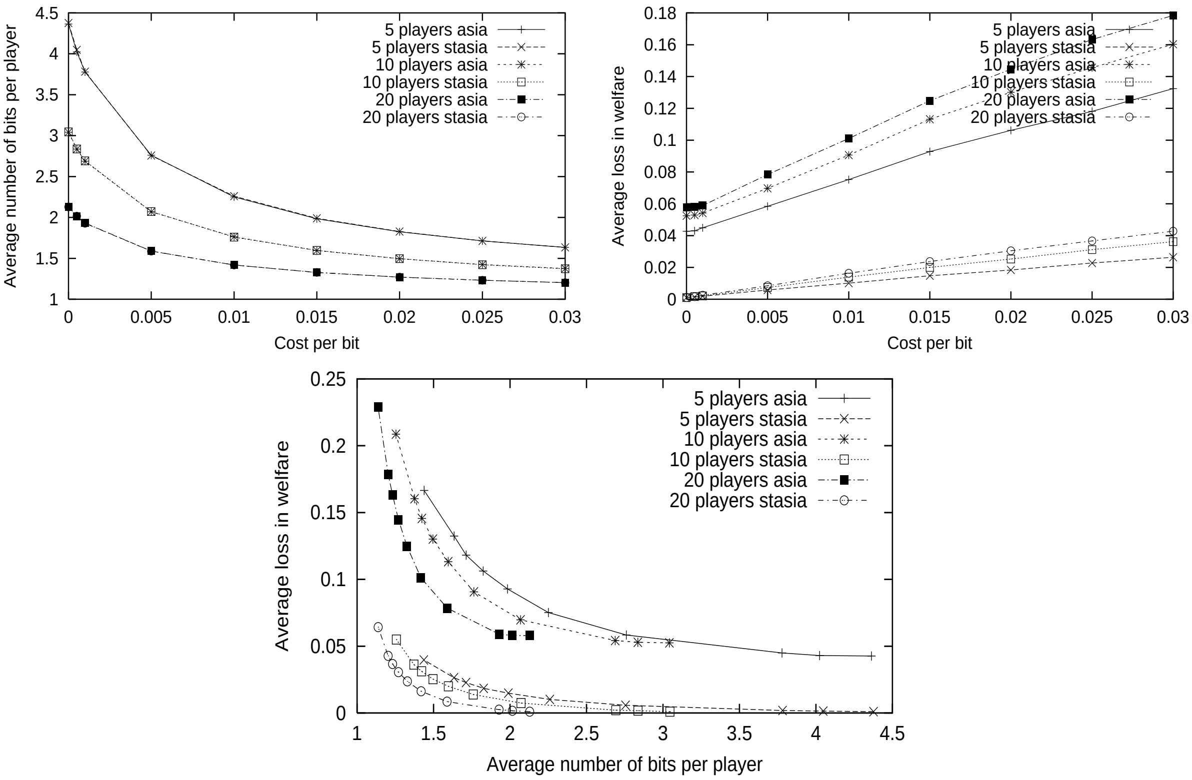

Figure 5.1: ASIA vs. ST-ASIA under the same price policies: (a) Required bits per player as costs vary; (b) Average welfare loss as costs vary; (c) Average welfare loss vs. bits per player.

munication of *ST-ASIA* is smaller than or equal to the total communication of the single-shot mechanism. This demonstrates that even with unoptimized prices, the types of mechanisms presented in this document offer advantages over the one-shot mechanisms, even the asymptotically optimal ones presented in [2]. Although we have presented the proof of this result only under the uniform distribution, we believe that the same will also occur under the other forms of distributions, the proof of this fact is beyond the scope of this document.

5.3 Evaluating the performance of Stochastic ASIA

In this section we empirically evaluate the performance of *ST-ASIA*.

As the first step, we will compare the performance of *ST-ASIA* to *ASIA*, with both auctions

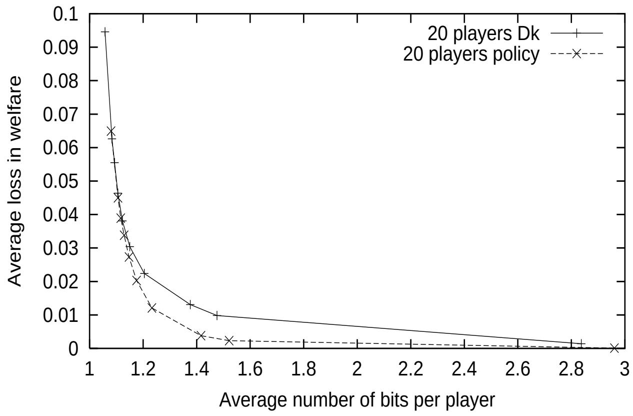

Figure 5.2: Comparison of loss in welfare in ST-ASIA under optimized policy vs. ST-ASIA under *Divide(k)* price update rule

operating under the same price policies, specifically the policies optimized for *ASIA*. To generate the results we ran each mechanism on each policy (10 policies for 10 different cost per bit values) over 100,000 runs with bidder valuations drawn randomly. Figure 5.1(a) demonstrates that the average number of bits communicated by each player is approximately the same for both mechanisms. Since both mechanisms use the same policies this is expected. Figure 5.1(b) shows that *ST-ASIA* achieves a lower loss in social welfare than *ASIA*. Figure 5.1(c) demonstrates the same result by presenting the loss in social welfare as a function of the average number of bits sent by each player. In both cases this result originates from the fact that *ST-ASIA* always allocates the good, and therefore the maximum loss in welfare it can incur is $p^t - p^{t-1}$, assuming the good was allocated immediately after iteration t.

We use the same approach as in the previous section, to produce the price policies optimized for *ST-ASIA*. As before, we will first compare the performance of these policies to *ST-ASIA* under *Divide(k)* price update rule. Figure 5.2 demonstrates that the optimized policy performs better that *Divide(k)*, however the difference is not as pronounced as with *ASIA*. This is primarily because even unoptimized *ST-ASIA* incurs a relatively small loss in social welfare, thus leaving a smaller gain due to optimization. Still, optimized policies should perform better; however the two policies have similar performance under low values of per player communication. This could be attributed the fact that the policies are optimized with respect to both social welfare and the cost of communication. This means that for any given limit on communication, it is potentially possible to achieve a smaller loss in welfare then we do here.

Finally, we will compare the performance of *ST-ASIA* under the price update policies optimized for *ASIA* to *ST-ASIA* with price update policies optimized for *ST-ASIA*. Figure 5.3(a) demonstrates that the policies are indeed different: policies optimized for *ST-ASIA* tend to use fewer bits of communication under the same cost. Figure 5.3(b) shows that the policies optimized for *ST-ASIA* do perform better than the policies optimized for *ASIA*.

Altogether, the results presented here demonstrate that ST-ASIA with optimization offers an improvement over ASIA with optimization, as well as ST-ASIA without optimization. More significantly we have shown that ST-ASIA can offer better performance than limited precision threshold based single-shot mechanisms while treating all players fairly and maintaining intuitive dominant strategy equilibria.

