Auction Theory in Computational Settings

Thesis submitted for the degree of

"Doctor of Philosophy"

By

Aharon Ron Lavi

Submitted to the Senate of the Hebrew University

August 2004

This work was carried out under the supervision of

Prof. Noam Nisan

Acknowledgments

It has been a great experience to work with Noam Nisan. Moral support, patience, help, understanding – Noam generously and open-heartedly awarded all these to me. It was so enriching to learn from Noam: the endless time he dedicated to show me how to shape the ideas expressed here, to improve results, and to overcome difficulties (all this with his amazing theoretical and conceptual abilities), are all truly appreciated. For all this, Noam, thanks!

I had the pleasure to learn from many people here at the Hebrew University. I would like to especially thank Yair Bartal, Daniel Lehmann, and Motty Perry, for always being willing to explain, answer questions, hear and comment on new ideas, and showing continuous interest in me and my work.

I will miss all the people here with whom I shared thoughts, ideas, and enjoyable (but also hard) times. A special warm thanks to Moshe Babaioff, Liad Blumrosen, Rica Gonen, and Ahuva Mu'alem, for countless hours of serious and foolish talks all over the globe, about computer science, and about the recent fresh gossip (please keep me updated). I really hope I will keep seeing you many more times.

My parents, Hanna and Yaacov, and my sister Vardit, supported me in the deepest sense of the word "support". I am sorry I have to go so far away... My dear Iris contributed to my life in so many ways, from which her help throughout my studies is just an example. Iris, without you, I could not have made it. All this is equally yours.

Abstract

This thesis studies distributed systems and distributed agent systems that are characterized by interactions among computers with different ownerships and incentives. We ask what happens when the input of an algorithm is kept by independent agents, acting selfishly to maximize their own utility. In such an environment, an agent might not fully cooperate with the algorithm, as a slight deviation from the behavior required by the protocol might increase the agent's own utility. The challenge is to design algorithms that have guaranteed performance with respect to the *true input* of the agents, despite the non-cooperative behavior.

One approach to this is to design a mechanism that is able to collect payments from the participants. The payments should be designed carefully, so that agents will indeed find it in their best interest to behave as required by the algorithm, according to their true input. Such algorithms are termed incentive compatible, or truthful.

A broad model that captures such questions is auctions. In an auction, a seller sells one, or several, items to a set of buyers. This is not just a pure economic model - it actually captures many traditional computational settings, like job scheduling, routing and bandwidth allocation, and much more. In this thesis I focus on the implications of the computational and algorithmic theory to the design of computerized auctions. I study possibilities and impossibilities that arise from the integration of new models and considerations from classic algorithmic theory with classic auction theory. Three groups of results are presented:

Competitive Analysis of Incentive Compatible Online Auctions.

We study auctions in a setting where the different bidders arrive at different times, and the auction mechanism is required to make decisions about each bid as it is received. We call such auctions, *on-line auctions*. While, in the economic literature, most auctions are analyzed using average-case analysis, we apply CS-related tools and give a competitive worst-case analysis of incentive compatible on-line auctions. We first characterize exactly those on-line auctions that are *incentive compatible*, i.e. where rational bidders are always motivated to bid their true valuation. We then embark on a competitive worst-case analysis of incentive compatible on-line auctions. We obtain several results, the cleanest of which is an incentive compatible on-line auction for a large number of identical items. We prove that this auction has the optimal competitive ratio, both in

terms of seller's revenue and in terms of the social welfare obtained. We also describe deterministic and randomized auctions for small numbers of items.

This set of results is based on a joint work with Noam Nisan.

Towards a Characterization of Truthful Combinatorial Auctions.

The classical VCG result, from the field of Mechanism Design, states that we can always ensure truthfulness when the algorithm reaches the exact optimal social welfare. Naturally, one wonders about other algorithmic goals. The motivation behind this question is twofold: First, in problems where the exact optimum is computationally hard, we would like to have approximation algorithms, instead of exact ones. Second, in many traditional computational models, our goal is not welfare maximization.

In the extreme case in which each player's input is just one secret number, then the answer to the above question is known: many recent positive examples for single parameter problems, that are not welfare maximizers, have been shown. However, in most interesting computational models players have multi-dimensional inputs. An excellent example, which we investigate in this study, is the model of combinatorial auctions. In combinatorial auctions, the seller has a set of different items, and bidders may value bundles of items. Since computing the exact optimum is NP-hard, Computer Scientists design approximation algorithms and practical heuristics. But none of these is known to be truthful.

We initiate an investigation whether it is at all possible to obtain truthful combinatorial auctions that are not welfare maximizers. We show that truthful mechanisms for combinatorial auctions (and related classes of problems) must be "almost welfare maximizers" if they also satisfy an additional requirement of "independence of irrelevant alternatives" (IIA). We also show that this condition is without loss of generality in some special cases. Along the way, we find a weak monotonicity condition that completely characterizes truthfulness for a large class of problems. Adding to it the extra condition of IIA results in an interesting "strong monotonicity" condition. Our main theorem has severe computational implications, as we show that "almost welfare maximizers" are as computationally hard as exact welfare maximizers. For combinatorial auctions, this implies that no truthful polynomial time combinatorial auction that satisfies IIA (and two more technical conditions) can reasonably approximate the welfare.

This set of results is based on a joint work with Ahuva Mu'alem and Noam Nisan.

Online Ascending Auctions for Gradually Expiring Items.

The growing intuitive understanding in the field is that the impossibility presented above is more robust than what we can currently prove. For many settings, it seems that truthful algorithms are rare. The need to find other suitable solutions therefore seems to be one of the important tasks of the community.

To investigate other solution concepts, we study a problem that provably forces us to take a different trail. From a conceptual point of view, instead of designing algorithms for which the players will be expected to take one specific behavior, we design auctions for which *many* selfish behaviors lead to an approximately optimal allocation. The new concept is that, although players are not expected to follow a specific behavior, but only one out of a set of behaviors, the outcome is still guaranteed to be close to optimal, for any choice the players make. From a game-theoretic point of view, we describe a complete hierarchy of "set equilibria" notions that aim to justify why will the players limit their choices to these sets of behaviors. We believe that these concepts will help to bypass the inherent difficulties of the truthfulness notion, in a way that suits the CS worst-case notions.

Specifically, we study online auction mechanisms for the allocation of M items that are identical to each other except for the fact that they have different "expiration times", and each item must be allocated before it expires. Players arrive at different times, and wish to buy one item before their deadline. In accordance with the main "theme" of this thesis, the main difficulty is that players act selfishly and may mis-report their values, deadlines, or arrival times. We begin by showing that the usual notion of truthfulness, where players follow a single dominant strategy, cannot be used in this case, since any (deterministic) truthful auction cannot obtain better than an M-approximation of the social welfare.

Therefore, instead of designing auctions in which players should follow a single strategy, we design two auctions that perform well under a wide class of selfish, "semi-myopic", strategies. For every combination of such strategies, the auction is associated with a different algorithm, and so we have a family of "semi-myopic" algorithms. We show that any algorithm in this family obtains a 3-approximation, and by this conclude that our auctions will perform well under any choice of such semi-myopic behaviors.

We then turn to provide a game-theoretic justification for acting in such a semi-myopic way. We suggest a new notion of "Set-Nash" equilibrium, where we cannot pin-point a single best-response strategy, but rather only a set of possible best-response strategies. We show that our auctions have a Set-Nash equilibrium which is all semi-myopic, hence guarantees a 3-approximation. We define a complete game-theoretic framework of similar "set equilibria" notions, and discuss their properties. As stated above, we believe that these notions will turn out to be useful in many settings where truthfulness does apply, and classic equilibria notions do not fit the CS framework.

This set of results is based on a joint work with Noam Nisan.

Contents

1 Introduction **1**

1.1 The Framework of Mechanism Design . . . 2

1.1.1 Social Choice Theory . . . 2

1.1.2 Non-Cooperative Games . . . 4

1.1.3 Mechanism Design . . . 7

1.1.4 Auction Theory . . . 11

1.2 Algorithmic Mechanism Design . . . 13

1.2.1 The Difference From Classical Mechanism Design . . . 13

1.2.2 Some Representative Results of the Field . . . 15

1.3 Overview of The Results of This Thesis . . . 16

2 Competitive Analysis of Incentive Compatible On-Line Auctions **21**

2.1 Exposition of Results . . . 21

2.2 On-Line Auctions . . . 23

2.2.1 The model . . . 23

2.2.2 Supply curves for on-line auctions . . . 25

2.3 Competitive Analysis . . . 28

2.3.1 A divisible good . . . 29

2.3.2 A randomized auction for k indivisible goods . . . 34

2.3.3 A deterministic auction for k indivisible goods . . . 35

2.4 Model Extensions . . . 37

2.5 Revenue Analysis for the Uniform Distribution . . . 38

3 Towards a Characterization of Truthful Combinatorial Auctions **41**

3.1 Introduction . . . 41

3.1.1 Motivation . . . 41

3.1.2 Characterizing Incentive Compatibility . . . 42

3.1.3 Our results . . . 44

3.2 Setting and Notations . . . 46

3.3 Truthfulness and Monotonicity . . . 48
3.3.1 Weak monotonicity . . . 48
3.3.2 W-MON characterizes truthfulness . . . 50
3.3.3 Strong monotonicity and IIA . . . 52
3.3.4 Equivalence of W-MON and S-MON . . . 54
3.4 Main Theorem . . . 56
3.4.1 Intuitive proof outline . . . 58
3.4.2 Proof of Theorem 3.3 . . . 60
3.5 The Implications for Combinatorial Auctions . . . 68
3.5.1 General Issues . . . 69
3.5.2 Approximation . . . 71
3.5.3 Polynomial-Time Computation . . . 73
3.6 Unrestricted Domains . . . 77
3.6.1 Connected references . . . 80
3.7 Deferred Proofs . . . 81
3.7.1 Proof of claim 3.2 . . . 81
3.7.2 Proof of claim 3.26 . . . 82
3.7.3 Additional Example for Theorem 3.4 . . . 82
3.7.4 The hardness of welfare maximization . . . 83

4 Online Ascending Auctions for Gradually Expiring Items 89
4.1 Overview of Results . . . 90
4.2 Model and Basic Definitions . . . 93
4.3 Two Online Ascending Auctions . . . 94
4.3.1 The Online Iterative Auction . . . 94
4.3.2 The Sequential Japanese Auction . . . 96
4.3.3 Semi-Myopic Algorithms . . . 97
4.4 The Impossibility of Truthful Approximations . . . 101
4.5 A Game-Theoretic Framework . . . 102
4.5.1 Set-Nash Equilibria . . . 103
4.5.2 Implementation in Set-Nash equilibria . . . 105
4.5.3 Ignorable Extensions of Games . . . 106
4.6 A Strategic Analysis of our Auctions . . . 107
4.6.1 Semi-Myopic Mechanisms . . . 107
4.6.2 Bad Examples . . . 110
4.6.3 The Online Iterative Auction has a Set-Nash Equilibrium . . . 111
4.6.4 The Sequential Japanese Auction has a Set-Nash Equilibrium . . . 114
4.7 Useful Properties of Offline Allocations and Matroids . . . 117

4.7.1 Some Useful Properties of Offline Allocations 118

5 Conclusions **121**

Bibliography **123**

Chapter 1

Introduction

With the growth of the Internet, a new kind of distributed systems and distributed agent systems have emerged: those are characterized by interactions among computers with different ownership and incentives. Many new algorithmic questions are being asked as a result of this new development. In contrast to the traditional assumption in Computer Science, that computers (or at least the "good" ones) will follow protocols and algorithm specifications, we now ask what happens when the input of the algorithm is kept by independent agents, acting selfishly to maximize their own utility? The algorithm cannot simply request the agents to reveal their private inputs in order to compute some global outcome, as some agent might yield some benefit from biasing his input, and by this slightly shifting the outcome in his favor. However, we would still like to design algorithms that have guaranteed performance with respect to the true input. Therefore, the algorithm should be carefully designed to motivate the agents to co-operate, while of-course still keeping the traditional considerations of performance and solution quality.

To answer such questions, computer scientists have begun to examine the effects of integrating micro-economic and game theoretic models with computational and algorithmic ones. General models of non-cooperative players (agents) are extensively studied in the fields of Microeconomics and Game Theory, and specifically in the theory of *Mechanism Design*. This theory aims to study how individuals with private preferences and information may be aggregated towards a "social-choice". However, much of this theory does not fit classic computational settings, because the desired goals and the studied models are different in large, and because crucial computational considerations are left out. Therefore, a new theoretical field which lies in the intersection of all these classical disciplines, commonly termed "algorithmic mechanism design" [80], has evolved. The first integration attempts were to simply "plug-in" known results and concepts from economy, and are now gradually taking the form of a more theoretical design and analysis of new models and algorithms. In the past few years we have seen new models and algorithms for multicast cost sharing [38], network routing [66, 96, 60, 90, 3], machine scheduling [80, 99, 2], combinatorial auctions and revenue maximizing auctions [93, 62, 39], and more.

One broad model, that captures many combinatorial questions, is auctions. In an auction, a seller sells one, or several, items to a set of buyers. Although it seems that this is a pure economic model, a crucial observation is that it actually captures and generalizes many traditional computational settings, like job scheduling, routing and bandwidth allocation, and much more. Although economists have studied auctions since the 60's, many new questions, and new models, arise when considering computational and algorithmic aspects.

In this thesis I study possibilities and impossibilities that result from this integration of algorithmic theory and auction theory. I suggest new models, new solution concepts, and describe both positive and negative results. I will begin by describing, in a slightly formal manner, the framework of mechanism design and auction theory, and, in a slightly informal manner, the new field of algorithmic mechanism design. The new results of this thesis are first briefly reviewed in the last section of this chapter, and are then fully presented in the next chapters.

1.1 The Framework of Mechanism Design

In order to fully understand the context of the contributions made by the field of algorithmic mechanism design, and, within it, the contribution of this thesis, we need to paint some global picture of Mechanism Design. To do this, we first need two more basic ingredients: social choice theory, and non-cooperative games. The relevant theory is vast and rich. I will describe only the essentials. For a more comprehensive description, see e.g. the textbooks [83, 67].

1.1.1 Social Choice Theory

In the Social Choice setting, the society has a set of alternatives, A. Each individual has a preference order over the alternatives in A. The set of all possible preferences of player i is denoted by $\mathcal{D}_i$, and $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$. The goal is to construct a "social preference" from the individual preferences. More accurately, the literature distinguishes between three aggregation types:

- *A social welfare function* $F : \mathcal{D} \to \mathcal{R}$, where $\mathcal{R}$ is the set of all possible preference orders over A. I.e. aggregating the individual preference orders to a single social preference order.

- *A social choice correspondence* $F : \mathcal{D} \to 2^A$. I.e., given the individual preferences, we determine a subset of the alternatives which are the socially most preferred alternatives.

- *A social choice function* $F : \mathcal{D} \to A$. I.e. the individual preferences exactly determine the chosen alternative.

To a computer scientist, I would guess that the most appealing type, at first glance, seems to be the third one, as it most resembles the concept of an algorithm: given individual preferences as input, we output an alternative. However, a careful re-examination will reveal that the first two types actually fit more to the intuition of computer science: for example, when one seeks an algorithm with an approximation ratio of c (an exposition to CS notions is given in section 1.2.1 below), the actual intention is a social choice correspondence that, given a specific input, identifies a subset of the alternatives from which the algorithm may choose one – these are all the alternatives that are far from optimum by a factor of at most c. In fact, this actually implies a social weak order over the alternatives, as an alternative which is c-far from optimum is preferred to an alternative which is c'-far from optimum, for $c' > c$.

Social Choice Theory is concerned with the possibilities and impossibilities of such social preference aggregations that will exhibit desirable properties. One of the first major results in the field was Arrow's impossibility theorem [4]. Arrow describes few desirable properties from a social welfare function, and shows that no social choice function can satisfy all:

Definition 1.1 (Arrow's desirable properties)

1. *A social welfare function satisfies "weak pareto" if whenever all individuals strictly prefer alternative a to alternative b then, in the social preference, a is strictly preferred to b.*

2. *A social welfare function is "a dictatorship" if there exists an individual for which the social preference is always identical to his own preference.*

3. *A social welfare function F satisfies the "Independence of Irrelevant Alternatives" property (IIA) if, for any preference orders $R, \tilde{R} \in \mathcal{R}$ and any $a, b \in A$,*

$$a >_{F(R)} b \text{ and } b >_{F(\tilde{R})} a \ \Rightarrow \exists\ i: \ a >_{R_i} b \text{ and } b >_{\tilde{R}_i} a$$

(where $a >_{R_i} b$ iff a is preferred over b in R_i). In other words, if the social preference between a and b was flipped when the individual preferences changed from R to $\tilde{R}$, then it must be the case that some individual flipped his own preference between a and b.

It seems that the last property, IIA, is the most questionable one. It is interesting to note, however, that it has a natural interpretation in the context of Mechanism Design, as will be described below.

Arrow's impossibility theorem holds for an *unrestricted* domain of preferences, i.e. all preference orders are possible:

Theorem 1.1 (Arrow [4]) *Assume $|A| \geq 3$. Any social welfare function over an unrestricted domain of preferences that satisfies both weak pareto and Independence of Irrelevant Alternatives must be a dictatorship.*

Similar impossibilities were later shown for social choice functions.

Under this basic impossibility, the social choice theory has evolved in two separate directions. The first examines aggregation methods that violate some of Arrow's conditions, and discusses their reasonability. An example for this is the Condorcet rule: socially prefer a to b iff more individuals prefer a to b. This rule satisfies weak pareto and IIA, and is not a dictatorship. However, it does not produce a social preference order – there are individual preferences for which the Condorcet rule results in a preferred to b, b to c, and c to a.

A second direction examines different restricted preference domains, establishing whether the same impossibility still holds. A class of restricted domains for which many possible social welfare functions that satisfy Arrow's conditions do exist is the class of Single-Peaked domains (first introduced by Black [18]). For example, if the set of alternatives is the real line, then a domain of preferences is single-peaked if for every preference R there exists a real number (an alternative) β, such that, for any two alternatives (real numbers) x, y, x is preferred to y in R if and only if $|x - \beta| < |y - \beta|$. On the other hand, many restricted domains exhibit the same impossibility that exists in the unrestricted domain. One of the first works that show this is the work of Kalai, Muller, and Satterthwaite on "saturated domains" [53]. Further details on these can be found in the excellent survey [25].

1.1.2 Non-Cooperative Games

Classical Game Theory deals with the interactions of selfish, payoff-maximizing entities. The basic concept is a non-cooperative game: There are n players, each player i chooses an action a_i from a space of possible actions $\mathcal{A}_i$, and, according to these, a payoff $u_i(a_1, ..., a_n)$ is determined for every player i. The underlying question is what will be the chosen actions.

Definition 1.2 (A Normal-Form Game) *A normal form game is a tuple* $(N, \{\mathcal{A}_i\}_{i \in N}, \{u_i\}_{i \in N})$*, where* N *is a finite set of players,* $\mathcal{A}_i$ *is a set of the possible actions of player* i*, and* $u_i : \mathcal{A} \rightarrow \Re$ *is the payoff function of player* i*.*

Notice that this models a game of *complete information*: Each player knows the action space of the other players, and their payoff functions.

In order to analyze what will be the strategic choices of the players, the basic assumption is that a player tries to maximize his payoff:

Assumption 1.1 (Players are payoff-maximizers) *The only goal of each player is to choose an action so as to maximize his payoff.*

This assumption is ofcourse problematic. For example, why can't a player take into consideration the "well-being" of a fellow player? One possible answer is that we implicitly assume that the payoff functions are designed to encompass such considerations. Another answer, that fits the framework of this thesis, is that we are analyzing automated entities that were programmed that way. Either way, the point is that under this assumption, the structure of the game can be analytically investigated.

Some games admit a very simple structure of *dominant actions*. An action is dominant if it guarantees the highest payoff for the player, for *any* choice of actions of the other players. In the following definition we use the notation $\mathcal{A}_{-i} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_{i-1} \times \mathcal{A}_{i+1} \times \cdots \times \mathcal{A}_n$.

Definition 1.3 (A weakly dominant action) *An action $a_i \in \mathcal{A}_i$ "weakly dominates" an action $a'_i \in \mathcal{A}_i$ if for any $a_{-i} \in \mathcal{A}_{-i}$, $u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$.*

Stronger definitions may require that the inequality will be strict for at least one specific s_{-i}, or even that it will always be strict.

A game that contains dominant actions is "easy to solve". Assuming that player i's only goal is to maximize his payoff, we can quite convincingly argue that each player will choose to play a dominant action, if he has one. A celebrated example for a game with dominant strategies is the "prisoner's dilemma" game.

Most games, however, does not have dominant actions. Another possibility to analyze games is to search for an "equilibrium point" – a tuple of actions from which no player can unilaterally profitably deviate. Formally,

Definition 1.4 (Nash-Equilibrium [76]) *A tuple of actions $(a^*_1, ..., a^*_n)$ is a Nash equilibrium if for any i, and any $a_i \in \mathcal{A}_i$, $u_i(a^*_i, a^*_{-i}) \geq u_i(a_i, a^*_{-i})$.*

A tuple of dominant actions is also a Nash equilibrium. The opposite is of-course not true: Nash equilibrium actions need not be dominant actions. Moreover, they might even be dominated.

Does there always exists a Nash equilibrium? Not necessarily. But if we allow players to randomize over actions, and assume they now aim to maximize their expected payoffs, then the answer changes: Define a *mixed strategy* of player i to be a probability distribution over $\mathcal{A}_i$, and a tuple of mixed strategies to be in Nash equilibrium if no player can unilaterlly increase his expected payoff by switching to another mixed strategy.

Theorem 1.2 ([76]) *Any finite game admits at least one mixed-strategy Nash equilibrium.*

Is it reasonable to expect players to *reach* a Nash equilibrium? Two explanations are usually offered. The first is an informal description of game dynamics, that will reach a steady state if and only if Nash equilibrium strategies are played (this explanation has an additional interesting meaning in the context of evolutionary games). The second explanation assumes the existence of an outside coordinator that "suggests" a specific Nash equilibrium to the players. This coordinator can be, in a computerized context, a software program that is programmed to play a specific equilibrium, and is freely accessible to every player.

The theory of games develops many more alternative solution concepts, which we will not survey here, e.g. a correlated equilibrium, iterative elimination of dominated strategies, rationalizable strategies, and more. Chapter 4 contains more discussion on these subjects. The theory also analyzes different types of games, e.g. extensive form games, in which players choose actions one after the other, and repeated games, in which a normal form game is played repeatedly. But the very basics we described here will suffice us to continue our exposition.

Bayesian Games: games with incomplete information

The first treatment of games with incomplete information, or bayesian games, is due to Harsanyi [48]. In such games, we introduce a set Ω that represents the possible "states of nature". The payoffs of the players now become dependent on the state of nature, as well as on the strategies of the players, i.e. $u_i : \Omega \times \mathcal{A} \rightarrow \Re$. This is termed "incomplete information" because it is assumed that a player does not know the exact state of nature – each player observes only a signal $\tau_i : \Omega \rightarrow T_i$ about the state of nature. It is convenient now to treat the action of player i as function of his signal, i.e. $s_i(\tau_i(\omega))$ denotes i's action when his type is $\tau_i(\omega)$. The function $s_i(\cdot)$ is termed *a strategy.* An example of such a game is an auction of a single item. The state of nature is a tuple of values, that denotes the worth of the item to each player. The signal of each player is the worth of the item to him (he does not know the values of the other players), and his strategy is, given his value, to decide how much to "bid" for the item.

As the results of this thesis are all in the special case of the *private value* model, it will be most convenient to switch context and define all future notions explicitly for this model:

Definition 1.5 (The private value model) *A game with incomplete information is said to be in the private value model if $\Omega = T_1 \times \cdots \times T_n$, for any $\omega = (t_1, ..., t_n) \in \Omega$, $\tau_i(\omega) = t_i$ for all i, and for any i, the payoff function depends on ω only through t_i, i.e. $u_i : T_i \times S \rightarrow \Re$.*

A player's signal $\tau_i(\omega) = t_i$ is usually referred to as his *type*, and T_i is his domain, or type space. The single item auction game from above is also an example for a game in the private value model.

The definition of a dominant strategy in a game with incomplete information is modified in the straightforward sense: a strategy $s_i(\cdot)$ is a dominant strategy if for any $t_i \in T_i$, any $t_{-i} \in T_{-i}$, any strategy tuple of the other players $s_{-i}(\cdot)$, and for any other strategy $s'_i(\cdot)$ of player i, it holds that $u_i(t_i, s_i(t_i), s_{-i}(t_{-i}) \geq u_i(t_i, s'_i(t_i), s_{-i}(t_{-i}))$, and there exists a specific combination of t_i, t_{-i}, and a strategy $s'_i(\cdot)$ for which $u_i(t_i, s_i(t_i), s_{-i}(t_{-i}) > u_i(t_i, s'_i(t_i), s_{-i}(t_{-i}))$.

In contrary to this, the definition of a Nash equilibrium does not follow through easily; few technicalities should be handled with care. Probably the most used notion is that of Bayesian-Nash equilibrium, where we additionally assume that each player has a *prior distribution* π_i over T. Then,

Definition 1.6 (Bayesian-Nash equilibrium) *A strategy tuple $s^*_1(\cdot), ..., s^*_n(\cdot)$ is a Bayesian-Nash equilibrium if, for any player i, any fixed type t_i, and any other strategy $s_i(\cdot)$ of i,*

$$E_{\pi_i}[u_i(t_i, s^*_i(t_i), s^*_{-i}(t_{-i})) \mid t_i] \geq E_{\pi_i}[u_i(t_i, s_i(t_i), s^*_{-i}(t_{-i})) \mid t_i].$$

I.e. it is assumed that a player chooses his strategy after receiving his own type, but before knowing the types of the other players.

In most works that consider Bayesian-Nash equilibria, it is usually assumed that players agree on their priors, in this case this distribution is called a *common prior*, and, furthermore, that this common prior is indeed the "true" distribution over the types. This is one of the major criticisms on classical Mechanism Design (see below).

An equilibrium notion that conceptually resides "in-between" dominant strategies and Bayesian-Nash equilibrium is an "ex-post" Nash equilibrium[1]:

Definition 1.7 *A strategy tuple $s_1^*(\cdot), ..., s_n^*(\cdot)$ is an "ex-post" Nash equilibrium if, for any player i, any type combination $t_i \in T_i$ and $t_{-i} \in T_{-i}$, and any other strategy $s_i(\cdot)$ of i,*

$$u_i(t_i, s_i^*(t_i), s_{-i}^*(t_{-i}))] \geq u_i(t_i, s_i(t_i), s_{-i}^*(t_{-i}))].$$

I.e. even after knowing the type realizations of the other players, player i will not regret he chose $s_i^*(t_i)$ (given that the other players played their equilibrium strategies).

The literature also discusses an "ex-ante" Nash equilibrium, in which a player chooses his strategy before obtaining his own type. Hence, the expectation is taken over i's own type as well.

1.1.3 Mechanism Design

The basis to Mechanism Design is a conceptual integration between social choice theory and the assumption of non-cooperative, "selfish" individuals. We start with a slightly informal description of general mechanism design, followed by a more concrete description of the specific model assumed in this thesis.

A social designer has a social choice correspondence $F : T \rightarrow 2^A$ in mind, where T is the domain of players' types/preferences, and A is the set of possible alternatives/outcomes. The problem is that, in order to obtain the "correct" alternative according to F, the designer needs to know the actual players' types. But, since players are selfish, they might gain by reporting false types instead. A *mechanism* is a game form with an outcome function $g : S \rightarrow A$. The mechanism *implements*[2] a social choice correspondence F if there exists an equilibrium strategies $s^*(\cdot)$ such that, for all $t \in T$, $g(s^*(t)) \in F(t)$. In other words, a mechanism is a game for which the result of an equilibrium strategies coincides with F. One may consider mechanisms with respect to several kinds of equilibria notions, e.g. equilibrium in dominant strategies, Nash equilibrium, Bayesian-Nash, etc.

[1]For direct mechanisms (see definition below) in the private value model, this reduces to dominant strategies equilibrium. Its importance is for indirect mechanisms as well for non-private value models.

[2]This implementation notion is the weakest one. Stronger notions require e.g. that *all* equilibria will lead to a result according to F.

Two sub-models are usually considered. In the *complete information* model, players know each other's type, but the social designer does not. In this model, Maskin [68] have shown that *any* social choice correspondence can be implemented in Nash equilibrium. Other works study implementation possibilities and impossibilities in other types of equilibria notions, and with other types of mechanisms (arguably more "reasonable" than Maskin's mechanism).

In this thesis, we focus on the *incomplete information* model, in which players do not know each other's type. The basic question is what social choice correspondences can be implemented, and with which solution concept?

As the most appealing equilibrium concept is that of dominant strategies, we can ask what functions F can be implemented in dominant strategies? Unfortunately, the Gibbard-Satterthwaite theorem [41, 94] shows that only dictatorial functions can be implemented in dominant strategies. More precisely, if $|A| \geq 3$, and the domain of player types is unrestricted, then a social choice function F is implementable if and only if it is dictatorial. The proof uses Arrow's theorem in a way that connects between implementation in dominant strategies and the condition of IIA: a social choice function over an unrestricted domain of preferences that can be implemented in dominant strategies can be "extended" to a social welfare function that satisfies IIA.

In order to bail out of this impossibility result, we have to consider restricted domains. Some works consider restricted domains for which Arrow's theorem does not hold (e.g. single-peaked domains). Most of the literature, however, uses the quasi-linear utilities assumption in order to restrict the domain of preferences. This is also assumed throughout this thesis.

Mechanism design with quasi-linear utilities in the private value model

In the quasi-linear utilities model, players are assumed to have monetary value for each alternative. Formally, there are n players, and a set A of alternatives. Each player has a type $v_i \in V_i$, where $v_i : A \rightarrow \Re$ describes the value that the player will obtain from each chosen alternative. Notice that this is a private value setting, as the value of a player depends solely on his own value function. We now formally define a mechanism:

Definition 1.8 (A mechanism) *A mechanism is a game form with an action space $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, an outcome function $g : \mathcal{A} \rightarrow A$, and price functions $p_i : \mathcal{A} \rightarrow \Re$ for every player i. The outcome of the mechanism, given that the players choose actions $a = a_1, ..., a_n$, is $g(a)$, and each player i pays the mechanism a monetary sum $p_i(a)$.*

In parallel to the "players are payoff-maximizers" assumption, we now assume that players are utility maximizers, and that their utility is quasi-linear:

Assumption 1.2 *Given a specific mechanism, the only goal of each player is to choose an action so as to maximize his own utility: $v_i(g(a)) - p_i(a)$, where a is the tuple of chosen actions.*

A player's *strategy* is a function $s_i : V_i \to \mathcal{A}_i$ that specifies his action for any realization of his type. In principle, the strategy of a player can capture her entire information set, which may include more than just her type. Simultaneous move one shot games usually eliminate such extra information, while other types of games may allow it. Here, the default assumption will be that the only information that a player has is her type, but more a carefull consideration will be provide for specific results, e.g. for the results of Chapter 4.

A strategy is dominant if the player will always maximize his utility by following this specific strategy, no matter what the others will do:

Definition 1.9 (A dominant strategy) *Given a mechanism $(\mathcal{A}, g, \{p_i\}_i)$, a strategy $s_i : V_i \to \mathcal{A}_i$ is dominant if for every $v_i \in V_i$, the action $s_i(v_i)$ is a weakly dominant action. I.e. if for any $a_{-i} \in \mathcal{A}_{-i}$ and any $a_i' \in \mathcal{A}_i$, $v_i(g(s_i(v_i), a_{-i})) - p_i(s_i(v_i), a_{-i}) \geq v_i(g(a_i', a_{-i})) - p_i(a_i', a_{-i})$.*

Our goal is to design a mechanism that implements the social choice function in dominant strategies:

Definition 1.10 (Implementation in dominant strategies) *A mechanism implements a social choice function $F : V \to A$ in dominant strategies if there exist dominant strategies $s_1(\cdot), ..., s_n(\cdot)$ for players $1, ..., n$, respectively, such that, for any $v_1 \in V_1, ..., v_n \in V_n$, $g(s_1(v_1), ..., s_n(v_n)) = F(v_1, ..., v_n)$.*

An extremely simple type of mechanisms are *direct revelation mechanisms*, in which players are required to simply report their type. It turns out that there is no loss of generality in using direct revelation mechanisms:

Definition 1.11 (direct revelation mechanisms) *A mechanism is a direct revelation mechanism if $\mathcal{A}_i = V_i$ for all i.*

Definition 1.12 (Truthfulness, or Incentive Compatability, or Strategy-proofness) *A direct revelation mechanism is "truthful" (or incentive-compatible, or strategy-proof) if the dominant strategy of each player is to reveal his true type. A social choice function is truthfully implementable if it can be implemented by a truthful mechanism.*

Proposition 1.1 (The direct revelation principle) *Any social choice function that is implementable in dominant strategies can also be truthfully implemented in dominant strategies.*

proof: Given a mechanism M that implements F, with dominant strategies $s_i^*(\cdot)$, we construct a direct revelation mechanism M' as follows: for any type tuple v, $g'(v) = g(s^*(v))$, and $p_i'(v) = p_i(s^*(v))$. Since $s_i^*(\cdot)$ is a dominant strategy in M, we have that for any fixed $v_{-i} \in V_{-i}$ and any $v_i \in V_i$, the action $a_i = s_i^*(v_i)$ is dominant when i's type is v_i. Hence declaring any other type $\tilde{v}_i$, that will "produce" an action $\tilde{a}_i = s_i^*(\tilde{v}_i)$, cannot increase i's utility. Therefore the strategy v_i in M' is dominant. ∎

For a given set of player types $v_1, ..., v_n$, the *welfare* obtained by an alternative $a \in A$ is $\sum_i v_i(a)$. A social choice function is termed *a welfare maximizer* if $F(v)$ is an alternative with maximal welfare, i.e. $F(v) \in \text{argmax}_{a \in A}\{ \sum_{i=1}^{n} v_i(a)\}$. Arguably the most impressive positive result of mechanism design for incomplete information is the class of Vickrey-Clarke-Groves mechanisms, that truthfully implement the welfare maximizing social choice function. These mechanisms fit *any* domain of player types. Vickrey, in his seminal paper [97], analyzed a special case of this family, which later led Clarke [27] and Groves [44] to the general formula.

Definition 1.13 (VCG mechanisms) *Given any set of alternatives A, and any domains of player types $V = V_1 \times \ldots \times V_n$, a VCG mechanism is a direct revelation mechanism such that, for any $v \in V$,*

1. $g(v) \in argmax_{a \in A}\{ \sum_{i=1}^{n} v_i(a) \}$.

2. $p_i(v) = -\sum_{j \neq i} v_j(g(v)) + h_i(v_{-i})$, *where $h_i : V_{-i} \rightarrow \Re$ is an arbitrary function.*

If we ignore for a moment the term $h_i(v_{-i})$ in the payment functions, we can describe the VCG mechanism as a mechanism that chooses an alternative with maximal welfare according to the reported types, and then, by making additional payments, equates the utility of each player to that maximal welfare level.

Theorem 1.3 (Groves [44]) *Any VCG mechanism truthfully implements the welfare maximizing social choice function.*

proof: We argue that $s_i(v_i) = v_i$ is a dominant strategy for i. Fix any $v_{-i} \in V_{-i}$ as the declarations (actions) of the other players, and any $v_i' \neq v_i$, and assume by contradiction that $v_i(g(v_i, v_{-i})) - p_i(v_i, v_{-i}) < v_i(g(v_i', v_{-i})) - p_i(v_i', v_{-i})$. Replacing $p_i(\cdot)$ with the specific VCG payment function, and eliminating the term $h_i(v_{-i})$ from both sides, we get: $v_i(g(v_i, v_{-i})) + \sum_{j \neq i} v_j(g(v_i, v_{-i})) < v_i(g(v_i', v_{-i})) + \sum_{j \neq i} v_j(g(v_i', v_{-i}))$. Therefore it must be that $g(v_i, v_{-i}) \neq g(v_i', v_{-i})$. Denote $g(v_i, v_{-i}) = a$ and $g(v_i', v_{-i}) = b$. The above equation is now $v_i(a) + \sum_{j \neq i} v_j(a) < v_i(b) + \sum_{j \neq i} v_j(b)$, or, equivalently, $\sum_{i=1}^{n} v_i(a) < \sum_{i=1}^{n} v_i(b)$, a contradiction to the fact that $g(v_i, v_{-i}) = a$. ∎

Clarke [27] suggests a specific form for the function $h_i(v_{-i})$, namely $h_i(v_{-i}) = \sum_{j \neq i} v_i(g(v_{-i}))$ (this is a slight abuse of notation, as g is defined for n players, but the intention is the straightforward one). In this case, if a player does not influence the social choice, his payment is zero, and, in general, a player pays the "monetary damage" that he caused to the others.

The major criticism on the VCG mechanisms in the classic economics literature is that they are not budget balanced – the sum of payments to the players may be strictly positive or strictly negative, meaning that the mechanism either collects money from the participants or endows money to them. A budget balanced mechanism has the property that the sum of payments is zero, i.e. it only causes money to shift hands, but does not change the total amount the players have. If we are willing to settle with an implementation in a Bayesian-Nash equilibrium, one can achieve a welfare maximizing implementation which is also budget balanced [29, 5]. For this, it is assumed that the player types are drawn from a commonly known distribution. However, this mechanism is not *individually rational*, i.e. a player may end up in some type realizations with a negative utility. The Myerson-Satterthwaite [75] impossibility result states that this is no accident – there is no implementation of the welfare maximizing social choice function in a Bayesian-Nash equilibrium that is both budget balanced and individually rational. Notice that a VCG mechanism with Clarke payments is individually rational.

1.1.4 Auction Theory

Auction theory is a special but elegant case of Mechanism Design. It has a particular importance to mechanism design, by at least two reasons. First, it seems that classical results from auction theory tend to generalize to Mechanism Design. For example, Vickrey's seminal paper [97], that is considered as one of the origins of mechanism design, analyzes auctions, mainly single-item auctions. The second reason to the particular importance of auctions is that they model a very general combinatorial structure, especially when extended to the case of several (different) items. Studying auctions therefore implies a better understanding of many classic combinatorial models, with the added flavor of a non-cooperative environment. In short, it seems that the model of auctions bears a "winning combination" of generality on the one hand, and concreteness, on the other.

Vickrey was the first to study auctions in a game-theoretic framework. In his seminal paper, he studied multi-unit auctions in general, but focused on single-item auctions: A single item should be sold to a set of n buyers, where each buyer has a private value for the item, and a quasi-linear utility: his value (if he wins the item) minus any price that he has to pay. As before, a non-cooperative environment is assumed. Vickrey's goal is to design a mechanism, an "auction", in order to allocate the item to the buyer with the highest value. Put in the mechanism design terminology, we want to implement the welfare maximizing social choice function.

The most simple (but yet extremely impressive) construction that Vickrey makes is the "2nd price auction": bidders submit their value, the highest bidder wins, and pays the second highest price. This auction has the remarkable property that it is a dominant strategy for a player to bid his true value. I.e. this auction is *truthful* (it is a direct revelation mechanism, that implements welfare maximization in dominant strategies). It should be pointed out that this formulation is exactly equivalent to a VCG mechanism with Clarke payments, as described above.

The second auction that Vickrey analyzes is the "first price auction": Each player submits a bid, the player with the highest bid wins, and pays his bid. Clearly, bidding one's true value is no longer a dominant strategy; in fact, it is a dominated strategy. However, assuming that player values are independently drawn from the same distribution which all bidders, it is shown that, in a Bayesian-Nash equilibrium, a bidder with value v bids $\frac{n-1}{n}v$. With these equilibrium strategies the bidder with the highest value still wins, and therefore a first price auction implements welfare maximization in Bayesian-Nash equilibrium. Vickrey also shows that the expected price in the two auctions, with equilibrium strategies, is equal.

This last observation turns out to be much more general. The "revenue equivalence theorem" of Myerson [74] shows that the expected revenue of the auction is solely determined by the allocation rule. Specifically, any two implementations in Bayesian-Nash equilibria, whose allocations are the same for every type realization, have the same expected revenue. The proof goes by showing that the prices are completely determined by the allocation rule. This theorem is originally stated for single item auction, but this turns out to be true for general mechanisms [50, 26]. As a result, Myerson [74] is able to find the auction that extracts the maximal revenue. It turns out that it is a 2nd price auction with an appropriate reservation price (a threshold price, beneath it no sale is made) that is determined by the underlying distribution.

The economic literature continues to study single item auctions in different settings, e.g. risk-averse and risk-taking buyers, interdependent and common values, and more. A comprehensive survey on Auction Theory may be found in [56].

We next turn to discuss auctions of several items. Two types are considered – multi-unit auctions, where several identical items are sold, and combinatorial auctions, where several different items are sold. In these case, the type of each player is a value function that assigns a real number to any subset of items. This implies the "no externality" property, i.e. a player does care about the allocation of the items that he did not receive; his value depends only on the items he received. Additional property that is usually assumed is the "free disposal" property, i.e. the value function is non-decreasing (we can freely dispose items).