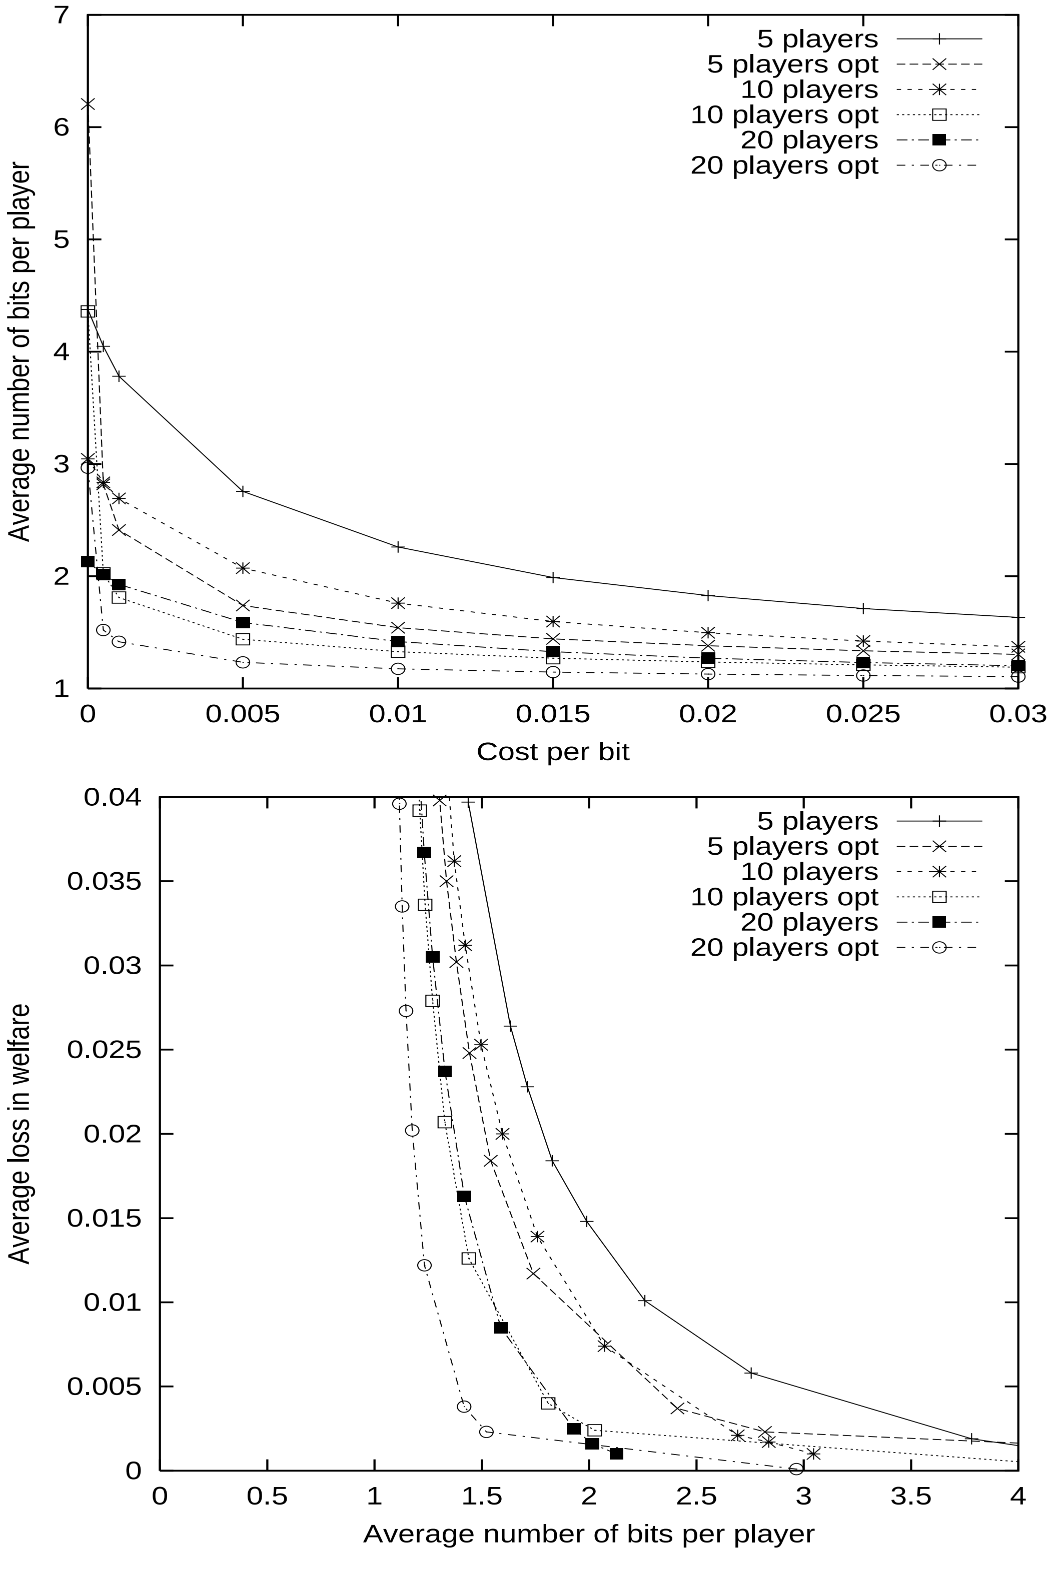

Figure 5.3: ST-ASIA under different price policies (a) Required bits per player as costs vary; (b) Average welfare loss vs. bits per player.

Chapter 6

Conclusion

6.1 Main contributions

In this dissertation we have looked at ways of conducting single item auctions under the restrictions of limited communication and (indirectly) limited computation and revelation.[1] In addition to these requirements we have also proposed that the desired auction mechanism must be fair and have an intuitive rules and dominant strategy equilibria. To develop these mechanisms we have taken the approach of using limited precision to minimize both the computational complexity faced by the participants and the amount of communication between them and the mechanism.

In Chapter 3 we developed a simple one-shot auction. This auction serves as an introduction to the main results. *LP-TIOLI* is clearly suboptimal, due to its symmetric nature. However it has an intuitive dominant strategy equilibrium and does not discriminate between players, features not often present in the limited precision auction mechanisms. Furthermore, *LP-TIOLI* can be optimized for any social choice objective, while maintaining all of these properties.

The main contributions of this document are centered in Chapters 4 and 5. We have conjectured that under the restriction of limited precision, incremental mechanisms will be superior

[1]The last two objectives were achieved through limiting the action space, a consequence of our approach to limited communication.

in performance to single-shot mechanisms. As before, we are interested in mechanisms with natural properties such as an intuitive dominant strategy equilibria and fairness. For the case of incremental auctions we have refined the concept of fairness to allow allocation only when a player with the unique highest valuation is known; that is, no decisions are left to chance. To further limit communication we have introduced a *limited participation* constraint, which states that only the players with sufficiently high valuations are allowed to participate. In Chapter 4 we have developed a number of necessary conditions that every mechanism under these restrictions has to satisfy. We have shown that all such mechanisms must belong to the class of *increasing price mechanisms*. We have proved that in these mechanisms the players' dominant strategies must be very similar, and even identical in some cases. This fact has a strong influence on the structure of the mechanisms of interest. It implies that in most cases, at each iteration, the mechanism only needs to provide the player with just two possible choices. This means that in most cases a players' communication is limited to sending one bit per iteration, the only time a player may face a larger set of possible answers is when she is about to be removed from the auction.