It is worth mentioning that if the social goal is to maximize the welfare then we can simply use the VCG mechanism[3]. For the case of multi-unit auctions with decreasing marginal values this is termed a "Generalized Vickrey Auction" as Vickrey itself describes it. Optimal multi-unit auctions, that aim to maximize revenue, were studied by Maskin and Riley [69]. They discuss the case where marginal values are decreasing, and the value function is drawn from a single-parameter distribution. The revenue equivalence theorem is extended to this case, and it is shown that the Vickrey auction with a reserve price is still optimal when the marginal values are step functions with a single step.

Much of the economic literature about multi-item auctions concentrates on designing *indirect* mechanisms, specifically iterative auctions. Indirect mechanisms are important, as "real life" auctions rarely require the players to simply reveal their true values. Vickrey itself draws parallels, in the single item case, between the 2nd price auction and the "open cry english auction" (in which the seller gradually raises the price, and the last remaining player wins), and between the first

[3] Although this is usually not computationally efficient. See details in the sequel.

price auction and the "open cry dutch auction" (in which the seller gradually lowers the price, and the first player to accept wins). The idea of an ascending-price auction that ends in a proper allocation and with prices for the buyers was extended by Demange, Gale, and Sotomayor [30] to the case of multiple different items, where each buyer has "unit-demand" (i.e. each buyer may have different values for different items, and the value of a set of two or more items is just the maximal value of an item in the set). Ausubel [6] has devised an ascending auction for the case of identical items, where buyers have decreasing marginal valuations. Both of these auctions implement welfare maximization in ex-post Nash. This follows from the strong property that both of these auctions posses: they both reach a welfare maximizing allocation, and along the way they extract the appropriate VCG prices (in equilibrium behavior). (for the unit-demand auction, this last property was observed by Gul and Stacchetti [45], relying also on a result of Leonard [63].) Ascending auctions for combinatorial auctions with general valuation functions has been recently suggested by Ausubel and Milgrom [7] ("proxy-bidding auctions"), and by Parkes [84] ("iBundle").

1.2 Algorithmic Mechanism Design

1.2.1 The Difference From Classical Mechanism Design

As we are located on the border of computer science and economics, it is natural to wonder what is the difference from classical economics/mechanism design? There are four factors that cause algorithmic mechanism design to significantly deviate from the lines of the classic theory, all drawn from the fact that a computer scientist implicitly carries with him a somewhat different approach towards problem modeling and solution criteria. These four factors are (a) the computational efficiency requirement, that completely changes the classic borders between possibilities and impossibilities, (b) the relaxation of allowing near optimal solutions, instead of exact ones, which is a classic relaxation in CS, but was rarely studied before in economics, hence opens the way to new exciting possibilities, (c) using worst-case analysis, that, again, is the gold-standard in CS but is quite new to economics, and (d) the introduction of new computational models, all bring with them new questions and understandings. The algorithmic theory is of-course benefited from all these as well, as the horizon opens up to new exciting explorations.

In what follows, I will try to expand a bit on each subject. This discussion will be mostly dedicated to a basic introduction to CS notions, in a somewhat detailed and a bit simplified way, for the sake of the unfamiliar reader.

Computational efficiency. Requiring the mechanisms to be computationally efficient completely re-opens many "solved" problems. Many classic positive results, like the VCG mechanism, sometimes require hard computation, and "fixing" this is a non-trivial task. Thus, under the computational requirement, new questions arise, and old answers become questionable.

Let me dedicate a few explanatory words about the notion of efficient computation for the unfamiliar reader. A common mistake is that computational efficiency is a "technical" condition that a good engineer can always solve. The definition of computational efficiency explicitly avoids this mistake by ignoring multiplicative constants (asymptotic analysis), and, more importantly, by differentiating between *polynomial* time and *non-polynomial* time computation – if a given algorithm/mechanism can be computed in time polynomial in the size of its input than it is considered computationally efficient, otherwise it is not efficient. An algorithm that is not efficiently computable will sometimes terminate only after a period longer than a human life time, even for moderate input sizes. "Usual" technological improvements will not solve this. This ofcourse has real and fundamental consequences on economics and mechanism design, as the implicit assumption, especially in mechanism design, is that the mechanisms are indeed computable.

A short word about terminology. If a problem is termed "NP-complete", or "NP-hard", then it is highly unlikely that it is polynomial time solvable. I will avoid exact definitions, but to please the ear, say that there is an extremely large class of such problems, and the interesting phenomena is that if one of them will turn out to be polynomial time solvable then they will all be, which, in turn, seems highly unlikely. This is perhaps the most important open question of computer science.

Approximately optimal solutions. Due to the computational difficulty of finding exact solutions, much concern is devoted to guaranteeing a "near-optimal" solution. Although this concept is fairly simple, it results in a complete conceptual shift. It is quite rare, in classic economics, to find procedures that do not achieve the exact goal of the designer, but only approximate it. It seems that this should open the way to many new possibilities, and to a better understanding of the border between the possibilities and impossibilities of a given problem.

A basic terminology explanation: we measure how "near" (to the optimal solution) the algorithm gets. Formally, given a specific measure of the "value" of a solution, an algorithm is a c-approximation if it *always* computes a solution with value at least the optimal value over c (for maximization problems; for minimization problems – if it always computes a solution with value at least the optimal value times c). A fully polynomial approximation scheme (FPAS) is an algorithm that obtains, for any given $\epsilon > 0$, a $(1 + \epsilon)$-approximation in running time polynomial in the size of the input and in $1/\epsilon$.

Worst-case analysis. All the above notions are from a worst-case point of view[4]. Much debate is devoted to the issue of the kind of analysis that should be performed. It somehow became basic grounds in computer science that strong distributional assumptions, under which a solution is "tailored", are strongly unfavored. Although the worst-case criteria is indeed an extreme one, practical evidence suggest that it correctly differentiates between solutions with different "real"

[4]The fact that an algorithm is not a polynomial time algorithm means that for at least one specific instance it will perform badly (not necessarily on average). The fact that an algorithm is a c-approximation means that for some specific instance the solution will be exactly c-far from the optimal solution (again, not on the average). And so on.

quality. Examining the economic and game theoretic models from a worst case point of view yields new understandings, and new types of solutions arise.

New models common in computational settings. The integration between classical computational models and game-theoretic considerations of a non-cooperative distributed environment yields many surprising and fascinating results. This contributes to new understandings and developments in both fields. A good example is the model of online auctions, first suggested as part of the work towards this thesis (see chapter 2).

1.2.2 Some Representative Results of the Field

The term Algorithmic Mechanism Design was coined in [80], where the connection between mechanism design and algorithmic theory was formally suggested (although some earlier work in this spirit exist, e.g. [62, 99, 96]). The main example studied in [80] is a job scheduling problem taken from classic computer science. The classic algorithmic solutions cannot be applied here because of the context of a non-cooperative environment, and the classic VCG mechanism cannot be applied because our goal here is different from welfare maximization. The paper presents inapproximability results, and a weak positive result, where a large gap is left as an open problem.

Later on, it has been realized that a representative problem with a richer structure is that of Combinatorial Auctions. This problem actually includes the above job scheduling problem, as well as many other combinatorial problems, as special cases. It has many inherent difficulties, and many interesting special cases. In this model, k different items are to be allocated among n players, where each players has a private valuation function that assigns a value to each subset of items. It turns out that even if our goal is to allocate the items so as to maximize the welfare, we cannot use the VCG mechanism as exactly solving the problem is NP-hard [62]. The early work first concentrated on understanding the computational possibilities of this problem, and a rich variety of allocation heuristics were suggested and investigated [64, 93, 43, 92, 82]. However, designing a mechanism with these algorithmic understandings turned out to be a problematic task (chapter 3 investigates the impossibility of this task). One special case for which possibilities are known is the case of *single minded bidders*, where each bidder is interested in only one single bundle. This is still computationally hard, but Lehmann, O'Callaghan, and Shoham [62] provide a truthful mechanism that approximates the social welfare. This was further studied by [72, 1]. In this context, [1] suggests to use *approximately truthful* mechanisms. [13] describe a truthful approximation mechanism for multi-parameter bidders: They assume a multi-unit combinatorial auction, where we have multiple copies from every item, and each bidder desires only a small fraction of the copies. [57] describe an approximately truthful multi-unit auction (where all items are identical) that is an FPAS.

Besides the computational difficulties, general combinatorial auctions are also problematic with respect to communication. How will the bidders communicate their valuation function to the auctioneer? The straight-forward way requires an exponential amount of communication. One approach is to design *bidding languages* [77], i.e. languages that represent the valuation of a player in a compact way. Another approach is to perform *preference elicitation* [28, 19]. In this approach, the mechanism elicits information by performing oracle queries, and the goal is to perform the smallest possible amount of queries. A third approach is to design indirect mechanisms, that reach an allocation through iterative bidding rounds, where the goal is to have only a few number of rounds [84]. However, all these will fit only to special cases of classes of valuation functions. [81, 78] show that reaching the optimal welfare (or to approximate it to a factor of $k^{0.5-\epsilon}$) must result in exponential communication.

Many other types of auctions were recently considered as well. In the setting of digital goods, there is an infinite number of goods to be sold, and each buyer desires one good. This model was suggested by [42, 39], where they design revenue maximizing auctions. The main result is a truthful 4-approximation of the revenue (with respect to the optimal fixed-price auction that is required to charge a flat price from all players). Revenue maximizing single item auctions were re-considered by [88, 89], where they apply the concept of an approximately optimal auction to this classic economic setting. The setting of online auctions (see chapter 2, where bidders arrive over time, and the auction is required to make decisions before seeing the entire sequence of bidders, also admits many new and interesting recent results [59, 21, 11, 13, 20, 40, 8, 55, 85, 47]. More types of auctions that I will just mention briefly include supply chains [9, 10], auctions with limited number of communication bits [22], bidding clubs [65], and auctions in which either the auctioneer or the players cheat [86].

It is important to stress that almost all the results in the field "insist" on dominant strategies, or almost dominant strategies. This is probably due to the fact that this notion resembles in spirit to the worst case arguments used in computer science. It seems that this has contributed much understanding to the possibilities of implementation in dominant strategies, an understanding that, quite surprisingly, is also new to classic mechanism design. We especially understand the possibilities of dominant strategy implementation in one parameter settings [2], where a complete characterization of truthful algorithms exist.

In the description above, I have concentrated on results related to auction theory. Other recently studied models on the border of computer science and economics, which I will not touch here, are, for example, distributed mechanism design (e.g. [38]), games in networks and "the price of anarchy" (e.g. [90]), and the computation of market (walrasian) equilibria (see e.g. [31, 32]). The efficient computation of Nash equilibria is a long standing question that attracts constant interest.

1.3 Overview of The Results of This Thesis

In this thesis I present three groups of results:

Competitive Analysis of Incentive Compatible Online Auctions. We study auctions in a setting where the different bidders arrive at different times, and the auction mechanism is required to make decisions about each bid as it is received. Such settings occur in computerized auctions of computational resources as well as in other settings. We call such auctions, *on-line auctions*. While, in the economic literature, most auctions are analyzed using average-case analysis, we apply CS-related tools and give a competitive worst-case analysis of incentive compatible on-line auctions. We first characterize exactly those on-line auctions that are *truthful* (or incentive compatible), i.e. where rational bidders are always motivated to bid their true valuation. We then design several truthful online auctions that have performance guarantees with respect to both the social welfare and the revenue of the auctioneer. Our main auction is an incentive compatible on-line auction for a large number of identical items. We prove that this auction has an optimal competitive ratio, in terms of both revenue and welfare. We also design deterministic and randomized truthful online auctions for a small number of indivisible goods.

Besides the technical contribution, we view this study as having two main conceptual contributions. First, we introduce online considerations, which are well-studied in algorithmic theory, to the setting of game theory and mechanism design. Second, we demonstrate how to design auctions with worst-case guarantees, which was a new approach at the time we conducted the research. Since then, these two conceptual contributions have been proven fruitful, and many follow-ups adopted this model of online auctions (see the citations in section 1.2.2 above).

This set of results is based on a joint work with Noam Nisan.

Towards a Characterization of Truthful Combinatorial Auctions. The class of VCG mechanisms describes a way to attach payments that induce truthfulness to any algorithm that computes the exact optimal social welfare in any given model. Naturally, one wonders about other algorithmic goals. The motivation behind this question is twofold: First, in problems where the exact optimum is computationally hard, we would like to have approximation algorithms, instead of exact ones. Second, in many traditional computational models, our goal is not welfare maximization. In two extreme cases, the answer to this question is known: on the one hand, if the input of the players consists of only one secret number, then many recent examples for incentive compatible algorithms that are not welfare maximizers are known (see section 1.2.2 above). On the other extreme, if no restrictions at all are put on the structure of the input, then a beautiful result by Roberts (1979) states that only welfare maximizers can be truthfully implemented. However, most interesting computational models lie in the intermediate range, for which nothing is known.

An excellent example, which we investigate in this work, is the model of combinatorial auctions. In a combinatorial auction, a seller has a set of different items, and bidders may value bundles of items. Since computing the exact optimum is hard, computer scientists design approximation algorithms and heuristics that work well in practice. But none of these are known to be truthful.

We show that truthful mechanisms for combinatorial auctions (and related restricted domains) must be "almost weighted welfare maximizers" if they also satisfy an additional requirement of "independence of irrelevant alternatives" (IIA). This condition is a natural analog in the quasi-linear case to Arrow's original IIA condition. We also show that this IIA requirement is without loss of generality in some special cases. We begin by describing a weak monotonicity condition that completely characterizes truthful combinatorial auctions (and related restricted domains). Although a monotonicity condition that characterizes truthfulness for single dimensional domains is well known, this is the first time that such a monotonicity condition is identified for multi dimensional domains. We then describe a slight, but significant, strengthening of this condition, to what we call strong monotonicity. We show that the difference between the two conditions exactly equals the IIA condition. We then prove that strong monotonicity (plus few technical conditions) implies "almost weighted welfare maximization".

This result has strong computational implications. We observe that weighted welfare maximization is as computationally hard as exact welfare maximization. From this we conclude that any polynomial time truthful combinatorial auction that satisfy our few additional conditions cannot achieve any polynomial approximation of the welfare (assuming $P \neq NP$, and a sufficiently powerful bidding language). For the case of multi-unit auctions among two players, where all items must always be allocated, we are able to show an even sharper result: no truthful auction for this case can obtain an approximation ratio better than 2, although without truthfulness there exists an FPAS for the problem.

This set of results is based on a joint work with Ahuva Mu'alem and Noam Nisan.

Online Ascending Auctions for Gradually Expiring Items. We consider online auction mechanisms for the allocation of M items that are identical to each other except for the fact that they have different "expiration times", and each item must be allocated before it expires. Players arrive at different times, and wish to buy one item before their deadline. The private information of the players contain all their parameters, i.e. their values, deadlines, and arrival times. We begin by showing that the usual notion of truthfulness cannot be used in this case, since any (deterministic) truthful auction cannot obtain better than an M-approximation of the social welfare. Therefore, instead of designing auctions in which players should follow a single, dominant, strategy, we design two natural ascending auctions that perform well under a wide class of selfish, "semi-myopic", strategies. For every combination of such strategies, the auction is associated with a different algorithm. Thus, our algorithmic construction is of a *family* of algorithms, each one corresponds to a specific combination of players' behaviors, and *all* of them obtain a near optimal outcome. The new concept is that, although players are not expected to follow a specific behavior, but only one out of a set of behaviors, the outcome is still guaranteed to be close to optimal, for any choice the players make. We call this family of algorithms "semi-myopic" algorithms, and show that any algorithm in this family obtains a 3-approximation. By this we conclude that our auctions will perform well under any choice of such semi-myopic behaviors. We next turn to provide a game-theoretic justification for acting in such a semi-myopic way. We suggest a new notion of "Set-Nash" equilibrium, where we cannot pin-point a single best-response strategy, but rather only a set of possible best-response strategies. We show that our auctions have a Set-Nash equilibrium which is all semi-myopic, hence guarantees a 3-approximation. We discuss some natural strengthenings of the equilibria notion we use, still keeping this general idea of "set equilibria". We believe that these concepts offer a new way to bypass the inherent difficulties of the truthfulness notion, in a way that suits the CS worst-case notions.

This set of results, coupled with the previous impossibility result, reflects the two central themes that this thesis tries to advocate: (I) the search for the exact "turning point" in which the computational model does not admit a solution in dominant strategies, and (II) the search for other solution concepts suitable for such models, that fit the worst case spirit of computer science.

This set of results is based on a joint work with Noam Nisan.

Chapter 2

Competitive Analysis of Incentive Compatible On-Line Auctions

2.1 Exposition of Results

This chapter[1] studies auctions in a setting where the different bidders arrive at different times and the auction mechanism is required to make decisions about each bid *as it is received*. This is in contrast to the traditional assumption (in theory and in practice) that the auction organizer must receive all the bids before determining the allocation. The traditional assumption implicitly assumes that all participants (including the auctioneer) are willing to wait for some amount of time (until all bids are gathered) before performing any trade. We argue that in many settings, especially computerized ones, players will not be willing to wait a long time for the allocation decision.

An example of such a setting is bandwidth allocation on a communication link. Consider a fixed communication link in some computer network. In cases where the demand for communication over this link exceeds the link's bandwidth, one approach for allocating the limited bandwidth is by auctioning it among all the possible uses [60, 66]. However, in such a setting one would expect the requests ("bids") for bandwidth to arrive at different times – each request needing an immediate answer. Similar situations arise in the allocation of other resources such as CPU time or cache space.

In our model, k identical items are sold in an auction. Each bidder has a (privately known) valuation for each quantity of the goods, where the marginal valuations of the bidders are non-increasing. The bidder learns this valuation at a certain time and must make a bid at that time. The auction mechanism must decide, as the bid is received (and before seeing future bids), how many items to allocate to this bidder and at what price. We term such an auction *on-line*. We also consider more general variants where the valuations as well as bids may be time-dependent –

[1]This chapter is based on a joint work with Noam Nisan [59]

all our results extend to the general variants.

Our main concern here is with the incentive compatibility – also called truthfulness or strategy-proofness – of the auction. As detailed in Chapter 1, an auction is called incentive compatible if participants are rationally motivated to reveal the truth about their valuations. Specifically, in game-theoretic terms, if the truth is a dominating strategy. This is a departure from the field of on-line algorithms (see [23]) which does not address any game-theoretic issues but only algorithmic ones.

Our first result in this chapter is a full characterization of incentive compatible on-line auctions: We define an on-line auction as "based on supply curves" if before receiving the i'th bid, $b_i(j)$ for $1 \leq j \leq k$, it fixes some function (supply curve) $p_i(j)$ based on previous bids, and,

1. The quantity q_i sold to bidder i is the quantity q that maximizes the sum $\sum_{j=1}^{q}(b_i(j) - p_i(j))$ (i.e. the bidder's utility).
2. The price paid by agent i is $\sum_{j=1}^{q_i} p_i(j)$.

Theorem: A deterministic on-line auction is incentive compatible if and only if it is based on supply curves.

We then employ a *worst case* analysis of on-line auctions. This is in the spirit of computer science and in sharp contrast to the usual Bayesian (average case) analysis employed in auction theory (as well as in other economic situations). We view this new way of analyzing an auction as one of the major contributions of this work. We strongly feel that as auction theory is increasingly applied to computational settings, the importance of worst case analysis increases.

Specifically, we assume here that bidders' valuations all belong to some range $[\underline{p}, \bar{p}]$ (without assuming any probability distribution), where $\underline{p}$ is also the seller's reservation price, i.e. each item is worth $\underline{p}$ to him. We compare the performance of an incentive compatible on-line auction to the performance of the standard *off-line* Vickrey auction – for one item, this is the sealed bid second price auction, or, equivalently, the popular open cry English auction with small bid increments. This auction is incentive compatible and obtains optimal *social welfare*, i.e. maximizes the sum of all players' valuations of the items they receive.

Similarly to the definition used in on-line analysis of algorithms, we focus on the, so called, competitive ratio: An on-line auction is called *c-competitive with respect to the revenue* (relative to the Vickrey auction) if for *every* sequence of valuations of bidders it obtains a revenue that is at least $1/c$ of the revenue obtained by the Vickrey auction for these valuations (where the Vickrey auction knows all bids in advance). Similarly we define *c-competitive with respect to the social welfare*. We note the dissimilarity between the economic meaning of the term "competitive" and its meaning in computer science, which is the one used here.

The tightest set of results is obtained when the number k of items is large, so it can be treated as a continuum. This can be viewed as the case of one divisible good, i.e. it can be divided to any number of small fractions. For this case we are able to find the optimal on-line incentive compatible auction. We define the *competitive on-line auction* by using the function suggested in [36] to construct the supply curves. Extending the results of [36] for on-line continuous one way trading, we prove the following upper and lower bounds. Let $\phi = \bar{p}/\underline{p}$, and let the constant c denote the solution of the equation $c = ln((\phi - 1)/(c - 1))$. It can be shown that $c = \Theta(ln\phi)$.

Theorem: The Competitive On-Line Auction is c-competitive with respect to the revenue as well as with respect to the social welfare of the off-line Vickrey auction. No other on-line auction has a better competitive ratio either with respect to the revenue or with respect to the social welfare.

For the case of a smaller values of k we obtain the following results. For one good the best competitive on-line auction achieves a competitive ratio of $\sqrt{\phi}$. For other values of k, we show a deterministic lower bound of $\phi^{\frac{1}{k+1}}$ and a deterministic upper bound of $k \cdot \phi^{\frac{1}{k+1}}$. We observe also that if *randomized* auctions are allowed then a better competitive ratio can be obtained. By using the supply curves of the previous theorem for probabilistic choices, a competitive ratio of c can be obtained (for any number of goods), where $c = \Theta(ln\phi)$ is as before. In this case, the on-line revenue is also competitive with respect to the *optimal social welfare*.

It should be noted that the competitive ratio is obtained in the worst case; in the average case the ratio is typically much better. As a demonstration, we also provide a "normal" Bayesian analysis of our competitive on-line auction for the divisible good, in the case of uniformly distributed valuations in the interval $[\underline{p}, \bar{p}]$. For example, for two bidders whose valuations are uniformly distributed in $[1, 2]$ this on-line auction achieves expected revenue of 1.31... as compared to 1.33.. for the Vickrey auction.

To the best of our knowledge, this is the first work to study competitive worst case analysis of online auctions. However, competitive analysis of offline auctions have been independently employed by [42, 39], where they analyze the (off-line) unlimited supply case, and design revenue maximizing randomized auctions. Other works have adopted the model of on-line auctions: [11] and [20] describe on-line auctions for unlimited supply. [21] describe on-line double auctions (where the auction mechanism is actually a market that matches buyers to sellers). Recent works on the subject are [13, 40, 8, 55, 85, 47].

The rest of this chapter is organized as follows. Section 2.2 describes our model and gives a full characterization of incentive compatible on-line auctions. In section 2.3 we describe the competitive on-line auction. Section 2.4 outlines an extended on-line model, and section 2.5 gives a distributional analysis of the competitive auction.

2.2 On-Line Auctions

2.2.1 The model

The goods: We consider an auction of k identical indivisible goods to a set of players. We distinguish the case of a very large k that can be treated as a continuum, viewing this case as auctioning one divisible good.

Players' valuations and utilities: Each player has some non-negative benefit (valuation) from receiving some quantity of the goods. This valuation is known only to the player himself. We denote the marginal valuation of player i as $v_i(j)$, for $1 \leq j \leq k$. I.e. $v_i(j)$ is the additional value gained from the j'th good. Thus, his total valuation for q goods is $\sum_{j=1}^{q} v_i(j)$. We assume that all players have downward sloping marginal valuation functions, i.e. $\forall i, j: \; v_i(j+1) \leq v_i(j)$.[2] When player i receives q goods and pays for them a total payment of P_i his utility is $U_i(q, P_i) = \sum_{j=1}^{q} v_i(j) - P_i$. We assume that each player aims to maximize his utility.

The on-line game and players' strategies: The on-line game has the following structure. Initially, the set of players is unknown to the auctioneer, and none of the players knows his valuation. At some point in time, t_i, player i determines his valuation and must declare his bid at that time. We focus on direct revelation mechanisms, in which the player declares his marginal valuation function. Thus, the bid is some non-increasing function $b_i(\cdot)$ of the form $b_i : [1 \ldots k] \rightarrow \mathcal{R}$. Of-course, a player may be motivated to lie, declaring some $b_i(\cdot) \neq v_i(\cdot)$, in order to increase his utility. The auctioneer must answer the bid immediately, before opening the next bid. In his answer, he determines the quantity to be sold and the total price to be paid for it. We assume that if a player does not receive any positive quantity then his total payment is zero.[3] The game ends when the auctioneer sells all the goods or when the last player announces his bid.

Incentive compatibility: We study truthful implementations in dominant strategies, as discussed in chapter 1, which we refer to as incentive compatible mechanisms. For easier readability, we briefly repeat the relevant definitions. A strategy (bid) $b_i(q)$ of player i is called dominant if for every other bid $\tilde{b}_i(q)$ and for every sequence of past and future bids of the other players, $U_i(q_i, P_i) \geq U_i(\tilde{q}_i, \tilde{P}_i)$, where $q_i, \tilde{q}_i$ are the quantities sold to player i when declaring $b_i(q), \tilde{b}_i(q)$, respectively, and $P_i, \tilde{P}_i$ are the total payments charged for each quantity. In other words, for every bid sequence of the other players the utility of player i is maximized by choosing the specific declaration $b_i(q)$. A direct revelation mechanism is incentive compatible if for every valuation $v_i(\cdot)$, declaring the true valuation is a dominant strategy.

Remark 1: This model explicitly limits the strategy space of the players, excluding any time considerations, i.e. player i must declare his bid at time t_i, and is not allowed to return. In addition, his strategy space does allow any considerations about past events to enter – the strategy of a player is a function from his type to a bid. All these assumptions are not necessary and are introduced for the purpose of clear exposition. In section 2.4 we remove this limitations and show that all our results still hold for an extended model with time considerations.

[2] This assumption is common in economics, and is assumed e.g. in Vickrey's original paper. Without it, finding an optimal allocation is NP-complete.

[3] This normalization ensures both participation constraints and no budget deficit.

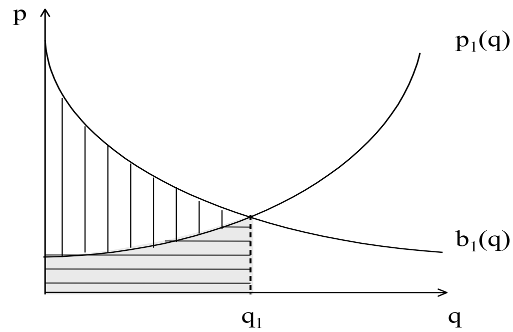

Figure 2.1: An example of supply curves based auction.

Remark 2: It is also possible to consider a *partially on-line* model, in which the set of players is known in advance to the auctioneer (but the valuation sequence is still revealed on-line). Although this approach weakens the on-line power, it is more close to regular game theory settings.

2.2.2 Supply curves for on-line auctions

We now characterize incentive compatible on-line auctions, as auctions that are *based on supply curves*. Intuitively, such auctions determine the prices for bidder i independently of i's bid, and then sell to i the quantity that maximizes his utility under these prices:

Definition 2.1 *(Supply Curves) An on-line auction is called "based on supply curves" if before receiving the i'th bid it fixes, based on previous bids*[4]*, some* $Q_i \subseteq \{1, ..., k\}$ *and a function (supply curve)* $p_i : Q_i \to \Re$*, such that:*

1. *The quantity* q_i *sold to bidder* i *is the quantity* $q \in Q_i$ *that maximizes the sum* $\sum_{j=1}^{q} b_i(j) - \sum_{j \in Q_i, j \leq q} p_i(j)$ *(ties may be broken arbitrarily).*
2. *The price paid by bidder* i *is* $\sum_{j \in Q_i, j \leq q_i} p_i(j)$*.*

A simpler form of supply curves, which we use below, is when Q_i is equal to all feasible quantities (i.e. $Q_1 = \{1, ..., k\}$, $Q_2 = \{1, ..., k - q_1\}$, and so on), and each supply curve $p_i(q)$ is non-decreasing. For this case, the quantity q_i becomes the largest quantity q such that $b_i(q) \geq p_i(q)$. For the divisible case, the price now becomes $\int_0^{q_i} p_i(q)dq$, and if both $b_i(q), p_i(q)$ are continuous then q_i is the unique solution to $b_i(q) = p_i(q)$.

For example, figure 2.1 illustrates a non-decreasing supply curve $p_1(q)$ and a bid $b_1(q)$. According to Def. 2.1, the quantity received by the player equals q_1, and the total price paid is the area below

[4]Here, the supply curves are determined deterministically. A possible extension of this definition, when allowing randomization, is described in section 2.3.2.

the supply curve, marked by the horizontal lines. The player's valuation of the quantity q_1 is the area below $b_1(q)$, and, thus, the resulting utility of the player is the area between the marginal valuation and the supply curve, marked by the vertical lines. This is the entire surplus, in economic terms. After the sale, the auction continues to the next player, presenting some new supply curve $p_2(q)$.

Theorem 2.1 *A deterministic on-line auction is incentive compatible if and only if it is based on supply curves.*

proof: We prove the two directions of the theorem by the following two lemmas.

Lemma 2.1 *An on-line auction that is based on supply curves is incentive compatible.*

proof: The utility of player i from receiving some quantity q is $U_i(q) = \sum_{j=1}^{q}(v_i(j) - p_i(j))$ (his valuation of the total quantity minus his total payment). Let $b_i(q) \neq v_i(q)$ be some bid and suppose the quantity sold for this bid is $\tilde{q}_i$, and for the truthful bid is q_i. Then it is the case that $U_i(q_i) \geq U_i(\tilde{q}_i)$, since this is explicitly verified in the first condition of the supply curves definition (when the bid is truthful then the term maximized there equals $U_i(q)$). Thus the claim follows. ∎

Lemma 2.2 *Any deterministic incentive compatible on-line auction is based on supply curves* [5].

proof: Fix any deterministic incentive compatible on-line auction A. We first argue that the total payment of player i is determined uniquely by the quantity sold to him (and by previous bids): Otherwise, there are two different bids $v(q), \tilde{v}(q)$ such that the quantity sold when declaring each one of them is the same but the total price paid is different. Let P be the total price when declaring $v(q)$ and $\tilde{P}$ the total price when declaring $\tilde{v}(q)$, and w.l.o.g suppose $P < \tilde{P}$. Thus a player with valuation $\tilde{v}(q)$ will increase his utility by declaring $v(q)$ since he will receive the same quantity and will pay a lower total payment, which is a contradiction since A is incentive compatible.

Now let Q_i denote the set of quantities q for which there is a bid $b(\cdot)$ such that the player will be allocated q items if declaring $b(\cdot)$. According to the previous paragraph we denote by $P_i(q)$, for any $q \in Q_i$, the total payment of player i when receiving the quantity q. We claim that i must be allocated the quantity q_i that maximizes $U_i(q) = \sum_{j=1}^{q} v_i(j) - P_i(q)$ (for any $q \in Q_i$). Otherwise, let $b(\cdot)$ be some bid for which A sells the quantity q_i to i. Then, if A sells a quantity $\tilde{q}_i \neq q_i$ for the truthful bid $v_i(\cdot)$, player i will increase his utility by declaring $b(\cdot)$ instead, which contradicts incentive compatibility.

For any $q \in Q_i$ let $q_{-1} = max\{q' \in Q_i \mid q' < q\}$. To conclude the argument we claim that $p_i(q) = P_i(q) - P_i(q_{-1})$ is the supply curve according to definition 2.1. Indeed, since $P_i(q) = P_i(0) + \sum_{j \in Q_i, j \leq q} p_i(j)$, and $P_i(0) = 0$, it follows from the argument above that the total quantity sold is the quantity q_i that maximizes $U_i(q) = \sum_{j=1}^{q} v_i(j) - P_i(q) = \sum_{j=1}^{q} v_i(j) - \sum_{j \in Q_i, j \leq q} p_i(j)$, and the total price paid is $P_i(q_i) = \sum_{j \in Q_i, j \leq q_i} p_i(j)$. ∎

[5] Notice, however, that the supply curves structure may be implicit in the formal auction description.

From the above two lemmas the theorem follows. ∎

Remark: Allowing any (non-increasing) marginal valuation functions may increase significantly the complexity of presenting the valuation function to the auctioneer. This problem can be solved by using a modified auction that, instead of receiving valuation functions, presents the (current) supply curve to all (interested) players. In this case each bid is simply a price-quantity coordinate on the supply curve. From the same considerations of incentive compatibility from lemma 2.1, declaring the truth (i.e. the maximal quantity according to Def. 2.1) is dominant. We note that the supply curves we give below can be presented easily.

An interesting special form of players' valuations is fixed marginal valuations. In this case, the marginal valuations are restricted to the form $v_i(q) = v_i$ for all i, q. This case is also useful since we use it for the lower bound we give below. For this case, it is possible to characterize the supply curves more precisely (this holds for the continuous case as well, with a similar proof).

Lemma 2.3 *Assume that all marginal valuations are of the form $v_i(q) = v_i$. Then any incentive compatible on-line auction is based on non-decreasing supply curves.*

proof: Fix some incentive compatible on-line auction A. According to Theorem 2.1, A is based on supply curves. Consider the sale to the i'th player. Denote A's total price function by $P_i(q)$, and A's allocation rule by $q_i(v)$, i.e. A sells a quantity $q_i(v)$ for a bid v (since the marginal valuation is fixed, each bid is simply a single value).

We first argue that $q_i(v)$ is non-decreasing. Otherwise, suppose there are two bids $\tilde{v} > v$ such that $q_i(\tilde{v}) < q_i(v)$. Denote $q_i(v) = q$, $q_i(\tilde{v}) = \tilde{q}$. Since the auction is incentive compatible, $\tilde{v} \cdot \tilde{q} - P_i(\tilde{q}) \geq \tilde{v} \cdot q - P_i(q)$, and therefore $P_i(q) - P_i(\tilde{q}) \geq \tilde{v}(q - \tilde{q}) > v(q - \tilde{q})$. Thus $v \cdot \tilde{q} - P_i(\tilde{q}) > v \cdot q - P_i(q)$ and according to the supply curves definition, A must sell a quantity of $\tilde{q}$ for a bid value of v, in contradiction. Now define

$$p_i(q) = \inf \{ \ v \mid q_i(v) \geq q \ \},$$

i.e. A sells at least q for any bid $v > p_i(q)$. Since $q_i(v)$ is non-decreasing then $p_i(q)$ is non-decreasing as well. We claim that A is based on $p_i(q)$. In other words, for every bid v, if A sells a quantity q then $P_i(q) = \sum_{j=1}^{q} p_i(j)$.

To see this, we argue that for any $l \geq 1$ and $q \geq l$ such that $p_i(q) = \ldots = p_i(q-l+1) > p_i(q-l)$, it is the case that $P_i(q) - P_i(q-l) = l \cdot p_i(q)$ (to clarify this and what follows, consider first the simpler case where $p_i(\cdot)$ is strictly increasing – then l always equals 1 and the claim becomes that $P_i(q) - P_i(q-1) = p_i(q)$). Denote $x = (P_i(q) - P_i(q-l))/l$, and suppose by contradiction that $x \neq p_i(q)$. If $x < p_i(q)$ then a bidder with marginal valuation $\tilde{v}$ such that $x, p_i(q-l) < \tilde{v} < p_i(q)$ will increase his utility by declaring $p_i(q)$ instead – he will now receive q units instead of $q-l$ (this follows from the definition of $p_i(\cdot)$), and will pay for each additional unit x, which is less than $\tilde{v}$, his value for this unit. In other words, his utility will change by $l(\tilde{v} - x) > 0$, since he will receive additional l units, and his additional payment will be $P_i(q) - P_i(q-l) = l \cdot x$. In a similar manner, if $x > p_i(q)$ then a bidder with marginal valuation $x > \tilde{v} \geq p_i(q)$ will increase his utility by declaring $p_i(q-l)$ instead, since his utility will change by $l(x - \tilde{v}) > 0$ [6].

We can now conclude that, if i bids v and receives q, then he pays $\sum_{j=1}^{q} p(j)$. To see this, notice that from the above, the sum $\sum_{j=1}^{q} p(j)$ becomes a telescopic sum that reduces to $P_i(q) - P_i(0) = P_i(q)$, as needed. ∎

In general, there is no specific relation between the different supply curves of an auction. However, a useful structure of supply curves, which we use in section 2.3 below, is when all the supply curves are derived from some global supply curve, as follows:

Definition 2.2 *(A Global Supply Curve) An on-line auction is called "based on a global supply curve $p(q)$" if it is based on supply curves and if $p_i(q) = p(q + \sum_{j=1}^{i-1} q_j)$, where q_j is the quantity sold to the j'th bidder.*

In other words, the i'th supply curve is a left shift of the $(i-1)$'st supply curve by q_{i-1}. Thus, the i'th bidder receives the quantity according to the first supply curve $p_1(q)$ minus the quantity that was sold previously.

2.3 Competitive Analysis

In this section we describe on-line auctions with worst-case performance guarantees, i.e. the on-line performance for every valuation sequence is not too far from the off-line performance for the same sequence. We first define our performance measure (revenue and social welfare) and the exact meaning of a performance guarantee (competitiveness).

For the worst-case analysis, we assume that all marginal valuations are taken from some known interval $[\underline{p}, \bar{p}]$, without assuming any distribution on them. We assume that $\underline{p} > 0$, and that it is also the reservation price of the auctioneer, i.e. the auctioneer has a value of $\underline{p}$ for any unit he did not sell [7].

Definition 2.3 *(Revenue and Social Welfare) The* **revenue** *of an auction A for a valuation sequence σ, denoted as $R_A(\sigma)$, is the resulting utility of the auctioneer, i.e. the total payment he received plus his valuation of the quantity he did not sell. More specifically, let q_i be the quantity sold to the i'th player in σ and P_i be the total price paid by the i'th player, then:*

$$R_A(\sigma) = \sum_i P_i + \underline{p}(k - \sum_i q_i).$$

[6]To be completely exact, all this holds only for q's that can be received, i.e. q s.t. there exists v with $q_i(v) = q$, but those are the only q's we need to worry about.

[7]This may be his manufacturing or shipping cost (that he may save for unsold units), an option to sell the units for a very low price, etc.

The **social welfare** *of an auction A for a valuation sequence σ, denoted as $E_A(\sigma)$, is the sum of the resulting utilities of all players, including the auctioneer. This is also equal to the sum of all the players' valuations of the quantity they possess (including the auctioneer). I.e:*

$$E_A(\sigma) = \sum_i \sum_{j=1}^{q_i} v_i(j) + \underline{p}(k - \sum_i q_i).$$

We compare the revenue and the social welfare obtained by on-line auctions to those obtained by the off-line Vickrey auction [97]:

Definition 2.4 *(The Vickrey auction) In the Vickrey auction, each player declares his (supposedly) marginal valuation function. The allocation chosen is the one that maximizes the social welfare (according to players' declarations). The price charged from player i for the quantity q_i he receives is the worth of this additional quantity to the other players, i.e. the additional value of the other players when dividing q_i optimally among them. Formally, denote by E_{-i} the optimal social welfare when player i is missing, and by E the actual optimal social welfare. Then, the price that i pays is $E_{-i} - (E - v_i(q_i))$.*

For example, if there is only one indivisible good, this auction becomes the well known second price auction, where the highest bidder wins and pays the second highest offer. This is approximately equivalent to the popular English auction, where increasing bids are announced until no bidder wishes to make any further higher bid [97] – As the bid increments become smaller, the price paid by the winner becomes closer to the second highest value.

The use of the Vickrey auction as our benchmark is not only due to its popularity but also since it is incentive compatible. Such a non-Bayesian equilibrium is required for the worst-case analysis we desire. In any case, the Vickrey auction is always optimal in terms of the social welfare. While the revenue is not necessarily optimal in a Bayesian setting, the revenue equivalence theorem [69] states that other auctions with equivalent outcomes extract the same revenue.

We compare our on-line auction to the Vickrey auction in the following worst-case sense:

Definition 2.5 *(Competitiveness) An on-line auction A is c-competitive with respect to the revenue if for every valuation sequence σ, $R_A(\sigma) \geq R_{vic}(\sigma)/c$. Similarly, A is c-competitive with respect to the social welfare if for every valuation sequence σ, $E_A(\sigma) \geq E_{vic}(\sigma)/c$.*

2.3.1 A divisible good

We first focus on the case of a divisible good, i.e. a good that can be divided to any number of small fractions (we assume w.l.o.g that we have *one* divisible good). We describe a global supply curve that is $\Theta(log(\bar{p}/\underline{p}))$-competitive with respect to both the revenue and the social welfare. For this purpose we use results of [36] for on-line continuous one way trading [8].

Let c be the unique solution to the equation:

$$c = ln\frac{(\bar{p}/\underline{p}) - 1}{c-1}. \tag{2.1}$$

It can be shown that $c = \Theta(ln(\bar{p}/\underline{p}))$. For example, if $(\bar{p}/\underline{p}) = 2$ then $c = 1.28$, and if $(\bar{p}/\underline{p}) = 8$ then $c = 1.97$ [36].

Definition 2.6 *(The Competitive On-Line Auction) Define the* **Competitive Supply Curve** *by:*

$$p(x) = \underline{p}(1 + (c-1)e^{cx}). \tag{2.2}$$

The **Competitive On-Line Auction** *has the Competitive Supply Curve as its global supply curve.*

In order to use the results of [36], we need to derive the following two functions from the global supply curve $p(x)$. Let $q(x) = p^{-1}(x)$ (the inverse function of $p(x)$) and let $r(x) = \int_0^{q(x)} p(y)dy$.[9] In our context, these functions can be interpreted as follows: $q(x)$ is the total quantity sold by the Competitive On-Line Auction when the last bid intersects the last supply curve at price x, and $r(x)$ is the total payment charged by the auction for such a sequence. [36] analyzes these functions (separately from their context to the supply curve), and shows that:

Lemma 2.4 (El-Yaniv, Fiat, Karp, and Turpin [36]) *The functions $q(x), r(x)$ preserve the following conditions:*

1. $\forall x \leq c \cdot \underline{p} \; : \; q(x) = 0, r(x) = 0$
2. $\forall x > c \cdot \underline{p} \; : \; r(x) + \underline{p} \cdot (1 - q(x)) = x/c.$
3. $q(\bar{p}) = 1.$

where c is as defined in Eq. 2.1.

The paper [36] also states the minimality of the constant c in the following sense (this lemma is implicit in [36]):

Lemma 2.5 (El-Yaniv, Fiat, Karp, and Turpin [36]) *For any constant $\tilde{c} < c$, there is no function $\tilde{q}(x)$ such that*

$$\forall x \in [\underline{p}, \bar{p}], \; \tilde{r}(x) + \underline{p} \cdot (1 - \tilde{q}(x)) \geq x/\tilde{c}$$

where $\tilde{r}(x) = \int_0^{\tilde{q}(x)} \tilde{p}(x)dx$ and $\tilde{p}(x) = \tilde{q}^{-1}(x)$ is the inverse function of $\tilde{q}(x)$.

[8] In this on-line model, a trader needs to convert dollars to yen. The exchange rate is unpredicted, and is determined by an adversary. [36] gives several algorithms to compete in such an environment. Here, we construct a global supply curve from a specific function developed in [36] for the purpose of describing the trader conversion behavior.

[9] The paper [36] uses $r(x) = \int_0^x yq'(y)dy$. It can be verified that both terms are equal.

Theorem 2.2 *The Competitive On-Line Auction is c-competitive with respect to the revenue and the social welfare.*

proof: We prove the following lemmas:

Lemma 2.6 *For any sequence of valuations σ, $R_{cola}(\sigma) \geq R_{vic}(\sigma)/c$, where "cola" is the Competitive On-Line Auction and "vic" is the Vickrey auction.*

proof: Fix some valuation sequence σ. For a player i let q_i be the quantity he received in the online auction, and denote $p_i = p_i(q_i)$. Let m be the last player that received a positive quantity q_m. For all i and $q > q_i$, $b_i(q) \leq p_i$. According to condition 2 of Lemma 2.4, the on-line revenue is p_m/c, It also follows that for all i, $p_i \leq p_{i+1}$ since $p(q)$ is non-decreasing. Thus, every player values any additional quantity (to that he already received), Δq, by no more than $p_m \cdot \Delta q$.

The price that the Vickrey auction determines for the quantity q_i^* it allocates to player i is the highest valuation of the other players for the additional quantity $\Delta q = q_i^*$ divided among them. Let us devise an upper bound to this price.

If there exists a player, say i', for which $q_{i'}^* > q_{i'}$, then this implies that there is some small quantity ϵ that is worth to i' at most p_m (i.e. $\int_{q_{i'}^*-\epsilon}^{q_{i'}^*} v_{i'}(x)dx \leq p_m \cdot \epsilon$). From the maximality of the Vickrey allocation it follows that for any other player, an addition of an ϵ quantity to the quantity he received does not worth more than $p_m \cdot \epsilon$. Since the marginal valuations are non increasing, this must be true for any Δq quantity addition. Thus, allocating q_i^* optimally among the other players will result in a total additional worth of at most $p_m \cdot q_i^*$, and the sum of Vickrey prices is therefore at most $\sum_i p_m \cdot q_i^* = p_m$. Hence the claim follows.

Otherwise, for all i, $q_i^* \leq q_i$. Since the Vickrey auction allocates all the quantity, it follows that $q_i^* = q_i$ for all i. From this, again, it follows that any quantity addition Δq to i will be worth at most $p_m \cdot \Delta q$, and the claim follows. ∎

Lemma 2.7 *For any sequence of valuations σ, $E_{cola}(\sigma) \geq E_{opt}(\sigma)/c$, where $E_{opt}(\sigma)$ is the optimal social welfare for σ.*

proof: Fix some valuation sequence σ and denote q_i, p_i, and m as in the previous lemma. Consider a new sequence σ^* as follows:

$$b_i^*(q) = \begin{cases} p_i & q \leq q_i \\ b_i(q) & \text{otherwise,} \end{cases}$$

I.e. player i has fixed marginal valuation up to $q = q_i$ and then as before. The on-line allocation for this sequence does not change since $b_i^*(q)$ intersects the supply curve at $p(q_i)$. Since $b_i(q) \leq p_i \leq p_m$ for all i, it follows that $E_{opt}(\sigma^*) \leq p_m$. It is also true that $E_{cola}(\sigma^*) \geq R_{cola}(\sigma^*) = p_m/c$ (the equality is due to condition 2 of Lemma 2.4). Thus, $E_{cola}(\sigma^*) \geq E_{opt}(\sigma^*)/c$.

Now consider moving from σ to σ^* in m steps. In each step i, if $b_i(q) > b_i^*(q)$ at some points, then $b_i(q)$ is decreased to $b_i^*(q)$. Let σ^i be σ after i such modifications. The on-line auction allocates to i

the entire quantity whose value decreased, and thus $E_{cola}(\sigma^i) - E_{cola}(\sigma^{i+1}) \geq E_{opt}(\sigma^i) - E_{opt}(\sigma^{i+1})$, i.e. the on-line welfare decrease is greater than the off-line decrease since it is the maximal possible. From this it follows that $E_{cola}(\sigma) - E_{cola}(\sigma^*) \geq E_{opt}(\sigma) - E_{opt}(\sigma^*)$, and we get:
$E_{cola}(\sigma) \geq (E_{opt}(\sigma) - E_{opt}(\sigma^*))/c + E_{cola}(\sigma^*) \geq$
$(E_{opt}(\sigma) - E_{opt}(\sigma^*))/c + E_{opt}(\sigma^*)/c = E_{opt}(\sigma)/c$. ∎
From the above two lemmas the theorem follows. ∎

A natural question to ask is whether the on-line revenue is competitive with respect to some higher revenue criteria. As it turns out, it can be shown that for the special case of fixed marginal valuations, the on-line revenue is c-competitive with respect to the *optimal social welfare*, i.e. with respect to the off-line auction that extracts the total surplus – clearly this is the best revenue we can hope for since no player will pay more than his value. This follows basically from the following argument: Let p_i denote the y-coordinate of the intersection point of the i'th bid with the supply curve (as in the proof above). Since players have fixed marginal valuations v_i, it follows that $v_i = p_i$. Let m be the last player that received a positive quantity. Since for all i, $p_i \leq p_{i+1}$, it follows that $v_m \geq v_i$ for any other player i. Therefore, the optimal social welfare is $v_m \cdot 1$, by allocating the entire quantity to player m. On the other hand, the online revenue is p_m/c (as shown in the proof above), which proves the claim.

In contrast, for general valuations, the on-line revenue is significantly lower than the optimal welfare in cases where the Vickrey revenue is significantly lower than the optimal welfare. For example, consider the following scenario of two players. Let p^* be some price and q^* be the quantity such that $p^* = p(q^*))$. The first player has a fixed marginal valuation of $\bar{p}$ up to q^*, and the second player has a fixed marginal valuation of p^*. The optimal welfare for this scenario is $q^* \cdot \bar{p} + (1 - q^*) \cdot p^*$. In the on-line auction, player 1 will receive a quantity of q^*, since this is the maximal quantity for which his valuation is higher then the supply curve. Player 2 will receive nothing, since the second supply curve is higher than p^* (as the auction is based on a global supply curve). Therefore, the on-line revenue is at most $q^* \cdot p^* + (1 - q^*) \cdot \underline{p}$. Thus, for example, when setting $p^* = \sqrt{\underline{p} \cdot \bar{p}}$ then the optimal welfare to on-line revenue ratio is larger than $\sqrt{\bar{p}/\underline{p}}$. It is interesting to observe that, if the arrival order of the players is reversed, then the on-line revenue increases significantly to $\bar{p}/c$ (although the Vickrey revenue remains the same).

If the players' valuations are drawn independently from a known probability distribution, then the Vickrey auction with an appropriate reservation price [10] is known to have optimal revenue for several special cases (e.g. when each player has unit demand) [69]. We note that our auction can be modified to be competitive with respect to the Vickrey auction with reservation price by simply taking, as the supply curve, the maximum between the original supply curve and the reservation price (at each point).

We now show that the competitive ratio of the Competitive On-Line Auction is the best we can

[10] A threshold price - no sale is performed for a lower price.

expect:

Theorem 2.3 *Any incentive compatible on-line auction must have a competitive ratio of at least c with respect to both the revenue and the social welfare, where c is the solution to Eq. 2.1.*

proof: We prove the claim for the special case of fixed marginal valuations (in other words, even if the adversary is restricted to use only fixed marginal valuations the claim holds). For this case we can assume w.l.o.g (according to lemma 2.3) that A is based on non-decreasing supply curves. We also assume w.l.o.g that $\underline{p} = 1$, and denote $\bar{p} = \phi$. Let f_n be the n'th root of ϕ, i.e. ${f_n}^n = \phi$, and $c_n = c/({f_n}^2)$.

The following lemma assumes only the more restricted partially on-line model, in which the number of players, n, is known in advance. For this case, it lower bounds the competitive ratio of any on-line auction by c_n, thus also implying that knowing the number n in advance may help significantly only for small values of n.

Lemma 2.8 *No on-line auction with n bidders achieves social welfare that is better than c_n competitive with respect to the revenue of the Vickrey auction.*

proof: Assume we have a better than c_n competitive auction, we will build a function $\tilde{q}(x)$ satisfying the condition of lemma 2.5 with a constant $\tilde{c} < c$, a contradiction.

Consider the behavior of the on-line auction on the sequence of bids of the n bidders: $p_1 = f_n, p_2 = {f_n}^2, ..., p_n = \phi$. Let q_i be the quantity allocated to bidder i. For all x in the range $1 \leq x \leq \phi$, define $\tilde{q}(x)$ as $\sum_{j=1}^{i} q_j$, where i is such that $p_{i-1} \leq x < p_i$ (for completeness, denote $p_0 = 1,\ p_{n+1} = \infty$). The function $\tilde{r}(x)$ (as defined in lemma 2.5) is now $\tilde{r}(x) = \sum_{j=1}^{i} q_j p_j$ for the same i. [11]

Now, for each i, consider the sequence of bids where the first i bids are $p_1, \ldots, p_i$, but the other $n - i$ bids are simply 1. The revenue of the Vickrey auction in this case is $p_{i-1} = p_i/f_n$. The social welfare of the on-line auction is given by $\tilde{r}(p_i) + (1 - \tilde{q}(p_i))$. Since we assumed better than c_n competitiveness, we have $\tilde{r}(p_i) + (1 - \tilde{q}(p_i)) > p_i/(c_n f_n) = p_i f_n/c$. It follows that for every x, if we let i be such that $p_{i-1} \leq x < p_i$, then we have $\tilde{r}(x) + (1 - \tilde{q}(x)) = \tilde{r}(p_i) + (1 - \tilde{q}(p_i)) > p_i f_n/c \geq (x/f_n) f_n/c = x/c$. This is exactly the condition of lemma 2.5, completing the contradiction. ■

From this lemma it follows that for every on-line auction A with n players there is a valuations sequence σ such that $R_A(\sigma) \leq E_A(\sigma) \leq R_{vic}(\sigma)/c_n \leq E_{vic}(\sigma)/c_n$. Therefore A is no less than c_n-competitive with respect to both the revenue and the social welfare. Since c_n approaches c as n grows to infinity, the theorem follows. ■

[11] Since $q(x)$ is not a one to one function there is no inverse function $p^{-1}(x)$, but it can be verified that the function $p(x) = p_i$ for i such that $q_{i-1} < x \leq q_i$ is the appropriate function to use for the definition of $r(x)$.

2.3.2 A randomized auction for k indivisible goods

We now discuss the discrete case. First we show that, by using randomization, it is possible to obtain an *expected* revenue and social welfare that are c-competitive (where c is as before), *both* with respect to the optimal social welfare. Thus, randomization enables us to improve the performance with respect to the revenue.

When allowing randomization, the definition of supply curves should be altered so that each supply curve may be chosen randomly according to some distribution (an auction that is based on such supply curves is incentive compatible in a strong sense, as detailed below). As noted in [37], the function $q(x) = p^{-1}(x)$ (i.e. the inverse function of the Competitive Supply Curve $p(x)$ of Eq. 2.2) may be viewed as a cumulative distribution function in the interval $[\underline{p}, \bar{p}]$. I.e., if we choose x randomly using $q(\cdot)$, we have that for any fixed $v \in [\underline{p}, \bar{p}]$, $Pr(x \leq v) = q(v)$ (notice that $q(\underline{p}) = 0$, $q(\bar{p}) = 1$, and that $q(\cdot)$ is non-decreasing).

Definition 2.7 *The Randomized On-Line Auction: Before receiving any bids, the auction first chooses some fixed price p_{on} randomly by using the cumulative distribution $q(\cdot)$. The supply curve is then simply $p(x) = p_{on}$. I.e. the auction sells, to each player, all the goods with value of at least p_{on} to him, with price p_{on} for each good (until all the goods are sold).*

This auction is incentive compatible in the following strong sense: For *any* result of the randomized choice, a player will maximize his utility by declaring his true valuation. This is because the randomized choice actually determines a supply curve independently of his bid. We note that it is possible to consider a weaker notion of incentive compatibility, in which a player will maximize his *expected* utility (with respect to the distribution of the randomized choice) by declaring his true valuation.

The following theorem shows that this auction is c-competitive with respect to its *expected* revenue and social welfare. That is, in some cases the on-line revenue and welfare will not be within a factor of $1/c$ of the optimal welfare. But, for *any* particular valuation sequence, the expected on-line revenue and welfare is within a factor of $1/c$ of the optimal welfare.

Theorem 2.4 *For any sequence of valuations σ, the expected revenue of the Randomized Auction is at least $1/c$ times the optimal social welfare, i.e. $E(R_{on}(\sigma)) \geq E_{opt}(\sigma)/c$.*

proof: Suppose OPT allocates the k goods to players with valuations $v_1, \ldots, v_k$ (possibly several goods to the same player) such that $v_i \geq v_{i+1}$. For convenience assume $v_0 = \bar{p}$, $v_{k+1} = \underline{p}$. Let p_{on} be the actual price determined by the on-line auction. For the specific i such that $v_{i+1} \leq p_{on} \leq v_i$, the on-line auction sells at least i goods, and thus its revenue is at least $i \cdot p_{on} + \underline{p}(k - i)$. Denote the density function that the auction uses by $f(x) = \frac{d[q(x)]}{dx}$. Therefore we have:

$$E(R_{on}|v_{i+1} \leq p_{on} \leq v_i) \geq \int_{v_{i+1}}^{v_i} i \cdot x \cdot \frac{f(x)}{Pr(v_{i+1} \leq p_{on} \leq v_i)} dx + \underline{p}(k - i)$$

and thus

$$E(R_{on}) \geq \sum_{i=0}^{k} [\ i \cdot \int_{v_{i+1}}^{v_i} xf(x)dx + \underline{p}(k-i) \cdot Pr(v_{i+1} \leq p_{on} \leq v_i)\ ] \ =$$
$$\sum_{i=0}^{k} [\ i \cdot \int_{v_{i+1}}^{v_i} xf(x)dx + \underline{p}(k-i)(q(v_i) - q(v_{i+1}))\ ]\ .$$

Changing the summation order of the right side we get

$$E(R_{on}) \geq \sum_{i=1}^{k} [\ \int_{\underline{p}}^{v_i} xf(x)dx\ ] \ + \sum_{i=1}^{k} [\ \underline{p}(q(v_0) - q(v_i))\ ] =$$
$$\sum_{i=1}^{k} [\ \int_{\underline{p}}^{v_i} xf(x)dx \ + \ \underline{p}(1 - q(v_i))\ ] \ = \sum_{i=1}^{k} \frac{v_i}{c} = E_{opt}(\sigma)/c\ ,$$

where the last equality follows from the specific character of the function $q(x)$, as stated in Lemma 2.4, and thus the claim follows. ∎

2.3.3 A deterministic auction for k indivisible goods

We next examine the deterministic case. First consider the case of $k = 1$. It follows from theorem 2.1 that for this case the on-line auction must fix some reservation price p_i for the i'th player, i.e. the good is sold to the i'th player for price p_i if $v_i(1) > p_i$. This is similar to the search algorithm of [37], where it is shown that a reservation price of $\sqrt{\bar{p} \cdot \underline{p}}$ is $\sqrt{\phi}$-competitive. It is not hard to verify that this is optimal. The general case for any $k \geq 1$ may be handled similarly:

Definition 2.8 *(The Discrete On-Line Auction) The Discrete On-Line Auction is based on the following global supply curve:*

$$p(j) = \underline{p} \cdot \phi^{\frac{j}{k+1}},\ \textit{for}\ j = 1, \ldots, k. \tag{2.3}$$

We next analyze the competitiveness of this auction, and give also a lower bound for this case.

Theorem 2.5 *The Discrete On-Line Auction is $k \cdot \phi^{\frac{1}{k+1}}$-competitive with respect to the revenue and to the social welfare. When $k \geq 2 \cdot ln\phi$ then the Discrete On-Line Auction is also $2 \cdot e \cdot (ln(\phi) + 1)$-competitive with respect to the revenue and to the social welfare.*

proof: Fix some scenario and suppose that the on-line auction sold q goods. We first prove the claim with respect to the revenue. Since the on-line auction sold q goods, the valuation of one additional good of any player is at most $p(q+1) = \underline{p} \cdot \phi^{\frac{q+1}{k+1}}$. Therefore the Vickrey auction may charge a unit price of at most $p(q+1)$, thus the Vickrey to on-line revenue ratio is at most:

$$\frac{k \cdot p(q+1)}{\sum_{j=1}^{q} p(j) + (k-q) \cdot \underline{p}} = \frac{k \cdot \phi^{\frac{q+1}{k+1}}}{\sum_{j=1}^{q} \phi^{\frac{j}{k+1}} + (k-q)} \leq \frac{k \cdot \phi^{\frac{k+1}{k+1}}}{\sum_{j=1}^{k} \phi^{\frac{j}{k+1}}}$$

where the inequality follows from the fact that for any $q,\ 0 \leq q \leq k-1$,

$$\frac{k \cdot \phi^{\frac{q+1}{k+1}}}{\sum_{j=1}^{q} \phi^{\frac{j}{k+1}} + (k-q)} \leq \frac{k \cdot \phi^{\frac{q+2}{k+1}}}{\sum_{j=1}^{q+1} \phi^{\frac{j}{k+1}} + (k-q-1)}$$

(since $k - q \geq 1 + \frac{k-q-1}{\phi^{\frac{1}{k+1}}}$). The first part of the claim follows since $\frac{k \cdot \phi^{\frac{k+1}{k+1}}}{\sum_{j=1}^{k} \phi^{\frac{j}{k+1}}} \leq \frac{k \cdot \phi^{\frac{k+1}{k+1}}}{\phi^{\frac{k}{k+1}}} = k \cdot \phi^{\frac{1}{k+1}}$. For the second part of the claim, let $l^* = \frac{k+1}{ln\phi} - 1$. If $ln\phi < 1$ then $2 \cdot e \cdot (ln(\phi) + 1) > e > \phi$ and the claim is trivial since any auction is ϕ-competitive. Otherwise $ln\phi \geq 1$, and $1 \leq l^* \leq k$ since $k \geq 2 \cdot ln\phi$. We claim that:

$$\sum_{j=1}^{k} \phi^{\frac{j}{k+1}} \geq l^* \cdot \phi^{\frac{k+1-l^*}{k+1}}.$$

Clearly this is true for any integer l. Let the right hand function be $f(l)$. It receives its maximum for $l^* + 1$ and it is increasing in $[l^*, l^* + 1]$. Thus for some integer $x \in [l^*, l^* + 1]$ it holds that $\sum_{j=1}^{k} \phi^{\frac{j}{k+1}} \geq f(x) \geq f(l^*)$. Thus:

$$\frac{k \cdot \phi}{\sum_{j=1}^{k} \phi^{\frac{j}{k+1}}} \leq \frac{k \cdot \phi}{l^* \cdot \phi^{\frac{k+1-l^*}{k+1}}} = \frac{k}{l^*} \cdot \phi^{\frac{l^*}{k+1}} = \frac{k \cdot ln\phi}{k+1-ln\phi} \cdot \phi^{\frac{1}{ln\phi} - \frac{1}{k+1}} \leq 2 \cdot e \cdot ln\phi$$

where the last inequality follows from the fact that $k \geq 2 \cdot ln\phi$, and $\phi^{\frac{1}{ln\phi}} = e$.

This proves the claim for the revenue. Now consider the case of social welfare. Given a valuation sequence $\sigma = (b_i(q))$, consider a new sequence $\sigma^* = (b_i^*(q))$ built in a similar manner to that of lemma 2.7, i.e.player i has fixed marginal valuation up to $q = q_i$ and then as before. By similar arguments to those of lemma 2.7, the off-line to on-line welfare ratio of the new sequence is an upper bound to the ratio of the original sequence (since the off-line welfare decrease is no more than the on-line decrease). Additionally, the on-line allocation for the two scenarios is identical. Let $\sum_i q_i = q \leq k$. Since $b_i^*(q) \leq p(q+1)$ then $E_{opt}(\sigma^*) \leq k \cdot p(q+1)$. Clearly $E_{on}(\sigma^*) \geq R_{on}(\sigma^*) \geq \sum_{j=1}^{q} p(j)$, where "on" is the Discrete On-Line Auction. Thus:

$$\frac{E_{opt}(\sigma^*)}{E_{on}(\sigma^*)} \leq \frac{k \cdot p(q+1)}{\sum_{j=1}^{q} p(j)} \leq 2 \cdot e \cdot ln\phi$$

where the last inequality was shown above for the revenue claim. ■

Theorem 2.6 *Any incentive compatible on-line auction of k goods has a competitive ratio of at least $m = max\{\phi^{\frac{1}{k+1}}, c\}$ with respect to the revenue and to the social welfare, where c is as defined in Eq. 2.1.*

proof: We prove the claim for the special case of fixed marginal valuations and assume, according to lemma 2.3, that A is based on non-decreasing supply curves. We prove each lower bound separately:

Lemma 2.9 *Any incentive compatible on-line auction of k goods has a competitive ratio of at least $\phi^{\frac{1}{k+1}}$ with respect to the revenue and to the social welfare.*

proof: Fix some incentive compatible on-line auction A. Consider the behavior of A for the sequence of players: $\phi^{\frac{1}{k+1}}, \phi^{\frac{1}{k+1}}, \phi^{\frac{2}{k+1}}, \phi^{\frac{2}{k+1}}, \ldots, \phi, \phi$ (i.e. $2(k+1)$ players). Let q be the first i such that both players with valuation $\phi^{\frac{i}{k+1}}$ does not receive any positive quantity (there is such q since there are k goods and $k+1$ pairs of players). Denote by σ the above sequence with only the first $2(q+1)$ players. Vickrey's revenue is $k \cdot \phi^{\frac{q}{k+1}}$, while A's welfare is at most $k \cdot \phi^{\frac{q-1}{k+1}}$. Thus $R_A(\sigma) \leq E_A(\sigma) \leq R_{vic}(\sigma)/(\phi^{\frac{1}{k+1}}) = E_{vic}(\sigma)/(\phi^{\frac{1}{k+1}})$, and the claim follows. ∎

Lemma 2.10 *Any incentive compatible on-line auction has a competitive ratio of at least c with respect to the revenue and to the social welfare.*

proof: The claim follows from the fact that for the special case of fixed marginal valuations the result of the Vickrey auction for the indivisible case is the same as for the divisible case (i.e. a single player receives all the good(s) and pays the second price). Thus, if there was an on-line auction for k indivisible goods with a competitive ratio $\tilde{c} < c$ it can be used for the divisible case (i.e. allocating quantity multiples of $1/k$), achieving the same competitive ratio $\tilde{c}$. This is in contradiction to Theorem 2.3, since the lower bound there holds for the special case of fixed marginal valuations. ∎

Remark: When considering a partially on-line model, in which the number of players, n, is known in advance, this lower bound weakens, becoming dependent in n. For example, consider the following auction of one good to two players: The price for the first bidder is $\underline{p} \cdot \phi^{\frac{2}{3}}$, while the price for the second bidder is $\underline{p} \cdot \phi^{\frac{1}{3}}$. It is easy to verify that this is $\phi^{\frac{1}{3}}$-competitive with respect to the revenue. As can be seen from the lower bound proof, as long as $n < k+2$ then a similar improvement (with respect to the revenue) is possible.

2.4 Model Extensions

We now discuss some natural extensions to our on-line model, incorporating the following time considerations:

1. Delayed bidding: Player i learns his valuation at time t_i, and his strategy space allows placing his bid at any time $t \geq t_i$.

2. Split bidding: Player i's strategy space allows placing several bids at any time $t_{i_1}, \ldots, t_{i_l} \geq t_i$.

3. The Players' valuations may be time-dependent (in a non-increasing way). Specifically, player i's valuation is given by $v_i(q,t)$, where $v_i(\cdot)$ is non-increasing both in q and in t. $v_i(q,t)$ is player i's marginal valuation of the q'th good at time t.

4. The strategy of a player may depend on past events.

A truthful bid is still considered as bidding the true valuation exactly once, at time t_i.

We note that even under any of these extensions, when the supply curves are non-decreasing over time there is no possible gain for a player from delaying his bid. Clearly, a non-decreasing global supply curve holds this property. Thus we conclude:

Theorem 2.7 *Any On-Line Auction that is based on a non-decreasing global supply curve is incentive compatible even in any of these extensions.*

Thus all our auctions remain truthful. Their competitiveness also remains, since the off-line Vickrey allocation is not affected by the on-line assumptions. The lower bounds we have shown still obviously remain true. In fact it turns out that they even generalize to partially on-line auctions (where the number of players is known in advance).

2.5 Revenue Analysis for the Uniform Distribution

We compare the expected revenue of the Competitive On-Line Auction to the expected revenue of the Vickrey off-line auction for a divisible good in the special case of fixed marginal valuations uniformly distributed in $[a, b]$. This is a simple example that demonstrates that the on-line revenue is similar to the Vickrey revenue in some cases.

Table 2.1 compares the revenue of the on-line auction to the revenue of the Vickrey auction for several values of n and $\bar{p}$, where $\underline{p} = 1$. From the table, we see that for small values of n and $\bar{p}$ the on-line revenue is close to the Vickrey revenue. When n increases, the Vickrey to on-line revenue ratio approaches c, the competitive ratio.

	On-Line revenue	Vickrey revenue
$\bar{p} = 1.5,\ n = 2$	1.15	1.17
$\bar{p} = 3,\ n = 2$	1.60	1.67
$\bar{p} = 10,\ n = 2$	3.33	4.00
$\bar{p} = 2,\ n = 2$	1.31	1.33
$\bar{p} = 2,\ n = 3$	1.37	1.50
$\bar{p} = 2,\ n = 100$	1.56	1.98

Table 2.1: On-line and Vickrey revenue in the average case.

The details of this exercise are as follows. We consider fixed marginal valuations uniformly distributed in $[a, b]$ (for $a > 0$). Let $f(x) = \frac{1}{b-a}, F(x) = \frac{x-a}{b-a}$ be the distribution function and the cumulative distribution function, respectively, and assume that the players' utilities are independent. For this case, the revenue of the on-line auction is determined by the maximal marginal utility. Its distribution function for n players is $g_n(x) = nf(x)(F(x))^{n-1}$. Let $c = c(b/a)$ be the appropriate competitive ratio. Then,

$$\int g_n(x)\frac{x}{c}dx = \frac{n}{c(b-a)^n}\int (x-a)^{n-1}xdx = \frac{n}{c(b-a)^n}(\frac{(x-a)^{n+1}}{n+1} + a\frac{(x-a)^n}{n})$$

$$E(R_{on}) = \int_a^{ca} g_n(x)adx + \int_{ca}^b g_n(x)\frac{x}{c}dx = \int_a^b g_n(x)\frac{x}{c}dx + \int_a^{ca} g_n(x)(a - \frac{x}{c})dx = \frac{n}{n+1}\frac{b}{c} + \frac{1}{n+1}\frac{a}{c} + \epsilon$$

The ϵ addition is relatively small, e.g. for $a = 1,\ b = 2,\ n = 2$ it is lower than 0.006.

The distribution function for the second maximal price is $h_n(x) = 2f(x)(F(x))^{n-2}(1 - F(x))$, which opens to:

$$n(n-1)(F(x))^{n-2}(1-F(x))\int h_n(x)xdx = \frac{n(n-1)}{(b-a)^{n-1}}(\int (x-a)^{n-2}xdx - \frac{1}{b-a}\int (x-a)^{n-1}xdx)$$

This is solved in a similar manner to the previous integral, thus the expected revenue of the Vickrey auction (see also [56]) is: $E(R_{vic}) = \int_a^b h_n(x)xdx = \frac{n-1}{n+1}b + \frac{2}{n+1}a$.

Chapter 3

Towards a Characterization of Truthful Combinatorial Auctions

3.1 Introduction

3.1.1 Motivation

This chapter[1] is concerned with the general search of implementations in dominant strategies. There is only one known general method for designing incentive compatible mechanisms: the Vickrey-Clarke-Groves (VCG) mechanisms [97, 27, 44]. As described in chapter 1, this method applies for cases where the social goal is to maximize the welfare: the sum of players' valuations. There are two main reasons that should motivate us to look for other types of mechanisms. First, the social goal may be different than welfare maximization. For example, one might desire other fairness criteria like rawls' max-min principle, minimizing the sum-of-squares of the values (or other norms of the valuation vector), or considering the trade-offs between the different criteria. In an auction setting, another common goal would be to ignore fairness and efficiency all together, and instead maximize the seller's revenue. Second, even if the social goal is the maximization of the welfare, in many cases this optimization problem is computationally infeasible. In such cases, it seems reasonable to settle in achieving an approximate optimum. The key difficulty is the fact that attaching VCG-payments to approximation methods, or to any other social goal, does *not* ensure incentive compatibility. This problem was shown to be essentially universal in [80]. Thus, the basic positive tool of mechanism design fits only to a very limited setting.

At this point we should mention that a widely accepted approach to this difficulty is to weaken the notion of the equilibrium, replacing dominant strategies with e.g. a Bayesian-Nash equilibrium. However, as dominant strategies equilibrium is significantly stronger, and certainly much more convincing, it seems we should try to justify the use of other significantly weaker notions in an

[1]This chapter is based on a joint work with Ahuva Mu'alem and Noam Nisan [58].

analytic manner. Indeed, for a very restricted class of "single dimensional" domains, we witness a recent surge of positive results with dominant strategies. These include e.g. scheduling with a min-max criteria [2], revenue maximization for digital goods [39, 89], auctioning with bounded communication [22], as well as combinatorial auctions with very restrictive bidders (see below). This naturally leads us to wonder whether there exist a variety of incentive compatible mechanisms for multi-dimensional domains.

A particularly central problem that captures all these difficulties is Combinatorial Auctions. In a combinatorial auction, k items are simultaneously auctioned among n bidders. Bidders value bundles of items in a way that may depend on the combination they win, i.e. each bidder has a valuation function v_i that assigns a real value $v_i(S)$ for each possible subset of items S that he may win. This model has many real world applications (e.g. the FCC spectrum rights auction), and, equally important, it generalizes many classic combinatorial problems like scheduling and allocation of network resources. Even if the social goal is to maximize the welfare, i.e. to find a partition $S_1 ... S_n$ of the items in a way that maximizes $\sum_i v_i(S_i)$, it is NP-hard to exactly solve it[2]. Experimental results have shown that reasonable heuristics can quickly obtain an approximate optimum for problems with thousands of items. Unfortunately, it is not known how to turn such non-fully-optimal heuristics or approximation algorithms into incentive compatible mechanisms. Thus, all the abstract discussion given above seems to boil down to a very concrete real problem: what types of *truthful* combinatorial auctions can we design?

3.1.2 Characterizing Incentive Compatibility

A general approach to the question of designing incentive compatible mechanisms would be to obtain a characterization of their powers. To do this, let us get slightly more formal about the basic model.

There is a set A of possible outcomes of the mechanism, and each player has a valuation function $v_i : A \to \mathcal{R}$ that specifies his value $v_i(a)$ for each possible outcome $a \in A$, where v_i is chosen from some possible domain of valuations V_i. For each n-tuple of valuations $v = (v_1, \ldots, v_n)$, the mechanism produces some outcome $f(v)$ that may be viewed as aggregating the preferences v_i of the n players. The function f is called the *social choice function.* Additionally, the mechanism hands out payments to the players. For example, in the case of combinatorial auctions, A is the set of all possible partitions $(a_1, \ldots, a_n)$ of the items, and each V_i is the set of valuations that depend only on a_i ("no externalities") and are monotone in a_i ("free disposal"). It turns out that for each social choice function f that may be the obtained by an incentive compatible mechanism there is essentially a single way to "implement it" (in dominant strategies), i.e. to set the payments needed as to ensure incentive compatibility. The basic question is *what social choice functions are implementable?*

[2]The optimum may be approximated to within a factor of $O(\sqrt{n})$ (but no better [62, 49]).

The VCG mechanism mentioned above implements the social choice function that maximizes the social welfare, i.e. the social choice function $f(v) = \text{argmax}_{a \in A} \sum_i v_i(a)$. Three generalizations may be applied to the VCG payment scheme, yielding generalizations to the implemented social choice function: (a) the range may be restricted to an arbitrary $A' \subset A$; (b) different non-negative weights ω_i can be given to the different players; (c) different additive weights γ_a can be given to different outcomes. All three generalizations can be combined, yielding an implementation for any social choice function that is an *affine maximizer* [3]:

Definition: *A social choice function f is an affine maximizer if for some $A' \subset A$, non-negative $\{\omega_i\}$, and $\{\gamma_a\}$, for all $v_1 \in V_1, \ldots, v_n \in V_n$ we have $f(v_1, \ldots, v_n) \in argmax_{a \in A'}(\sum_i \omega_i v_i(a) + \gamma_a)$.*

What other social choice functions can be implemented? A classic negative result of Roberts [87] shows that if the domain of players' valuations is unrestricted, and the range is non-trivial, then nothing more:

Theorem (Roberts, 1979): *If there are at least 3 possible outcomes, and players' valuations are unrestricted ($V_i = \mathcal{R}^{|A|}$), then any implementable[4] social choice function is an affine maximizer.*

The requirement that the valuations are unrestricted is very restrictive. In almost all interesting scenarios the domain of valuations is restricted. E.g., as mentioned, for the combinatorial auction problem the valuations are restricted in two ways: "free disposal" and "no externalities", and thus $V_i \neq \mathcal{R}^{|A|}$. Indeed, some assumption about the space of valuations is also necessary: In the extreme opposite case, the domain is so restricted as to become single dimensional, for which truthful non affine maximizers exist, as mentioned above. Interesting examples in the context of combinatorial auctions involve "single-minded" bidders, where the valuation function is given by a single value v_i offered for a single set of items S_i [62]. While the optimization problem in this case is still NP-hard and thus affine maximization is not efficiently computable, [62] presented computationally efficient *truthful* approximation mechanisms for it. Additional mechanisms for this single-minded case were presented in [72, 1].

However, most interesting computational problems are not single dimensional either – they lie somewhere between the two extremes of "unrestricted" and "single dimensional". This intermediate range includes combinatorial auctions and many of their interesting special cases such as, multi-unit (homogeneous) auctions, or unit-demand auctions (matching). It also includes most examples of other combinatorial optimization problems such as various variants of scheduling and routing problems. Almost nothing is known about this intermediate range. The only positive example of a non-VCG mechanism for non-single-dimensional domains is for a special case of multi-unit combinatorial auctions where each bidder is restricted to demand at most a fraction of the number of units of each type [13].

[3]This term was coined by Meyer-ter-Vehn and Moldovanu [70].

[4]Roberts, as we do here, only discusses implementation in private-value environments. See [70] for a generalization to environments with inter-dependent valuations.

It is interesting to draw parallels with the non-quasi-linear case, i.e. the model where player preferences are given by order relations $\succeq_i$ over the possible outcomes. The classic Gibbard-Satterthwaite result [41, 94] shows that, in this case, the dictatorial social choice function over an unrestricted domain is implementable. The proof shows that any implementable social choice function must essentially satisfy Arrow's condition of "Independence of Irrelevant Alternatives", and thus Arrow's impossibility result [4] applies. On the other hand, in this non-quasi-linear case, there exists much literature for various interesting restricted domains. For example, over "single peaked domains" [24, 71], many non-dictatorial social choice functions are implementable, and over "saturated domains" [53], only dictatorial functions are implementable.

3.1.3 Our results

In this paper we initiate an analysis of implementable social choice functions over restricted domains in quasi-linear environments. It is widely known that certain monotonicity requirements characterize implementable social choice functions. E.g. Roberts starts by defining a condition of "positive association of differences" (PAD) that characterizes implementable social choice functions over unrestricted domains. It turns out that this condition is usually meaningless for restricted domains. We start with a formulation of a "weak monotonicity" condition (W-MON), that provides this characterization for "usual" restricted domains (exact definitions are given below)[5]. We also demonstrate that other natural notions are not appropriate.

Theorem: *Every implementable social choice function over every domain must satisfy W-MON. Over "usual" domains, W-MON is also a sufficient condition.*

As opposed to the case of unrestricted domains, it turns out that, for restricted domains, W-MON by itself does not imply affine maximization! A key contribution of this paper is the identification of a key additional property, *Independence of Irrelevant Alternatives (IIA)*, that will provide this implication. This property is a natural analog, in the quasi-linear setting, of Arrow's similarly named property in the non-quasi-linear setting. This condition states that if the social choice function changes its value from one outcome a to another outcome b, then this is due to a change in some player's preference between a and b.

Definition: *A social choice function f satisfies IIA if for any $v, u \in V$, if $f(v) = a$ and $f(u) = b \neq a$ then there exists a player i such that $u_i(a) - u_i(b) \neq v_i(a) - v_i(b)$.*

For example, in a combinatorial auction that satisfies IIA, the effect of some player increasing his value for the bundle that contains all items will be either that this player will now receive all items, or that the same allocation will still be chosen.

[5]Bikhchandani, Chatterji,and Sen [17] independently study the same condition for a restricted class of Multi-Unit Auctions. Later on, Muller and Vohra [73] provided different proofs and some generalizations for Combinatorial Auctions.

We show that the IIA property is equivalent to a slight, but significant, strengthening of the W-MON condition, termed "strong monotonicity". We further show that in unrestricted domains IIA may be assumed without loss of generality. This is also true in a class of domains that includes the case of combinatorial auctions with two players in which all items are always allocated. In other domains we demonstrate that IIA may not be assumed without loss of generality.

We then get to our main result: incentive compatible mechanisms that also satisfy IIA must be "almost" affine maximizers. The theorem is proved in a general setting and requires certain technical conditions.

Main Theorem: *In "auction-like" domains, any implementable social choice function that additionally satisfies IIA and certain technical conditions must be an "almost" affine maximizer.*

The proof of this theorem is different from the one Roberts provides for unrestricted domains, and uses ideas suggested, in a somewhat different context, by Archer and Tardos [3]. This theorem applies to combinatorial auctions as well as to multi-unit (non-combinatorial) auctions. It even applies to the case of "known double minded bidders", i.e. where each bidder has only two bundles on which he may bid – showing that the mechanisms of [62, 72] regarding single-minded bidders cannot be generalized this way (if one additionally requires IIA to be satisfied). For unrestricted domains, the IIA condition may be assumed without loss of generality, and therefore this yields a new proof of Roberts' theorem (the qualifications in the theorem statement all disappear in this case). For two-player auctions where all items must always be allocated, the IIA condition can similarly be dropped. We also show that in this two-player case, the requirement that all items must always be allocated is necessary – without it, there exist implementable social choice functions that are not almost affine maximizers (and do not satisfy IIA)[6].