We have developed two auctions based on these ideas, *ASIA* and *SPSIA*. Both of these auctions have an intuitive set of rules and an easy to understand dominant strategy equilibrium. They are also both optimizable for various social choice functions; the optimizations process does not disrupt any of the beneficial properties of these mechanisms. We have used a Markov Decision Process formulation to construct the optimization procedures. To demonstrate the performance of these type of protocols we ran a number of simulations on *ASIA*. These simulations show that *ASIA* provides superior performance with more players, they also demonstrate the benefits of optimization.

In Chapter 5 we provided more evidence pointing to the advantages of incremental auctions. The main idea behind the results of this chapter was the partial relaxation of the determinism assumption. The auctions were still required to treat players equally, however, now only in expectation, which permitted the stochastic outcomes. This allowed the development of a stochastic variant of *ASIA*. We showed that under the assumption of uniformly and independently distributed valuations stochastic *ASIA* can be significantly better than any limited precision fixed structure threshold based single-shot auction. We have also provided a MDP based optimization and demonstrated that stochastic *ASIA* performs better than the original *ASIA*.

6.2 Future research directions

There are a number of possible extensions to the results presented in this dissertation. The most obvious and potentially the most difficult is the generalization of the results of Chapter 4 to multi-unit and combinatorial auctions. These auctions have the most to gain from a reduction in communication and computational complexity, however, whether it is even possible to design similar protocols with the same intuitive forms of dominant strategy equilibria is an open question.

A somewhat less ambitious undertaking would be to show that single item incremental auctions are superior to limited precision one-shot mechanisms (not only threshold based) under all possible distributions over players' valuations. The empirical evidence suggests that this is indeed the case, however until proven this evidence is not a sufficient justification.

Another interesting line of research would be to prove a set of necessary conditions similar to ones in Chapter 4 with the relaxed determinism condition, that is under the assumptions of Chapter 5. The question there is whether the auctions satisfying these conditions must also be in some sense similar or belong to some restrictive class.

Finally, another important direction would be to take the approach similar to that of Blumrosen and Nisan and develop the optimal incremental auctions. It is very interesting to see whether these optimal mechanisms would also have dominant strategy equilibria or any other properties of the mechanisms presented in this thesis.

6.3 Final remarks

The topic of this thesis dealt with one, although arguably somewhat important, problem faced in designing systems for facilitating interactions of a large number of agents. As the world becomes even more dependent on software agents for handling the familiar tasks such as finding the best vacation, getting your favorite seat in a concert or more critical applications such as stock trading or contract procurement, these problems will become more pronounced and the inefficiencies caused by them could cause a significant drag on market performance. Fortunately, a recent surge in research in this area gives hope that the most important of these problems will be solved or at least identified before they are encountered in real business or personal applications. It is the hope of this author that this work would give some, no matter how small, contribution to that objective.

Appendix A

Proofs and Derivations for Chapter 3

Proposition A.1 In *LP-TIOLI* it is a weakly dominant strategy of any player *i* with valuation v_i to bid the *least* price $p \in \{p_1, \ldots, p_{k-1}, p_k = 1\}$ which is at least as great as v_i. Formally, $\sigma_i(v_i) = \arg\min_p \{p \in \{p_1, \ldots, p_{k-1}, p_k = 1\} | p \geq v_i\}$. In case of a tie it is dominant to always refuse the offer.

Proof: The proof is similar to that of the dominant strategy in Vickrey auction, although slightly longer because of the possibility of ties.

Take any set *I* of players and select a player $i \in I$. Suppose that v_i is the valuation of player *i* and let $\sigma_i(v_i)$ be player *i*'s strategy (as defined above). Let $\sigma_i'(v_i) \neq \sigma_i(v_i)$ be any other strategy of *i*. Denote the strategies (not necessarily dominant) of all the other players as σ_{-i} and let p_{-i} be the maximum bid among these players, $p_{-i} = \max_{j \neq i} \sigma_j(v_j)$. We will show that it is suboptimal to follow $\sigma_i'(v_i)$.

There are three possible cases.

Case 1 $\sigma_i(v_i) > p_{-i}$, by following $\sigma_i(v_i)$ *i* pays p_{-i}, receives the good and obtains non negative utility. If *i* follows $\sigma_i'(v_i)$ then

(a) $\sigma_i'(v_i) > p_{-i}$, *i* wins with the same payment and allocation

(b) $\sigma_i'(v_i) < p_{-i} < \sigma_i(v_i)$, *i* loses, receives utility of 0

(c) $\sigma_i'(v_i) = p_{-i} < \sigma_i(v_i)$, *i* enters a tie, receives utility of at most 0

Case 2 $\sigma_i(v_i) < p_{-i}$, by following $\sigma_i(v_i)$ *i* loses and receives utility of 0. If *i* follows $\sigma_i'(v_i)$ then

(a) $\sigma_i'(v_i) > p_{-i} > \sigma_i(v_i)$, *i* wins and pays $p_{-i} > v_i$, obtains negative utility

(b) $\sigma_i'(v_i) = p_{-i} > \sigma_i(v_i)$, *i* enters a tie, receives utility of at most 0

(c) $\sigma_i'(v_i) < p_{-i}$, *i* loses, receives utility of 0

Case 3 $\sigma_i(v_i) = p_{-i}$, by following $\sigma_i(v_i)$ *i* enters a tie, receives utility of 0. If *i* follows $\sigma_i'(v_i)$ then

(a) $\sigma_i'(v_i) > \sigma_i(v_i)$, *i* wins and pays $p_{-i} > v_i$, obtains negative utility

(b) $\sigma_i'(v_i) < \sigma_i(v_i)$, *i* loses, receives utility of 0

This shows that the strategy $\sigma_i'(v_i) \neq \sigma_i(v_i)$ will always result in a utility which is less than or equal to the utility from following $\sigma_i(v_i)$, independent of the moves of other players. Therefore, $\sigma_i(v_i)$ is a dominant strategy of player *i*.

Lemma A.1 For any distribution $\Phi(x)$ it is always possible to find the set of threshold prices for a 2-bidder ***LP-TIOLI*** such that $ER_2^L(p_1, ..., p_k) > ER_2^V$.

Proof: Recall that $G(x)$ was defined as $G(x) = x * (1 - \Phi(x))$ and p_r is the smallest number in (0,1) such that $G(p_r)$ is a global maximum of $G(x)$ on [0,1].

$$
\begin{aligned}
ER_2^V &= \int_0^1 2 * x * (1 - \Phi(x)) * \phi(x)dx \\
< \int_0^{p_r} 2 * G(p_r) * \phi(x)dx &+ \int_{p_r}^1 2 * x * (1 - \Phi(x)) * \phi(x)dx \\
= 2 * G(p_r) * (\Phi(p_r) - \Phi(0)) &+ \int_{p_r}^1 2 * x * (1 - \Phi(x)) * \phi(x)dx \\
&= \lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k)
\end{aligned}
$$

(assuming $p_2, \cdots, p_{k-1}$ are uniformly distributed on $(p_r, 1)$)

For any $\epsilon > 0$ there always exists k', $k' > 0$ and $k' - 1$ price thresholds $p'_2, p'_3, ..., p'_{k'}$ such that $\lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k) \le ER_2^L(p_r, p'_2, ..., p'_{k'}) + \epsilon$.