The major open problem we leave is whether the IIA condition is necessary:

Main Open Problem: *Are there incentive compatible combinatorial auctions that are not "essentially" affine maximizers?*

The meaning of "essentially" in this open problem is soft, as we demonstrate that various "minor" variations from affine maximization are possible. The question is really whether anything *useful* is possible, e.g. can any non-trivial welfare approximation be achieved.

Our results has important implications to the existence of computationally efficient incentive compatible approximation mechanisms. Formally, a mechanism has an approximation ratio of c (or is a c-approximation) if it always produces outcomes with a social welfare of at least the optimal social welfare divided by c. We observe that essentially any affine maximizer is as computationally hard as exact social welfare maximization. This implies that if exact computation of the optimal allocation is computationally hard, then incentive compatible mechanisms that satisfy

[6]We note that the incentive compatible mechanism of [13] also does not always allocate all items.

IIA are essentially powerless. For an exact statement of computational hardness we must first fix an input format, i.e. a "bidding language" [77] that is powerful enough to make the exact optimization problem computationally intractable[7]. We say that a combinatorial auction mechanism is *unanimity-respecting* if whenever every bidder values only a single bundle, and furthermore, these bundles compose a valid allocation, then this allocation is chosen[8]. This condition essentially ensures that all allocations are possible outcomes, ruling out "bundling" auctions[9].

Theorem: *(Assuming $P \neq NP$ and a sufficiently powerful bidding language) Any unanimity-respecting truthful polynomial-time combinatorial (or multi-unit) auction that satisfies IIA cannot obtain any polynomially-bounded approximation ratio.*

An especially crisp result is obtained for the case of two-player multi-unit auctions. This case is still computationally hard, but has a $1+\epsilon$ approximation for any $\epsilon > 0$ (where the computation time depends on ϵ). However, this approximation is not incentive compatible. Indeed, [57] who considered this problem were only able to show "almost incentive compatibility" [10]. Our results show that this is no accident. Exact incentive compatibility directly collides with an approximation scheme.

Corollary: *(Assuming $P \neq NP$ and a sufficiently powerful bidding language) No polynomial time incentive compatible mechanism for a multi-unit auction between two players that always allocates all units can achieve an approximation factor better than 2.*

The rest of the chapter is organized as follows. In section 3.2 we describe our model. In section 3.3 we discuss the connection between truthfulness and monotonicity. Section 3.4 gives our main theorem and its proof. Section 3.5 discusses the implications to computationally efficient combinatorial auctions. Section 3.6 provides our alternative proof of Roberts' theorem.

3.2 Setting and Notations

We study a general model of a social choice function $f : V_1 \times ... \times V_n \rightarrow A$. The interpretation is that f gets as its input a vector of *players' preferences* and chooses an alternative among a finite set of possible alternatives A. We denote $|A| = m$, and assume w.l.o.g that f is onto A.

[7]E.g.: for general combinatorial auctions any complete bidding language that can succinctly express single-minded bids is enough; if the number of players is a fixed constant, the language must allow OR-bids; for multi-unit (non-combinatorial) auctions, the bidding language must allow specifying the number of items in binary.

[8]This is essentially equivalent to the property of a "reasonable" auction of [79].

[9]E.g., where all items are sold as a single bundle in a simple auction – this clearly gives a factor $\min(n,k)$-approximation. Slightly better approximations in polynomial time are possible by partitioning the items into a constant number of bundles [51].

[10]A somewhat similar notion of "almost incentive compatibility" for an approximation scheme for a different problem was also obtained in [1].

Each player i ($1 \le i \le n$) assigns a real value $v_i(a)$ to each possible alternative from A. The vector $v_i \in R^m$ is called the player's *type* and is interpreted as specifying the player's preferences. The set $V_i \subseteq R^m$ is the set of possible valuations v_i. We denote $V = V_1 \times ... \times V_n$. We use the notation $v = (v_1, ..., v_n) \in R^{nm}$, and $v(a) = (v_1(a), ..., v_n(a)) \in R^n$. We also use the notation $v_{-i} = (v_1 ... v_{i-1}, v_{i+1} ... v_n) \in R^{n-1}$. For $v_i \in V_i$, we denote by $u_i = v_i|^{a+=\delta}$ the following type: $u_i(a) = v_i(a) + \delta$, and for all $b \neq a$, $u_i(b) = v_i(b)$. Similarly, $u_i = v_i|^{a=\delta}$ denotes the type $u_i(a) = \delta$, and for all $b \neq a$, $u_i(b) = v_i(b)$. We use 1^m to denote the vector $(1, \ldots, 1) \in \mathcal{R}^m$.

The main point is that V_i may be a proper subset of R^m. Here are some of the domains that we are concerned with in this paper:

- **Unrestricted Domains.** We say that the domain is *unrestricted* if $V_i = R^m$. In other words, the value of alternative a for player i does not place any restrictions upon i's values for the other alternatives.

- **Combinatorial Auctions (CA).** In a combinatorial auction, a set Ω of k items are auctioned between n bidders. The "alternatives" that the auction chooses among are allocations of items to bidders. That is, an alternative a is an allocation $a = (a_1 ... a_n)$, where $a_i \subseteq \Omega$ is the set of items allocated to player i, and $a_i \cap a_j = \emptyset$ for $i \neq j$ (each item can be allocated to at most one player). The valuations are assumed to satisfy three conditions:

  1. No externalities: v_i should only depend on i's allocated bundle a_i. I.e. $v_i(a) = v_i(a_i)$.
  2. Free disposal: v_i should be non-decreasing with the set of allocated items. I.e. For every $a_i \subseteq b_i$, we have that $v_i(a_i) \le v_i(b_i)$.
  3. Normalization: $v_i(\emptyset) = 0$.

- **Multi Unit Auctions.** This is the special case of combinatorial auctions where the items are homogeneous. In this case an allocation $(a_1 ... a_n)$ is simply a vector of nonnegative integers, subject to the restriction that $\sum_i a_i \le k$, and the valuation functions v_i can be represented as non-decreasing non-negative functions $v_i : \{1 ... k\} \to R_+$.

- **Order-Based Domains.** We will phrase our results in this paper in terms of a general family of domains termed "order-based", which contains all the previous examples, as well as others. These are domains where each V_i is defined by a (finite) family of inequalities and equalities of the form $v_i(a) \le v_i(b)$, $v_i(a) < v_i(b)$, $v_i(a) = v_i(b)$ or $v_i(a) = 0$. Thus for example an unrestricted domain is defined by the empty family, while the domain of valuations for combinatorial auctions is defined by the following set of inequalities: for all $a, b \in A$ such that $a_i = b_i : v_i(a) = v_i(b)$ (no externalities); for all $a, b \in A$ such that $a_i \subseteq b_i$: $v_i(a) \le v_i(b)$ (free disposal); for all $a \in A$ such that $a_i = \emptyset$: $v_i(a) = 0$.

  We denote by $R_i(a, b)$ the relation of player i between alternatives a, b, and use $R_i(a, b) = null$ to denote that there is no such relation. We also use $0_i = \{a \in A \mid v_i(a) = 0 \}$.

- **Strict Order-Based Domains.** A subset of order-based domains for which we can prove strong statements is those defined only by strict inequalities $v_i(a) < v_i(b)$ (i.e. $R_i(a,b) \in \{>, <, \text{"}null''\}$), as well as at most a single equality of the form $v_i(a) = 0$. Examples of strict order-based domains are two-players combinatorial auctions, or two-player multi-unit auctions, where all items must be allocated, i.e. $a_1 \cup a_2 = \Omega$ (this is discussed in details in section 3.5). Trivially, unrestricted domains are also strict order based.

We assume that players' valuations are private information. Thus, a player might be motivated to declare a different type than his true type, in order to shift the social choice in some direction desirable for him. The solution we consider is to construct a *mechanism*, which is allowed to charge payments ($p_i : V \rightarrow \mathcal{R}$) from the players, in addition to producing the chosen alternative. We assume that players are quasi-linear and rational in the sense of maximizing their total utility: $u_i = v_i(f(v)) - p_i(v)$. *Truthful* mechanisms are direct revelation mechanisms in which the dominant strategy of a player is to declare his true type, v_i, rather than a different type, u_i. A mechanism *truthfully implements* a social choice function if it is truthful, and its outcome function coincides with the social choice function. The only known general class of truthfully implementable social choice functions over multi-dimensional domains are *affine maximizers*, which can be implemented using a simple generalization of VCG payments:

Definition 3.1 (Affine maximization) *A social choice function f is an* affine maximizer *if there exist constants $\omega_1, \ldots, \omega_n \geq 0$ and $\{\gamma_a\}_{a \in A}$ such that for any $v \in V$:*

$$f(v) \in argmax_{a \in A}\{\sum_{i=1}^{n} \omega_i v_i(a) + \gamma_a\}.$$

It can be verified that, in this case, f is implemented by the payments $p_i = -\omega_i^{-1}(\sum_{j \neq i}^{n} \omega_j v_j(a) + \gamma_a)$.

3.3 Truthfulness and Monotonicity

It is well known that truthfulness is related to some notions of monotonicity. In this section we derive these relationships which serve as the embarking point towards our main characterization.

3.3.1 Weak monotonicity

In simple "one parameter" domains, monotonicity is usually the property of "still winning when raising my value". In general domains, we must examine value *differences*. Roberts [87] used a definition of monotonicity called PAD: f satisfies PAD if for every $v, u \in V$, $f(v) = a$ and $u_i(a) - v_i(a) > u_i(b) - v_i(b)$ for all $i = 1, \ldots, n$ and all $b \in A$ implies that $f(u) = a$. However, PAD has no real meaning for most restricted domains: Suppose there exists a player i and two alternatives a, b s.t. $v_i(a) = v_i(b)$ for all $v_i \in V_i$ (e.g. in CA, when i gets the same bundle in a and b). Then the condition of PAD is never satisfied if $f(v) = a$. One can make several attempts to "fix" this. Below we describe several natural "candidates" for a more general monotonicity condition, and demonstrate that they fail to be necessary for truthfulness. We first identify the "correct" notion of monotonicity:

Definition 3.2 (Weak Monotonicity (W-MON)) *A social choice function f satisfies W-MON if for any $v \in V$, player i, and $u_i \in V_i$: $f(v) = a$ and $f(u_i, v_{-i}) = b$ implies that $u_i(b) - v_i(b) \geq u_i(a) - v_i(a)$.*

In other words, if player i caused the outcome of f to change from a to b by changing his valuation from v_i to u_i, then it must be that i's value for b has increased at least as i's value for a. W-MON implies PAD on every domain but makes sense also in domains where PAD does not.

Claim 3.1 *If f satisfies W-MON then f satisfies PAD.*

proof: Fix any $v, u \in V$. Suppose $f(v) = a$, and $u_i(a) - v_i(a) > u_i(b) - v_i(b)$ for all $i = 1, \ldots, n$ and $b \in A$. Let $v^0 = v,\ v^1 = (u_1, v_2, \ldots, v_n),\ v^2 = (u_1, u_2, v_3, \ldots, v_n),\ v^n = (u_1, \ldots, u_n) = u$. Now, $f(v^0) = a$ and $f(v^i) = a$ implies by W-MON that $f(v^{i+1}) = a$. ∎

For restricted domains, it turns out that W-MON is crucially important, as it is essentially equivalent to truthfulness:

Theorem 3.1 *Every implementable social choice function in any domain satisfies W-MON. If V is an order based domain then W-MON is also a sufficient condition for truthfulness.*

A proof of this theorem is given in subsection 3.3.2 below.

The condition that the domain is order-based is needed (although it may be relaxed) to ensure that W-MON is a sufficient condition. The following example, inspired by [91], shows that W-MON by itself is not a sufficient condition for truthfulness.

Example 3.1 *Consider a single player with $A = \{a, b, c\}$ and a domain of three possible types v_a, v_b, v_c, as follows: $v_a = (0, 1, -2)$; $v_b = (-2, 0, 1)$; $v_c = (1, -2, 0)$, where the first coordinate in each type is a's value, the second is b's value, and the third c's value.*

The function f has $f(v_x) = x$, for every $x \in A$. f satisfies W-MON since $v_x(x) - v_y(x) > v_x(y) - v_y(y)$ for any $x, y \in A$.

Suppose by contradiction that there are truthful prices. Therefore: $-1 = v_c(c) - v_c(a) \geq p(c) - p(a)$. Similarly, $-1 = v_a(a) - v_a(b) \geq p(a) - p(b)$, and $-1 = v_b(b) - v_b(c) \geq p(b) - p(c)$. But the last two inequalities imply $p(c) - p(a) \geq 2$, a contradiction.

We next describe two natural "candidates" for a more general monotonicity condition, and show by an example that they fail to be necessary for truthfulness.

Strong PAD: For every $v, u \in V$, where $f(v) = a$), if for all $i = 1, \ldots, n$ and $b \in A$: $u_i(a) - v_i(a) \geq u_i(b) - v_i(b)$ then $f(u) = a$.

Generalized W-MON: For every $v, u \in V$, if $f(v) = a$ and $f(u) = b$ then there exists a player i such that: $u_i(b) - v_i(b) \geq u_i(a) - v_i(a)$.

To contradict both types of monotonicity, consider the following example:

Example 3.2 *Suppose there are two players, and four alternatives: $A = \{YY,\ YN,\ NY,\ NN\}$. A player type is determined by one positive value v_i, as follows. For any $a \in A$ (denote $a = a_1a_2$ where $a_i \in \{Y, N\}$): if $a_i = N$ then $v_i(a) = 0$, and if $a_i = Y$ then $v_i(a) = v_i$. Define $f(v) = a_1a_2$, where $a_i = Y$ if $v_i > 2v_j - 10$, otherwise $a_i = N$. It is easy to verify that f is truthful, with the payments $p_i(N, v_j) = 0$ and $p_i(Y, v_j) = 2v_j - 10$.*

Now, suppose $v_1 = v_2 = 9$, and $u_1 = u_2 = 11$. Then $f(v) = YY$, but $f(u) = NN$! [11]

3.3.2 W-MON characterizes truthfulness

In this subsection we prove theorem 3.1. To show the first direction we start with a basic known claim, that essentially states that the prices of player i do not depend on i's type, and that f always chooses an alternative that maximizes i's utility, under these prices (for completeness, we provide the proof in appendix 3.7.1):

Claim 3.2 *Any truthful function f has (price) functions $p_i : A \times V_{-i} \to \mathcal{R} \cup \{\infty\}$ such that, for any $v \in V$ and any player i, $f(v) \in argmax_{a \in A}\{v_i(a) - p_i(a, v_{-i})\}$.*

Lemma 3.1 *Every truthful social choice function satisfies W-MON.*

proof: Let $p_i : A \times V_{-i} \to \mathcal{R} \cup \{\infty\}$ be the price functions according to claim 3.2. Suppose that $f(v) = a$ and $f(u_i, v_{-i}) = b$. Therefore $v_i(a) - p_i(a, v_{-i}) \geq v_i(b) - p_i(b, v_{-i})$, and $u_i(b) - p_i(b, v_{-i}) \geq u_i(a) - p_i(a, v_{-i})$. Thus $u_i(b) - u_i(a) \geq p_i(b, v_{-i}) - p_i(a, v_{-i}) \geq v_i(b) - v_i(a)$, and the claim follows. ∎

For the second direction of the theorem, we assume that V is ordered based, and use the following definitions. Fix any player i. For any $a, b \in A$, let $E_i(a) = \{d \in A \mid R_i(a, d) = \text{“} = \text{”} \text{ or } d = a\}$, and define:

$$\delta_{ab}(v_{-i}) = \inf\ \{v_i(a) - v_i(b) \mid v_i \in V_i \text{ and } f(v_i, v_{-i}) \in E_i(a)\}.$$

Claim 3.3 *For any $a, b, c \in A$, and $v_{-i} \in V_{-i}$:*

1. *W-MON implies that $\delta_{ab}(v_{-i}) \geq -\delta_{ba}(v_{-i})$.*

[11] Note that PAD trivially holds – its condition is never satisfied. This example is also “far from affine maximization”, and can be extended to more than two players.

2. *If* $R_i(a,b) \in \{=,\leq\}$ *then* $\delta_{cb}(v_{-i}) \leq \delta_{ca}(v_{-i})$.

proof: Suppose by contradiction that $\delta_{ab}(v_{-i}) < -\delta_{ba}(v_{-i})$. Take $v_i \in V_i$ such that $v_i(a) - v_i(b) = \delta_{ab}(v_{-i}) + \epsilon$ and $f(v) = \tilde{a}$, where $\tilde{a} \in E_i(a)$, and $u_i \in V_i$ such that $u_i(b) - u_i(a) = \delta_{ba}(v_{-i}) + \epsilon$ and $f(u_i, v_{-i}) = \tilde{b}(\tilde{b} \in E_i(b))$. Since $R_i(a,\tilde{a}) = R_i(b,\tilde{b}) =$ " $=$ " it follows that $v_i(\tilde{a}) - v_i(\tilde{b}) < u_i(\tilde{a}) - u_i(\tilde{b})$. But by W-MON, since $f(v) = \tilde{a}$ it follows that $f(u_i, v_{-i}) \neq \tilde{b}$, a contradiction.

For the second part, assume by contradiction that $\delta_{cb}(v_{-i}) > \delta_{ca}(v_{-i})$, and choose some v_i such that $v_i(c) - v_i(a) = \delta_{ca}(v_{-i}) + \epsilon < \delta_{cb}(v_{-i})$ and $f(v) \in E_i(c)$. Since $v_i(a) \leq v_i(b)$ it follows that $v_i(c) - v_i(b) \leq v_i(c) - v_i(a) < \delta_{cb}(v_{-i})$, contradicting the definition of δ_{cb}. ∎

We now describe a price function $p_i : A \times V_{-i} \to \mathcal{R}$ that induces truthfulness, for all $v \in V$: $f(v) \in argmax_{a \in A}\{v_i(a) - p_i(a, v_{-i})\}$. For this, fix some alternative $c \in A$ such that for any other $d \in A$, $R_i(c,d) \notin \{\leq, <\}$ (there always exists such alternative since the R_i relations depict partial order over A) [12], and set:

$$p_i(a, v_{-i}) = \begin{cases} 0 & a \in E_i(c) \\ -\delta_{ca}(v_{-i}) & \text{otherwise} \end{cases} \quad (3.1)$$

Claim 3.4 *For any* $a \in A$, $\tilde{c} \in E_i(c)$, *and* $v \in V$:

1. *If* $v_i(a) - p_i(a, v_{-i}) < v_i(\tilde{c}) - p_i(\tilde{c}, v_{-i})$ *then* $f(v) \neq a$.

2. *If* $v_i(a) - p_i(a, v_{-i}) > v_i(\tilde{c}) - p_i(\tilde{c}, v_{-i})$ *then* $f(v) \neq \tilde{c}$.

proof: By definition, $p_i(\tilde{c}, v_{-i}) = 0$ and $v_i(\tilde{c}) = v_i(c)$. First suppose that $v_i(a) - p_i(a, v_{-i}) < v_i(c)$. By definition and by claim 3.3, $v_i(a) - v_i(c) < -\delta_{ca}(v_{-i}) \leq \delta_{ac}(v_{-i})$, and therefore $f(v) \neq a$. In the other direction, $v_i(c) - v_i(a) < -p_i(a, v_{-i}) = \delta_{ca}(v_{-i})$, and therefore $f(v) \neq \tilde{c}$. ∎

We can now finish the proof.

Lemma 3.2 *If* V *is an order based domain then W-MON is a sufficient condition for truthfulness.*

proof: Suppose that f satisfies W-MON. We will show that the prices of equation 3.1 induce truth-telling. Suppose by contradiction that there exists $v \in V$ such that $f(v) = a$, but $v_i(a) - p_i(a, v_{-i}) < v_i(b) - p_i(b, v_{-i})$. By claim 3.4 it follows that $a, b \notin E_i(c)$, and that $v_i(c) - p_i(c, v_{-i}) \leq v_i(a) - p_i(a, v_{-i})$. Choose some small enough $\epsilon > 0$ and some δ such that $v_i(a) + \epsilon - p_i(a, v_{-i}) < v_i(c) + \delta - p_i(c, v_{-i}) < v_i(b) - p_i(b, v_{-i})$. Define $T_i = \{a\} \cup \{d \in A \mid v_i(d) = v_i(a) \text{ and } R_i(a,d) \in \{\leq, =\}\ \}$, and let $u_i = v_i|^{E_i(c) += \delta,\ T_i += \epsilon}$. Notice that $u_i \in V_i$ (we can raise $E_i(c)$ as we wish, and raise T_i by some small enough ϵ) [13].

[12] We also assume that $c \notin 0_i$. This is w.l.o.g since we can always "normalize" the domain with respect to any other alternative a, as follows: we convert any original type v_i to a new type $u_i = v_i - v_i(a) \cdot 1^m$. It is not hard to verify that this maintains truthfulness.

[13] It is possible that there exists some $\tilde{c} \in E_i(c) \cap T_i$. In this case we raise $\tilde{c}$ by δ

By claim 3.3, for any $d \in T_i$, $p_i(a, v_{-i}) \leq p_i(d, v_{-i})$, and therefore $v_i(d) - p_i(d, v_{-i}) \leq v_i(a) - p_i(a, v_{-i})$. From this we conclude that $b \notin T_i$, and also that for any $d \in T_i$, $u_i(d) - p_i(d, v_{-i}) < u_i(c) - p_i(c, v_{-i})$. Thus, by claim 3.4, $f(u_i, v_{-i}) \neq d$. Similarly, for any $\tilde{c} \in E_i(c)$, $u_i(\tilde{c}) - p_i(\tilde{c}, v_{-i}) < u_i(b) - p_i(b, v_{-i})$, and so $f(u_i, v_{-i}) \neq \tilde{c}$. But, by W-MON, since $f(v) = a$ it must be the case that $f(u_i, v_{-i}) \in E_i(c) \cup T_i$, a contradiction. ∎

3.3.3 Strong monotonicity and IIA

So far we have seen that weak monotonicity is almost equivalent to truthfulness. We identify the following slightly stronger monotonicity condition, where the inequality in the definition is strict, as being of particular importance. We require this stronger condition for our main result.

Definition 3.3 (Strong Monotonicity (S-MON)) *A social choice function f satisfies S-MON if for any $v \in V$, player i, and $u_i \in V_i$: $f(v) = a$ and $f(u_i, v_{-i}) = b \neq a$ imply that $u_i(b) - v_i(b) > u_i(a) - v_i(a)$.*

In both definitions, we have the situation that i's valuation changed from v_i to u_i and this caused the outcome of f to change from a to b. S-MON asserts that this implies that i's valuation of b had to increase more than did the valuation of a. W-MON only requires that it did not increase less. While this seems like a slight change, it is in fact crucial. S-MON is not a necessary condition for truthfulness – we give several counter examples for this in section 3.5, in the context of Combinatorial Auctions. The following definition, inspired by Arrow's notion for non-quasi-linear environments [4], essentially characterizes the difference between W-MON and S-MON:

Definition 3.4 (Independence of Irrelevant Alternatives (IIA)) *f satisfies IIA if for any $v, u \in V$, if $f(v) = a$ and $f(u) = b \neq a$ then there exists a player i such that $u_i(a) - u_i(b) \neq v_i(a) - v_i(b)$.*

In other words, if the social choice function on some valuations clearly prefers a over b, as a is chosen, and no player changes his preference of a with respect to b, then it cannot be the case that the social choice function would now choose b. For example, imagine some setting of a combinatorial auction, and an initial valuation declaration that causes some allocation to be chosen. Suppose now that player 1 raises his value for the bundle that contains all items, and that nothing else is changed. Then, a combinatorial auction that satisfies IIA would have to now choose either the previous allocation, or the allocation that hands in all items to player 1. Any other allocation violates IIA.

We would like to explicitly state the connection between W-MON, S-MON, and IIA. As we will show, W-MON plus IIA always implies S-MON. The other direction is not always true – the following example demonstrates that S-MON does not always imply IIA:

Example 3.3 *Suppose there are four alternatives ($A = \{a, b, c, d\}$) and two players, each one with two possible types v_i, u_i such that: $u_1(c) - v_1(c) > u_1(a) - v_1(a) = u_1(b) - v_1(b) > u_1(d) - v_1(d)$, and $u_2(d) - v_2(d) > u_2(a) - v_2(a) = u_2(b) - v_2(b) > u_2(c) - v_2(c)$. Define f as follows: $f(v_1, v_2) = a$, $f(u_1, u_2) = b$, $f(u_1, v_2) = c$, and $f(v_1, u_2) = d$. It is not hard to verify that S-MON holds (there are four inequalities to check, all of them follow from the way the types are defined). IIA does not hold since $f(v) = a$, $f(u) = b$, but $u(a) - u(b) = v(a) - v(b)$.*

However, for order based domains, IIA exactly characterizes the difference between W-MON and S-MON:

Proposition 3.1 *If f satisfies W-MON and IIA then it satisfies S-MON. In the other direction, if V is order based and f satisfies S-MON then f satisfies W-MON and IIA.*

Remark: We actually show that, for order based domains, S-MON implies the following "generalized S-MON": $f(v) = a$ and $f(u) = b \Rightarrow \exists i : u_i(b) - u_i(a) > v_i(b) - v_i(a)$. This clearly implies IIA.

We prove the proposition using several claims:

Claim 3.5 *If f satisfies W-MON and IIA then f satisfies S-MON.*

proof: Fix any $v \in V$, player i, and $u_i \in V_i$. Suppose $f(v) = a$ and $f(u_i, v_{-i}) = b$. We need to show that $u_i(b) - v_i(b) > u_i(a) - v_i(a)$. By W-MON it follows that $u_i(b) - v_i(b) \geq u_i(a) - v_i(a)$. Suppose by contradiction that $u_i(b) - v_i(b) = u_i(a) - v_i(a)$. But then, denote $u = (u_i, v_{-i})$, and we have $f(v) = a$, $f(u) = b$, and for any player j, $v_j(a) - v_j(b) = u_j(a) - u_j(b)$, thus contradicting IIA. ∎

For the other direction, we first claim that we can assume w.l.o.g that V is not normalized, i.e. $0_i = \emptyset$ for all i:

Claim 3.6 *If V is normalized then there exists a non-normalized order based domain $\tilde{V}$ and a function $\tilde{f} : \tilde{V} \to A$ such that:*

1. *If f satisfies S-MON then $\tilde{f}$ satisfies S-MON as well.*
2. *$V \subseteq \tilde{V}$, and for any $v \in V$, $f(v) = \tilde{f}(v)$.*
3. *If $\tilde{f}$ satisfies IIA then f satisfies IIA as well.*

proof: Define $\tilde{V}$ as the order based domain defined by exactly the same relations $R_i(a, b)$ but with $0_i = \emptyset$, for all i. $V \subseteq \tilde{V}$ since for any $v \in V$, all the relations $R_i(a, b)$ hold, and therefore $v \in \tilde{V}$. Define $\tilde{f} : \tilde{V} \to A$ as follows: For every i, choose some $a^i \in 0_i$. For any $\tilde{v} \in \tilde{V}$, let $v_i = \tilde{v}_i - \tilde{v}_i(a^i)$, and define $\tilde{f}(\tilde{v}) = f(v)$ ($v \in V$ since all inequalities hold after a translation, and for any $b \in 0_i$, $\tilde{v}_i(b) - \tilde{v}_i(a) = 0$ since $R_i(a,b) = \text{``} = \text{''}$).

To see that $\tilde{f}$ satisfies S-MON, suppose $\tilde{f}(\tilde{v}) = a$ and $\tilde{f}(\tilde{u}_i, \tilde{v}_{-i}) = b$. Let $v_j = \tilde{v_j} - \tilde{v_j}(a^j)$ (for $j = 1, \ldots, n$), and $u_i = \tilde{u_i} - \tilde{u_i}(a^i)$. By definition, $f(v) = a$ and $f(u_i, v_{-i}) = b$. Since f satisfies S-MON, $u_i(b) - v_i(b) > u_i(a) - v_i(a)$. Therefore $\tilde{u}_i(b) - \tilde{v}_i(b) > \tilde{u}_i(a) - \tilde{v}_i(a)$, and thus $\tilde{f}$ satisfies S-MON.

Since $V \subseteq \tilde{V}$, contradicting IIA for f implies contradicting IIA for $\tilde{f}$, and the claim follows. ∎

Claim 3.7 (Dependence on Differences (DOD)) *Suppose V is order based and non normalized, and f satisfies S-MON. Then for any $v \in V$ and $\delta \in \mathcal{R}$: $v_i + \delta \cdot 1^m \in V_i$, and $f(v) = f(v_i + \delta \cdot 1^m, v_{-i})$.*

proof: $v_i + \delta \cdot 1^m \in V_i$ since all inequalities hold after a translation. Since $[v_i(b) + \delta] - v_i(b) = [v_i(a) + \delta] - v_i(a)$ for any $a, b \in A$, it follows from S-MON that $f(v) = f(v_i + \delta \cdot 1^m, v_{-i})$. ∎

Claim 3.8 (Generalized S-MON) *Suppose V is order based, and f satisfies S-MON. Then for any $u, v \in V$, if $f(v) = a$ and $f(u) = b$ then there exists a player i such that $u_i(b) - v_i(b) > u_i(a) - v_i(a)$.*

proof: By claim 3.6 we can assume w.l.o.g that V is not normalized: otherwise, move to $\tilde{f}$, and then, contradicting generalized S-MON for f implies contradicting generalized S-MON for $\tilde{f}$. By claim 3.7 we can assume w.l.o.g that $u_i(a) = v_i(a)$: otherwise let $\tilde{v}_i = v_i + [u_i(a) - v_i(a)] \cdot 1^m$, then $f(\tilde{v}) = a$, and finding i such that $u_i(b) - \tilde{v}_i(b) > u_i(a) - \tilde{v}_i(a) = 0$ implies that $u_i(b) - v_i(b) > u_i(a) - v_i(a)$.

Now, we "move" from v to u by L "elementary steps" $v = v^1, v^2, \ldots, v^L = u$, such that: (1) for any index j there exists a player i and $d \in A$ such that $v_i^{j+1} = v_i^j|^{d+=u_i(d)-v_i(d)}$, (2) every pair (i, d) appears only once in the sequence, and (3) there exists an index l^* such that for any $l \leq l^*$, $u_i(d) - v_i(d) < 0$, and for any $l > l^*$, $u_i(d) - v_i(d) > 0$ (since V is order based, we can construct such a sequence of types). By S-MON, $f(v^{l^*}) = a$, and for any $l > l^*$, if $f(v^l) = c$ then $f(v^{l+1}) \in \{c, d\}$ (where d is the alternative that changes from v^l to v^{l+1}). Therefore, if $f(v^L) = b$ it follows that there exists i such that $u_i(b) - v_i(b) > 0$, as claimed. ∎

Clearly, generalized S-MON implies IIA, and S-MON implies W-MON, hence the second direction of the proposition follows. ∎

3.3.4 Equivalence of W-MON and S-MON

For some domains, S-MON can be assumed without loss of generality for our main purpose of proving affine maximization. Intuitively, in such domains, the only possibility of having W-MON but violating S-MON is due to "tie-breaking" rules, which cannot harm the affine maximization property. The formal statement is:

Theorem 3.2 *If V is an open set* [14] *then for every $f : V \to A$ there exists $\tilde{f} : V \to A$ such that:*

1. *If f satisfies W-MON then $\tilde{f}$ satisfies S-MON.*

2. *If $\tilde{f}$ is affine maximizer then f is affine maximizer.*

By this theorem, proving that S-MON implies affine maximization exactly implies that W-MON implies affine maximization: Using the first step of the theorem we "generate" from f that satisfies W-MON an $\tilde{f}$ that satisfies S-MON. We then show that this $\tilde{f}$ is an affine maximizer using the main theorem. Finally, by the second step of theorem 3.2 we conclude that our original f is also an affine maximizer.

Proof of theorem 3.2: We use the notation $v + \epsilon 1_{i,b} = (v_i|^{b+=\epsilon}, v_{-i})$, and $v + \epsilon 1_b = v + \epsilon 1_{1b} + \ldots + \epsilon 1_{nb}$. For any $v \in V$, define:

$$T(v) \;=\; \{\; b \in A \mid \exists \epsilon^* > 0 \text{ s.t. } \forall \epsilon \in (0, \epsilon^*) : f(v + \epsilon 1_b) = b \;\}$$

Claim 3.9 *For any $v \in V$, i, and $u_i \in V_i$: if $a \in T(v)$, $b \in T(u_i, v_{-i})$, and $u_i(a) - v_i(a) \geq u_i(b) - v_i(b)$, then $a \in T(u_i, v_{-i})$.*

proof: For any (small enough) $\epsilon > 0$, since $a \in T(v)$, $f(v + \epsilon 1_a) = a$. By W-MON, it follows that:

$$f(v_i + \epsilon 1_{i,a},\; v_{-i} + 2\epsilon 1_{-i,b} + 4\epsilon 1_{-i,a}) = a \tag{3.2}$$

(this follows by changing the player types one at a time). Similarly, since $b \in T(u_i, v_{-i})$, we get that $f(u_i + \epsilon 1_{i,b},\; v_{-i} + \epsilon 1_{-i,b}) = b$. By W-MON (changing the player types one at a time):

$$f(u_i + 2\epsilon 1_{i,b} + 4\epsilon 1_{i,a},\; v_{-i} + 2\epsilon 1_{-i,b} + 4\epsilon 1_{-i,a}) \in \{a, b\} \tag{3.3}$$

Since $u_i(a) - v_i(a) \geq u_i(b) - v_i(b)$ it follows that $[u_i(a) + 4\epsilon] - [v_i(a) + \epsilon] > [u_i(b) + 2\epsilon] - v_i(b)$. Therefore, comparing Eq. 3.3 to Eq. 3.2, and by W-MON, we conclude that $f(u_i + 2\epsilon 1_{i,b} + 4\epsilon 1_{i,a},\; v_{-i} + 2\epsilon 1_{-i,b} + 4\epsilon 1_{-i,a}) = a$. Thus also $f(u_i + 5\epsilon 1_{i,a},\; v_{-i} + 5\epsilon 1_{-i,a}) = a$, hence $a \in T(u_i, v_{-i})$, and the claim follows. ∎

We can now define $\tilde{f}$. Fix any complete order $\succ$ on A, and then:

$$\tilde{f}(v) = \max_{\succ} T(v)$$

Claim 3.10 *$\tilde{f}$ satisfies S-MON.*

[14] V is open if for any $v \in V$ there exists $\epsilon_v > 0$ such that for any $u \in \mathcal{R}^{m \times n}$, if $|u_i(a) - v_i(a)| < \epsilon_v$ for all i, a then $u \in V$ as well.

proof: Suppose that $\tilde{f}(v) = a$ and $\tilde{f}(u_i, v_{-i}) = b$. Therefore $a \in T(v)$ and $b \in T(u_i, v_{-i})$. Assume by contradiction that $u_i(a) - v_i(a) \geq u_i(b) - v_i(b)$. By claim 3.9 it follows that $a \in T(u_i, v_{-i})$, and thus $b \succ a$. On the other hand, it is also the case that $v_i(b) - u_i(b) \geq v_i(a) - u_i(a)$, and so, by claim 3.9 again (changing variable names), we get that $b \in T(v)$ and therefore $a \succ b$, a contradiction. ∎

Claim 3.11 *If $\tilde{f}$ is an affine maximizer, then f is an affine maximizer as well.*

proof: Assume that for any $v \in V$, $\tilde{f}(v) \in argmax_{a \in A}\{\sum_i \omega_i v_i(a) + \gamma_a\}$, and suppose that $\tilde{f}(v) = a$ but $f(v) = b$. We first claim that for any (small enough) $\epsilon > 0$, $\tilde{f}(v + \epsilon 1_b) = b$: otherwise, suppose it equals c. By definition, this implies that $f(v + \epsilon 1_b + (\epsilon/2) 1_c) = c$, contradicting PAD (claim 3.1), since $f(v) = b$ and b was raised strictly more than all other alternatives for all players. Since $\tilde{f}$ is affine maximizer it follows that $\sum_i \omega_i[v_i(b) + \epsilon] + \gamma_b \geq \sum_i \omega_i[v_i(a) + \epsilon] + \gamma_a$. This is true for any (small enough) $\epsilon > 0$, so it follows that $f(v)$ chooses a maximal alternative as well, as claimed. ∎

This concludes the proof of theorem 3.2. ∎

Since an unrestricted domain is an open set, this theorem immediately applies to it. The theorem also applies to strict order based domains:

Corollary 3.1 *If V is strict order based then for every $f : V \to A$ there exists $\tilde{f} : V \to A$ such that, if f satisfies W-MON then $\tilde{f}$ satisfies S-MON, and then, if $\tilde{f}$ is an affine maximizer, then f is an affine maximizer as well.*

proof: If V is not normalized (i.e. $0_i = \emptyset$ for all i then it is an open set, by definition, and the corollary immediately follows. Otherwise, we expand V to a non normalized $\hat{V}$, exactly as in claim 3.6. Then $\hat{f}$ also satisfies W-MON, and if $\hat{f}$ is affine maximizer then f is affine maximizer as well. Since $\hat{V}$ an open set, there exists $\tilde{f}$ that satisfies S-MON, and if $\tilde{f}$ is affine maximizer then $\hat{f}$ is affine maximizer, which in turn implies that f is affine maximizer as needed. ∎

In order to use all this for our main theorem, we have to verify that all the translations from f to $\tilde{f}$ also preserve the other requirements of the theorem. It is not hard to verify that the player decisiveness and the non-degeneracy conditions are indeed preserved. As for the "conflicting preferences" requirement, the removal of the normalization in the translation from f to $\tilde{f}$ does not harm it, since the structure of "top" and "bottom" alternatives is not affected.

3.4 Main Theorem

Our main theorem shows that, under certain conditions, social choice functions that satisfy S-MON are "almost" affine maximizers. Let us first explain these conditions and qualifications:

- **The Domain:** The theorem holds for a family of restricted domains which we call *order-based domains with conflicting preferences* – These are essentially order based domains in which the most preferred alternative of player i is the least preferred alternative of all other players:

  Definition 3.5 (top and bottom alternatives of player i) *Suppose V_i is order based. The alternative $a \in A \setminus 0_i$ is a top alternative if its value is never smaller than the value of any other alternative. I.e. if for all other $b \in A$, $R_i(a,b) \in \{ > , \geq , null \}$. Similarly, the alternative $a \in A$ is a bottom alternative if for all other $b \in A$, $R_i(a,b) \notin \{ > , \geq \}$.*

  Definition 3.6 (Conflicting preferences) *An order based domain has conflicting preferences if:*

  1. *Any player i has at least one top alternative (denoted c^i).*
  2. *For all i and $j \neq i$, c^j is a bottom alternative for player i, and $c^j \in 0_i$* [15].

  Note that $c^j \neq c^i$ for all $i \neq j$ as $c^j \notin 0_j$ and $c^i \in 0_j$. Combinatorial Auctions and Multi-Unit Auctions have conflicting preferences: the allocation of all the goods to player i is a top alternative for i, and is indeed a bottom alternative (with a value of zero) for all other players. Matching, however, does not have conflicting preferences, since there is no top alternative – every alternative is coupled with many other alternatives (all the ones that match i to the same person).

- **The Range:** The actual range of the social choice function must be *non-degenerate*:

  Definition 3.7 (Non-degenerate range) *A is non-degenerate if for any player $i > 1$ there exists $a \in A$ such that $a \notin 0_1$ and $a \notin 0_i$.*

  For combinatorial auctions or multi-unit auctions this means that there exists some player (w.l.o.g player 1) such that, for every other player i, the range includes an allocation a with $a_1 \neq \emptyset$ and $a_i \neq \emptyset$. Without this condition, the problem may essentially be reduced to a single-dimensional setting (e.g. when the range contains only the allocations that allocate all items to one player), in which case many truthful non affine maximizers exist.

- **The Social Choice Function:** We require *player decisiveness*. This means that a player can ensure that his top alternative is chosen if he bids high enough on it:

  Definition 3.8 (Player decisiveness) *f is player-decisive if for any $v \in V$ and any player i there exist $u_i = v_i|^{c^i += \delta}$ for some $\delta > 0$ such that $f(u_i, v_{-i}) = c^i$.*

[15] This normalization is for convenience. We can instead just assume that for any i, j, l, $R_i(c^j, c^l)$ is "=", and use S-MON to normalize the domain.