Since $ER_2^V < \lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k)$, set $\epsilon < \lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k) - ER_2^V$.

$$
\begin{aligned}
&\lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k) \le ER_2^L(p_r, p'_2, ..., p'_{k'}) + \epsilon. \\
&\lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k) < ER_2^L(p_r, p'_2, ..., p'_{k'}) + \lim_{k\to\infty} ER_2^L(p_r, p_2, ..., p_k) - ER_2^V. \\
&\qquad\qquad ER_2^V < ER_2^L(p_r, p'_2, ..., p'_{k'})
\end{aligned}
$$

Therefore there exists k' and $k' - 1$ price thresholds such that ***LP-TIOLI*** has a higher expected revenue than standard Vickrey.

Lemma A.2 To optimize the seller's revenue with 2-bidders, the lowest price threshold p_1 of the ***LP-TIOLI*** mechanism should be set to at least p_r (set p_1 such that $p_1 \ge p_r$).

Proof: We will show that for any k, any $l < p_r$ and any set of $k-1$ price thresholds $p_2, \cdots p_k$, $ER_2^l(p_r, p_2, \cdots, p_k) > ER_2^l(l, p_2, \cdots, p_k)$

$$
\begin{aligned}
& ER^{l}{}_{2}(l, p_2, \cdots, p_k) - ER^{l}_{2}(p_r, p_2, \cdots, p_k) \\
= \quad & 2 * l * (1 - \Phi(l)) * \Phi(l) + 2 * p_2 * (1 - \Phi(p_2)) * (\Phi(p_2) - \Phi(l)) \\
& - 2 * p_r * (1 - \Phi(p_r)) * \Phi(p_r) - 2 * p_2 * (1 - \Phi(p_2)) * (\Phi(p_2) - \Phi(p_r)) \\
= \quad & 2 * G(l) * \Phi(l) - 2 * G(p_r) * \Phi(p_r) \\
& + 2 * G(p_2) * (\Phi(v_2) - \Phi(l) - \Phi(p_2) + \Phi(p_r)) \\
= \quad & 2 * G(l) * \Phi(l) - 2 * G(p_r) * \Phi(p_r) + 2 * G(p_2) * (\Phi(p_r) - \Phi(l)) \\
< \quad & 2 * G(p_r) * \Phi(l) - 2 * G(p_r) * \Phi(p_r) + 2 * G(p_r) * (\Phi(p_r) - \Phi(l)) \\
& \qquad \text{Since } G(p_r) > G(l),\ G(p_r) \geq G(p_2) \text{ and } \Phi(p_r) \geq \Phi(l) \\
= \quad & 2 * G(p_r)(\Phi(l) - \Phi(p_r)) + 2 * G(p_r) * (\Phi(p_r) - \Phi(l)) \\
= \quad & 0
\end{aligned}
$$

Appendix B

Proofs and Derivations for Chapter 4

Proposition B.1 All single-good, iterative auctions with a *dominant strategy equilibrium* satisfying the conditions listed in section 4.2 are *increasing price mechanisms*.

Proof: Suppose there exists some iterative single-good mechanism that satisfies the requirements of section 4.2 and has a *dominant strategy equilibrium* but does not belong to *increasing price mechanisms*. Then for some player *i* there exists an opponents' strategy profile σ_{-i} such that:

1. The earliest iteration at which *i* can win against σ_{-i} is t.
2. If *i* wins at iteration t she is required to make a payment $p_1 < v_i^{max}$.
3. It is possible to *i* to win against σ_{-i} at some iteration $k > t$.
4. If *i* wins at iteration k she is required to make a payment $p_2 < p_1$.

Suppose *i* has a dominant strategy $\sigma_i(v_i)$, for any possible valuation v_i. Then as long as player *i*'s valuation v_i is greater than p_1 it is optimal for *i* to win at some iteration greater than t. Which means that playing against σ_{-i} the dominant strategy $\sigma_i(v_i)$ would dictate a bid at iteration t that would not result in immediate win, but instead would move the game into the next iteration.

Construct an alternative opponents strategy σ'_{-i} such that $\forall r \leq t$, $\sigma'_{-i} = \sigma_{-i}$ and for every iteration r greater than t it must be the case that $\forall j \neq i, \forall v_j$, $\sigma_j(v_j)' = \max_b b \in M^r$. This implies that σ'_{-i} is equivalent to σ_{-i} up to iteration t, but after t all opponents of *i* are bidding the greatest allowed message at each iteration.

If player *i* with any valuation $v_i > p_1$ uses her dominant strategy $\sigma_i(v_i)$ against σ'_{-i} she would not be allocated the object, since she would either lose or be eternally tied. Therefore, her utility from using $\sigma_i(v_i)$ would equal 0, alternatively if she deviates from $\sigma_i(v_i)$ and wins at iteration t (which is possible since by construction $\forall r \leq t\, \sigma'_{-i} = \sigma_{-i}$) player *i* would receive the good and obtain positive utility, since $v_i > p_1$. Therefore, $\sigma_i(v_i)$ is not a dominant strategy. This means that there does not exist an iterative single-good mechanism with *dominant strategy equilibrium*, satisfying the requirements of section 4.2 that does not belong to *increasing price mechanisms*.

Definition B.1 As was stated in Section 4.2.2 the sequences of messages received from the players are comparable. Therefore, the allocation and payment decisions have to be made based on the full sequences of messages (and not individual messages). To formalize that we will use $x(s_1[t], s_2[t], \cdots, s_n[t]) \to I \cup \{NI, NA\}$ to denote the allocation function, which, based on the message sequence, either assigns the object to one of the players, moves the game to the next iteration (*NI*), or stops the game without allocating the object (*NA*). In this definition $s_i[t]$ is a length *t* message sequence of player *i*. If the object is allocated the function $\tau_t(s_1[t], s_2[t], \cdots, s_n[t])$ is used to calculate the payment, here *t* is the iteration at which the allocation is made (that is the payment function can be different depending on the iteration). Given that the object is allocated to player *i* at iteration *t* we will use a simpler notation to denote player *i*'s payment, $\tau_t(s_i[t], s_{-i}[t])$, where $s_{-i}[t]$ is the set of message sequences of all players except *i*.

Proposition B.2 Let player i have dominant strategy σ_i and valuation v_i. Suppose there exist opponents' strategies σ_{-i} and valuations v_{-i} such that if $\sigma_i(v_i)$ is played against $\sigma_{-i}(v_{-i})$, the mechanism terminates at iteration t with i winning and paying $p_i < v_i$. Then for any $v_i' \geq v_i$ and any dominant strategy σ_i', if i plays $\sigma_i'(v_i')$ against $\sigma_{-i}(v_{-i})$, it must be the case that i wins and pays p_i (as with $\sigma_i(v_i)$)

Proof: We will use the notation as described in Definition B.1.