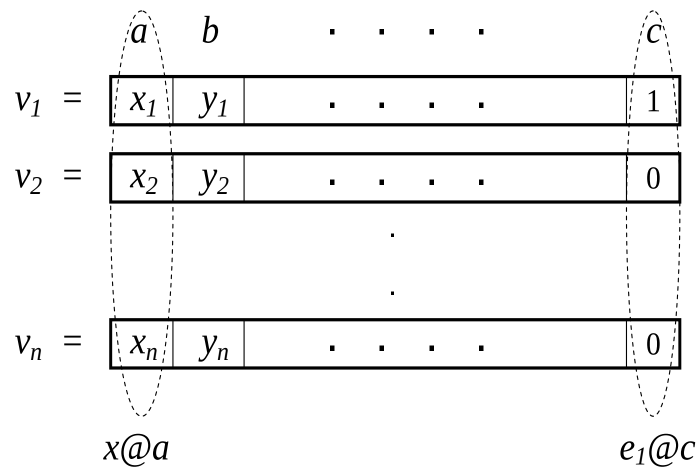

Figure 3.1: The structure of the valuation vector, and the notion $x@a$.

For CAs and MUAs, this means that a player can always receive all goods if he bids high enough on them. We note the difference between this requirement and the decisiveness requirement of [70], where it is required that some player will be able to cause *any* alternative to be chosen, when declaring appropriately. For CAs, this is very strong – for example, it requires that player 1 will be able to decide whether player 2 or player 3 will receive all goods.

- **Almost Affine Maximizer:** The theorem only shows that the social choice function must be an affine maximizer for large enough input valuations. I.e. there exists a threshold M s.t. the function is an affine maximizer if $v_i(a) \geq M$ for all a and i (except from inherently zero alternatives). We believe that this restriction is a technical artifact of the current proof, although we were not able to remove it.

Theorem 3.3 *Every social choice function over an order-based domain with conflicting preferences and onto a non-degenerate range, that is player decisive and satisfies S-MON, must be an almost affine maximizer.*

3.4.1 Intuitive proof outline

We now provide an intuitive outline of the proof. Full details appear below. It will be first useful to visualize the valuation vector v as a matrix, like the one in Figure 3.1. The i'th row contains the valuation vector of player i, and each column represents an alternative. Thus, i's value for alternative a, $v_i(a)$, is the first number in the i'th row. In the proof we extensive use the notation $x@a$ (x at a), which simply denotes the fact that $x = v(a) = (v_1(a), ..., v_n(a))$.

Our first step in the proof is to infer some order that f induces on the domain. Specifically, if for some vector v of valuations the choice is $a = f(v)$ then we may say that the vector of values

$v(a) = (v_1(a), \ldots, v_n(a))$ has more weight than the vector $v(b)$. This leads us to the following definition:

Definition 3.9 ("x at a" is larger than "y at b") *For $a, b \in A$ and $x, y \in \mathcal{R}^n$ we say that $x@a > y@b$ if there exists $v \in V$ such that $v(a) = x$, $v(b) = y$, and $f(v) = a$.*

This notation certainly suggests that ">" is an order. In unrestricted domains this is indeed the case. However, in restricted domains, it is not generally so. The requirements of the theorem imply "just enough" of the properties of an order to proceed with the proof.

Once such a "near-order" is defined, we can compare every $x@a$ to multiples of some fixed reference $z@c$. This is inspired by the "min-function" model of Archer and Tardos [3]. We would expect that for small values of α we would have $x@a > (\alpha z)@c$, while for large values of α we would have $x@a < (\alpha z)@c$. The value of α where the change happens somehow summarizes the "size" of $x@a$. To proceed we need to find such c and z where this holds for "enough" x and a. From now on, let such appropriate c and z be fixed.

Definition 3.10 *The "measure of x at a" is defined as:*

$$m(x@a) = inf \ \{ \ \alpha \mid x@a < (\alpha \cdot z)@c \ \}$$

This measure captures the choice function, as the following property shows:

Claim: *Under the conditions of the theorem, if $m(v(a)@a) < m(v(b)@b)$ then $f(v) \neq a$.*

This claim basically shows that $f(v) \in \text{argmax}_a\{m(v(a)@a)\}$. What remains to show is that $m(x@a)$ is in fact an affine function (in x) on R^n. (And, that it does not depend on a, up to an additive constant.) To get this result let us, informally, consider the partial derivative $\partial m(x@a)/\partial x_i$. A key observation is that this partial derivative must be equal to $\partial m(y@b)/\partial y_i$ for any other "compatible" y and b. Let us see the intuition for this: consider some v such that $v(a) = x$ and $v(b) = y$. Since the S-MON requirement only looks at differences $v_i(a) - v_i(b)$ when "choosing between a and b", we would expect that adding a constant δ to both $x_i = v_i(a)$ and to $y_i = v_i(b)$ will also leave $m(x@a) - m(y@b)$ unchanged. This is indeed the case:

Claim: *Under the conditions of the theorem, for all (appropriate) a, b, x, y and δ we have that $m(x@a) - m(y@b) = m((x + \delta \cdot e_i)@a) - m((y + \delta \cdot e_i)@b)$.* [16].

But now we claim that this means that $\partial m(x@a)/\partial x_i$ is independent of x. To see this, fix some y and denote $h_i(\delta) = m((y + \delta \cdot e_i)@b) - m(y@b)$. The previous claim states that $m((x + \delta \cdot e_i)@a) - m(x@a) = h_i(\delta)$. It is a simple exercise to verify that such a condition on $m(x@a)$ implies that it is

[16] Here e_i is the i'th unit vector

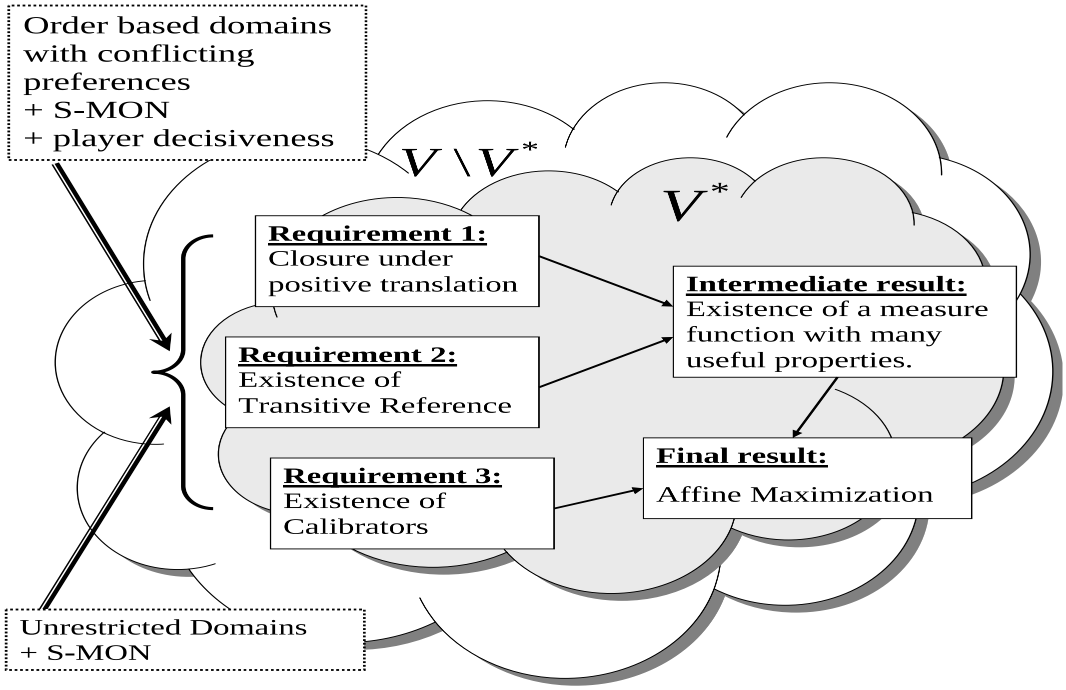

Figure 3.2: The structure of the full proof

linear in x. Specifically it must have the form $m(x@a) = \sum_i h_i(1) \cdot x_i + \gamma_a$, where γ_a is an arbitrary constant. This is (almost) the required result. (Delicate difficulties enter when we are unable to choose a single $y@b$ in an appropriate way for all $x@a$ – those are treated in the full proof.)

3.4.2 Proof of Theorem 3.3

We describe the proof in a somewhat abstract way: we define several requirements for f, V, and A, and show in parallel that: (1) these requirements are satisfied for order based domains with conflicting preferences, and, (2) under these requirements, f must be an almost affine maximizer. This abstraction will enable us to show that the proof, without most of the qualifiers, holds for unrestricted domains (this is described in section 3.6 below).

As we cannot measure $x@a$ if there is no $v \in V$ such that $f(v) = a$ and $v(a) = x$, we refer to the abstract sets $V^a \subseteq \{x \mid \exists v \in V \text{ such that } v(a) = x\}$, for all $a \in A$, and $V^* = \{\ v \in V \mid \forall a : v(a) \in V^a\ \}$. The proof will show that f is affine maximizer for any $v \in V^*$.

The proof structure is demonstrated in figure 3.2: we define three abstract requirements from a sub-domain V^* of V, and show that they imply affine maximization over V^*. We also of-course show that these abstract requirements are satisfied in our case.

For order based domains with conflicting preferences, we define the V^a's (and by this V^*) as follows:

Definition of V^a, for all $a \in A$: Denote $T = \{c^1, c^2, \ldots, c^n\}$ and $S = \{a \in A \mid a_1 \notin 0_1 \text{ and } a \neq c^1\}$, then

1. $V^{c^1} = \{ v(c^1) \mid v \in V \}$.
2. For any $a \in S$: $V^a = \{ x \mid \exists v \in V \text{ s.t. } v(a) = x \text{ and } f(v) = a \}$.
3. For any $c^i, i > 1$: $V^{c^i} = \{ y \mid \exists a \in S,\ a \notin 0_i, \text{ and } \exists x \in V^a \text{ s.t. } y@c^i > x@a \}$.
4. For any $a \in 0_1 \setminus T$: $V^a = \{ x \mid \text{ for all } i \text{ s.t. } a \notin 0_i \text{ there exists } y \in V^{c^i} \text{ s.t. } x@a > y@c^i \}$.

We will show below that all the V^a's are non-empty.

Requirement 1: Closure under positive Translation. For any $a, b \in A$, player i such that $a, b \notin 0_i$, and $\delta > 0$:

1. $x@a > y@b$ implies $(x + \delta \cdot e_i)@a > (y + \delta \cdot e_i)@b$.
2. $x \in V^a$ implies $(x + \delta \cdot e_i) \in V^a$.

Proposition 3.2 *If V is order based with conflicting preferences, and f satisfies S-MON, then Requirement 1 is satisfied.*

proof: We prove this in several steps, as follows:

Definition 3.11 (Closure under minimum) *A domain V is "closed under minimum" if for any $v, v' \in V$: $min(v, v') \in V$, where the minimum is taken componentwise.*

Claim 3.12 *Any order based domain is closed under minimum.*

proof: Suppose $v_i, v_i' \in V_i$, and let $u_i = min(v_i, v_i')$. For any $a \in 0_i$, since $v_i(a) = v_i'(a) = 0$ then $u_i(a) = 0$. For any $a, b \in A$, if $R_i(a, b) \in \{ =, \text{ null} \}$ then trivially "$u_i(a)\ R_i(a, b)\ u_i(b)$". Suppose $R_i(a, b)$ is "$\geq$", and suppose $u_i(a) = v_i(a)$. Then $u_i(b) \leq v_i(b) \leq v_i(a) = u_i(a)$, as needed. The other cases are similar. ∎

Claim 3.13 *Suppose V is closed under minimum and f is strongly monotone. Then for any $x@a$, $y@b$, $\tilde{x} \geq x$, and $\tilde{y} \leq y$: $x@a > y@b \Rightarrow \neg(\tilde{x}@a < \tilde{y}@b)$.*

proof: Since $x@a > y@b$, $\exists v \in V$ such that $v(a) = x$, $v(b) = y$, and $f(v) = a$. Suppose by contradiction that $\tilde{x}@a < \tilde{y}@b$. Therefore $\exists v' \in V$ such that $v'(a) = \tilde{x}$, $v'(b) = \tilde{y}$, and $f(v') = b$. Let $u = min(v, v')$. By S-MON, since $u(a) = v(a)$ and $u \leq v$, $f(v) = a$ implies $f(u) = a$. And since $u(b) = v'(b)$ and $u \leq v'$, $f(v') = b$ implies $f(u) = b$, a contradiction. ∎

Claim 3.14 *Suppose V is order based with conflicting preferences, and f is strongly monotone. Then for any $a, b \in A$, if $x@a > y@b$ and $a, b \notin 0_i$ then $(x + \delta \cdot e_i)@a > (y + \delta \cdot e_i)@b$ for any $\delta > 0$.*

proof: Since $x@a > y@b$ there exists $v \in V$ with $v(a) = x$, $v(b) = y$, and $f(v) = a$. Define u_i as follows: $u_i(d) = v_i(d) + \delta$ for any $d \in A \setminus 0_i$, and $u_i(d) = v_i(d)$ for any $d \in 0_i$. Since V is order based, and the alternatives in 0_i are all bottom alternatives (this is a side-effect of the second condition of conflicting preferences), raising all non-zero coordinates by a constant does not violate any relation, and so $u_i \in V$. Since $u_i(a) - v_i(a) \geq u_i(d) - v_i(d)$ for any $d \neq a$ it follows from S-MON that $f(u_i, v_{-i}) = a$, so $(x + \delta \cdot e_i)@a > (y + \delta \cdot e_i)@b$. ∎

Claim 3.15 *For any $a \in A$, V^a is non-empty, and $x \in V^a \Rightarrow (x + \delta \cdot e_i) \in V^a$, for all i such that $a \notin 0_i$, and for any $\delta > 0$.*

proof: For any $a \in S$ take any $v \in V$ such that $f(v) = a$, thus $v(a) \in V^a$, and, similarly to claim 3.14, $(v(a) + \delta \cdot e_i) \in V^a$, whenever $\delta > 0$ and $a \notin 0_i$.

For any c^i, choose some $a \in S \setminus 0_i$ (it is not empty since A is non-degenerate) and $x \in V^a$. By player decisiveness there exists some $y@c^i$ such that $y@c^i > x@a$, and thus $y \in V^{c^i}$. By S-MON and since c^i is a top alternative, $(y + \delta \cdot e_i)@c^i > x@a$, and thus $(y + \delta \cdot e_i) \in V^{c^i}$ as well.

For any $a \in 0_1 \setminus T$, take any $v \in V$ such that $f(v) = a$, thus $x@a > y@c^i$ (where $y = v(c^i)$). Notice that, since V is normalized, $y = y_i \cdot e_i$. Therefore there exists $\delta \geq 0$ such that $(y + \delta \cdot e_i) \in V^{c^i}$, and, by claim 3.14, $(x + \delta \cdot e_i)@a > (y + \delta \cdot e_i)@c^i$, thus $(x + \delta \cdot e_i) \in V^a$. If $x \in V^a$ then $x@a > y@c^i$ and, similarly to claim 3.14, $(x + \delta \cdot e_j)@a > y@c^i$ for any distinct i, j ($a \notin 0_i, 0_j$). ∎

We can therefore conclude that Requirement 1, closure under positive translation, follows from claims 3.14 and 3.15. ∎

We now continue to Requirement 2. An affine maximizer attaches a measure to every alternative (its weighted welfare), and chooses the one with the highest measure. Essentially, we show that every social choice function with the characteristics of the theorem must do the same. It has an underlying measure, it always chooses the alternative with the highest measure, and this measure is an affine function. In order to actually calculate the measure of alternative a, for some type v, we use a *reference* alternative, c. Intuitively, we start with some type profile where $x@a$ is chosen, and raise c's welfare until c is chosen. If we do this "slowly", then at some point the measure of the two alternatives will become equal, and thus, by knowing the measure of c, we obtain the measure of a. For this to succeed, we need an alternative with several technical properties:

Definition 3.12 *(Comparable) $x@a$ and $y@b$ are comparable if either $x@a > y@b$ or $y@b > x@a$.*

Requirement 2: A Transitive Reference $z@c$ for the set $M_c \subseteq A \setminus \{c\}$.

$z@c$ (where $c \in A$ and $z \in \mathcal{R}_+^n, z \neq 0$) is termed a transitive reference for the set $M_c \subseteq A \setminus \{c\}$ if, for any $a, b \in M_c$, $x \in V^a$, and $y \in V^b$, we have:

1. Measurability: There exist $\alpha, \beta \in \mathcal{R}$, $\alpha < \beta$, such that $(\alpha \cdot z)@c < x@a < (\beta \cdot z)@c$, and for any $\alpha' \in (\alpha, \beta)$: $x@a$ and $(\alpha' \cdot z)@c$ are comparable.

2. (semi) Transitivity: If $x@a < (\alpha \cdot z)@c$ and $\neg(y@b < (\alpha \cdot z)@c)$ then $\neg(x@a > y@b)$.

3. R-monotonicity: $x@a > (\alpha \cdot z)@c \;\Rightarrow\; \neg(x@a < (\beta \cdot z)@c)$ for any $\beta \leq \alpha$.

4. L-monotonicity: For $\delta > 0$, $(x + \delta \cdot e_i)@a < (\alpha \cdot z)@c \;\Rightarrow\; x@a < (\alpha \cdot z)@c$.

Proposition 3.3 *If V is order based with conflicting preferences, and f satisfies S-MON and player decisiveness, then Requirement 2 is satisfied by taking $e_1@c^1$ to be a transitive reference for $A \backslash \{c^1\}$.*

proof: We show that $e_1@c^1$ is a transitive reference by several claims:

Claim 3.16 *If $x@a < (\alpha \cdot e_1)@c^1$ and $\neg(y@b < (\alpha \cdot e_1)@c^1)$ then $\neg(x@a > y@b)$.*

proof: Suppose by contradiction that $x@a > y@b$, and take $v \in V$ such that $v(a) = x$, $v(b) = y$, and $f(v) = a$. Since V is normalized, $v(c^1) = \beta \cdot e_1$. Thus $x@a > (\beta \cdot e_1)@c^1$ and by claim 3.13 it follows that $\beta \leq \alpha$. Denote $u_1 = v_1|^{c^1=\alpha}$. By S-MON, $f(u_1, v_{-1}) \in \{a, c^1\}$. If it is a then $x@a > (\alpha \cdot e_1)@c^1$, contradicting the assumption $x@a < (\alpha \cdot e_1)@c^1$ (by claim 3.13). Thus it is c^1, and therefore $y@b < (\alpha \cdot e_1)@c^1$, a contradiction. ∎

Claim 3.17 *$\forall a \in A$ and $x \in V^a$, if $(x + \delta \cdot e_i)@a < (\alpha \cdot e_1)@c^1$, for some $\delta > 0$, then $x@a < (\alpha \cdot e_1)@c^1$.*

proof: Choose some v such that $v(a) = (x + \delta \cdot e_i)$, $v(c^1) = \alpha \cdot e_1$, and $f(v) = c^1$, and v' such that $v'(a) = x$ (there is one since $x \in V^a$). Let $u = min(v, v')$ (thus $u(a) = x$ and $u_1(c^1) \leq \alpha$) and $u' = u|^{c^1=\alpha \cdot e_1}$. Thus $u'(a) = x$, $u'(c^1) = \alpha \cdot e_1$, and $u' \leq v$. By S-MON, $f(u') = c^1$, as needed. ∎

Definition 3.13 (Compatible) *$x@a$ and $y@b$ are "compatible" if there exists $v \in V$ such that $v(a) = x$ and $v(b) = y$.*

Claim 3.18 *For any $a \in A$, $c^i \in T$, $x \in V^a$, and $(\alpha \cdot e_i)@c^i$, if $x@a$ and $(\alpha \cdot e_i)@c^i$ are compatible then they are comparable.*

proof: Since $x \in V^a$ there exists $v \in V$ such that $v(a) = x$ and $f(v) = a$. Since $x@a$ and $(\alpha \cdot e_i)@c^i$ are compatible then there exists $v' \in V$ such that $v'(a) = x$ and $v'(c^i) = \alpha \cdot e_i$. Let $u = min(v, v')$. Thus $f(u) = a$ and $u_i(c^i) \leq \alpha$. Therefore $f(u|^{c^i=\alpha \cdot e_i}) \in \{a, c^i\}$, and the claim follows. ∎

We can now conclude that $e_1@c^1$ is a transitive reference for $A \setminus \{c^1\}$: Transitivity is by claim 3.16, R-monotonicity follows from claim 3.13, and L-monotonicity is by claim 3.17. Measurability follows since, by the definition of the V^a's, there exists $(\alpha \cdot z)@c$ such that $x@a > (\alpha \cdot z)@c$, by player decisiveness there exists $(\beta \cdot z)@c$ such that $x@a < (\beta \cdot z)@c$, and by claim 3.18, since for any $\alpha' \in (\alpha, \beta)$, $x@a$ and $(\alpha' \cdot e_1)@c^1$ are compatible, they are comparable. ∎

Using Requirements 1 and 2, we can measure $x@a$ in a way that has several important properties:

Definition: *The "measure of x at a", for $a \in M_c$ and $x \in V^a$, is defined as:*

$$m(x@a) = inf \ \{ \ \alpha \mid x@a < (\alpha \cdot z)@c \ \}$$

and also: $m((\alpha \cdot z)@c) = \alpha$.

The following claims specify some important properties of the measure function.

Claim 3.19 *For any $a, b \in M_c$, $x \in V^a$, and $y \in V^b$: $m(x@a) = sup \ \{ \ \alpha \mid x@a > (\alpha \cdot z)@c \ \}$, and, as a conclusion, $-\infty < m(x@a) < \infty$.*

proof: By measurability, there exist $\alpha, \beta \in \mathcal{R}$ such that $(\alpha \cdot z)@c < x@a < (\beta \cdot z)@c$. Hence the infimum is at most β, and the supremum is at least α. By R-monotonicity, the infimum is not smaller than the supremum. Therefore they both reside in the interval $[\alpha, \beta]$. Since for any $\alpha' \in [\alpha, \beta]$, $x@a$ and $(\alpha' \cdot z)@c$ are comparable, the claim follows. ■

Claim 3.20 *For any $a, b \in M_c$, $x \in V^a$, and $y \in V^b$: $m(x@a) < m(y@b) \Rightarrow \neg(x@a > y@b)$ (and thus $x@a > y@b \Rightarrow m(x@a) \geq m(y@b)$).*

proof: Choose some α, $m(x@a) \leq \alpha < m(y@b)$, such that $x@a < (\alpha \cdot z)@c$. Since $\neg(y@b < (\alpha \cdot z)@c)$ it follows, by Transitivity, that $\neg(x@a > y@b)$. ■

Claim 3.21 *For any $a \in M_c \setminus 0_i$, $x \in V^a$, and $\delta > 0$: $m((x + \delta \cdot e_i)@a) \geq m(x@a)$.*

proof: Suppose by contradiction that $m(x@a) > m((x+\delta \cdot e_i)@a)$, and choose some α, $m(x@a) > \alpha \geq m((x + \delta \cdot e_i)@a)$, such that $(x + \delta \cdot e_i)@a < (\alpha \cdot z)@c$. But since $\neg(x@a < (\alpha \cdot z)@c)$, this contradicts L-monotonicity. ■

Definition 3.14 (The "support" of $z@c$) $S_c = \{ \ a \in M_c \mid a \notin 0_i$ *for all i s.t.* $z_i > 0 \ \}$.

The following claim implies that, for the alternatives in the support of $z@c$, we can "inflate" the measure in a very simple way.

Claim 3.22 *For any $a \in S_c$, $x \in V^a$, and $\beta > 0$: $m((x + \beta \cdot z)@a) = m(x@a) + \beta$.*

proof: Suppose by contradiction that $m((x+\beta \cdot z)@a) > m(x@a)+\beta$. Choose some α, $m((x+\beta \cdot z)@a) - \beta > \alpha \geq m(x@a)$, such that $x@a < (\alpha \cdot z)@c$. By positive translation, $x@a < (\alpha \cdot z)@c \ \Rightarrow (x + \beta \cdot z)@a < (\alpha + \beta) \cdot z@c$. But since $\alpha + \beta < m((x + \beta \cdot z)@a)$, we have a contradiction.

Similarly, suppose by contradiction that $m((x+\beta \cdot z)@a) < m(x@a)+\beta$. Choose some α, $m((x+\beta \cdot z)@a) - \beta < \alpha \leq m(x@a)$, such that $x@a > (\alpha \cdot z)@c$ (such α exists by claim 3.19). By positive translation, $x@a > (\alpha \cdot z)@c \ \Rightarrow \ (x + \beta \cdot z)@a > (\alpha + \beta) \cdot z@c$. But since $\alpha + \beta > m((x + \beta \cdot z)@a)$, we have a contradiction to claim 3.19. ■

Claim 3.23 *For any* $a, b \in M_c \setminus 0_i$, $x \in V^a$, *and* $y \in V^b$*: if* $x@a$ *and* $y@b$ *are comparable, and* $m(x@a) < m(y@b)$*, then* $m((x + \delta \cdot e_i)@a) \leq m((y + \delta \cdot e_i)@b)$*, for any* $\delta > 0$.

proof: $m(x@a) < m(y@b)$ implies $\neg(x@a > y@b)$ by claim 3.20. Since they are comparable it follows that $x@a < y@b$. Thus $(x+\delta\cdot e_i)@a < (y+\delta\cdot e_i)@b$ by the closure under positive translation property, and thus by claim 3.20 again, $m((x + \delta \cdot e_i)@a) \leq m((y + \delta \cdot e_i)@b)$, as claimed. ∎

The next claim argues that, in some sense, f will choose an alternative with maximal measure:

Claim 3.24 *For any* $a, b \in M_c$*, and* $v \in V$ *such that* $v(a) = x \in V^a$*,*

1. *If* $v(b) = y \in V^b$*, and* $m(x@a) < m(y@b)$*, then* $f(v) \neq a$.
2. *If* $v(c) = \alpha \cdot z$ *(for some* $\alpha \in \mathcal{R}$*) and* $m(x@a) < \alpha$*, then* $f(v) \neq a$.
3. *If* $v(c) = \alpha \cdot z$ *(for some* $\alpha \in \mathcal{R}$*) and* $m(x@a) > \alpha$*, then* $f(v) \neq c$.

proof: (1) If, by contradiction, $f(v) = a$ then $x@a > y@b$ by definition, contradicting claim 3.20.
(2) Suppose $f(v) = a$, thus $x@a > (\alpha \cdot z)@c$, and by claim 3.19, $m(x@a) \geq \alpha$, a contradiction.
(3) Suppose $f(v) = c$, thus $x@a < (\alpha \cdot z)@c$ and, by definition, $m(x@a) \leq \alpha$, a contradiction. ∎

Using these properties of the measure function, we next show that it is affine.

Definition 3.15 $x@a$ *"calibrates"* $y@b$ *using* $w \in \mathcal{R}^n_+$ *if* $x@a < y@b$*, and if* $(x+\alpha\cdot w)@a$ *and* $y@b$ *are comparable for any* $\alpha > 0$.

In a descriptive manner, $x@a$ calibrates $y@b$ using w if we can "inflate" $x@a$ to $(x+\alpha\cdot w)@a$ while still keeping it comparable to $y@b$. The following claim implies that, if we can inflate the measure function as well, then the derivatives of $m(x@a)$ and $m(y@b)$ are identical:

Claim 3.25 *Suppose there exist* $\omega \in \mathcal{R}$, $\omega > 0$ *and* $w \in \mathcal{R}^n_+$ *such that* $m((x + \alpha \cdot w)@a) = m(x@a) + \omega \cdot \alpha$ *for all* $x \in V^a$ *and* $\alpha > 0$*. Fix any* $x \in V^a$ *and* $y \in V^b$ *such that* $x@a$ *calibrates* $y@b$ *using* w*. Then for any* i *such that* $a, b \notin 0_i$ *and any* $\delta > 0$*:* $m((x + \delta \cdot e_i)@a) - m(x@a) = m((y + \delta \cdot e_i)@b) - m(y@b)$.

proof: Since $x@a < y@b$ then $m(x@a) \leq m(y@b)$. Let β be such that $\omega \cdot \beta = m(y@b) - m(x@a)$ (thus $\beta \geq 0$). If $\beta > 0$, then for any $0 < \alpha < \beta$, $m((x+\alpha\cdot w)@a) = m(x@a)+\omega\cdot\alpha < m(y@b)$, and thus by claim 3.23, $m((x+\delta\cdot e_i+\alpha\cdot w)@a) \leq m((y+\delta\cdot e_i)@b)$. Therefore $m((x+\delta\cdot e_i+\beta\cdot w)@a) \leq m((y + \delta \cdot e_i)@b)$. If $\beta = 0$ then since $x@a < y@b$, $(x + \delta \cdot e_i)@a < (y + \delta \cdot e_i)@b$ and so $m((x + \delta \cdot e_i)@a) \leq m((y + \delta \cdot e_i)@b)$.

For any $\alpha > \beta$, $m((x + \alpha \cdot w)@a) = m(x@a) + \omega \cdot \alpha > m(y@b)$, and thus by claim 3.23, $m((x+\delta\cdot e_i+\alpha\cdot w)@a) \geq m((y+\delta\cdot e_i)@b)$. Therefore $m((x+\delta\cdot e_i+\beta\cdot w)@a) = m((y+\delta\cdot e_i)@b)$. Since $m((x+\delta\cdot e_i+\beta\cdot w)@a) = m((x+\delta\cdot e_i)@a)+\omega\cdot\beta$ and $\omega\cdot\beta = m(y@b)-m(x@a)$, the claim follows. ∎

To use this property, we need the domain to have "calibrators":

Requirement 3: A calibrator. An alternative $d \in M_c$ is a "calibrator" for player i if,

1. $d \notin 0_i$, and there exist $y \in V^d$, $(\alpha \cdot z)@c$, and $\epsilon, \delta > 0$ such that $y@d < (\alpha \cdot z)@c$, and $(y + \delta \cdot e_i)@d > ((\alpha + \epsilon) \cdot z)@c$.

2. For all $y, \tilde{y} \in V^d$ there exists $a \in S_c \setminus 0_i$ and $x \in V^a$ such that $x@a$ calibrates both $y@d$ and $\tilde{y}@d$ using z.

3. For all $a \in S_c \setminus 0_i$ and $x \in V^a$ there exists $y \in V^d$ such that $x@a$ calibrates $y@d$ using z.

4. For all $b \in (M_c \setminus S_c) \setminus 0_i$, and $x \in V^b$, there exists $y \in V^d$ such that $y@d$ calibrates $x@b$ using e_i.

Proposition 3.4 *If V is order based with conflicting preferences, f satisfies S-MON and player decisiveness, and A is non-degenerate, then Requirement 3 is satisfied: alternative c^i is a calibrator for i (for any $i > 1$).*

proof: The first calibrator requirement follows immediately from player decisiveness. For the second requirement, notice that S is the support of $e_1@c^1$. Now fix any $y, \tilde{y} \in V^{c^i}$, and suppose $y \leq \tilde{y}$. Since A is non-degenerate, there exists some $a \in S \setminus 0_i$. By definition of V^{c^i}, there exists $x \in V^a$ such that $y@c^i > x@a$. Thus also $\tilde{y}@c^i > x@a$. Since $i > 1$, for any $\alpha > 0$, $y@c^i$ and $(x + \alpha \cdot e_1)@a$ are compatible. Since $(x + \alpha \cdot e_1) \in V^a$ it follows from claim 3.18 that they are comparable (the same is true for $\tilde{y}@c^i$), and so the second requirement holds. The third and fourth requirements follows from essentially the same arguments. ∎

To show that Requirement 3 implies affinity, we use this basic technical fact (for completeness, we provide a proof in appendix 3.7.2):

Claim 3.26 *Fix some $m : X \to \mathcal{R}$, where $X \subseteq \mathcal{R}^n$ has the property that $x \in X$ and $y \geq x$ implies $y \in X$. Suppose also that m is monotonically non-decreasing.*

1. *If there exists $h_i : \mathcal{R}_+ \to \mathcal{R}_+$ such that $m(x + \delta \cdot e_i) - m(x) = h_i(\delta)$ for any $x \in X$ and $\delta > 0$, then there exist $\omega_i \in \mathcal{R}$ such that $h_i(\delta) = \omega_i \cdot \delta$.*

2. *If there exist $\omega_1, \ldots, \omega_n$ such that $m(x + \delta \cdot e_i) - m(x) = \omega_i \cdot \delta$ for any i, $x \in X$, and $\delta > 0$ then $m(x) = \sum_{i=1}^{n} \omega_i \cdot x_i + \gamma$ (for some constant $\gamma \in \mathcal{R}$).*

Claim 3.27 *If there exists a calibrator d for player i then there exists some $\omega_i \in \mathcal{R}$, $\omega_i > 0$ such that for any $b \in M_c \setminus 0_i$, for any $x \in V^b$ and for any $\delta > 0$: $m((x + \delta \cdot e_i)@b) - m(x@b) = \omega_i \cdot \delta$.*

proof: By the second calibrator property, for any $y, \tilde{y} \in V^d$ there exists $x@a$ that satisfies the conditions of claim 3.25 (recall that by claim 3.22, $m((x + \alpha \cdot z)@a) = m(x@a) + \alpha$). Thus:

$$\begin{aligned} &m((y + \delta \cdot e_i)@d) - m(y@d) = \\ &m((x + \delta \cdot e_i)@a) - m(x@a) = \\ &m((\tilde{y} + \delta \cdot e_i)@d) - m(\tilde{y}@d) \end{aligned}$$

so we conclude that there exists some $h : R_+ \to R_+$ such that $m((y + \delta \cdot e_i)@d) - m(y@d) = h(\delta)$ for all $y \in V^d$ and $\delta > 0$. By claim 3.26 there exists some $\omega_i \in \mathcal{R}_+$ such that $h(\delta) = \omega_i \cdot \delta$.

By the first calibrator property we can conclude that $\omega_i > 0$, since, for $y \in V^d$ that is specified by the property, $m(y@d) \leq \alpha$, $m((y+\delta \cdot e_i)@d) \geq \alpha + \epsilon$, and thus $\omega_i \cdot \delta = m((y+\delta \cdot e_i)@d) - m(y@d) > 0$.

By the third and fourth calibrator properties, for any $b \in M_c \setminus 0_i$, for any $x \in V^b$ and for any $\delta > 0$, there exists $y \in V^d$ that satisfies the conditions of claim 3.25 (recall we already know that $m((y + \delta \cdot e_i)@d) = m(y@d) + \omega_i \cdot \delta$). Thus:

$$\begin{aligned} &m((x + \delta \cdot e_i)@b) - m(x@b) = \\ &m((y + \delta \cdot e_i)@d) - m(y@d) = \omega_i \cdot \delta \end{aligned}$$

and the claim follows. ∎

Claim 3.28 *Requirements 1 to 3 imply that there exist constants* $\omega_1, \ldots, \omega_n$ *and* $\{\gamma_a\}_{a \in A}$ *such that:*

1. $m(x@a) = \sum_{i=1}^{n} \omega_i \cdot x_i + \gamma_a$ *for all* $a \in A$ *and* $x \in V^a$.

2. $\forall a, b \in A,\ v \in V:$ *if* $v(a) \in V^a$, $v(b) \in V^b$, *and* $\sum_{i=1}^{n} \omega_i \cdot v_i(a) + \gamma_a < \sum_{i=1}^{n} \omega_i \cdot v_i(b) + \gamma_b$, *then* $f(v) \neq a$.

proof: For any $i \neq 1$, $a \neq c$ and $x \in V^a$ there exists $\omega_i \in \mathcal{R}$ such that $m((x+\delta \cdot e_i)@a) - m(x@a) = \omega_i \cdot \delta$ for any $\delta > 0$ (by claim 3.27 if $a \notin 0_i$, and trivially if $a \in 0_i$). For $i = 1$, if $a \notin 0_i$ then $a \in S_c$, and thus by claim 3.22 $m((x + \delta \cdot e_1)@a) - m(x@a) = \delta$ so we take $\omega_1 = 1$. By claim 3.26 we conclude that $m(x@a) = \sum_{i=1}^{n} \omega_i \cdot x_i + \gamma_a$ for $a \neq c$. For c, since $x = v(c) = \alpha \cdot e_1$ then $m(x@c) = \alpha = \sum_{i=1}^{n} \omega_i \cdot x_i$ (so we take $\gamma_c = 0$). Therefore the first part of the claim follows. The second part is exactly claim 3.24, when replacing $m(x@a) = \sum_{i=1}^{n} \omega_i \cdot x_i + \gamma_a$ as shown in the first part of the proof. ∎

We can now immediately conclude:

Theorem 3.3 *Suppose* V *is order based with conflicting preferences,* f *is strongly monotone and player-decisive, and* A *is non-degenerate. Then* f *is affine maximizer for any* $v \in V^*$.

Corollary 3.2 *Suppose V is order based with conflicting preferences, f is strongly monotone and player-decisive, and A is non-degenerate. Then there exist $\omega_1, \ldots, \omega_n$, $\{\gamma_a\}_{a \in A}$, and a constant $M \in \mathcal{R}$ such that:*

$$f(v) \in argmax_{a \in A}\{ \sum_{i=1}^{n} \omega_i \cdot v_i(a) + \gamma_a \}$$

for all $v \in V$ such that $v_i(a) > M$ for all i and $a \notin 0_i$.

proof: (of corollary) Take representatives $v^a \in V^a$ and denote $M = max_{i,a}\{v_i(a)\}$. By the closure under positive translation of the V^a's, if $v_i(a) > M$ for all i and $a \notin 0_i$ then $v \in V^*$. ∎

Corollary 3.3 *Suppose V is order based with conflicting preferences, f is strongly monotone and player-decisive, and A is non-degenerate. For any $I' \subseteq \{1, \ldots, n\}$, denote $B_{I'} = \{a \in A | a \in 0_i \Leftrightarrow i \in I'\}$. Then there exist $\omega_1, \ldots, \omega_n$ and $\{\gamma_a\}_{a \in A}$ such that:*

$$f(v) \in \cup_{I' \subseteq \{1,\ldots,n\}} argmax_{b \in B_{I'}} \{ \sum_{i=1}^{n} \omega_i v_i(b) + \gamma_b \}$$

proof: Fix the constants implied by claim 3.28. Notice that $\cup_{I' \subseteq \{1,\ldots,n\}} B_{I'} = A$. Fix any $v \in V$. Suppose $f(v) = b \in B_{I'}$, but, by contradiction, there exists $a \in B_{I'}$ such that $\sum_{i=1}^{n} \omega_i v_i(b) + \gamma_b < \sum_{i=1}^{n} \omega_i v_i(a) + \gamma_a$ Choose some large enough δ so that $v(b) + \delta \cdot 1^n \in V^b$ (this is somewhat an abuse of notation, since only the non-zero coordinates are raised) and $v(a) + \delta \cdot 1^n \in V^a$. For every player $i \notin I'$ (i.e. $a \notin 0_i$) let $u_i = v_i + \delta \cdot 1^m$, and for player $i \in I'$, $u_i = v_i$. By S-MON, $f(u) = f(v) = a$, contradicting claim 3.28. ∎

Among the three conditions on f, it seems that the crucial one is the strong monotonicity (indeed, in section 3.5 we show examples of truthful CAs with non-degenerate range, that satisfy player decisiveness, but do not satisfy S-MON, and are not affine maximizers). On the other hand, for one parameter domains, it is not hard to construct strongly monotone functions that are not (almost) affine maximizers. The main question remained is, for exactly what domains is S-MON the main characterization of affine maximization:

Open Problem 1: *Is there a weaker condition than S-MON that implies affine maximization for combinatorial auctions?*

Open Problem 2: *Does S-MON imply affine maximization for order based domains that do not have conflicting preferences (e.g. matching) ?*

3.5 The Implications for Combinatorial Auctions

In this section we discuss the applicability of the main theorem to the main motivating problem: truthful mechanisms for approximating the optimal allocation in combinatorial auctions (CAs) and multi-unit auctions (MUAs). For this application most of the technical issues in the main theorem can be dropped. We start dealing with general issues, proceed with those implied by approximation factors, and conclude with the computational ones.

3.5.1 General Issues

CAs and MUAs satisfy all the requirements on the domain of theorem 3.3. Thus any CA or MUA that satisfies S-MON and player decisiveness, onto a non-degenerate domain, must be almost affine maximizer. In fact, a non-degenerate domain captures even the case where each bidder is interested in only two, known in advance, bundles ("known double minded bidders"), where one of bundles is the set of all goods.

Let us now look at the different requirements of the theorem. First notice that if the range is degenerate, then as discussed above, the social choice function need not be an almost affine maximizer[17]. As for the strong monotonicity, the following example demonstrates a truthful CA that does not satisfy S-MON, and indeed is *not* an affine maximizer:

Example 3.4 *Assume at least three players. Define A to be all possible allocations (where all goods are allocated). Define constants $\gamma_a = 0$ if $a_1 \neq \emptyset$, and $\gamma_a = 1$ if $a_1 = \emptyset$. The function f is as follows. For player 1, choose some allocation a that maximizes $\sum_{i=1}^{n} v_i(a) + \gamma_a$ and allocate a_1 to player 1. (clearly this is truthful for 1, e.g. with a price $\sum_{i=2}^{n} v_i(a) + \gamma_a$). For the others, if $v_1(c^1) \geq 1$, choose the allocation a from before. If $v_1(c^1) < 1$, choose the allocation that maximizes $\sum_{i=2}^{n} \omega_i v_i(a)$ (for some fixed ω_i's) (clearly this is also truthful for the other players, from the same reason as before, and since the choice between the two different affine maximizers depends only on player 1's declaration). Notice that f always chooses a feasible allocation: if $v_1(c^1) < 1$ then it must be the case that player 1 gets the empty set.*

Notice that f is player decisive and A is non-degenerate. To see that f does not satisfy S-MON, consider the following two types of 1: at first, $v_1(c^1) = 1 + \epsilon$, but the others declare high enough, so 1 gets nothing. Now, if 1 lowers all his values by ϵ, the allocation changes since now, f maximizes $\sum_{i=2}^{n} \omega_i v_i(a)$.