Suppose there exists a valuation $v_i' \geq v_i$ and a dominant strategy $\sigma_i'(v_i')$ of player *i* such that $\sigma_i'(v_i')$ wins against $\sigma_{-i}(v_{-i})$ at some iteration *k* with some payment $p_i' \neq p_i$. Therefore, $\tau_t(s_i[t], s_{-i}[t]) \neq \tau_k(s_i'[k], s_{-i}'[k])$ where $s_i[t]$ and $s_{-i}[t]$ are the message sequences resulting from using $\sigma_i(v_i)$ against $\sigma_{-i}(v_{-i})$, similarly $s_i'[k], s_{-i}'[k]$ are the message sequences resulting from using $\sigma_i'(v_i')$ against $\sigma_{-i}(v_{-i})$.

Case 1 $\tau_t(s_i[t], s_{-i}[t]) < \tau_k(s_i'[k], s_{-i}'[k])$

Therefore $u_i(x(s_i[t], s_{-i}[t]), \tau_t(s_i[t], s_{-i}[t]), v_i') > u_i(x(s_i'[k], s_{-i}'[k]), \tau_k(s_i'[k], s_{-i}'[k]), v_i')$ since the utility functions are assumed to be quasi-linear (Definition 2.18). This means that it is better for player *i* with valuation v_i' to follow $\sigma_i(v_i)$ and not $\sigma_i'(v_i')$. Therefore, $\sigma_i'(v_i')$ can't be a dominant strategy of player *i*. Contradiction.

Case 2 $\tau_t(s_i[t], s_{-i}[t]) > \tau_k(s_i'[k], s_{-i}'[k])$

Therefore $u_i(x(s_i[t], s_{-i}[t]), \tau_t(s_i[t], s_{-i}[t]), v_i) < u_i(x(s_i'[k], s_{-i}'[k]), \tau_k(s_i'[k], s_{-i}'[k]), v_i)$. This implies that for player *i* with valuation v_i it is better to follow $\sigma_i'(v_i')$ and not $\sigma_i(v_i)$. Therefore, $\sigma_i(v_i)$ can't be a dominant strategy of player *i*. Contradiction.

Proposition B.3 Let player i have dominant strategy σ_i and valuation v_i. Suppose there exist opponents' strategies σ_{-i} and valuations v_{-i} such that if $\sigma_i(v_i)$ is played against $\sigma_{-i}(v_{-i})$, the mechanism terminates at iteration t with i winning and paying $p_i < v_i$. Then for any $v_i' \geq v_i$ and any dominant strategy σ_i', if i plays $\sigma_i'(v_i')$ against $\sigma_{-i}(v_{-i})$, it must be the case that the mechanism terminates at iteration t (as with $\sigma_i(v_i)$).

Proof: We will use the notation as described in Definition B.1.

Suppose there exists a valuation $v_i' \geq v_i$ and a dominant strategy $\sigma_i'(v_i')$ of player *i* such that $\sigma_i'(v_i')$ wins against $\sigma_{-i}(v_{-i})$ at some iteration $k \neq t$. From Proposition B.2 know that the payment in this case will equal p_i.

Case 1 $k > t$

Construct the opponents strategy profile σ'_{-i} such that $\forall r \leq t,\ \sigma'_{-i} = \sigma_{-i}$ and for every iteration r greater than t it must be the case that $\forall j \neq i, \forall v_j,\ \sigma_j(v_j)' = \max_b b \in M^r$.

Since, $\sigma_i'(v_i')$ played against $\sigma_{-i}(v_{-i})$ wins at some iteration greater than *t*, it must be the case that $\sigma_i'(v_i')$ can't win against $\sigma'_{-i}(v_{-i})$ at any iteration earlier than $t+1$. Therefore, if $\sigma_i'(v_i')$ is played against $\sigma'_{-i}(v_{-i})$ it can never win (and get the good). This means that the utility *i* receives from this game is 0, which is less than *i* receives if it follows $\sigma_i(v_i)$, since $p_i < v_i \leq v_i'$, by assumption. Therefore, $\sigma_i'(v_i')$ can't be a dominant strategy of player *i* with valuation v_i'. Contradiction.

Case 2 $k < t$

By Proposition 4.1 in order to have a dominant strategy the mechanism must belong to the space of *increasing price mechanisms*. Noting this fact the remainder of the proof is very similar to Case 1, construct an appropriate σ'_{-i} to show that $\sigma_i(v_i)$ can't be a dominant strategy of player *i* with valuation v_i, which leads to a contradiction.

Proposition B.4 Let player *i* have dominant strategy σ_i and valuation v_i. Suppose there exists some σ_{-i} and v_{-i} such that, if *i* plays $\sigma_i(v_i)$ against $\sigma_{-i}(v_{-i})$ (inducing history h_i^{t-1}), then $LPI(v_i, h_i^{t-1}) = \mathit{future}$ and player *i* is active at iteration *t*. Then for any other $v_i' \geq v_i$ and any dominant strategy σ_i' it must be the case that if *i* plays $\sigma_i'(v_i')$ against $\sigma_{-i}(v_{-i})$ (inducing history $\hat{h}_i^{t-1}$), we must have $h_i^r = \hat{h}_i^r$ and $\sigma_i(v_i, h_i^r) = \sigma_i'(v_i', \hat{h}_i^r)$, for all $r \leq t$.

Proof: It is assumed that $LPI(v_i, h_i^{t-1}) = \mathit{future}$. Therefore by the definition of LPI there exists some opponents' strategy profile $\widehat{\sigma}_{-i}$ and valuations $\widehat{v}_{-i}$ such that playing $\sigma_i(v_i)$ against

$\widehat{\sigma}_{-i}(\widehat{v}_{-i})$ induces the history identical to h_i^{t-1} up to iteration *t*. Furthermore, $\sigma_i(v_i)$ wins against $\widehat{\sigma}_{-i}(\widehat{v}_{-i})$ at some iteration $k > t$ with payment p_i smaller than v_i.

Take any other valuation $v_i' \geq v_i$ and any dominant strategy σ_i' for player *i*. Since $v_i' \geq v_i$ it is profitable for player *i* with this valuation to win against $\widehat{\sigma}_{-i}(\widehat{v}_{-i})$. Therefore, $\sigma_i'(v_i')$ wins against $\widehat{\sigma}_{-i}(\widehat{v}_{-i})$ since σ_i' is a dominant strategy. By Proposition B.2 and Proposition B.3 $\sigma_i'(v_i')$ will win at iteration k with payment p_i (same payment and iteration as with $\sigma_i(v_i)$).

Player *i* gets identical outcomes if it uses $\sigma_i(v_i)$ or $\sigma_i'(v_i')$ against $\widehat{\sigma}_{-i}(\widehat{v}_{-i})$. Therefore, with both strategies player *i* has to remain active but not win at all iterations r, $0 \leq r < k$. By the *limited participation* assumption, this implies that player *i* has to report a bid equal to the highest bid of her opponents, at each of these iterations.

Consider the following inductive argument on the number of iterations. Let h_{-i}^{r-1} be the history known to player *i*'s opponents before iteration *r* (with some facts possibly being privately known only by some opponents). Let $m_{-i}^r(h_{-i}^{r-1})$ be the maximum message submitted by player *i*'s opponents at iteration *r* based on history h_{-i}^{r-1}.

Base case: before the first iteration the history h_l^{-1} is empty for all $l \in I$. Therefore, for player *i* to remain active after iteration 0, she had to submit a message equal to $m_{-i}^r(\varnothing)$, which means that $\sigma_i(v_i, \varnothing) = \sigma_i'(v_i', \varnothing) = m_{-i}^r(\varnothing)$, furthermore $\forall l \in I$ $h_l^0 = \hat{h}_l^0$ where h_l^0 and $\hat{h}_l^0$ are the histories known to player *l* after $\sigma_i(v_i)$ or $\sigma_i'(v_i')$ is used at iteration 0 against $\widehat{\sigma}_{-i}(\widehat{v}_{-i})$, respectively.

Inductive step: assume that at some iteration $r < k-1$, $\sigma_i(v_i, h_i^{r-1}) = \sigma_i'(v_i', \hat{h}_i^{r-1}) = m_{-i}^{r-1}(h_{-i}^{r-1})$ and $h_l^r = \hat{h}_l^r$ $\forall l \in I$, which means that player *i* remains active after iteration *r*. For player *i* to remain active after iteration $r+1$ she needs to submit a bid equal to $m_{-i}^{r+1}(h_{-i}^r)$. Therefore, $\sigma_i(v_i, h_i^r) = \sigma_i'(v_i', h_i^r) = m_{-i}^{r+1}(h_{-i}^r)$ and $\forall l \in I,\ h_l^{r+1} = \hat{h}_l^{r+1}$.

This argument shows that both strategies when used against the same opponents will produce the same bids and therefore the same histories for all iterations less than *k*, therefore this will also holds for all iterations r, $0 \leq r \leq t$, since $t < k$.

Proposition B.5 Suppose i has a dominant strategy. Suppose that for some history h_i^{t-1} and valuation v_i, $LPI(v_i, h_i^{t-1}) = now$. Then there is a dominant strategy σ_i in which $\sigma_i(v_i, h_i^k) = minbid$ for any $k > t+1$ and history h_i^k s.t. $h_i^k[t-1] = h_i^{t-1}$.

Proof: The proof comes directly from the definition of *LPI* (Definition 4.4). Since $LPI(v_i, h_i^{t-1}) = now$, if player *i* wins at any iteration greater than *t*, she will pay a price at least as high as her valuation. It is therefore weakly dominant for *i* to ensure the object is never won at any iteration later than *t*.

The *termination* assumption, Section 4.2.3, states that the object is allocated only to a *unique* highest bidder, therefore bidding $minbid$ ensures that *i* never wins. This means that if there exists a dominant strategy σ_i' for player *i*, then it is possible to construct an alternative dominant strategy σ_i for *i* which is identical to σ_i', except σ_i submits a bid of $minbid$ at every iteration past *i*'s point of *LPI* ($LPI(v_i, h_i^{t-1}) = now$).

Proposition B.6 Given a *fully symmetric* auction mechanism with dominant strategy equilibrium, if σ_i is a dominant strategy of player *i*, then it is also a dominant strategy of any other player $j \neq i$.

Proof: Suppose this is not true. Then there must exist some opponents' strategies σ_{-j} and a player *j* with valuation v and a dominant strategy $\sigma_j \neq \sigma_i$ such that $u_j(\sigma_j, \sigma_{-j}, v) > u_j(\sigma_i, \sigma_{-j}, v)$ (we use a simplified notation where $u_i(\sigma_i, \sigma_{-i}, v_i)$ denotes $u_i(x(\sigma_i, \sigma_{-i}), \tau(\sigma_i, \sigma_{-i}), v_i)$)

The utilities are assumed to be quasi-linear, so a utility for an outcome of player *i* is a difference between *i*'s valuation for the object (if she receives it) and the payment *i* has to make. Since the auction mechanism is *fully symmetric* this implies that $u_i(\sigma_j, \sigma_{-j}, v) = u_j(\sigma_j, \sigma_{-j}, v)$ and $u_i(\sigma_i, \sigma_{-j}, v) = u_j(\sigma_i, \sigma_{-j}, v)$.

Therefore,

$$\begin{aligned} u_i(\sigma_j, \sigma_{-j}, v) = u_j(\sigma_j, \sigma_{-j}, v) > \\ u_j(\sigma_i, \sigma_{-j}, v) = u_i(\sigma_i, \sigma_{-j}, v) \end{aligned}$$

This implies that $u_i(\sigma_j, \sigma_{-j}, v) > u_i(\sigma_i, \sigma_{-j}, v)$, which is impossible since σ_i is a dominant strategy for player *i*. Contradiction.

Proposition B.7 Under the rules of ASIA, it is a weakly dominant strategy of any player *i* with valuation v_i to bid 1 at iteration *t* as long as $v_i > p^t$, and to bid 0 otherwise.

Proof: Suppose that at some iteration *t* the set of active players is A (that is all players in A reported 1 at iteration $t-1$) and the mechanism announces some price p_t. Take any player $i \in A$.

Suppose *i* does not follow the above described strategy, that is she will report 1 when the above strategy dictates 0, and/or 0 when it dictates 1. Denote the alternative strategy of *i* by σ_i' and the proposed dominant strategy by σ_i. As before σ_{-i} denotes the strategy profile of other players, h_i^{t-1} the history known to player *i* after the end of iteration $t-1$ and h_{-i}^{t-1} the history known to all the other players. We will prove that σ_i is indeed the dominant strategy for player i.

Case 1 $v_i > p_t$, player *i* reports 0, when following σ_i'.

Given the message 0, player *i* drops out of the auction and $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0, \forall \sigma_{-i}$. However, there exists a sequence of moves of other players σ_{-i}^* where $\forall \sigma_j^* \in \sigma_{-i}^*, \sigma_j^*(v_j, h_j^{t-1}) = 0$ (all opponents of *i* bid 0 at *t*), therefore since $v_i > p_t$, $u_i(\sigma_i(v_i, h_i^{t-1}), \sigma_{-i}^*, v_i) > 0$. Given any other $\sigma_{-i} \neq \sigma_{-i}^*$ if the player reports 1 the auction will move into the next iteration with player *i* active.

Case 2 $v_i \leq p_t$, player *i* reports 1, when following σ_i'.

Given the message 1, player *i* either wins the auction or stays active at the next iteration. Similarly to above, there exists a sequence of moves of other players σ_{-i}^* where $\forall \sigma_j^* \in \sigma_{-i}^*, \sigma_j^*(v_j, h_j^{t-1}) = 0$. Given this sequence, $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}^*, v_i) \leq 0$, since $v_i \leq p_t$. Furthermore, since *ASIA* is an increasing price mechanism it is impossible for i to win with non negative utility in any future iteration. With σ_i however, $u_i(\sigma_i(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0, \forall \sigma_{-i}$.

Cases 1 and 2 show that at any iteration σ_i archives the maximum possible utility for player *i* independent of the actions of player *i*'s opponents, therefore σ_i is the dominant strategy for player *i*.

Proposition B.8 It is a weakly dominant strategy of any player *i* in SPSIA to bid 1 at iteration *t* as long as $v_i > b_t$ (equivalent to $LPI(h_i^t, v_i) = \textit{future}$). Otherwise, the player should bid 0. If at any point the player is asked to choose from some set of prices she should pick the (announced) price which is just above her valuation.