When there are two players, and all the goods are always allocated, then S-MON is no longer a burden: in this case, for any distinct allocations a and b we have that $a_i \neq b_i$. Thus, V is very close to being strict order domain, so we expect to be able to use Corollary 3.1 to reduce S-MON to W-MON. Specifically, define the *interior* of V to be $\mathring{V} = \{v \in V \mid v_i(a) < v_i(b) \text{ for all } a, b \in A \text{ s.t. } a_i \subsetneq b_i\}$ and define $\mathring{f} : \mathring{V} \to A$ by $\mathring{f}(v) = f(v)$.

[17] A simple class of examples is the "bundled auction" that allocates all items to the player with maximum value of $t_i(v_i(\Omega))$, where each t_i is an arbitrary monotone real function.

Theorem 3.4 *Fix any truthful CA or MUA f for two players, that always allocates all the goods. Suppose that $\mathring{f}$ is player decisive and onto a non-degenerate range* [18]*. Then:*

1. *f must be almost affine maximizer in the interior of V.*

2. *If the $\{\gamma_a\}_{a\in A}$'s are all zero then f is almost affine maximizer in all of V.*

proof: (1) $\mathring{V}$ is strict order based, with conflicting preferences. Since f is truthful, it satisfies W-MON. Thus $\mathring{f}$ satisfies W-MON as well. Therefore, by theorem 3.1, we can assume w.l.o.g that $\mathring{f}$ satisfies S-MON. Therefore, by theorem 3.3, $\mathring{f}$ is almost affine maximizer, and, therefore, so is f.

(2) We show that, if there exists $v \in V$ such that $f(v) \notin argmax_{a\in A}\{\sum_i \omega_i v_i(a)\}$ then there exists $u \in \mathring{V}$ such that $f(u) \notin argmax_{a\in A}\{\sum_i \omega_i u_i(a)\}$, thus a contradiction. We do this in two steps, moving from v_i to $u_i \in \mathring{V}_i$ (so suppose w.l.o.g that $i = 1$).

Let $f(v) = d$. Take any $v_1' \in \mathring{V}_1$ such that $|v_1'(a) - v_1(a)| < \epsilon$ for all $a \in A$. Define $D = \{a \in A \mid d_1 \subseteq a_1 \text{ and } v_1(d) = v_1(a)\ \}$. Let $u_1 = v_1'|^{D+=2\epsilon}$. Choose ϵ small enough so that $u_1 \in \mathring{V}_1$ and $argmax_{a\in A}\{\omega_1 u_1(a) + \omega_2 v_2(a)\} \subseteq argmax_{a\in A}\{\sum_i \omega_i v_i(a)\}$. By W-MON, $f(u_1, v_2) \in D$ since $d \in D$, and all alternatives not in D were raised by at most ϵ, while d was raised by 2ϵ. We claim that for any $b \in argmax_{a\in A}\{\omega_1 u_1(a) + \omega_2 v_2(a)\}$, $b \notin D$ and the claim follows. To see this suppose by contradiction that $b \in D$. Therefore $d_1 \subset b_1$, and so $b_2 \subset d_2$. But also $v_1(d) = v_1(b)$, and since $\omega_1 v_1(d) + \omega_2 v_2(d) < \omega_1 v_1(b) + \omega_2 v_2(b)$ (since $b \in argmax_{a\in A}\{\sum_i \omega_i v_i(a)\}$ and $d \notin argmax_{a\in A}\{\sum_i \omega_i v_i(a)\}$) it follows that $v_2(d) < v_2(b)$, contradicting $b_2 \subset d_2$. ∎

If we drop the assumption of always allocating all goods, then S-MON cannot be assumed without loss of generality – here is a specific CA that sometimes leaves unallocated goods, is player decisive and onto a full range, but does not satisfy S-MON, and is not almost affine maximizer:

Example 3.5 *For any $X \subseteq \Omega$, let $\tilde{p}_i(X, v_j) = v_j(\Omega) - v_j(\Omega \setminus X)$ (these are the Clarke prices). Suppose f allocates to each player the bundle that maximizes his utility under these prices (breaking ties for the two players in a consistent manner). It is not hard to verify that this is truthful and always chooses an optimal allocation. Now, change the prices of player 2 to be $p_2(X, v_1) = \tilde{p}_2(X, v_1) + v_1(\Omega)/2$ for any $X \neq \emptyset$, and $p_2(\emptyset, v_1) = \tilde{p}_2(\emptyset, v_1) = 0$. Clearly, f is still truthful. To see that it chooses a legal allocation, suppose that the Clarke function chooses $\tilde{a}$, and that $f(v) = a$. $a_1 = \tilde{a}_1$, since the prices of 1 did not change. a_2 is either the empty set or $\tilde{a}_2$: since we added a constant to the price of all the non-empty bundles, the difference between the utility of $\tilde{a}_2$ and any other non-empty bundle remains the same as in the Clarke function. Since we break ties in the same manner, f chooses either the empty set or $\tilde{a}_2$.*

To see that f is not almost affine maximizer (and does not satisfy S-MON), take any $g \in \Omega$, and fix $v_1(\Omega) = 16$, $v_1(\Omega \setminus \{g\}) = 10$, $v_2(\{g\}) = 9$ the rest of v is just small perturbations of the values,

[18] It is not enough to require that f has a non-degenerate domain (i.e. we must require that $\mathring{f}$ has that), as example 3.6 in appendix 3.7.3 demonstrates.

so that v will be in the interior of V). The optimal allocation is $(\Omega \setminus \{g\}, \{g\})$, so, f allocates $\Omega \setminus \{g\}$ to 1. But $p_2(\{g\}, v_1) = 13 > v_2(\{g\})$, so 2 gets nothing. If player 1 lowers his value for Ω to 12, then $p_2(\{g\}, v_1) = 8$, and so the allocation changes, contradicting S-MON.

3.5.2 Approximation

Since exact welfare optimization in CAs is computationally hard (see also below), we ask whether there exist truthful welfare approximations. A social choice function is a c-approximation of the optimal welfare if, for any type v, the alternative $f(v)$ has welfare of at least $1/c$ times the optimal welfare for v. For this class of functions, we are able to show that most of the qualifiers of the main theorem can be dropped.

Specifically, we define an auction to be *unanimity-respecting* (essentially equivalent to the notion of "reasonable" in [80]) if, whenever every player values only a single bundle a_i, and $a_i \cap a_j = \emptyset$ for all i, j, then f chooses the allocation $a = (a_1, \ldots, a_n)$. Using these, the "almost" qualifier and the player decisiveness property are dropped from the main theorem:

Lemma 3.3 *Any unanimity-respecting truthful CA or MUA that satisfies IIA and achieves a c-approximation must be an affine maximizer. Furthermore, the weights must satisfy $\gamma_a = 0$ for all alternatives a and and $(1/c) \le (\omega_i/\omega_j) \le c$ for all players i, j.*

proof: First notice that, by the unanimity-respecting property, it follows that f has full range, since every possible allocation is obtained when the players are unanimous for it. Also notice that, since f is a c-approximation, it must be player decisive.

We now show that V^* from the proof of the main theorem now becomes $V^* = \{v \in V \mid v_i(a) > 0 \forall \text{i and } a \in A \setminus 0_i\}$. This will immediately follow from the claim that, for any $a \in A$: $V^a = \{v(a) | v_i(a) > 0 \text{ for all i s.t. } a \notin 0_i\}$, where the V^a's are defined in the proof of the main theorem. For $a \in S$ ($S = \{a \in A | a_1 \notin 0_1\}$), then by definition, $x \in V^a$ iff there exists $v \in V$ s.t. $f(v) = a$ and $v(a) = x$. Therefore, take v_i so that player i is interested in a_i with $v_i(a_i) = x_i$ (and if $a_i \in 0_i$ then i has a value of zero for all bundles), and so $f(v) = a$, since f is unanimity-respecting. For the c^i alternatives, $y \in V^{c^i}$ iff there exists some $a \in S$ and $x \in V^a$ such that $y@c^i > x@a$. Take some allocation a s.t. $a_1, a_i \neq \emptyset$. Thus $a \in S$. For any $\epsilon > 0$, let $x = (\epsilon, \ldots, \epsilon)$. As shown before, $x \in V^a$. Let v be some type in which all players are interested in the single bundle a_i with a value of ϵ, and player i, in addition, has a value of $2nc\epsilon$ for c^i. Since f is a c-approximation, it follows that $f(v) = c^i$, and so $y@c^i > x@a$ (where we choose ϵ so that $y_i = 2nc\epsilon$). For the other alternatives a, $x \in V^a$ iff there exists $y \in V^{c^i}$ s.t. $x@a > y@c^i$ (for all i s.t. $a \notin 0_i$). For this, start with a type v in which all players are unanimous for a with value x, except player i who has value $x_i - \epsilon$. Thus $f(v) = a$. Now, if we raise all non-zero coordinates of i by ϵ, then by S-MON f will still choose a, and so $x@a > y@c^i$, where $y \in V^{c^i}$, as needed.

Therefore, by our main theorem, f is affine maximizer for any $v \in V^*$. We now show that the γ_a constants are all zero. Assume w.l.o.g that $\gamma_a \ge 0$ for any $a \in A$. Let b an alternative with

$\gamma_b = 0$. Suppose all players are unanimous with value δ for b. Then $f(v) = b$. Suppose that $b_j \neq \emptyset$. Then, if j raises all his non-zero values by ϵ, b is still chosen. Since $v(b) \in V^b$ and $v^{c^j} \in V^{c^j}$, then by claim 3.28, $\gamma_{c^j} \leq \gamma_b + \sum_i \omega_i \delta$. Since this is true for any $\delta > 0$ it follows that $\gamma_{c^j} = 0$. Now suppose by contradiction that $\gamma_a > \gamma_{c^j}$ for some alternative a. Consider the type where j values all the goods for a value of $c\epsilon$, and all players value a by some ϵ' s.t. $\sum_i \omega_i \epsilon' < \epsilon$. Then, for small enough ϵ, ϵ', j will not receive all the goods since $\gamma_a > \gamma_{c^j}$. But this contradicts the approximation guarantee.

To verify that for any i, j, $(\omega_i/\omega_j) \leq c$, suppose i, j are interested only in the bundle that contains all the goods, for a value of 1 and $c + \epsilon$, respectively. From the approximation ratio it follows that j wins, and therefore $\omega_i \cdot 1 \leq \omega_j(c + \epsilon)$.

We are left to show that f is an affine maximizer for any $v \in V$. Suppose by contradiction that there exists a type v s.t. $f(v) = a$ but $\sum_i \omega_i v_i(a) < \sum_i \omega_i v_i(b)$ for some alternative b. We first verify that $v(a) \in V^a$ by adding ϵ to all non-zero coordinates of any player i with $a_i \neq \emptyset$ (by S-MON the result remains a). By claim 3.28, it follows that for every c^i, $\omega_i v_i(c^i) \leq \sum_i \omega_i v_i(a)$ (as shown above, if $v_i(c^i) > 0$ then $v(c^i) \in V^{c^i}$). We turn this inequality to be strict by choosing two players i, j with $a_i, a_j \neq \emptyset$ and raise all their non-zero values by ϵ. [19] We now move to some $u \in V^*$ in which the measure of a is still smaller than that of b: For every player i, increase all non-zero coordinates by ϵ, and c^i's value by 2ϵ. Let u_i denote this new type of i. By S-MON, $f(v_{-i}, u_i)$ is either a or c^i, and since $\omega_i u_i(c^i) + \gamma_{c^i} < \sum_i \omega_i v_i(a) + \gamma_a$ (we choose a small enough ϵ) we have that $f(v_{-i}, u_i) = a$. By induction, $f(u) = a$. But now $u \in V^*$, and we still have $\sum_i \omega_i u_i(a) + \gamma_a < \sum_i \omega_i u_i(b) + \gamma_b$, thus a contradiction. ∎

For two players, where all the goods are always allocated, we can drop even the remaining qualifiers:

Lemma 3.4 *Any truthful CA or MUA for two players that always allocates all items and achieves an approximation factor of $c < 2$ must be an affine maximizer. Furthermore, it must have a full range, and the weights must satisfy $\gamma_a = 0$ far all a and $0.5 < (\omega_i/\omega_j) < 2$ for all i, j.*

proof: We observe the following:

1. Any c-approximation algorithm must satisfy player decisiveness: Fix any player i and v_{-i}. If $v_i(\Omega) = (c+1)\max_{j \neq i} v_j(\Omega)$ then $f(v_i, v_{-i})$ must allocate all goods to i in order to c-approximate the optimal welfare. (In fact this is true for any number of players).

2. Any $(2 - \epsilon)$-approximation that always allocates all the goods must have a full range, even in its interior: Fix any allocation $a = (a_1, a_2)$ where $a_2 = \Omega \setminus a_1$. If player i wants a_i with some value x (for $i = 1, 2$) then a has welfare of $2x$ and any other allocation has a value at most x (if player i is

[19] The possibility that only one player, say j, receives a non-empty bundle in a is handled by performing the move from v_j to u_j last. Then, neither a nor c^j can be chosen since both measures are strictly smaller than b's measure – and now $v(b) \in V^b$.

allocated a bundle that contains a_i then the bundle of player j is partial to a_j). This type is on the boundary of V, but we can easily shift it to the interior by choosing a small enough δ (w.r.t. ϵ and x) and "space" the values for other bundles with δ jumps: a bundle $X \not\supseteq S_i$ has value $l\delta$, where $l = |X|$, and a bundle $X \supset S_i$ has value $x + l\delta$. We need to choose δ so that $2x/(x + L\delta) > 2 - \epsilon$ (where L is the number of goods).

3. For the case of a $(2-\epsilon)$-approximation CA (or MUA) for two players (that always allocates all the goods), V^* from the main theorem's proof equals V, as follows. Notice that definition of V^a's for this case become: $V^{c^1} = V|_{c^1}$, $V^a = \{v(a)|f(v) = a\}$ (for any $a \neq c^i$), and $V^{c^2} = \{y|\exists a, x \in V^a \text{ s.t. } y@c^2 > x@a\}$. For any $a \in A$ s.t. $a_i \neq \emptyset$ for $i = 1, 2$, for any $x > 0$ we have seen in the last section that there exists $v \in V$ such that $v_1(a) = v_2(a) = x$ and $f(v) = a$. Thus $(x, x) \in V^a$. For any $y \geq (x, x)$ it follows from the closure under positive translation, that $y \in V^a$. Since this is true for any $x > 0$ it follows that $V^a = \mathcal{R}^2_+$. For the alternative c^2, and any $y@c^i$, since $x = (y_2/4, y_2/4) \in V^a$ for some $a \in A$, it follows that $y@c^i > x@a$ (since this is a $(2-\epsilon)$-approximation), and so $y \in V^{c^i}$. Thus $V^{c^i} = \mathcal{R}^2_+$, and so $V^* = \overset{\circ}{V}$.

Since all the goods are always allocated, we conclude, by theorem 3.4, that f is affine maximizer in its interior.

We now claim that the γ_a constants are all equal to zero: Otherwise, suppose there are two allocations a, b s.t. $\gamma_b > \gamma_a$. Then, if we choose $x = (\gamma_b - \gamma_a)/4$ to be the x of observation 2 (as the value of a_i), we get that a will not be chosen, contradicting the fact that a must be chosen in order to be a $(2-\epsilon)$-approximation (as shown there). Therefore, by theorem 3.4 again, f is affine maximizer.

We are left to show that $\omega_i \leq 2\omega_j$. Otherwise suppose $\omega_i > 2\omega_j$, and consider the case where player i is interested only in Ω, for a value of 1, and player j is also interested only in Ω, for a value of 2. Then, f will allocate Ω to i, contradicting the $(2-\epsilon)$-approximation. ∎

3.5.3 Polynomial-Time Computation

All treatment of mechanisms so far assumed a fixed number of players n and a fixed number of items k. When formalizing the notion of computational running time we must let these parameters (or at least the number of items k) grow, and consider the running time as a function of them. A mechanism whose running time we wish to analyze would apply to all k and, if n is not fixed, for all n, i.e. would really be a uniform family of mechanisms. The characterization as affine maximizer above would then only apply to each mechanism in the family separately (with no explicit relationship across the different values of n and k.) This implies that, for a given k and n, the constants ω_i, γ_a, and the range A, may all depend on k and n. We denote these by the superscript n, k, i.e. $\omega_i^{n,k}, \gamma_a^{n,k}, A^{n,k}$ (we sometimes drop the n if it is clear from the context). Notice that, if these constants are large (w.r.t. n and k), then this may limit the range of the auction in a way that will enable it to become polynomial (e.g. if ω_i is much larger than the input size and the other constants, this depicts that player i will always receive all goods). This motivates the following definition:

Definition 3.16 *An affine maximizer CA or MUA has* **polynomially bounded constants** *if there exists a constant c such that* $(\omega_i^{n,k}/\omega_j^{n,k}), \gamma_a^{n,k} \leq 2^{(n \cdot \log k)^c}$ *for all number of goods k, for any number of players n and any players* $i, j \in \{1, \ldots, n\}$, *and for any* $a \in A^{n,k}$.

Note that $\omega_i^k/\omega_j^k, \omega_j^k/\omega_i^k, \gamma_a^k$ are real numbers with possibly infinitely many digits. The only consideration about these numbers is that they are not too small or too large.

In order to represent the mechanisms' running time as a function of its input size, we must fix an input representation for the valuations, i.e. a *bidding language* [77]. Our results apply to any such choice of a bidding language as long as it is complete (i.e. can represent all valuation) and sufficiently powerful. In fact, for claiming that affine maximization is as computationally hard as exact maximization, we only need the bidding language to have the following two elementary properties:

Definition 3.17 *A bidding language L is* **elementary** *if,*

1. *For any bid* $b \in L$ *that implicitly represents some valuation v, there exists a polynomial time procedure to construct a bid* $b' \in L$ *that represents the valuation* $\alpha \cdot v$, *i.e. multiplying all values of all bundle by some constant* $\alpha > 0$.

2. *There exists a valid bid in which all bundles except* Ω *are valued as* 0, *and* Ω *is valued as* α, *for any* $\alpha \geq 0$.

For example, OR bids and XOR bids (see details below) are elementary: the first property is satisfied by just going over all the bid's blocks and multiplying their value by α.

We can now state formally that affine maximizers CAs and MUAs are as hard to compute as exact welfare maximizers:

Lemma 3.5 *Any affine maximizer CA or MUA with an elementary bid language, with polynomially bounded constants, and with the additive constants being equal to zero, is as computationally hard as the exact welfare maximization problem (with the same bidding language and the same range A).*

proof: Denote by AM the affine maximizer CA or MUA, and by EM the exact welfare maximizer. To prove the claim, we need to show a reduction from EM to AM. Before showing this, we need a method to calculate a close enough bound on the constants ω_i. We assume w.l.o.g that $\omega_1 = 1$ (any affine maximizer with constants $\{\omega_1, \ldots, \omega_n\}$ is also an affine maximizer with constants $\{\omega_1/\omega_1, \ldots, \omega_n/\omega_1\}$). Suppose the input bid is of size l (i.e. it contains l bits), let $M = 2^l$ be an upper bound on the value of any bundle, and $1/R = 1/2^l$ be a lower bound on the precision of the bundle values, i.e. if $v_i(X) > v_i(Y)$ then $v_i(X) \geq v_i(Y) + (1/R)$.

Claim 3.29 *There is a polynomial time procedure that computes* $\tilde{\omega}_i$ *such that* $1 \leq (\omega_i/\tilde{\omega}_i) \leq 1 + 1/(2nMR)$.

proof: We describe a simple iterative procedure: maintaining an interval I that contains ω_i, while reducing its size half until it is sufficiently small. We use a bid $b(\alpha_1, \alpha_i)$, which represents n players, where players 1 and i are interested only in the bundle Ω for a value of α_1, α_i, respectively, and the other players have a value of zero for all bundles.

The procedure works as follows. Initially, find some α s.t. $AM(b(\alpha, 1)) = 1$ (i.e. the auction allocates all goods to player 1). This is done by starting with $\alpha = 1$ and doubling it until the desired allocation is achieved. Since ω_i is polynomially bounded, i.e. $\omega_i \leq 2^{(n \cdot \log k)^c}$, this requires at most $2(n \cdot \log k)^c$ steps. Since the auction choose the allocation with maximal weighted welfare, we have that $\omega_i \cdot 1 \leq \omega_1 \alpha = \alpha$. Then we find c_1 such that $AM(b(\alpha, c_1)) = i$, using the same doubling method. This again takes polynomial time in the number of players and the input size. Therefore we now have that $\omega_i \in [(\alpha/c_1), (\alpha/c_0)]$, where $c_0 = 1$. We now set $c^* = (c_1 + c_0)/2$. And test $AM(b(\alpha, c^*))$. If this equals 1 then we set $c_0 = c^*$, otherwise this equals i and we set $c_1 = c^*$. Thus we maintain $\omega_i \in [(\alpha/c_1), (\alpha/c_0)]$. We repeat this until $c_1 - c_0 \leq 1/(2nMR)$, and then determine $\tilde{\omega}_i = \alpha/c_1$. Therefore $1 \leq \omega_i/\tilde{\omega}_i \leq c_1/c_0$. Since $c_0 \geq 1$ it follows that $c_1/c_0 = (c_1 - c_0)/c_0 + 1 \leq 1 + 1/(2nMR)$. This binary search procedure takes $\log(\beta(2nMR))$, where β is the initial length of the interval. This is again polynomial in the number of players and the input size. ∎

We can now describe a reduction from EM to AM:

1. Given an input bid $b = (b_1, \ldots, b_n)$ for EM, first compute the bounds $\{\tilde{\omega}_i\}_i$ according to claim 3.29.

2. Create a bid $\tilde{b}$ such that $\tilde{b}_i$ represents the valuation $\tilde{v}_i = v_i/\tilde{\omega}_i$ (where v_i is the valuation that b_i represents) – there is an efficient method to compute $\tilde{b}$ from b since the bid language is elementary.

3. Return the allocation $AM(\tilde{b})$ (as the allocation that EM outputs).

The correctness of this reduction immediately follows from the following claim:

Claim 3.30 *For any two allocations* $a, b \in A$, *if* $\sum_i \omega_i \tilde{v}_i(a) \geq \sum_i \omega_i \tilde{v}_i(b)$ *then* $\sum_i v_i(a) \geq \sum_i v_i(b)$.

proof: We show that $\sum_i v_i(a) < \sum_i v_i(b)$ implies $\sum_i \omega_i \tilde{v}_i(a) < \sum_i \omega_i \tilde{v}_i(b)$. First note that $\sum_i \omega_i(\tilde{v}_i(b) - \tilde{v}_i(a)) = \sum_i (\omega_i/\tilde{\omega}_i)(v_i(b) - v_i(a)) = \sum_i (v_i(b) - v_i(a)) + \sum_i ((\omega_i/\tilde{\omega}_i) - 1)(v_i(b) - v_i(a))$. Since $\sum_i v_i(b) > \sum_i v_i(a)$ then $\sum_i (v_i(b) - v_i(a)) \geq 1/R$. Since $0 \leq ((\omega_i/\tilde{\omega}_i) - 1) \leq 1/(2nMR)$ and $(v_i(b) - v_i(a)) \geq -M$, it follows that, for every i, $((\omega_i/\tilde{\omega}_i) - 1)(v_i(b) - v_i(a)) \geq (-M)(1/(2nMR)) = -(1/(2nR))$. Therefore: $\sum_i ((\omega_i/\tilde{\omega}_i) - 1)(v_i(b) - v_i(a)) \geq n(-1/(2nR)) = -1/(2R)$. So we can conclude that $\sum_i \omega_i(\tilde{v}_i(b) - \tilde{v}_i(a)) \geq 1/R - 1/(2R) > 0$. ∎

This concludes the proof of the lemma. ∎

Our interest is in cases where the bidding language is sufficiently powerful as to make exact welfare maximization NP-complete. If the bid language forces the input to be long, e.g. the value of all possible bundles must be specified, then clearly we can construct an affine maximizer that will take linear time in the size of this input. Therefore, we need to allow short inputs. In particular, [62] show that as long as even single-minded bids are possible then the CA problem with n players is NP-complete (where n is not fixed). We observe that this is true for MUAs as well, as long as the number of desired items may be given in binary (rather than unary). When the number of players is fixed, then single-minded bids (as well as XOR-bids) may be handled in polynomial time, but we show that allowing OR bids results in an NP-complete optimization problem. More formally:

Definition 3.18 (Single Minded Bids) *A single minded bid of player i has the form $(q^i,\ v^i)$, which implies the following valuation: for MUA, any quantity not smaller than q^i has a value v^i, and, for CA, any bundle that contains the bundle q^i has value v^i. All other bundles have value* 0.

Definition 3.19 (OR Bids) *Player i's valuation is represented by OR bids if it is a collection of pairs $(q^i_1,\ v^i_1), (q^i_2,\ v^i_2), \dots, (q^i_l,\ v^i_l)$, where each v^i_j is the value of i for the bundle q^i_j – for MUA q^i_j specifies just the number of items in the bundle, where in CA it identifies uniquely some bundle. From this representation, it is implicit that the value of any bundle X is: $v_i(X) = \max\ \{\ \sum_{j\in I} v^i_j \mid I \subseteq \{1,\dots,l\}$ s.t. $\cup_{j\in I} q^i_j \subseteq X$ and for all $j, j' \in I,\ q^i_j \cap q^i_{j'} = \emptyset\}$* [20].

Claim 3.31 *Any welfare maximizing CA or MUA for n players (where n is not fixed), with full range, is NP-hard, even with single minded bids. If the number of players is fixed, then the above holds with OR bids as the bidding language.*

proof: We give the proof in appendix 3.7.4. ∎

To integrate our main characterization with this computational hardness, we need a bidding language that will be rich enough to express all possible valuations, since the characterization does not assume any limitations on the possible valuation of the players. Notice that single minded bids and OR bids are not rich enough (OR bids can express only super-additive valuations).

Definition 3.20 *A bidding language L* **generalizes** *the bidding language L' if,*

1. *L contains all valid bids of L'.*
2. *L can express all possible player valuations.*

For example, XOR bids generalize single minded bids. And, OR bids with dummy items, and XOR of ORs, both generalize OR bids.

We can now integrate the above claims with our characterization of truthful welfare approximations:

[20]In MUA, X and the q^i_j are number of goods, and so the condition becomes $\sum_{j\in I} q^i_j \leq X$.

Theorem 3.5 *Any Unanimity-respecting truthful polynomial-time combinatorial (or multi-unit) auction, with a bidding language that generalizes single minded bids, and that satisfies IIA, cannot obtain* $poly(n, k)$ *welfare approximation (unless* $P = NP$*).*

proof: By Lemma 3.3, any truthful CA or MUA that satisfies Unanimity-respecting and IIA, and is a $poly(n, k)$ welfare approximation is an affine maximizer with polynomially bounded constants and the additive constants are zero. By Lemma 3.5, the affine maximization problem is as computationally hard as the exact maximization problem, and by claim 3.31, this problem is NP-hard. Therefore the auction cannot be polynomial (unless $P = NP$). ∎

For the case of two-player auctions, we can omit the "unanimity-respecting" and "IIA" assumptions:

Corollary 3.4 *Any truthful polynomial-time multi-unit (or combinatorial) auction between two players, with a bidding language that generalizes OR bids, and that always allocates all goods, cannot obtain a welfare approximation better than* 2 *(unless* $P = NP$*).*

proof: Follows from essentially the same arguments as above, replacing Lemma 3.3 with Lemma 3.4. ∎

In contrast, for MUA without the truthfulness requirement there exists an FPAS [81]! Also notice that a truthful 2-approximation can be easily obtained using a simple auction of the bundle of all goods.

3.6 Unrestricted Domains

In this section we give an alternative proof to Roberts' Theorem, using the notions of our main theorem. Notice that an unrestricted domain is (trivially) an order based domain with conflicting preferences. Therefore, our main theorem applies for it (with its qualifiers). However, if we want to remove the qualifiers, we can prove requirements one to three in a different (and in fact easier) way. We choose $V^* = V$, and show that for any alternative $c \in A$, $\vec{1}@c$ is a transitive reference for $A \setminus \{c\}$, and that any alternative $a \in A$ is a calibrator for any player i. For this we assume only that f satisfies S-MON:

Proposition 3.5 *If* V *is unrestricted and* f *is strongly monotone, then Requirements 1 to 3 are satisfied.*

proof: We show this using several claims: the first one summarizes some nice properties of an unrestricted domain:

Claim 3.32 *Suppose* V *is unrestricted and* f *is strongly monotone. Then:*

1. *$x@a > y@b$ implies $(x + \delta \cdot e_i)@a > (y + \delta \cdot e_i)@b$ for any i and $\delta > 0$.*

2. *For any $x@a$ there exists $v \in V$ such that $f(v) = a$ and $v(a) = x$.*

3. *Any $x@a$ and $y@b$ are comparable.*

proof: (1) Take $v \in V$ such that $f(v) = a, v(a) = x$, and $v(b) = y$. By S-MON, $f(v_i + \delta \cdot \vec{1}, v_{-i}) = a$ [21], and the claim follows.

(2) Fix some $u \in V$ such that $f(u) = a$. Let $\delta_i = x_i - u_i(a)$, and $v_i = u_i + \delta_i \cdot \vec{1}$. By S-MON, $f(v_i, u_{-i}) = a$, and thus also $f(v) = a$. Since $v(a) = x$, the claim follows.

(3) Fix any $u \in V$ such that $u(a) = x$ and $f(u) = a$. Let us construct some $v \in V$ as follows: $v(a) = x, v(b) = y$, and for any $c \neq a, b$: $v(c) = u(c)$. Since $x@a > u(c)@c$ then by claim 3.13, it follows that $f(v) \neq c$ (otherwise $x@a < u(c)@c$). Therefore $f(v) \in \{a, b\}$, and the claim follows. ∎

Requirement 1 follows from the part 1 of the claim.

Claim 3.33 *For any $c \in A$, $\vec{1}@c$ is a transitive reference for $A \setminus \{c\}$. Therefore, requirement 2 is satisfied.*

proof:

Measurability: Fix any $v \in V$ such that $f(v) = a$ and $v(a) = x$. Therefore $x@a > y@c$ (where $y = v(c)$). By claims 3.13 and 3.32, $x@a > (\alpha \cdot \vec{1})@c$ for any α such that $\alpha \cdot \vec{1} \leq y$. For the other direction, fix any $u \in V$ such that $f(u) = c$. By claim 3.32 we can assume w.l.o.g that $u(a) = x$. Therefore, for any β such that $\beta \cdot \vec{1} \geq u(c)$ it follows that $x@a < (\beta \cdot \vec{1})@c$.

Transitivity: Assume $x@a < (\alpha \cdot \vec{1})@c$ and $\neg(y@b < (\alpha \cdot \vec{1})@c)$, but, by contradiction, $x@a > y@b$. Choose $v \in V$ "for" $x@a > y@b$. By claim 3.13 it must be the case that $v(c) < (\alpha \cdot \vec{1})$. By S-MON, $f(v|^{c=\alpha \cdot \vec{1}} \in \{a, c\}$. But if this equals a it contradicts $x@a < (\alpha \cdot \vec{1})@c$, and if this equals c it contradicts $\neg(y@b < (\alpha \cdot \vec{1})@c)$.

R-monotonicity: follows immediately from claim 3.13.

L-monotonicity: Suppose by contradiction that $(x + \delta \cdot e_i)@a < (\alpha \cdot z)@c$ but $\neg(x@a < (\alpha \cdot z)@c)$. Therefore $x@a > (\alpha \cdot z)@c$, contradicting to claim 3.13.

∎

Since a transitive reference cannot measure itself, we need some "rich enough" structure of transitive references:

Definition 3.21 (A zero player) *Fix any reference $z@c$. Player i is a zero player w.r.t. M_c if for any $a \in M_c$ and any $x \in V^a$: $x@a < (\alpha \cdot z)@c \Rightarrow (x + \delta \cdot e_i)@a < (\alpha + \epsilon)z@c$, for any $\epsilon, \delta > 0$.*

[21] By definition, $\vec{1} = (1, \ldots, 1)$.

Claim 3.34 *For any $c \in A$ and any non-zero player i (w.r.t. $A \setminus \{c\}$) there exists a calibrator for i. Therefore, requirement 3 is satisfied.*

proof: Since i is a non-zero player, there exists $b \in A \setminus \{c\}, y@b, \alpha \cdot z@c$, and $\epsilon, \delta > 0$ such that the first calibrator requirement holds. For the second requirement, first notice that $0_i = \emptyset$, and $S_c = A \setminus \{c\}$. For any $y@b$ choose some a and $x@a$ such that $y@b > x@a$. Similarly, for $\tilde{y}@b$ choose some $\tilde{x}@a$ such that $\tilde{y}@b > \tilde{x}@a$. Then, for $x' = min(x, \tilde{x})$ it follows that $x'@a$ calibrates both $y@b$ and $\tilde{y}@b$. The third and fourth requirements are also immediate. ∎

This concludes the proof of the proposition. ∎

Requirement 4: Connected references. f has (a set of) "connected references" if there exists some $R \subseteq A$ such that, for every $c \in R$, $z@c$ is a transitive reference for M_c (for some $z \in \mathcal{R}^n,\ z \geq 0$), and,

1. For any $a, b \in A$ there exists $c \in R$ such that $a, b \in M_c$.
2. For any $c, d \in R$ and any player i there exists $a \in A \setminus 0_i$ such that $a, c \in M_d$, and $a, d \in M_c$.
3. For any $c \in R$ and any player i, either i is a zero player w.r.t. M_c, or there exists a calibrator $c^i \in M_c$ for i (the calibrators c^i, c^j are not necessarily distinct).

Proposition 3.6 *If V is unrestricted, f is strongly monotone, and $|A| \geq 3$, then f has connected references.*

proof: Fix any three alternatives as R, the set of connected references. For any $c \in R$, it follows from claim 3.33 that $\vec{1}@c$ is a transitive reference for $A \setminus \{c\}$. The first two requirements from R immediately follow from this, and the third requirement follows from claim 3.34. ∎

Lemma 3.6 *Suppose f has connected references. Then there exist constants $\omega_1, \ldots, \omega_n$ and $\{\gamma_a\}_{a \in A}$ such that:*

$$f(v) \in argmax_{a \in A}\{\ \sum_{i=1}^{n} \omega_i \cdot v_i(a) + \gamma_a\ \}$$

for all $v \in V^*$.

We prove this Lemma in sub-section 3.6.1 below. From all this we can immediately conclude:

Theorem 3.6 *If V is unrestricted, f is strongly monotone, and $|A| \geq 3$, then f is affine maximizer.*

Corollary 3.5 *Any truthful social choice function on an unrestricted domain, with at least three alternatives, must be affine maximizer.*

proof: By theorem 3.1, f satisfies W-MON. By theorem 3.1, since an unrestricted domain is an open set, there exists a function $\tilde{f}$ that satisfies S-MON, and that if $\tilde{f}$ is affine maximizer, then so is f. By theorem 3.6, $\tilde{f}$ must be affine maximizer, and the claim follows. ∎

3.6.1 Connected references

Since there are several references, we denote by $m_c(x@a)$ the measure of $x@a$ according to the reference $z@c$. Before proving the lemma, we need some useful claims:

Claim 3.35 *For any reference $c \in R$, $a \in M_c$ and $x \in V^a$,*

1. *If i is a zero player w.r.t. M_c, then $m_c((x + \delta \cdot e_i)@a) - m_c(x@a) = 0$, for any $\delta > 0$.*

2. *There exist $\omega_1^c, \ldots, \omega_n^c, \{\gamma_{a,c}\}_{a \in M_c}$ such that: $m_c(x@a) = \sum_{i=1}^n \omega_i^c \cdot x_i + \gamma_{a,c}$.*

3. *$\sum_{i=1}^n \omega_i^c \cdot z_i = 1$.*

proof: (1) From the definition of a zero player it follows that $m_c((x + \delta \cdot e_i)@a) \leq m_c(x@a)$. By claim 3.21, $m_c((x + \delta \cdot e_i)@a) \geq m_c(x@a)$, and the claim follows.

(2) Similarly to claim 3.28, since for every $c \in R$ and every player i, either there exists a calibrator for i or i is a zero player, it follows that $m_c((x + \delta \cdot e_i)@a) - m_c(x@a) = \omega_i^c \cdot \delta$, and the claim follows.

(3) Fix any $a \in S_c$, $x \in V^a$, and some $\beta > 0$. By claim 3.22, $m_c((x + \beta \cdot z)@a) = m_c(x@a) + \beta$. Therefore $\sum_{i=1}^n \omega_i^c \cdot z_i \cdot \beta = \beta$, and the claim follows. ∎

Claim 3.36 *For any $c, d \in R$:*

1. *For any $a \in M_c \cap M_d$, and $x \in V^a$: $m_c(x@a) = m_d(x@a) - \gamma_{c,d}$.*

2. *For any i, $\omega_i^c = \omega_i^d = \omega_i$.*

proof: First suppose by contradiction that $m_c(x@a) < m_d(x@a) - \gamma_{c,d}$, and choose some α such that $x@a < (\alpha \cdot z)@c$, and $\alpha < m_d(x@a) - \gamma_{c,d}$. Since $c \in M_d$, and by claim 3.35, $m_d((\alpha \cdot z)@c) = \alpha + \gamma_{c,d} < m_d(x@a)$. By claim 3.20, this contradicts $x@a < (\alpha \cdot z)@c$. Similarly, suppose by contradiction that $m_c(x@a) > m_d(x@a) - \gamma_{c,d}$, and choose some α such that $x@a > (\alpha \cdot z)@c$, and $\alpha > m_d(x@a) - \gamma_{c,d}$. Thus $m_d((\alpha \cdot z)@c) = \alpha + \gamma_{c,d} > m_d(x@a)$, a contradiction.

The second claim follows from the first one by taking some $a \in M_c \cap M_d \setminus 0_i$, $x \in V^a$, and $\delta > 0$, and therefore $\omega_i^c \cdot \delta = m_c((x + \delta \cdot e_i)@a) - m_c(x@a) = m_d((x + \delta \cdot e_i)@a) - m_d(x@a) = \omega_i^d \cdot \delta$. ∎

Claim 3.37 *For any $c, d \in R$, and any $a \in M_c \cap M_d$,*

1. *$\gamma_{c,d} = -\gamma_{d,c}$.*

2. *$\gamma_{a,c} = \gamma_{a,d} + \gamma_{d,c}$.*

proof: Fix any $x \in V^a$. By claim 3.36, $m_c(x@a) = m_d(x@a) - \gamma_{c,d}$, and also $m_d(x@a) = m_c(x@a) - \gamma_{d,c}$. Therefore $\gamma_{c,d} = -\gamma_{d,c}$. Since $m_c(x@a) = \sum_{i=1}^n \omega_i \cdot x_i + \gamma_{a,c}$, $m_d(x@a) = \sum_{i=1}^n \omega_i \cdot x_i + \gamma_{a,d}$, and $m_c(x@a) = m_d(x@a) - \gamma_{c,d}$, the second claim follows. ∎

.6

Lemma 3.3 *Suppose f has connected references. Then there exist constants $\omega_1, \ldots, \omega_n$ and $\{\gamma_a\}_{a\in A}$ such that:*

$$f(v) \in argmax_{a\in A}\{ \sum_{i=1}^{n} \omega_i \cdot v_i(a) + \gamma_a \}$$

for all $v \in V^*$.

proof: Fix any $c^* \in R$. For any $a \in A$ and $x \in V^a$, choose any $c \in R$ such that $a \in M_c$, and define $w(x@a) = m_c(x@a) - \gamma_{c^*,c}$ (where $\gamma_{c^*,c^*} = 0$ by definition, also notice that for any $c \neq c^*$, $c^* \in M_c$ by the second property of R, and so $\gamma_{c^*,c}$ is well defined).

We first claim that for any $c \in R$, $a, b \in M_c$, $x \in V^a$, and $y \in V^b$, $w(x@a) - m_c(x@a) = w(y@b) - m_c(y@b)$. Let $d \in R$ (not necessarily distinct from c) be the reference that determined $w(x@a)$. Therefore: $w(x@a) = m_d(x@a) - \gamma_{c^*,d} = m_c(x@a) - \gamma_{d,c} - \gamma_{c^*,d} = m_c(x@a) - \gamma_{c^*,c}$ (where the second equality follows from claim 3.36, and the third equality follows from claim 3.37, since $c^* \in M_c \cap M_d$). Similarly, let $e \in R$ (not necessarily distinct from c, d) be the reference that determined $w(y@b)$. Therefore: $w(y@b) = m_e(y@b) - \gamma_{c^*,e} = m_c(y@b) - \gamma_{e,c} - \gamma_{c^*,e} = m_c(y@b) - \gamma_{c^*,c}$, and so $w(x@a) - m_c(x@a) = w(y@b) - m_c(y@b)$.

From this it follows that $f(v) \in \text{argmax}_{a\in A, x=v(a)}\{ w(x@a)\}$ for any $v \in V^*$: by contradiction, let $v \in V^*$ be such that $f(v) = a$, but $w(x@a) < w(y@b)$ (for $x = v(a), y = v(b)$). Let $c \in R$ be such that $a, b \in M_c$. Since $w(x@a) - m_c(x@a) = w(y@b) - m_c(y@b)$, it follows that $m_c(x@a) < m_c(y@b)$, contradicting claim 3.24. From this the Lemma immediately follows. ■

3.7 Deferred Proofs

3.7.1 Proof of claim 3.2

Claim 3.2 Any truthful function f has (price) functions $p_i : A \times V_{-i} \to \mathcal{R} \cup \{\infty\}$ such that, for any $v \in V$ and any player i, $f(v) \in argmax_{a\in A}\{v_i(a) - p_i(a, v_{-i})\}$.

proof: Since f is truthful it has price functions $\tilde{p}_i : V \to \mathcal{R}$. Suppose by contradiction that there exists $v \in V$ and $u_i \in V_i$ such that $f(v) = f(u_i, v_{-i}) = a$, but $\tilde{p}_i(v) \neq \tilde{p}_i(u_i, v_{-i})$. W.l.o.g $\tilde{p}_i(v) > \tilde{p}_i(u_i, v_{-i})$. Thus when the other players declare v_{-i}, and the true type of player i is v_i, she will increase her utility by declaring u_i, a contradiction. Therefore we can define the price functions $p_i : A \times V_{-i} \to \mathcal{R} \cup \{\infty\}$, as follows. For any i, $v_{-i} \in V_{-i}$, and $a \in A$, if there exists $v_i \in V_i$ such that $f(v) = a$ we set $p_i(a, v_{-i}) = \tilde{p}_i(v)$, otherwise $p_i(a, v_{-i}) = \infty$.

To see that $f(v) \in argmax_{a\in A}\{v_i(a) - p_i(a, v_{-i})\}$, suppose by contradiction that there exists $v \in V$ such that $f(v) = a$, and $v_i(a) - p_i(a, v_{-i}) < v_i(b) - p_i(b, v_{-i})$. Let $u_i \in V_i$ be the type that determined $p_i(b, v_{-i})$. Therefore if i will declare u_i instead of v_i, when his true valuation is v_i, she will increase his utility, a contradiction. ■

3.7.2 Proof of claim 3.26

We split claim 3.26 to two claims:

Claim 3.38 *Suppose $m : \mathcal{R}_+ \to \mathcal{R}$ is monotonically non-decreasing and there exists $h : \mathcal{R}_+ \to \mathcal{R}_+$ such that $m(x+\delta) - m(x) = h(\delta)$ for any $x, \delta \in \mathcal{R}_+$. Then there exist $\omega \in \mathcal{R}_+$ such that $h(\delta) = \omega \cdot \delta$.*

proof: Let $\omega = h(1)$ (note that $\omega \geq 0$ since m is non-decreasing). First we claim that for any two integers p, q, $h(p/q) = \omega \cdot (p/q)$. Note that $h(1) = m(1) - m(0) = \sum_{i=0}^{q-1} m((i+1)/q) - m(i/q) = q \cdot h(1/q)$. Thus $h(1/q) = (1/q) \cdot h(1)$. Similarly, $h(p/q) = m(p/q) - m(0) = \sum_{i=0}^{p-1} m((i+1)/q) - m(i/q) = p \cdot h(1/q) = (p/q) \cdot h(1) = (p/q) \cdot \omega$. Now we claim that for any real δ, $h(\delta) = \delta \cdot \omega$. Notice that since m is monotonically non-decreasing then h must be monotonically non-decreasing as well. Suppose by contradiction that $h(\delta) > \delta \cdot \omega$. Choose some rational $r > \delta$ close enough to δ such that $h(\delta) > r \cdot \omega$. Since h is monotone and $r > \delta$ then $h(r) \geq h(\delta)$, but since r is rational, $h(r) = r \cdot \omega < h(\delta)$, a contradiction. A similar argument holds if $h(\delta) < \delta \cdot \omega$. ∎

Claim 3.39 *Suppose that $X \subseteq \mathcal{R}^n$ has the property that $x \in X$ and $y \geq x$ implies $y \in X$. Let $m : X \to \mathcal{R}$ be monotonically non-decreasing, and suppose there exist $\omega_1, \ldots, \omega_n$ such that $m(x + \delta \cdot e_i) - m(x) = \omega_i \cdot \delta$ for any i, $x \in X$, and $\delta > 0$. Then there exist $\gamma \in \mathcal{R}$ such that $m(x) = \sum_{i=1}^{n} \omega_i \cdot x_i \ + \gamma$.*

proof: First we claim that for any $x, y \in X$ such that $y_i \geq x_i$ for all i, it is the case that $m(y) = m(x) + \sum_{i=1}^{n} \omega_i \cdot (y_i - x_i)$. Notice that $(y_1, x_2, \ldots, x_n) \in X$, and $m(y_1, x_2, \ldots, x_n) = m(x) + h_1(y_1 - x_1)$. Repeating this step n times we get $m(y) = m(x) + \sum_{i=1}^{n} \omega_i \cdot (y_i - x_i)$.

Now fix some $x^* \in X$. We claim that for any $x \in X$, $m(x) = m(x^*) + \sum_{i=1}^{n} \omega_i \cdot (x_i - x_i^*)$. To see this, choose some y such that $y_i \geq x_i, x_i^*$ for all i. Thus $m(y) = m(x) + \sum_{i=1}^{n} \omega_i \cdot (y_i - x_i)$ and also $m(y) = m(x^*) + \sum_{i=1}^{n} \omega_i \cdot (y_i - x_i^*)$, therefore the claim follows. ∎

3.7.3 Additional Example for Theorem 3.4

This example shows that it is not enough to require that f has a non-degenerate domain – we must require that $\overset{\circ}{f}$ has that:

Example 3.6 *Suppose a CA for two players, who considers three alternatives: c^i: allocating all the goods to player i, and a: allocating half the goods to player 1 and half to player 2. The allocation rule is as follows:*

$$f(v) = \begin{cases} c^2 & v_2(c^2) > \sqrt{v_1(c^1)} \\ c^1 & v_2(c^2) < \sqrt{v_1(c^1)} \\ a & v_2(c^2) = \sqrt{v_1(c^1)} \text{ and } v_2(c^2) = v_2(a) \\ & \quad \text{and } v_1(c^1) = v_1(a) \\ c^1 & \text{otherwise} \end{cases}$$

If a player gets nothing he pays zero. If player 1 is allocated some non-empty bundle he pays $(v_2(c^2))^2$*, and if player 2 is allocated some non-empty bundle he pays* $\sqrt{v_1(c^1)}$*. To verify that this is truthful, notice that the prices of each player does not depend on his deceleration, and that a player always receives an allocation that maximizes his utility under these prices.*

This auction is player decisive and has a non-degenerate domain, but is not almost affine maximizer - indeed, it does not satisfy S-MON: if a *is chosen and player* i *raises his value for every non empty bundle by some* $\delta > 0$*, then* c^i *is chosen, thus contradicting S-MON.*

In the interior of the domain, however, the situation changes. First, $\overset{\circ}{f}$ *satisfies S-MON (and this is not accidental, as theorem 3.1 implies). But, clearly,* $\overset{\circ}{f}$ *is not almost affine maximizer, and this is since the range of* $\overset{\circ}{f}$ *becomes degenerate.*

3.7.4 The hardness of welfare maximization

In this section we prove that CA or MUA that is an exact welfare maximizer (with the appropriate bid language) is NP-hard. For two players, we prove this for any affine maximizer, even with additive constants not equal to zero (this claim is stronger then proving NP-hardness for exact welfare maximization and using Lemma 3.5, since Lemma 3.5 requires the additive constants to be zero). For n players, we prove this for exact welfare maximizers.

Lemma 3.4 *An affine maximizer CA or MUA for two players, with OR bids as the input, that has polynomially bounded constants and full range* [22]*, is an NP-complete problem.*

proof: We show this in two parts. First, we show how to calculate polynomial bounds on the constants ω_i and γ_a (we omit the superscript k when it is clear from the context), in polynomial time. We then use these bounds to describe a reduction of exact-subset-sum to MUA, and of independent-set to CA.

We also assume w.l.o.g that $\omega_1 = 1$ and $\gamma_a \geq 0$ for all $a \in A$ (since f is also an affine maximizer with all the constants multiplied by $1/\omega_1$, and with all the γ_a constants increased by the same value). Denote by c the constant implied from the polynomially bounded constants definition.

By an abuse of notation, we denote by k the alternative that allocates all goods to player 1, by $k-1$ the alternative that allocates $k-1$ goods to player 1 and 1 good to player 2 (for CA, there are several such alternatives - we define below exactly to which one we refer), and by 0 the alternative that allocates all goods to player 2. We need three bounds on the constants, according to the following three claims:

Claim 3.40 *There exists a polynomial time procedure to calculate a bound* $\bar{\gamma} > \max\{(\gamma_a - \gamma_k),\ (\gamma_a - \gamma_{k-1})\}$*, for all* $a \in A$*.*

[22] In fact it is enough to assume that the range contains the following three allocations: allocating all goods to player 1, allocating $k-1$ goods to player 1, and one goods to player 2, and allocating all goods to player 2.

proof: We first show how to find a bound on $\gamma_a - \gamma_k$. Assume that player 1 has the single OR bid $(k:2)$, and 2 has a single OR bid $(1:1)$. We double 1's price (for k) l times, until k is chosen. We denote this as $f(k:2^l \mid 1:1) = k$. Since the auction is affine maximizer with $\omega_1 = 1$, we have: $2^l + \gamma_k \geq \omega_2 + \gamma_a$ for every $a \in A$, therefore $2^l \geq \gamma_a - \gamma_k$, so we can take the bound to be 2^l. To verify that l is polynomial, notice that $2^{l-1} \leq \omega_2 + \gamma_a$ (where $f(k:2^{l-1} \mid 1:1) = a \neq k$), and so $l \leq 4(\log k)^c$, i.e. the number of bits and iterations l is linear in $(\log k)^c$.

To bound $\gamma_a - \gamma_{k-1}$, we use a similar procedure: we iteratively find the minimal r s.t. $f(k-1: 2^r \mid 1:2^r) = k-1$. Notice first that such an r exists: if $2^r > \gamma_a - \gamma_{k-1}$ this implies that $2^r + \omega_2 2^r + \gamma_{k-1} > \omega_2 2^r + \gamma_a$ and so any $a \neq k, k-1$ cannot be chosen, and if $\omega_2 2^r > \gamma_a - \gamma_{k-1}$ then this implies that $2^r + \omega_2 2^r + \gamma_{k-1} > 2^r + \gamma_k$, and so k cannot be chosen. Since $f(k-1:2^r \mid 1:2^r) = k-1$ it follows that $2^r + \omega_2 2^r + \gamma_{k-1} > \gamma_a$, and so $\gamma_a - \gamma_{k-1} \leq 2^r(\bar{\omega}+1)$, where $\bar{\omega}$ is the upper bound on ω_2 that is calculated in the next claim. To verify that r is polynomial, notice that either $2^{r-1} \leq \gamma_a - \gamma_{k-1}$ or $\omega_2 2^r \leq \gamma_a - \gamma_{k-1}$, and hence r is linear in $(\log k)^c$. ∎

Claim 3.41 *There exists a polynomial time procedure to calculate a bound $\bar{\omega} \geq \omega_2$ in polynomial time.*

proof: We start with $f(k:1 \mid k:2)$, and double 2's bid until $f(k:1 \mid k:2^{r_2}) = 0$, that is 2 wins all the goods. Thus, $\omega_2 2^{r_2} + \gamma_0 \geq 1 + \gamma_k$, and so $2^{r_2}\omega_2 \geq \gamma_k - \gamma_0$.

We continue with $f(k:2 \mid k:1+2^{r_2})$ and double 1's bid until $f(k:2^{r_1} \mid k:1+2^{r_2}) = k$. Now, $2^{r_1} + \gamma_k \geq \omega_2(1+2^{r_2}) + \gamma_0$. In particular $\omega_2(1+2^{r_2}) \leq 2^{r_1} + \gamma_k - \gamma_0 \leq 2^{r_1} + 2^{r_2}\omega_2$. We conclude that $\omega_2 \leq 2^{r_1} = \bar{\omega}$.

To verify that r_1 and r_2 are polynomial, notice first that $\omega_2 2^{r_2-1} - \gamma_0 \leq 1 + \gamma_k$, and therefore r_2 is linear in $(\log k)^c$. Similarly, r_1 is polynomial since $2^{r_1-1} + \gamma_k \leq \omega_2(1+2^{r_2}) + \gamma_0$, in fact r_1 is $O(\log k)^{2c}$. ∎

Claim 3.42 *There exists a polynomial time procedure to calculate a bound $\underline{\omega} \leq \omega_2$ in polynomial time.*

proof: We start by iteratively finding $\hat{\beta}$ s.t. $f(k:1 \mid k:\hat{\beta}) = 0$, and then m s.t. $f(k:m \mid k:\hat{\beta}) = k$. Therefore $\gamma_a \leq \omega_2\hat{\beta} + \gamma_0 \leq m + \gamma_k$ (for all $a \in A$), i.e. $0 \leq \omega_2\hat{\beta} - (\gamma_k - \gamma_0) \leq m$. Define $\beta = \hat{\beta} + 1$. It follows that $f(k:1 \mid k:\beta) = 0$. We note that finding m takes $O(\log k)^{2c}$ time (as detailed in the proof of claim 3.41).

Consider the interval $I = [\omega_2\hat{\beta} - (\gamma_k - \gamma_0) \ , \ \omega_2\beta - (\gamma_k - \gamma_0)]$. The length of I is ω_2 but we do not have exactly its two ends. We shall find 2 distinct points in I, then the distance between these 2 points is a lower bound for ω_2. However we have an interval $[1, m]$ that contains I. From this we can find a point $\alpha \in I$, using binary search as follows: Set $l_0 = 1$, $l_1 = m$. Iteratively, let $\alpha = (l_0 + l_1)/2$. (notice that $f(k:\alpha \mid k:\hat{\beta})$ may only be either 0 or k, and the same for β instead of $\hat{\beta}$, since $f(k:1 \mid k:\hat{\beta}) = 0$). Test if $f(k:\alpha \mid k:\hat{\beta}) = 0$: If so, $\alpha \leq \omega_2\hat{\beta} - (\gamma_k - \gamma_0)$.

Therefore set $l_0 = \alpha$ and start another iteration. Otherwise, test if $f(k : \alpha \mid k : \beta) = k$: If so, $\alpha \geq \omega_2\beta - (\gamma_k - \gamma_0)$. Therefore set $l_1 = \alpha$ and start another iteration. Otherwise we have that $\alpha \in I$. Let L be the number of iterations performed. Thus, after $L - 1$ iterations we still have that $I \subseteq [l_0, l_1]$. Therefore $m/2^{L-1} = l_1 - l_0 \geq |I| = \omega_2$, and so the procedure will iterate at most $O(\log k)^{3c}$ times.

To find a second point in I we find ϵ s.t. either $(\alpha + \epsilon) \in I$ or $(\alpha - \epsilon) \in I$. We start with $\epsilon = 1$, and check if either $(\alpha+\epsilon) \in I$ or $(\alpha-\epsilon) \in I$, using the test described above. If not, we decrease ϵ by half and continue. When we stop, we will have $\epsilon > \omega_2/4$ - the distance between α to one of I's ends must be at least $|I|/2 = \omega_2/2$. Thus, since $\omega_2 > 1/(\log k)^c$ (this follows since $\omega_1/\omega_2 < (\log k)^c$), we conclude that finding ϵ took polynomial time as well (we assume that to represents a polynomial proper fractional number we separately store its denominator and numerator as integers occupying together polynomial number of bits). Now, we have an interval of size ϵ that is contained in I. Therefore $\omega_2 = |I| \geq \epsilon$. On the other hand, we have $\epsilon \geq \omega_2/4 \geq 1/4(\log k)^c$, so we can take $\underline{\omega} = \epsilon$.
∎

Reducing Exact Subset Sum to MUA: In order to show that an affine maximizer MUA (with OR bids, denoted below as AMOR) is NP complete we show that it is harder from the following NP-complete problem:

The Problem **Exact Subset Sum** *denoted below as* **Exact***:*

1. *Input: a finite collection of positive integers $S, r_1, r_2, \ldots, r_d$.*
2. *Output: "yes" if there is a sub-collection $I \subseteq \{1, \ldots, d\}$ of r_i's that amounts to S, that is $\Sigma_{i \in I} r_i = S$, and "no" otherwise.*

The reduction: *Given an input $J = (S, r_1, r_2, \ldots, r_d)$ for EXACT, the reduction constructs the following input $\tau(J)$ for AMOR:*

1. *$k = S$.*
2. *compute the bounds $\bar{\gamma}, \bar{\omega}$, and $\underline{\omega}$ (notice that these are computed for this specific k).*
3. *The OR bids for player 1 are: $(r_1,\ c_1 r_1), \ldots, (r_d,\ c_1 r_d)$, where $c_1 = \frac{2 \cdot \bar{\gamma} \cdot \bar{\omega}}{\underline{\omega}}$*
4. *The OR bids for player 2 is: $(1,\ c_2)$, where $c_2 = \frac{\bar{\gamma}}{\underline{\omega}}$.*

And then, answer "yes" if and only if all goods are allocated to player 1.

c_1 can be viewed as the average price of one item for player 1 $c_1 > \omega_2 c_2$ implies that the donation of player 2 to the welfare is always smaller in case both compete the same item. Intuitively, the prices of both players are factored by $\bar{\gamma}$ and so the γ_a's never affect the chosen allocation.

Claim 3.43 *If Exact(J) is "yes" then AMOR($\tau(J)$) allocates all the S items to player 1.*

proof: If Exact(J) is "yes" then there is $I \subseteq \{1,\ldots,d\}$ such that $\Sigma_{i\in I} r_i = S$. The weighted welfare of allocating all items to player 1 is then $c_1 \cdot S + \gamma_S$. We show that in this case any other allocation achieves a sub optimal weighted welfare. The following is an upper bound for the weighted welfare achieved whenever at least one item is allocated to player 2: $c_1(S-1)+\omega_2 c_2+\gamma_{x_0}$, where $\gamma_{x_0} \geq \gamma_a$ for all alternatives $a \neq k$. We argue that $c_1 \cdot S + \gamma_S$ is greater than this upper bound and hence AMOR($\tau(J)$) would allocate all the items to player 1.

$c_1 \cdot S + \gamma_S > c_1(S-1) + \omega_2 c_2 + \gamma_{x_0}$ if and only if $c_1 > \omega_2 c_2 + \gamma_{x_0} - \gamma_S$. Now, $c_1 = \frac{2\cdot\bar{\gamma}\cdot\bar{\omega}}{\underline{\omega}} \geq \frac{2\cdot\bar{\gamma}\cdot\omega_2}{\underline{\omega}} = 2\omega_2 c_2$. Thus, it is suffice show that $2\omega_2 c_2 > \omega_2 c_2 + \gamma_{x_0} - \gamma_S$. This is true since $\omega_2 c_2 = \omega_2 \frac{\bar{\gamma}}{\underline{\omega}} \geq \bar{\gamma} > \gamma_{x_0} - \gamma_S$. ∎

Claim 3.44 *If Exact(J) is "no" then AMOR($\tau(J)$) allocates at least one item to player* 2.

proof: Assume by contradiction that AMOR($\tau(J)$) allocates all the items to player 1. We argue that the allocation of $S-1$ items to player 1 and one item to player 2 has a higher welfare. That is, we argue that $v_1(S)+\gamma_S < v_1(S-1)+\omega_2 v_2(1)+\gamma_{S-1}$. Note that in this case $v_1(S) = v_1(S-1)$, since otherwise this implies that there exists $I \subseteq \{1,\ldots,d\}$ s.t. $\Sigma_{i\in I} r_i = S$. Thus it is suffice to show that $\gamma_S < \omega_2 c_2 + \gamma_{S-1}$, or equivalently $\gamma_S - \gamma_{S-1} < \omega_2 c_2$. But, $\gamma_S - \gamma_{S-1} < \bar{\gamma} \leq \omega_2 c_2$. ∎

This completes the proof w.r.t. MUA.

We now show a reduction from Independent Set to an affine maximizer CA for two players. This is in the spirit of [62], but for two players, and where the CA obtains the weighted optimum, and not simply the optimum.

The Problem **Max. Independent Set***:*

1. *Input: An undirected graph $G = (V, E)$.*
2. *Output: The size of the maximal independent set*[23] *of G.*

The reduction: *Given a graph G, choose some node $u_0 \in V$, and define the graph $G_{-u_0} = G \setminus \{u_0\}$* [24]*. Let x, x_{-u_0} be the size of the max. independent set of G, G_{-u_0} respectively. It is either the case that $x = x_{-u_0}$ or that $x = x_{-u_0}+1$. We first determine which case is it, using the following procedure. We then compute recursively x_{-u_0} and, by that, determine x.*

1. *Construct a set of items s.t. each edge becomes an item. Define the specific bundles (for any $u \in V$): $B_u = \{(u,u') \in E | u' \in V\}$ (i.e. all the edges of u).*
2. *Compute the bounds $\bar{\gamma}, \bar{\omega}$, and $\underline{\omega}$ for this problem instance. Here, the allocation termed $k-1$ is the allocation where player* 1 *receives $E \setminus B_{u_0}$, and player* 2 *receives B_{u_0}.*

[23] a set of vertices $I \subseteq V$ s.t. for any $u, v \in I$, $(u,v) \notin E$

[24] I.e. $V_{-u_0} = V \setminus \{u_0\}$ and $E_{-u_0} = \{(u,u') \in E | u, u' \neq u_0\}$

3. *Define* $c_1 = (2\bar{\omega}\bar{\gamma})/\underline{\omega}$, *and* $c_2 = \bar{\gamma}/\underline{\omega}$.

4. *Construct The OR bids for player* 1*:* $(B_u, c_1)_{u \neq u_0}$ *– i.e.* 1 *values any bundle* B_u *(except* B_{u_0}*) by* c_1*. And the OR bids for player* 2*:* (B_{u_0}, c_2) *(i.e.* 2 *only wants the bundle* B_{u_0}*).*

5. *Execute the CA. If player* 2 *receives all the items in* B_{u_0} *then* $x = x_{-u_0} + 1$*, otherwise* $x = x_{-u_0}$*.*

Before proving the correctness of the reduction, it is useful to notice that the value of player 1 for the entire set of goods is $v_1(E) = x_{-u_0} \cdot c_1$, since there are x_{-u_0} (but no more) disjoint bundles that player 1 is interested in, and each has a value of c_1.

Claim 3.45 *If* $x = x_{-u_0}$ *then player* 2 *will not receive* B_{u_0}*.*

proof: If $x = x_{-u_0}$ then every max. IS for G_{-u_0} contains a neighbor of u_0 (otherwise, we can take this set, add u_0, and get an IS for G with size $x_{-u_0} + 1$). Thus, $v_1(E \setminus B_{u_0}) \leq (x_{-u_0} - 1)c_1 < xc_1 = v_1(E)$. Therefore, if 2 receives all the items of B_{u_0}, then the maximal weighted welfare that can be achieved is $v_1(E \setminus B_{u_0}) + \omega_2 v_2(B_{u_0}) + \gamma_a \leq (x-1)c_1 + \omega_2 c_2 + \gamma_a$ (for some γ_a). Allocating all items to 1 will result in a weighted welfare of $v_1(E) + \gamma_k = xc_1 + \gamma_k$. We claim that the latter term is strictly larger, and by that the claim is proved. But this follows since $c_1 - \omega_2 c_2 > \bar{\gamma}$. ∎

Claim 3.46 *If* $x = x_{-u_0} + 1$ *then player* 2 *will receive all the items of* B_{u_0}*.*

proof: Since $x = x_{-u_0} + 1$, every max. IS for G contains u_0 (if we had a max. IS that does not contain u_0, it will be also a max. IS for G_{-u_0}, contradicting $x_{-u_0} < x$). Let S be any max. IS for G. Since S contains u_0 it does not contain any of its neighbors. Thus it contains $x - 1$ nodes s.t. none of them has an edge in B_{u_0}. Therefore, the set of goods $E \setminus B_{u_o}$ contains $x - 1$ disjoint bundles that player 1 values, and so $v_1(E \setminus B_{u_o}) \geq (x-1)c_1$. Suppose by contradiction that player 2 does not receive B_{u_0}. Then the maximal weighted welfare is at most $v_1(E) + \gamma_a = x_{u_0}c_1 + \gamma_a = (x-1)c_1 + \gamma_a$. But the following allocation has a larger weighted welfare: $v_1(E \setminus B_{u_o}) + \omega_2 v_2(B_{u_o}) + \gamma_{k-1} \geq (x-1)c_1 + \omega_2 c_2 + \gamma_{k-1}$ (this follows since $\omega_2 c_2 > \gamma_a - \gamma_{k-1}$), a contradiction. ∎

The correctness of the reduction now follows from these two claims: if player 2 receives B_{u_0} then it cannot be the case that $x = x_{u_0}$, and therefore $x = x_{u_0} + 1$. And if player 2 does not receive B_{u_0} then it cannot be the case that $x = x_{u_0} + 1$, and therefore $x = x_{u_0}$. ∎

Lemma 3.5 *A exact welfare maximizer MUA or CA for n players, with single minded bids as the input and a full range is an NP-hard problem.*

proof: For CA this was proved by [62]. We prove this for MUA.

Reducing Exact subset sum to MUA: *Given an input* $J = (S, r_1, r_2, \ldots, r_d)$ *for EXACT, the reduction constructs the following input* $\tau(J)$ *for Multi Unit Auction with Single Minded Bids and* n *players (denoted as MU-SMB):*

1. $k = S$*,* $n = d + 1$*.*

2. The Single minded bids for players $i = 1, \ldots, d$ *are:* $(r_i,\ 2 \cdot r_i)$*, i.e. every player desires* r_i *items for a value of* $2 \cdot r_i$*.*

3. The Single minded bid for player $d + 1$ *is:* $(1,\ 1)$*.*

And then, answer "yes" if and only if none of the items is allocated to player $d + 1$*.*

Claim 3.47 *If Exact*(J) *is "yes" then MU-SMB*$(\tau(J))$ *allocates none of the items to player* $d+1$*.*

proof: If Exact(J) is "yes" then there is $I \subseteq \{1, \ldots, d\}$ such that $\Sigma_{i \in I} r_i = S$. Thus allocating r_i items to players $i = 1, \ldots, d$ has total welfare of $2S$. Allocating one item to player $d + 1$ means that at least one player from $i = 1, \ldots, d$ will now be allocated a quantity less than r_i, and thus his value will be zero. Since this player has a value of at least 2 for r_i items, we have that the total welfare when allocating player $d + 1$ a non-empty bundle is at most $2S - 2 + 1 < 2S$. Therefore MU-SMB will not allocate any item to player $d + 1$. ∎

Claim 3.48 *If Exact*(J) *is "no" then MU-SMB*$(\tau(J))$ *allocates at least one item to player* $d + 1$*.*

proof: Let I be the set of players that received a non empty bundle. Suppose by contradiction that $d + 1 \notin I$. Since Exact(J) is "no" then $\Sigma_{i \in I} r_i < S$. Therefore there exists an item who is not allocated, or is allocated to someone that is indifferent to not having it. Delivering this item to player $d + 1$ will increase the welfare by 1, a contradiction. ∎

Chapter 4

Online Ascending Auctions for Gradually Expiring Items

In the model studied throughout this thesis, in order to design algorithms that will lead to desirable outcomes, one needs two ingredients. The first is a protocol, handed in to the participants. The second is a prediction of expected player behaviors (who tune their actions with respect to the given protocol). With these, we get a precise algorithmic description, and we are able to analyze its performance. In principle, for a given protocol, one may expect several different player behaviors. However, the most studied approach is to avoid this difficulty by designing truthful auctions, in which players follow a very specific, simple behavior. By this, the problem reduces to an algorithmic construction that should satisfy one more important requirement – truthfulness. The difficulty of analyzing several different player behaviors is therefore avoided.

Unfortunately, for many settings, truthful algorithms are rare, and a need to find other suitable solutions rises. This difficulty was demonstrated by the impossibility result of the last chapter, and is commonly believed to be much severe. In this chapter[1] we study a problem that forces us to take a different trail, instead of truthfulness. As we (provably) cannot design algorithms for which the players will be expected to take one single behavior, we design auctions for which *many* selfish behaviors lead to an approximately optimal allocation. Thus, our algorithmic construction is of a *family* of algorithms, each one corresponds to a specific combination of players' behaviors, and *all* of them obtain a near optimal outcome. The new concept is that, although players are not expected to follow a specific behavior, but only one out of a set of behaviors, the outcome is still guaranteed to be close to optimal, for any choice the players make. We also provide a game-theoretic rational why will the players limit their choice to this set of actions. We discuss some natural strengthenings of the equilibria notion we use, still keeping this general idea of "set equilibria". We believe that these concepts offer a new way to bypass the inherent difficulties of the truthfulness notion, in a

[1] This chapter is based on a joint work with Noam Nisan.

way that suits the CS worst-case notions.

4.1 Overview of Results

The problem we study is the online allocation of M items that are all identical except that they "expire" at different times: the first item expires at time 1, the second at time 2, and so on. Players arrive over time, and items must be allocated at or before their expiration time. Each player j desires any single item between his arrival time, r_j, and his deadline, d_j, and has a value v_j for receiving the item. All information r_j, d_j, v_j is private to player j, and players act rationally to maximize their utility: the value v_j, if they are allocated an item, minus any payment that they must pay. Our goal is to design a mechanism that maximizes the "social welfare", i.e. to allocate the items so that the sum of values of players that receive an item is maximized.

This model seems applicable to many scenarios in which items are sequentially allocated as time progresses, where both items and players have a finite "life-time". In a computational setting, this model is equivalent to online scheduling of unit length jobs with deadlines. Focusing on the algorithmic question only, and ignoring incentive issues, it is known that the offline problem can be solved exactly in polynomial time and that the online problem has a simple greedy algorithm that achieves a 2-approximation [54, 12], but no online algorithm can achieve an approximation ratio better than the "golden ratio" [46]. Of course, this algorithm requires each player to reveal his true type (value, arrival time, and deadline) and ignores players' strategic considerations.

To incorporate the strategic considerations of the players, our first attempt was to design a truthful mechanism for this problem, in which players are motivated to reveal their true input by incorporating some sophisticated payment scheme. However, this cannot be achieved, as the following strong impossibility result shows:

Theorem: *Any truthful deterministic online mechanism cannot obtain an approximation ratio better than M.*

One could approach this difficulty by adding more assumptions about the players. E.g. one can assume that player values are taken from some known interval $[v_{min}, v_{max}]$. We assumed this in chapter 2, and it was later turned out to be a useful assumption for many online auctions [59, 21, 11, 13, 20, 40, 8, 55, 85, 47]. With this assumption, one can construct a randomized truthful auction with an approximation ratio of $O(log(v_{max} - v_{min}))$ (this can be obtained as a special case of the market clearing algorithm of [21], or by using the general method of [8] to convert online algorithms to truthful online algorithms). To our view, this is a too heavy toll to pay in this model for truthfulness, as a deterministic 2-approximation without any assumptions on players exists once truthfulness is dropped. In addition, the resulting truthful auctions sometimes appear somewhat artificial (e.g. without real competition among the different bidders – all items are sold in a fixed price, determined before the first bidder arrives).

Instead, we will not assume any additional assumptions on players, but will relax the required notion of equilibria: instead of specifying a single tuple of strategies (the equilibrium point) that provides a good approximation ratio, we will specify a large set of strategies with the property that the mechanism will perform well on *any* of them. Our strategic analysis will not be able to pinpoint exactly which of these strategies will be rationally chosen, but rather only that one of them will be – this is enough to guarantee good performance.

At this point we embark with the algorithmic analysis of two (variants of) classic ascending auctions, under a wide class of possible player behaviors. The first auction we consider is a natural adaptation of the iterative auction of Demange, Gale, and Sotomayor [30] (similar offline scheduling auctions were also considered by [99]). *The Online Iterative Auction* constantly maintains a current price p_t and a current winner win_t for every item t. At each time t, each player (in his turn) may place his name as the temporary winner of some item t' (bid on t'), deleting the previous temporary winner, and increasing the price by some fixed small δ (a player can be a temporary winner only for one item). When none of the players wishes to bid, the time t phase ends: item t is sold to player win_t for a price of $p_t - \delta$. At time $t + 1$ the prices and temporary winners from time t are kept, and the auction continues similarly.

In the offline setting, where all players arrive at time 1, [30] show that if all players behave *myopically*, i.e. always bid on the item with the lowest price among those that interest them, then the auction will reach the optimal allocation. Moreover, such myopic behavior is indeed the player's best interest [45]. But what will the player choose, facing this auction in the online setting? This depends on the beliefs of the player about the future: if he fears that new competitive bidders will arrive in the future, he may bid aggressively for earlier items, offering a higher price for them but reducing his risk of future competition. To incorporate such considerations, we call a player *semi-myopic* if he always bids on some item with price lower than his value (not necessarily the item with the lowest price, as the myopic behavior requires). Thus there exist many semi-myopic behaviors, that represent different beliefs. The point is that the auction will obtain near optimal allocation under *any* combination of such behaviors:

Theorem: *If all players are semi-myopic then the Online Iterative Auction achieves a* 3-*approximation of the welfare.*

We prove this by analyzing a family of semi-myopic algorithms, where each such algorithm corresponds to a specific combination of semi-myopic behaviors. One algorithm in the family is the greedy algorithm[2], but the analysis of the entire family is completely different and non-trivial (even the analysis of the "myopic algorithm", which results from the myopic behaviors of the players, is completely different).

[2] which corresponds to the completely aggressive behavior that continues to bid on the current item as long as its price is lower than the player's value (even if the next item has a price of zero).

The second auction we consider is *The Sequential Japanese Auction*: Item t is sold at time t using a classic one-item ascending auction (exact details appear in the sequel). Surprisingly, we show that this auction has a similar structure to the previous one (in our setting). We define a myopic behavior that leads to the optimal allocation in the offline case (when all players arrive at time 1), and, similarly to above, a family of semi-myopic behaviors aimed to capture players' uncertainties about the future. These semi-myopic behaviors again exactly correspond to our family of semi-myopic algorithms, hence a 3-approximation is obtained for every combination of semi-myopic behaviors.

But why should the players play as we expect? We now turn to give a more accurate game-theoretic analysis of the players' behaviors. As truthful auctions do not exist, we instead seek an equilibrium notion that will capture the idea advocated above, i.e. that the best we can do is recommend on a set of strategies, and not on a specific, single strategy.

In the game-theoretic setting, each player is required to choose a strategy $s_i \in S_i$. The resulting payoff of each player i is $u_i(s_1 ... s_n)$[3]. Player i's strategy $r_i \in S_i$ is a best response to a specific combination of strategies of the other players $s_{-i} \in S_{-i}$ if for any $s_i \in S_i$, $u_i(r_i, s_{-i}) \geq u_i(s_i, s_{-i})$. Our notion of "set equilibria" captures the situation where we describe a subset $R_i \subseteq S_i$ of "recommended strategies" (best response strategies) to choose from, instead of describing a single strategy r_i as the equilibrium point:

Definition: *The sets R_i are in* **Set-Nash equilibrium** *if for any player i, and any strategy combination of the other players $s_{-i} \in R_{-i}$, player i has a best response to s_{-i} in R_i.*

Thus the definition requires that a best response to any tuple of recommended strategies of the others may be found within the recommended strategies of player i. This becomes equivalent to regular Nash equilibrium when $|R_i| = 1$ for all i. It should be pointed out that there always exists a trivial Set-Nash equilibrium in which the recommended strategies are the entire set of strategies. Therefore this notion is interesting only when one can guarantee some performance bound whenever players play any one of their recommended strategies, as we do.

Although Set-Nash equilibrium is a weak notion, certainly weaker than regular Nash equilibrium, it seems to us that it carry some weight, especially in computerized environments, in which appropriate protocols and software programs that act "as recommended" are available, and so a deviation would seem to require some effort. In such cases players can be realistically expected to act as recommended unless they have clear incentives to deviate. Such a clear incentive would seem to be absent when the recommended strategies are in Set-Nash equilibrium.

We also provide some discussion on ways to strengthen the basic definition. We describe a hierarchy of four "set equilibria" notions, with a growing strength. While, for our motivating

[3] We actually need the framework of games with incomplete information, in which each player additionally has a type (in our case, a triplet {arrival time, value, and deadline}), and both his strategy and his payoff may depend on this type; this adds few technicalities that we deal with in the sequel, but does not change the spirit of the results, as we describe here.

example, we were able to use only the basic definition, we believe that the complete hierarchy will turn out to be useful for other models, in which truthfulness does not exist, and one wishes to remain within the worst-case framework and to avoid any strong distributional assumptions.

Returning to our model, we show that both our online ascending auctions have Set-Nash equilibria that are all semi-myopic. The main point we arrive at is that players do not have a clear incentive to deviate outside of these sets of recommended strategies; and when they do stay inside the set of recommended strategies, the mechanism obtains a 3-approximation.

Main Theorem: *The Online Iterative Auction and the Sequential Japanese Auction both have a Set-Nash equilibrium which is all semi-myopic, hence results in a* 3*-approximation of the welfare.*

The rest of the chapter is organized as follows. The model and basic definitions are given in section 4.2. In section 4.3 we describe the two online ascending auctions, and show their algorithmic properties by characterizing a family of 3-approximations. Section 4.4 returns to the strategic setting, showing that no truthful auction can achieve an approximation ratio better than M. In section 4.5 we define the new notion of Set-Nash equilibrium and a hierarchy of three stronger notions. We discuss their properties and the relevant literature. In section 4.6 we analyze our auctions accordingly. Section 4.7 describes some useful observations about the offline allocation problem, used as basic building blocks in our proofs.

4.2 Model and Basic Definitions

Items: We wish to sell M identical items with different expiration times. W.l.o.g. we assume that the first item expires at time 1, the second at time 2, and so on. Each item must be sold (and received by the buyer) at or before its expiration time.

Players: The potential buyers of the items (players/bidders) arrive over time. Player i arrives to the market at time $r(i)$, and stays in the market for some fixed period of time, until his deadline $d(i)$. We assume w.l.o.g. that the arrival and departure times are integers[4]. Each player desires only one item (unit demand), that expires no earlier than his arrival time. He must receive it at or before his departure time[5]. Player i obtains a value of $v(i)$ from receiving such an item, otherwise his value is 0. We assume w.l.o.g. that different players have different values[6].

We assume the standard game-theoretic setting: Player i privately obtains his variables $r(i)$, $d(i)$, and $v(i)$, at time $r(i)$. He acts selfishly in order to maximize his own utility: his obtained value minus his price. I.e., a player may arrive at or after his true arrival time, and declare or act as if he has any value, and any deadline.

[4]as actions in a non-integral time point can be deferred to the next integral point with no affect.

[5]Our auctions also fit the more severe restriction that player i cannot get an item $t > d(i)$. E.g., player i cannot attend Saturday's show if he is leaving on Friday, even if he receives the ticket before Friday.

[6]I.e. fix some arbitrary order over players, and set $v(i) \succ v(j)$ iff $v(i) > v(j)$ or $v(i) = v(j)$ and $i \succ j$.

Our goal: is to maximize the social welfare: the sum of (true) values of players that receive an item.

Basic definitions: Player i is *active* at time t if $r(i) \leq t \leq d(i)$, and i did not win any item before time t. Let A_t be the set of all active players at time t. An *allocation* is a mapping of items to players such that, if player i receives item t, then $r(i) \leq t \leq d(i)$. Let X_t be an allocation of items $t, ..., M$. $X_t[d]$ denotes the player that receives item d according to X_t, and $X_t[d_1, d_2] = \cup_{d=d_1}^{d_2} X_t[d]$, the set of players that receive items d_1 through d_2. By a slight abuse of notation we also use X_t as the set of players $X_t[t, M]$. The *value* of X_t is $v(X_t) = \sum_{d=t}^{M} v(X_t[d])$, i.e. the welfare obtained by X_t. A set S of players is *independent* with respect to items $t, ..., M$ if there exists an allocation of (part of) the items $t, ..., M$ such that every player in S receives an item.