Proof: The above strategy describes a player's actions for both stages of the mechanism. The only time a player can be asked to pick among a set of prices (the second stage) is when she announced 0 (a winner is never asked to refine her bid). At this point the game becomes fully equivalent to *LP-TIOLI* auction. Therefore, by Proposition 3.1 it is weakly dominant in this case for each player to bid the price level just above her valuation. Given this fact we only need to prove that the other portion of the above strategy also constitutes a weakly dominant strategy.

Suppose that at some iteration *t* the set of active players is A (that is all players in A reported 1 at iteration $t-1$) and the mechanism announces some price $[a_t, b_t]$. Take any player $i \in A$.

Suppose *i* does not fully follow the above described strategy, that is she will report 1 when the above strategy dictates 0, and/or 0 when it dictates 1 (but will follow it in the second stage of the game). Denote the alternative strategy of *i* by σ_i' and the proposed dominant strategy by σ_i. As before σ_{-i} denotes the strategy profile of other players, h_i^{t-1} the history known to player *i* after the end of iteration $t-1$ and h_{-i}^{t-1} the history known to all other players. We will prove that σ_i is indeed the dominant strategy for player *i*.

Case 1 σ_{-i} is such that there exists at least one player in A_{-i} that bids 1 at iteration t

Part A $v_i > b_t$, player *i* reports 0, when following σ_i'.

Since i reports 0, and there is at least one other reporting 1, $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0 \; \forall \sigma_{-i}$. Following σ_i would result in i moving to the next iteration, with non negative utility, since σ_i dictates quitting only when $v_i \leq b_k$ for any $k > t$.

Part B $v_i \leq b_t$, player *i* reports 1, when following σ_i'.

This combination of messages moves i into the next iteration. By the design of the mechanism, $a_{t+k} \geq b_t, \forall k > 0$, therefore $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) \leq 0 \; \forall \sigma_{-i}$, but $u_i(\sigma_i(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0 \; \forall \sigma_{-i}$

Case 2 σ_{-i} is such that all players in A_{-i} bid 0 at iteration t

Part A $v_i > b_t$, player *i* reports 0, when following σ_i'.

Since $v_i > b_t$ and we assumed that in this situation σ_i' will act according to the dominant strategy of $LP - TIOLI$ auction, the utility from following σ_i' would, in this case, equal to the utility obtained from following σ_i. That is, $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = u_i(\sigma_i(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = v_i - m_2 > 0$ if $m_2 < \sigma_i(v_i, h_i^{t-1})$ and 0 otherwise. Where m_2 is the highest bid from players in A_{-i} at the second stage of the mechanism.

Part B $v_i \leq b_t$, player *i* reports 1, when following σ_i'.

If $v_i < b_t$ then there exists a set of moves of other players and a set of prices such that player *i* will win the auction and pay $m_2 > v_i$, with $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) < 0$. Following σ_i, on the other hand, will always insure a non negative utility, since *i* is allowed to select her bid, instead of having a mechanism enter it for her. Finally, since the payment of player *i* (if she wins) does not depend on her actions, a strategy that always enters the same bid can't do better than the strategy that provides a set of possibilities including the fixed bid of the other strategy.

Cases 1(A,B) and 2(A,B) show that at any iteration σ_i archives the maximum possible utility for player *i* independent of the actions of player *i*'s opponents, therefore σ_i is the dominant strategy for player *i*.

Appendix C

Proofs and Derivations for Chapter 5

Proposition C.1 In the ST-ASIA mechanism, it is a weakly dominant strategy of any player *i* with valuation v_i to bid 1 at iteration *t* as long as $v_i > p^t$, and to bid 0 otherwise.

Proof: The proof of this proposition is very similar to the proof of Proposition 4.6

Suppose that at some iteration *t* the set of active players is A (that is all players in A reported 1 at iteration $t-1$) and the mechanism announces some price p_t. Take any player $i \in A$.

Suppose *i* does not follow the above described strategy, that is she will report 1 when the above strategy dictates 0, and/or 0 when it dictates 1. Denote the alternative strategy of *i* by σ_i' and the proposed dominant strategy by σ_i. As before σ_{-i} denotes the strategy profile of other players, h_i^{t-1} the history known to player *i* after the end of iteration $t-1$ and h_{-i}^{t-1} the history known to all other players. We will prove that σ_i is indeed the dominant strategy for player *i*. Recall that at each iteration the mechanism randomly selects one active player and initially elicits a bid from every player except this one.

Case 1 Player *i* is not the selected player, $v_i > p_t$, player *i* reports 0, when following σ_i'.

Given the message 0, player *i* drops out of the auction and $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0, \forall \sigma_{-i}$. However, there exists a sequence of moves of other players σ_{-i}^* where $\forall \sigma_j^* \in \sigma_{-i}^*, \sigma_j^*(v_j, h_j^{t-1}) = 0$ (all opponents of *i* bid 0 at *t*), so since

$v_i > p_t$, $u_i(\sigma_i(v_i, h_i^{t-1}), \sigma_{-i}^*, v_i) > 0$. Given any other $\sigma_{-i} \neq \sigma_{-i}^*$ the auction will move into the next iteration with player *i* active.

Case 2 Player *i* is not the selected player, $v_i \leq p_t$, player *i* reports 1, when following σ_i'.

Given the message 1, player *i* either wins the auctions or stays active at the next iteration. Similarly to above, there exists a sequence of moves of other players σ_{-i}^* where $\forall \sigma_j^* \in \sigma_{-i}^*, \sigma_j^*(v_j, h_j^{t-1}) = 0$. Given this sequence, $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}^*, v_i) \leq 0$, since $v_i \leq p_t$. Furthermore, since *ASIA* is an increasing price mechanism it is impossible for *i* to win with positive utility in any future iteration. With σ_i however, $u_i(\sigma_i(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0, \forall \sigma_{-i}$.

Case 3 Player *i* is the selected player, $v_i > p_t$, player *i* reports 0, when following σ_i'.

The fact that player *i* is asked for a bid implies that at least one other player announced 1. Given message 0, player *i* drops out of the auction and $u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0 \; \forall \sigma_{-i}$. However as in Case 1, if *i* reports 1 then at the next iteration the utility of *i* would be $\forall \sigma_{-i},\ u_i(\sigma_i(v_i, h_i^t), \sigma_{-i}, v_i) \geq 0$, assuming *i* continues to follow σ_i.

Case 4 Player *i* is the selected player, $v_i \leq p_t$, player *i* reports 1, when following σ_i'.

Announcing 1 moves player to the next iteration. If the player is not selected at the next iteration and she bids 1 again than as in Case 2, there exists a sequence of moves of other players such that $u_i(\sigma_i'(v_i, h_i^t), \sigma_{-i}^*, v_i) < 0$, for any other opponents' moves $u_i(\sigma_i'(v_i, h_i^t), \sigma_{-i}, v_i) = 0$. Similarly, if she is selected then for any set of opponent moves σ_{-i}, $u_i(\sigma_i'(v_i, h_i^t), \sigma_{-i}, v_i) \leq 0$. In any case this is inferior to bidding 0 which assures $\forall \sigma_{-i},\ u_i(\sigma_i'(v_i, h_i^{t-1}), \sigma_{-i}, v_i) = 0$

Cases 1, 2, 3 and 4 show that at any iteration σ_i archives the maximum possible utility for player *i* independent of the actions of player *i*'s opponents, therefore σ_i is the dominant strategy for player *i*.

Lemma C.1 Under *Divide(k)* the price at iteration t can be expressed as $p_t = 1 - (\frac{k-1}{k})^{t+1}$.

Proof: *Divide(k)* is defined as $p_0 = \frac{1}{k}$ at iteration 0 and $p_t = p_{t-1} + \frac{1-p_{t-1}}{k}$ at any iteration $t > 0$.

It is possible to rewrite p_t as $p_t = \frac{1}{k} + p_{t-1} * \frac{k-1}{k} = p_0 + z * p_{t-1}$ where $z = \frac{k-1}{k}$. Expanding this get $p_t = p_0 + z * p_{t-1} = p_0 + z * (p_0 + z * p_{t-2}) = p_0 + z * p_0 + z^2 * (p_0 + z * p_{t-3}) = p_0 + z * p_0 + z^2 * p_0 + \cdots + z^t * p_{t-t} = p_0 * (1 + z + z^2 + \cdots + z^t)$

Therefore, since $z < 1$ it is the case that $p_0 * \sum_{r=0}^{t} z^t = p_0 * \frac{1-z^{t+1}}{1-z} = \frac{1}{k} * \frac{1-(\frac{k-1}{k})^{t+1}}{1-\frac{k-1}{k}} = 1 - (\frac{k-1}{k})^{t+1}$.

Proposition C.2 Suppose that the players' valuations are IID according to uniform distribution. The expected total number of bits sent by each player participating in *ST-ASIA* with unrestricted number of iterations and price update rule *Divide(k)* is upper bounded by k.

Proof: From Lemma C.1 know that the price at iteration t under *Divide(k)* can we written as $p_t = 1 - (\frac{k-1}{k})^{t+1}$ or $p_t = 1 - z^{t+1}$ if we let $z = \frac{k-1}{k}$.

Suppose there are n players initially participating in the auction. Due to the structure of *ST-ASIA* a player sends a bit at iteration t only if she is active at t. Therefore, since a player can send at most one bit at any iteration, the expected number of bits sent by a player at iteration t is equal to probability the player is active at iteration t.

$$\begin{aligned}
\mathit{Pr}(\text{player } i \text{ active at } t) &= \mathit{Pr}(v_i > p_{t-1} \text{ and } \exists j \neq i \text{ s.t. } v_j > p_{t-1}) \\
&= \mathit{Pr}(v_i > p_{t-1}) * \mathit{Pr}(\exists j \neq i \text{ s.t. } v_j > p_{t-1}) \\
&= \mathit{Pr}(v_i > p_{t-1}) * (1 - \mathit{Pr}(\forall j \neq i\ v_j \leq p_{t-1})) \\
&= (1 - p_{t-1}) * (1 - (p_{t-1})^{n-1}) \\
&= 1 - p_{t-1} - (1 - p_{t-1}) * (p_{t-1})^{n-1} \\
&< 1 - p_{t-1} = 1 - (1 - z^t) = z^t
\end{aligned}$$

To calculate the total expected amount of communication need to sum up the above over all iterations.

$$\begin{aligned}\sum_{t=0}^{\infty} z^t &= 1 + \sum_{t=1}^{\infty} z^t \\ &= 1 + \frac{z}{1-z} = 1 + \frac{\frac{k-1}{k}}{1 - \frac{k-1}{k}} \\ &= 1 + \frac{\frac{k-1}{k}}{\frac{1}{k}} = k\end{aligned}$$

Note that this value is independent of n, the total number of players.

Bibliography

[1] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communication. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS 02), Vancouver, Canada*, pages 406–416, November 2002.

[2] Liad Blumrosen, Noam Nisan, and Ilya Segal. Multi-player and multi-round auctions with severely bounded communication. In *11th Annual European Symposium on Algorithms (ESA 03), Budapest, Hungary*, September 2003.

[3] Craig Boutilier. A pomdp formulation of preference elicitation problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence, Edmonton, AB*, pages 239–246, 2002.

[4] Craig Boutilier. Solving concisely expressed combinatorial auction problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-2002), Edmonton, AB*, pages 359–366, 2002.

[5] Craig Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.

[6] Craig Boutilier and Holger H. Hoos. Bidding languages for combinatorial auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01), Seattle*, pages 1211–1217, 2001.

[7] Wolfram Conen and Tuomas Sandholm. Partial-revelation vcg mechanism for combinatorial auctions. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, pages 367–372, Edmonton, 2002.

[8] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI 02), Edmonton, AB*, pages 103–110, August 2002.

[9] Vincent Conitzer and Tuomas Sandholm. Computational criticisms of the revelation principle. In *Proceedings of the AAMAS-03 Workshop on Agent-Mediated Electronic Commerce V (AMEC V), Melbourne, Australia*, July 2003.

[10] S. DeVries and R. Vohra. Combinatorial auctions: A survey, 2000.

[11] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–602, 1973.

[12] Jerry R. Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

[13] Carlos Guestrin, Daphne Koller, and Ronald Parr. Multiagent planning with factored mdps. In *Proceeding of the 14th Neural Information Processing Systems (NIPS-14), Vancouver, BC*, pages 1523–1530, December 2001.

[14] David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.

[15] Vijay Krishna. *Auction Theory*. Academic Press, 2002.

[16] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[17] Kate Larson and Tuomas Sandholm. An alternating offers bargaining model for computationally limited agents. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 135–142, Bologna, 2002.

[18] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce, Minneapolis*, pages 233–241, October 2000.

[19] Peter Marbach. Analysis of a static pricing scheme for priority services. *IEEE/ACM Transactions on Networking*, 12(2):312 – 325, 2004.

[20] Andreu Mas-Colell, Micheal D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[21] Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive biding. *Econometrica*, 50:1089–1122, 1982.

[22] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[23] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce, Minneapolis*, pages 1–12, October 2000.

[24] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[25] David C. Parkes. ibundle: An efficient ascending price bundle auction. In *Proceedings of the ACM Conference on Electronic Commerce, Denver*, pages 148–157, November 1999.

[26] David C. Parkes. Optimal auction design for agents with hard valuation problems. In *Proceedings of the Agent Mediated Electronic Commerce (IJCAI Workshop), Stockholm, Sweden*, pages 206–219, July 1999.

[27] David C. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and AI*, 2003. to appear.

[28] Anatol Rapoport. *Decision Theory and Decision Behaviour*. Macmillan Press Ltd, 2 edition, 1998.

[29] Tim Roughgarden. Stackelberg scheduling strategies. In *Proceedings of ACM Symposium on Theory of Computing, Crete, Greece*, pages 104–113, July 2001.

[30] Tim Roughgarden and Eva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[31] Tuomas Sandholm. Automated mechanism design: A new application area for search algorithms. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP 03), Kinsale, County Cork, Ireland*, October 2003.

[32] Tuomas Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *Proceedings of the AAMAS-03 Workshop on Agent-Mediated Electronic Commerce V (AMEC V), Melbourne, Australia*, July 2003.

[33] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.