The offline allocation problem: The offline problem, in which all players arrive at time 1, is a matroid: a set of *players* is independent if there exists an allocation of (part of the items) to these players. This is known [52] for the unit-demand scheduling problem, which is equivalent to ours. This matroid structure is used extensively in our proofs. See Appendix 4.7 for more details.

4.3 Two Online Ascending Auctions

We first describe online adaptations of two well-known ascending auctions. These have the property that players do not have to choose specific actions for the auction to perform well: a 3-approximation is obtained for a large, reasonable family of behaviors that we term "semi-myopic". Under any such player behaviors, each of our auctions belongs to a general family semi-myopic algorithms, that we characterize. We then show that any semi-myopic algorithm obtains a 3-approximation, and therefore conclude that our auctions lead to a near optimal allocation for any choice of semi myopic behaviors of the players.

In this section, we focus on the algorithmic side. Therefore we give only intuitive justifications for the player behaviors that we assume. For the same reason, we also omit few technicalities about prices and tie-breaking rules from the definitions. All these are detailed when we analyze the strategic properties of our auctions, below.

4.3.1 The Online Iterative Auction

We consider an online adaptation of the iterative auction of Demange, Gale, and Sotomayor [30]:

Definition 4.1 (The Online Iterative Auction (intuitive version)) *The Online Iterative Auction constantly maintains a current price p_t and a current winner win_t for every item t. These are initialized to zero at $t = 0$, and updated according to players' actions at each time t, as follows:*

- *Each player, in his turn, may place his name as the temporary winner of some item t', causing the previous winner to be deleted, and the price to increase by some fixed small δ. A player* *cannot perform this action, and must relinquish his turn, if he is already a temporary winner.*

- *When none of the players that are not temporary winners wishes to place their names somewhere, the time t phase ends: item t is sold to the player win_t for a price of $p_t - \delta$.*
- *At time $t+1$ the prices and temporary winners from time t are kept. If additional players arrive then the auction continues according to the above rules.*

Before analyzing the online auction, it is useful to take a glimpse at the offline case, in which all players arrive at time 1. This is a special case of the unit-demand model studied by [30], [45]:

Definition 4.2 ([30]) *Player i has a* **myopic strategy** *in the iterative auction if, in his turn, he always places his name on the item $t \leq d(i)$ with the minimal price, unless the minimal price $\geq v(i)$, in which case he does not place his name at all.*

Lemma 4.1 ([30], [45]) *If all players are myopic and arrive at time 1 then the online iterative auction obtains the optimal allocation. Furthermore, if all other players are myopic then player i will maximize his utility by playing myopically.*

In the online setting, however, a player might not be completely myopic, depending on his beliefs about the future. For example, he may bid aggressively for the current item, not placing his name on future items at all. This is reasonable if he anticipates tight competition from players that will arrive later on. Viewing this behavior as one extreme, and the completely myopic behavior as the other, it seems that any combination of the two cannot be "ruled-out". On the other hand, a player might choose not to participate at all for some time units – if, for example, there are M high valued players that desire any item 1 through M, but they all do not participate up to time M, then the resulting welfare will be low. As it turns out, this is the only type of behavior we need to exclude:

Definition 4.3 *Player i is* **semi-myopic** *if, in his turn, i places his name on some item t with $p(t) \leq v(i)$ and $r(i) \leq t \leq d(i)$ (not necessarily the one with the lowest price). If there is no such item, i stops participating.*

Theorem 4.1 *If all players are semi-myopic then the online iterative auction achieves almost a 3-approximation: $v(OPT) \leq 3 \cdot v(ON) + 2 \cdot M \cdot \delta$, where OPT, ON are the optimal, online allocations.*

The proof is given in section 4.3.3 below, where we show that, under any semi-myopic behavior, the online iterative auction is a semi myopic algorithm, hence obtains the desired approximation.

4.3.2 The Sequential Japanese Auction

A different possibility is to sell item t at time t using a simple one item ascending auction:

Definition 4.4 (A Japanese Auction) *The (classic, one item) Japanese auction operates as follows: An auctioneer gradually raises a price, starting from* 0*. Each participating player should decide whether to drop out or to stay (once a player drops out, he cannot join again), as the price ascends. The price stops increasing exactly when all players, besides one, have dropped out. The winner is the player that did not drop out, and he pays the price that was reached.*

A natural adaptation of this to the online case is:

Definition 4.5 (The Sequential Japanese Auction (intuitive version)) *The Sequential Japanese Auction sells each item t at time t, separately, using a Japanese auction with one modification: the participants are allowed to observe how many drop-outs occur as the price ascends (and to incorporate this into their drop-out decision).* [7]

As before, it is useful to first consider this auction for the offline case, in which a rather surprising notion of myopic behavior leads to the optimal allocation:

Definition 4.6 *Player i is* **myopic** *in the Sequential Japanese Auction if, in the auction of any time t, (for $r(i) \leq t \leq d(i)$), he drops exactly when either the price reaches $v(i)$, or when there are exactly $d(i) - t$ other players that did not drop yet.*

The logic for dropping when $d(i) - t$ players remain is that at this point the player is assured that there are enough items before his deadline to be allocated to all bidders who are willing to pay the current price.

Lemma 4.2 *If all players are myopic and arrive at time* 1 *then the Sequential Japanese Auction obtains the optimal allocation.*

Our assumption that player have different values is important here. It is not hard to verify that this lemma is actually a special case of theorem 4.6 from the online strategic setting (specifically, it follows from claim 4.7). In this case, a myopic behavior (in the offline case) is a best response when all others are myopic only when using the modified prices of section 4.6.4.

In the online setting, again, players might not play myopically, and may insist on closer items (i.e. stay longer in the auction) if they anticipate much competition in the future. In the extreme, when every player remains in the auction until the price reaches his true value, we actually simulate the simple greedy algorithm, which is a 2-approximation. As before, any behavior between the two

[7] Prices are also modified. The time-t-winner pays the highest price among all time-t'-auctions in which he tied the time-t'-winner. Defining "a tie" is delicate, and requires the players to drop simultaneously. See section 4.6.4.

extremes can cause only a minor performance degradation. All we wish is that players will not drop out "too soon". Indeed, dropping out early in the auction also have disadvantages, as future auctions might be much more competitive, due to new arriving players.

Definition 4.7 *Player i's strategy is* **semi-myopic** *(for the Sequential Japanese Auction) if, at every time t, he drops no later than when the price reaches his value, $v(i)$, and no earlier than when only $d(i) - t$ other players remain in the auction.*

Theorem 4.2 *If all players play semi-myopic strategies then the Sequential Japanese Auction achieves a 3-approximation.*

In a similar manner to the iterative auction above, this theorem is proved by showing that, under any semi-myopic behavior, the Sequential Japanese Auction is a semi myopic algorithm. The proof is given is section 4.3.3 below.

4.3.3 Semi-Myopic Algorithms

For each combination of player strategies, the above auctions are associated with a different algorithm. In order to analyze their performance for a family of strategies, we therefore need to characterize a family of algorithms, that we call semi-myopic algorithms. The main point is that *any* semi myopic algorithm obtains a 3-approximation of the welfare.

Specifically, the *current best schedule* at time t, S_t, is the allocation with maximal value among all allocations of items $t, ..., M$ to the active players, A_t[8]. Define

$$f_t = \{\ j \in S_t \mid S_t \setminus j \text{ is independent w.r.t items } t+1, ..., M\ \}, \tag{4.1}$$

The set f_t contains all players that can receive item t, when one plans to allocate items $t, ..., M$ to the players of S_t (i.e. these are all the potentially *first* players). Now define the critical value at time t, v_t^*, as:

$$v_t^* = \begin{cases} 0 & S_t \text{ is independent w.r.t items } t+1, ..., M \\ \min_{j \in f_t}\{v(j)\} & \text{otherwise} \end{cases}$$

All active players with value larger than v_t^* must belong to S_t, because of its optimality (w.l.o.g the first player in S_t has value v_t^*, and if there was a higher valued player outside of S_t, we could switch between them and increase the value of S_t). Thus, it seems reasonable not to allocate item t to a player with value less than v_t^*, as this player cannot belong to any optimal allocation. Surprisingly, this condition is enough to obtain approximately optimal allocations:

[8]There exists one such allocation, by the matroid structure, and since different players have different values.

Definition 4.8 (A semi myopic algorithm) *An algorithm is semi myopic if every item t is sold at time t to some player j with $v(j) \geq v_t^*$.* [9]

Lemma 4.3 *The Online Iterative Auction with semi-myopic players and the Sequential Japanese Auction with semi-myopic players are both semi myopic algorithms*[10].

proof: We first show the claim for the Online Iterative Auction. If $v_t^* = 0$ then, trivially, $v(win_t) \geq v_t^* - \delta$. Thus assume that $v_t^* > 0$. Let Y_t be the allocation of items to the temporary winning players at the end of time t iterations. According to claim 4.17 in section 4.7.1, f_t is independent w.r.t items $t+1, ..., M$ if and only if $v_t^* = 0$. Therefore f_t is not independent, so there exists some player $j \in f_t$ such that $j \notin Y_t[t+1, M]$. Since $j \in f_t$ then $v(j) \geq v_t^*$. if $j = Y_t[t]$ $(= win_t)$ then we are done. Otherwise, j is not a temporary winner at the end of time t iterations. Since j is semi-myopic, this implies that $v_t^* \leq v(j) < p(t)$. Let $i = win_t$. Since i is also semi-myopic then $v(i) \geq p(t) - \delta$. Therefore $v(win_t) \geq v_t^* - \delta$, as needed. This concludes the claim for the Online Iterative Auction.

For the Sequential Japanese Auction, we show that the winner has value at least v_t^*. Let $j \in f_t$ be the first player in f_t that dropped. If he dropped because the price reached v_j then the winner has value at least v_j, which is at least v_t^*. Otherwise there were at most $d(j) - t + 1$ players that did not drop yet, including j. By claim 4.12[11], $d(j) - t + 1 \leq |f_t|$. Since no player in f_t dropped yet, it follows that every player that did not drop yet belongs to f_t, hence the winner belongs to f_t and has value at least v_t^* by definition. ∎

The family of semi myopic algorithms can be viewed as the entire range between the following two extremes: the first is the greedy algorithm that always chooses the player with maximal value[12], and the second is the "myopic" algorithm that always chooses the player that determined v_t^*. These two extremes are 2-approximations (both were studied in the context of online scheduling [54, 14]). The entire family has only a slightly larger approximation ratio:

Lemma 4.4 *Any semi myopic algorithm is a 3-approximation of the welfare (and this is tight).*

proof: The lemma follows immediately from claim 4.1, as any semi-myopic algorithm has $v(ON[t]) \geq v_t^*$, and therefore $v(OPT) = v(OPT \setminus ON) + V(ON) \leq 2\sum_{t=1}^{M} v_t^* + v(ON) \leq 2 \cdot v(ON) + v(ON) = 3 \cdot v(ON)$. This ratio is tight, as demonstrated by example 4.1. ∎

Claim 4.1 *Any online allocation algorithm that produces an allocation ON have $v(OPT \setminus ON) \leq 2\sum_{t=1}^{M} v_t^*$, where OPT is the optimal allocation.*

proof: We first prove two useful claims:

---

[9] A worst-case approximation cannot sell item t before time t, as a player with high value only for t may appear.

[10] For the online iterative auction, we show that $v(win_t) \geq v_t^* - \delta$, hence it is "almost" a semi-myopic algorithm.

[11] We can assume that there are no ϵ players in f_t, otherwise $v_t^* = 0$ and the claim trivially holds.

[12] Interestingly, this is a special case of the greedy algorithm of [61] for combinatorial auctions with sub-modular valuations. They study the offline case, but it is easy to verify that their algorithm actually works online.

Sub-Claim 4.1 *Let A, B be sets of players, where $A \subset B$. Let S_A, S_B be the allocation with optimal value for A, B, respectively (both are over the same set of items). Then if $j \in A$ but $j \notin S_A$ then $j \notin S_B$*

proof: Assume by contradiction that there exists $j \in S_B \cap A$ but $j \notin S_A$. Notice that S_A and S_B are both independent sets of the matroid over players in B. Notice also that, by the contradiction assumption, $S_A \not\subseteq S_B$, otherwise also $S_A \cup j \subseteq S_B$, implying that $S_A \cup j$ is independent, with players only from A, contradicting the maximality of S_A. Therefore, since $j \in S_B \setminus S_A$, there exists $j' \in S_A \setminus S_B$ such that $S_A \setminus j' \cup j$ and also $S_B \setminus j \cup j'$ are both independent. From the maximality of S_A and since $j \in A$, the first condition implies that $v(j') > v(j)$. But then we obtain a contradiction to the maximality of S_B. ∎

Sub-Claim 4.2 *Let S be the allocation with maximal value over the set of players A and the set of items $t, ..., M$. Assume that S is not independent w.r.t items $t+1, ..., M$. Let $j \in S$ be the player with minimal value such that $S \setminus j$ is independent w.r.t items $t+1, ..., M$. Then $S \setminus j$ has maximal value among all independent sets w.r.t items $t+1, ..., M$ and players in A.*

proof: Denote $S' = S \setminus j$. Suppose by contradiction that the maximal allocation X over items $t+1, ..., M$ has $v(X) > v(S')$. If $j \notin X$ then this contradicts the maximality of S, as $X \cup j$ is independent w.r.t items $t, ..., M$. Otherwise $j \in X \setminus S'$. $S' \not\subset X$, since otherwise $S = S' \cup j \cup X$ contradicting the fact that S is not independent w.r.t items $t+1, ..., M$. Hence there exists $j' \in S' \setminus X$ such that $X \setminus j \cup j'$ and $S' \setminus j' \cup j$ are independent w.r.t items $t+1, ..., M$. Therefore $S \setminus j'$ is independent w.r.t items $t+1, ..., M$, and from the choice of j it follows that $v(j) < v(j')$, contradicting the maximality of X. ∎

We now prove claim 4.1. Fix some scenario, and let OPT and ON be the optimal and online allocations for this scenario. We describe $f : OPT \setminus ON \to \{1, ..., M\}$ such that f is 2 to 1 and $v(j) \leq v^*_{f(j)}$ for any $j \in OPT \setminus ON$. From this, claim 4.1 immediately follows. The function f is defined as follows. Let X_t be the optimal allocation of items $t+1, ..., M$ among players in $OPT[1, t] \setminus ON$. For any $j \in OPT \setminus ON$ (say $j = OPT[t']$), let $t^*_j = min\{t \geq t' \mid j \notin X_t\}$. Then we fix $f(j) = t^*_j$.

Sub-Claim 4.3 *For any $j \in OPT \setminus ON$, $v^*_{f(j)} \geq v(j)$.*

proof: Let $t = f(j)$. First notice that $j \in A_t$: $j \notin ON$, $r(j) \leq t$ as $j \in OPT[1, t]$, and $d(j) \geq t$ since either $j \in X_{t-1}$ or $j = OPT[t]$. Let $m_t \in S_t$ be the player who determined v^*_t, (if $v^*_t = 0$ then set $m_t = null$, so $S_t \setminus m_t = S_t$). We first show that, by subclaim 4.1, $j \notin S_t \setminus m_t$: define A as $OPT[1, t] \setminus ON$ minus all players with deadline $< t$, and $B = A_t$. Clearly $A \subseteq B$. By definition, X_t is optimal for A (over items $t+1, ..., M$). $S_t \setminus m_t$ is optimal for B (over items $t+1, ..., M$): if $m_t = null$ this follows from the optimality of S_t, and if $m_t \neq null$ this follows from sub-claim 4.2. Therefore, since $j \notin X_t$ then $j \notin S_t \setminus m_t$. If $j \neq m_t$ then $j \notin S_t$, and since $j \in A_t$ it follows from the optimality of S_t that $v(j) \leq v(m_t)$. If $j = m_t$ then this trivially holds. Therefore $v(j) \leq v(m_t) = v^*_{f(j)}$, and the claim follows. ■

Sub-Claim 4.4 *f is 2 to 1.*

proof: Fix any time t. We need to show that f maps at most two players to t. Let $j_1 \in X_{t-1}$ be the player with minimal value such that $X_{t-1} \setminus j_1$ is an allocation of items $t+1, ..., M$, and denote $Y = X_{t-1} \setminus j_1$ (if X_{t-1} itself is independent w.r.t items $t+1, ..., M$ then set $Y = X_{t-1}$). If $X_t \subseteq Y$ then by the optimality of X_t it follows that $X_t = Y$ and the claim follows: by definition, f maps only j_1 and $OPT[t]$ to t. Otherwise, $X_t \setminus Y \neq \emptyset$. We first show that $X_t \setminus Y = \{OPT[t]\}$. This is implied by sub-claim 4.1: set $A = OPT[1, t-1] \setminus ON$, and $B = OPT[1, t] \setminus ON$. Since Y is optimal for A (by sub-claim 4.2) and X_t is optimal for B (by definition) it follows that, if $j \in OPT[1, t-1]$ but $j \notin Y$ then $j \notin X_t$, i.e. that $X_t \setminus Y = \{OPT[t]\}$. To conclude, we observe that X_t is a base in the matroid over items $t+1, ..., M$ and players $OPT[1, t] \setminus ON$, and that Y is an independent set of that matroid. Therefore $|Y \setminus X_t| \leq |X_t \setminus Y| = 1$, and thus $|X_{t-1} \setminus X_t| \leq 2$. Since $OPT[t] \in X_t$ then, by definition, the players mapped to t are exactly those in $|X_{t-1} \setminus X_t|$, and the claim follows. ■

This concludes the proof of Claim 4.1. ■

Corollary 4.1 *The Online Iterative Auction is "almost" a 3-approximation of the welfare: for any scenario, $v(OPT) \leq 3 \cdot v(ON) + 2 \cdot M \cdot \delta$.*

The Sequential Japanese Auction is a 3-approximation of the welfare: for any scenario, $v(OPT) \leq 3 \cdot v(ON)$.

proof: For the Online Iterative Auction, for any fixed δ, it follows from claim 4.1 that $v(OPT) = v(OPT \setminus ON) + v(ON) \leq 2 \sum_{t=1^M} v^*_t + v(ON) \leq 2(v(ON) + M \cdot \delta) + v(ON)$, as needed. For the Sequential Japanese Auction, this follows immediately from Lemma 4.3 and Lemma 4.4. ■

The following example shows that the 3-approximation factor analysis is tight:

Example 4.1 *Consider the following scenario for three items. At time 1 arrive two players, j_1 has value ϵ and deadline 1 and j_2 has value 1 and deadline 2. It is easy to verify that $v^*_1 = 0$, and so the online algorithm allocates item 1 to j_1. At time 2 arrive two additional players, j_3 has deadline 2 and j_4 has deadline 3, and both have a value of 1. Therefore $v^*_2 = 1$ and the online algorithm chooses j_4. At time 3 no new players arrive, so item 3 remains unallocated by the online algorithm. Therefore its welfare is $1 + \epsilon$. The optimal welfare is, however, 3, as needed.*

4.4 The Impossibility of Truthful Approximations

We now move from algorithmic considerations to game-theoretic ones, in order to analyze player strategies. Since our goal is to find approximately optimal allocations with respect to the *true* variables of the players, we would prefer to design a truthful mechanism, i.e. an allocation rule with price functions such that, regardless of how the other players act, player i will maximize his utility by declaring his true variables. Let us briefly formally re-state the definition of truthfulness. Let T_i be the domain of all valid player i types/bids $(r(i), v(i), d(i))$, and let $T_{-i} = \times_{j \neq i} T_j$. Consider the allocation constructed by the mechanism upon receiving the type $b_i \in T_i$ from player i and $b_{-i} \in T_{-i}$ from the other players, and let $v(i, b)$ be the value that player i obtains from this allocation, i.e. $v(i)$ if i receives one of his desired items, and 0 otherwise.

Definition 4.9 (Truthfulness) *A mechanism is truthful if there exist price functions $p_i : T_1 \times \cdots \times T_n \to \Re$ such that, for any i, any $b_{-i} \in T_{-i}$, any true type $b_i \in T_i$, and any $\tilde{b}_i \neq b_i$[13]:*

$$v(i, b_i, b_{-i}) - p_i(b_i, b_{-i}) \geq v(i, \tilde{b}_i, b_{-i}) - p_i(\tilde{b}_i, b_{-i}).$$

Such a property is highly desirable, as it guarantees that each player will be motivated to reveal his true type, by an argument similar to the traditional worst-case arguments of Computer Science. Indeed, many recent examples show truthful mechanisms for various models. However, for our model, no such algorithm performs well:

Theorem 4.3 *Any truthful deterministic algorithm for our online allocation problem cannot obtain an approximation ratio better than M.*

proof: Assume w.l.o.g. that a player that does not win any item pays 0. This implies that i's price must not be higher than his value.

Claim 4.2 *Fix some truthful deterministic mechanism with some fixed approximation ratio. Then, for any player i with $r(i) = 1$ there exists a price function $p_i : T_{-i} \to \Re$ such that, for any combination of players that arrive at time 1, b_{-i}:*

- *If $v(i) > p_i(b_{-i})$ then i wins item 1 and pays $p_i(b_{-i})$ (regardless of his deadline).*
- *If $v(i) < p_i(b_{-i})$ then i does not win any item.*

proof: Fix any combination of players that arrive at time 1, b_{-i}. Suppose first that i has deadline equal to 1. For this case, the player becomes one parameter, and by truthfulness there exist a price function according to the claim [2][14].

[13] We actually restrict the possible $\tilde{b}_i$'s such that $\tilde{r}_i \geq r_i$.

[14] The argument essentially states that, if i wins for some $v(i)$ then he ins with any higher value, and pays the same. Therefore there exists a threshold value $p_i^* = p_i^*(b_{-i})$, such that i wins and pays p_i^* if $v(i) > p_i^*$, and looses otherwise.

We now show that this function p_i satisfies the conditions of the claim, regardless of i's deadline. Fix any deadline $d(i)$ of i. If $v(i) > p_i(b_{-i})$ then i must win some item until his deadline, otherwise he can declare $\tilde{d}_i = 1$ and have strictly better utility. But then, if i does not win item 1, the adversary will produce players with higher and higher values, forcing the mechanism not to allocate any item to i in order to maintain the approximation ratio, thus contradicting truthfulness. Therefore i will receive item 1. He will pay $p_i(b_{-i})$ as otherwise, if he pays a higher price, he will declare $\tilde{d}_i = 1$ and will reduce his price, and if he pays less, then if i will have deadline equals 1 he will declare $d(i)$ instead, thus still winning item 1 but paying less. Therefore the function p_i satisfies the first condition.

Suppose now that $v(i) < p_i(b_{-i})$, and suppose there exists a scenario in which i wins one of his desired items. His price must be at most $v(i) < p_i(b_{-i})$. But then, if i had some value larger than $p_i(b_{-i})$ he would have been better off declaring $v(i)$ instead, by this still winning but paying less. Therefore i cannot win any item at all, and the claim follows. ∎

We can now quickly finish the proof of theorem. Fix any price functions $p_i : T_{-i} \rightarrow \Re$. For any $\epsilon > 0$ we will show that there exist player types $b_1, ..., b_M$ such that, for all i, $r(i) = 1$, $d(i) = M$, $1 \leq v(i) \leq 1 + \epsilon$, and $v(i) \neq p_i(b_{-i}$. By the above claim, it follows that the mechanism can obtain welfare of at most $1 + \epsilon$, while the optimal allocation is at least M, and the theorem follows. To verify that such types exist, fix $L > M$ real values in $[1, 1 + \epsilon]$. Choose M values $v(i)$ uniformly at random from these L values. Then, for any given i, $Pr(v(i) = p_i(v(-i))) \leq 1/L$, as the values were drawn i.i.d. Thus, $Pr(\exists i,\ v(i) = p_i(v(-i))) \leq M/L < 1$, hence there exist a choice of values with $v(i) \neq p_i(v(-i))$ for all i. ∎

Remark 1: Although the proof utilizes an extreme scenario with players with very large values, the worst case ratio occurs in common, simple scenarios, as the proof demonstrates. I.e., since the algorithm defends itself against such extremes, it must make wrong decisions even in simple cases.

Remark 2: A simple truthful deterministic M-approximation exists: For any player i, set p_i to be the highest bid received in time slots $1, ..., t$, excluding i's own bid. Sell item t to player i if and only if $v(i) > p_i$, for a price of p_i.

4.5 A Game-Theoretic Framework

Our main motivation at this point is to justify the assumption that players will behave "as expected". We desire a rational justification, i.e. one that shows that expected strategies are, in some sense, utility maximizers for the players. The settings that we are interested in are ones in which "recommended" strategies are indeed to be intuitively expected, and deviating from them would seem to require some effort. In such cases, even rather weak notions of rational justification carry some weight. Such settings include, in particular, situations where computer protocols are announced and appropriate software that acts "as expected" is available. From the onset, we should note that our notions are intended for cases where the existing standard notions of games with incomplete information do not apply: ex-post Nash equilibria do not exist, and no reasonable common prior can be assumed (i.e. we seek "worst-case" notions as in computer science rather than Bayesian notions common in economics).

4.5.1 Set-Nash Equilibria

We first describe the set equilibrium notions for games with complete information, and then explain how to extend them to a setting of incomplete information, which suits our needs here. There are n players, where each player i has a strategy space S_i. The outcome of the game is given by the n utility functions $u_i : S \rightarrow \Re$ where $u_i(s_i, s_{-i})$ denotes i's payoff he plays strategy s_i and the others play the strategy tuple s_{-i}. The basic assumption is that, given that the other players play s_{-i}, player i will choose a strategy $s_i \in \text{argmax}\{u_i(s_i, s_{-i})\}$.

In our setting, a set of recommended strategies, R_i, is defined for each. The motivating scenario is where it is known that if all players i play recommended strategies then the outcome is "good" in some sense. E.g., in our case, the obtained social welfare approximates the optimal one (therefore we do not put any emphasis on the minimality of the sets; see the discussion on related literature below for details). We would like to capture the notion that the sets R_i are in equilibrium. In other words, formalize when can it be said that given that other players $j \neq i$ all play strategies in R_j, then player i also rationally plays some strategy in R_i.

We give four definitions below, all maintain the spirit of this "set equilibrium" notion, in order of increasing strength. Some of these notions have been defined before in the literature in the context of complete information games – we discuss this below in subsection 4.5.1. All of the following definitions behave the same on the two extreme cases: When each R_i is a singleton set ($\forall i | R_i = 1$) then they are equivalent to Nash equilibrium. When R_i is the entire strategy space ($R_i = S_i$) then they are trivially satisfied.

Definition 4.10

1. *We say that R_i are in "Set-Nash equilibria" (in the pure sense) if for every i, every $s_{-i} \in R_{-i}$, and every $s_i \in S_i$ there exists $r_i \in R_i$ such that $u_i(r_i, s_{-i}) \geq u_i(s_i, s_{-i})$. I.e. for very tuple of recommended strategies there exists a best response strategy in the recommended set.*

2. *We say that R_i are in "Set-Nash equilibria" (in the mixed sense) if for every i, for every series of distributions π_j on R_j for all $j \neq i$, and every $s_i \in S_i$ there exists $r_i \in R_i$ such that $u_i(r_i, s_{-i}) \geq E_{\{\pi_j\}_{j \neq i}}[u_i(s_i, s_{-i})]$. I.e for every series of distributions on the recommended strategies of the other players there exists a best response in the recommended set. This definition captures an expected-utility scenario,* over all possible priors.

3. *We say that $\{R_i(\cdot)\}$ are in "Set-Nash equilibria" (in the mixed-correlated sense) if for every i, for every π on $s_{-i} \in R_{-i}$, and every $s_i \in S_i$, there exists $r_i \in R_i$ such that $u_i(r_i, s_{-i}) \geq E_\pi[u_i(s_i, s_{-i})]$. This definition extends the previous one in the sense of allowing the other players to correlate strategies.*

4. *We say that R_i are in "Set-Domination equilibria" if for every i, and every $s_i \in S_i$ there exists $r_i \in R_i$ such that for every $s_{-i} \in R_{-i}$, we have that $u_i(r_i, s_{-i}) \geq u_i(s_i, s_{-i})$. I.e. for every unrecommended strategy, there is a recommended strategy that that is not worse-off, as long as others act as recommended.*

These definitions extend to games with incomplete information in a straightforward way. Each player i has a privately known type (input) $t_i \in T_i$. No probability distribution is assumed on $T = T_1 \times ... \times T_n$. The utility functions now depend on the player's type, as well ($u_i : T_i \times S \to \Re$, where $u_i(t_i, s_i, s_{-i})$ denotes i's payoff when his type is t_i, he plays strategy s_i and the others play the strategy tuple s_{-i}). The set of recommended strategies may now depend on the player's type, i.e. $R_i : T_i \to 2^{S_i}$, and we denote also $R_i(*) = \cup_{t_i \in T_i} R_i(t_i)$. All four definitions are modified so that the condition specified should now hold for all possible types t_i. In addition, the best response r_i must exist in player i's recommended set *according to his true type*, $R_i(t_i)$, and this r_i should be a best response to any tuple of strategies out of $R_{-i}(*)$ (i.e. the requirement holds for all possible type realizations of the other players). For example, the first definition is altered so that the set functions $R_i(\cdot)$ are in "Set-Nash equilibria" (in the pure sense) if for every i, every t_i, every $s_{-i} \in R_{-i}(*)$, and every $s_i \in S_i$ there exists $r_i \in R_i(t_i)$ such that $u_i(t_i, r_i, s_{-i}) \geq u_i(t_i, s_i, s_{-i})$.

In all definitions, we require the existence of a pure recommended strategy $r_i \in R_i(t_i)$. One can in principle relax the definition to allow r_i to be a mixed strategy, i.e. a probability distribution on $R_i(t_i)$. It is easy to verify that this does not change the first three definitions (the best mixed strategy is always a pure one), while for the Set-Domination definition, this will weaken it to become equivalent to Set-Nash for correlated strategies (using von-Neuman's max-min principle, in the sense of Yao showing equivalence between distributional complexity and probabilistic one).

The first three definitions suffer from the same caveats of regular Nash-equilibria, in particular noting that inequalities are not strict. Thus for example one can have any of these equilibria in strictly dominated strategies. More refined notions may require that strategies in $R_i(t_i)$ are undominated, or even that all undominated best-responses are in $R_i(t_i)$.

Another refinement is to show that the best response is in $R_i(t_i)$ even when other players' strategies reside in a wider class than $R_{-i}(*)$ (this may be interesting also when i assumes only partial rationality of the other players). One may formally define the wider set of acceptable strategies $A_i \subseteq S_i$, where $R_i(*) \subseteq A_i$, and replace the quantification of $s_{-i} \in R_{-i}(*)$ in the definition with $s_{-i} \in A_{-i}$.

In this work we use the basic definition (and drop the qualifier "in the pure sense" hereafter). In addition, all our Set-Nash strategies are undominated, and one can show that they are best response to a set of acceptable strategies wider than $R_{-i}(*)$. Indeed, we feel that an interesting problem we leave open is to find Set-Nash equilibria that contain best responses for mixed recommended strategies as well.

Related notions in the Game-Theory literature

The game theory literature defines and discusses similar notions to the above set equilibria notions for games with complete information. Most study is done on the existence and uniqueness of minimal such equilibria. We are not aware of any study in the setting of implementation theory, that examines such notions with respect to the quality of the outcome they yield.

Shapley [95] defines a notion of "a saddle" for two-person zero-sum games, which is almost the same as the Set-Domination notion (but the inequalities are strict). Shapley shows that there always exists a unique minimal saddle in a zero-sum game (the strictness of the inequalities are crucial for this), but does not address the quality of the obtained outcome. Duggan and Le Breton [33, 35] define a "mixed saddle", which allows mixed strategies in the definition. As we noted above, this is actually equivalent to the definition of Set-Nash in the correlated sense. Their results are again for the complete information case (mainly for zero-sum games, and for voting procedures). Duggan and Le Breton [34] develop a general approach to construct "choice sets". They require both an "outer stability", which resembles our logic of constructing a set equilibria, and also require an "inner stability", in order to have a minimal choice set. We replace this inner stability with a requirement on the quality of the outcome. This of-course can be done in our context of implementation theory, but not in their context of normal form games with complete information. Basu and Weibull [15] study sets of strategies that contain all their best replies (a "curb" set). Voorneveld [98] defines a "prep-set", which is equivalent to our definition of Set-Nash in the mixed sense. He studies the existence of a minimal such set in games with complete information. Bernheim [16] considers "point rationalizability", in which only pure strategies are considered. Although the motivation behind rationalizable strategies is different than the motivation of set equilibria, it is interesting to parallel the shift from rationalizable strategies (in which mixed strategies are allowed) to point rationalizabilty, to the shift from Set-Nash in the mixed sense to Set-Nash in the pure sense, that we make.

4.5.2 Implementation in Set-Nash equilibria

As our context is the framework of implementation theory, we wish to formally specify how the notion of Set-Nash equilibria fits in, in parallel to classical results. We do this for the basic definition of Set-Nash, but the entire discussion follows through for all four definitions in an immediate way. The setting contains a set of outcomes/alternatives, A, from which we have to choose one outcome. The choice depends on the players types $t \in T$, according to some social choice correspondence $F : T \rightarrow 2^A$. In our example, A is the set of all valid allocations of items to players, and $F(t)$

outputs all allocations that are 3-approximations w.r.t t. This social correspondence represents the fact that our goal is to obtain a 3-approximation of the welfare, and any allocation that obtains this will satisfy us. All the classic definitions from implementation theory can be adapted to our Set-Nash definition:

Definition 4.11 *Given $F : T \to 2^A$, an implementation in Set-Nash equilibrium is a mechanism with strategy sets $S_1, ..., S_n$, and an outcome function $g(s_1, ..., s_n) \in A$, such that there exists a Set-Nash equilibrium $\{R_i(\cdot)\}_i$ that satisfies that $g(s) \in F(t)$ for* **all** *$s \in R(t)$.*

Notice that we cannot hope to require that *all* equilibria will produce results according to F, as there always exists the trivial set-equilibria that contains all strategies.

The celebrated revelation principle states that whenever we can implement a social function in some equilibrium, we can also implement it using a direct revelation implementation, in which the strategy space of the players is simply to reveal their type. For our "set equilibrium" notion, we can have an "extended direct revelation" implementation which is "extended truthful":

Definition 4.12 *An implementation is an "extended direct revelation implementation" if the strategies of the players are of the form (t_i, l_i), where $t_i \in T_i$, and l_i represents any additional information.*

An extended direct revelation implementation is "extended truthful" (in Set-Nash equilibrium) if there exists a Set-Nash equilibrium in which $R_i(t_i) = (t_i, *)$, i.e. the player declares his true type in every one of his recommended strategies.

Proposition 4.1 (An extended revelation principle) *Every function $F : T \to 2^A$ that can be implemented in Set-Nash equilibrium can be implemented by an extended truthful implementation.*

proof: Given an implementation M to F in Set-Nash equilibrium, we build an extended truthful implementation M', that encapsulates M, as follows. Let $R_i(t_i)$ be the recommended strategies of M. Then the strategy space of a player in M' is to specify his type t_i, and a strategy in $R_i(t_i)$. The mechanism then uses M' with the specified strategies to determine the result. It is immediate to verify that the sets $R'_i(t_i) = \{(t_i, s_i) \mid s_i \in R_i(t_i)\ \}$ are indeed a Set-Nash that fits the definition. ∎

It is worth pointing out that our auctions, which are not direct revelation, have an interesting extended direct revelation counterpart – we describe this in section 4.6.1 below.

4.5.3 Ignorable Extensions of Games

This section formalizes a concept used in our proof of main theorem, below. In the proof, we first describe an extended truthful mechanism that implements a 3-approximation, and then show that each of our ascending auctions has "inside" it a semi-myopic mechanism. In this section, we describe this type of a building block more generally.

When one actually attempts to implement a game as a software protocol, it often turns out that the set of strategies that is available to players has grown: the protocol that allows a player to play any strategy in S_i turns out to enable also other strategies, informally ones that are "locally" like a valid strategy $s_i \in S_i$, but that do not correspond to any single valid strategy. These new strategies may open up new strategic behaviors. We will specify the requirements from such implementations needed to maintain Set-Nash equilibria.

Formally, given a game with incomplete information $G = (T, S, u)$ (where T, S, u are the players' type space, the players' strategies, and the players' utility functions, as described in section 4.5.1 above) we say that $\bar{G} = (T, \bar{S}, \bar{u})$ is an extension of G if $S_i \subseteq \bar{S}_i$ for all i and $\bar{u}_i(t_i, s) = u_i(t_i, s)$ for all $t_i \in T_i$ and $s \in S$ (i.e. $\bar{u}$ when restricted to S is identical to u).

Clearly a strategy that was best response in G need not be a best response in $\bar{G}$ since the new strategies $\bar{S}_i \setminus S_i$ may be better. "Ignorable" extensions of G will not allow such better strategies:

Definition 4.13 *We say that $\bar{G}$ is an* **ignorable extension** *if for all i, all $t_i \in T_i$, all $s_{-i} \in S_{-i}$ and all $\bar{s}_i \in \bar{S}_i$ there exists $s_i \in S_i$ such that $u_i(t_i, s_i, s_{-i}) \geq u_i(t_i, \bar{s}_i, s_{-i})$. I.e. if others play an original strategy then I have an original strategy which is a best response.*

Proposition 4.2 *If $\{R_i(\cdot)\}$ are a Set-Nash equilibrium of G and $\bar{G}$ is an ignorable extension of G then $\{R_i(\cdot)\}$ are a Set-Nash equilibrium of $\bar{G}$.*

We point out that, although these notions where related to the notion of Set-Nash equilibria, we can, in an immediate and similar way, define ignorable extensions to any one of the other three definitions of equilibria.

4.6 A Strategic Analysis of our Auctions

4.6.1 Semi-Myopic Mechanisms

We now devise an extended direct revelation auction with our two basic building blocks: it has a Set-Nash equilibrium, and, for these equilibrium strategies, the auction is a semi-myopic algorithm. We will later use this to show that both our online ascending auctions also have such a structure.

Definition 4.14 *We define the* **semi-myopic mechanism** *as follows:*

Strategy space: *Each player declares, as he arrives, his value, his deadline, and a tentative deadline between his arrival time and his deadline. The variable $d(i, t)$ holds i's tentative deadline if t is not larger than his tentative deadline, otherwise $d(i, t)$ equals his final deadline.*

Winner determination at time t: *Let A_t, S_t, and f_t be the natural parallels of the notions in definition 4.8, where the deadline of each player in A_t is $d(i, t)$. The mechanism allocates item t to some player in f_t (this choice may depend on the contents and structure of A_t, S_t, and f_t).*

Prices: *For each player i, the mechanism maintains a tentative price for each time t, $p_t(i)$, as follows: If $i \notin S_t$ then $p_t(i) = 0$. For any $i \in S_t$, let*

$$c_t(i) = max\{v(j) \mid j \in A_t \setminus S_t,\ S_t \setminus i \cup j \text{ is independent w.r.t items } t, ..., M\}. \tag{4.2}$$

For any $i \in f_t$, the mechanism sets $p_t(i) = c_t(i))$. For any $i \in S_t \setminus f_t$, the mechanism may set any price $p_t(i) \in [0, c_t(i)]$. The winner i of time t pays $max_{r(i) \leq t' \leq t} p_{t'}(i)$.

The recommended strategies: *In a recommended strategy, i declares his true value and deadline at time $r(i)$, and may declare any tentative deadline.*

Lemma 4.5 *When all players play recommended strategies according to their true types then the semi-myopic mechanism is a semi-myopic algorithm.*

proof: Fix any time t. We need to show that the mechanism chooses a player with value at least v_t^*. Let f_t^{true} be the "true" one, i.e. the relevant set computed with the true player deadlines, and let S_t, f_t be the actual sets computed by the mechanism according to the declared tentative deadlines. If $f_t^{true} \subseteq S_t$ then, by the prefix construction process described in section 4.7.1, and since tentative deadlines are not larger than true ones, $f_t \subseteq f_t^{true}$, and the claim follows. Otherwise there is some $j \in f_t^{true} \setminus S_t$, and so for every $i \in f_t$, $v(i) \geq v(j) \geq v_t^*$ (recall that the declared values are the true ones), as claimed. ∎

Lemma 4.6 *For any player i, and any $s_{-i} \in R_{-i}(*)$, i has a best response to s_{-i} in $R_i(t_i)$.*

proof: Let σ be the scenario in which all players besides i play s_{-i}, and i does not show up at all. Let

$$t^* = argmin_{r_i \leq t \leq d_i}\{v_t^*(\sigma)\}. \tag{4.3}$$

Notice that player i can win and pay exactly $v_{t^*}^*$ by arriving at time t^*, declaring any value larger than $v_{t^*}^*$, and a deadline equals to t^*.

Claim 4.3 *t^* and $v_{t^*}^*$ does not depend on the choice of the winner $i \in f_t$ of time $t \in [r(i), d(i)]$ (where the winners prior to time $r(i)$ are fixed).*

proof: By contradiction, assume that there exist two different scenarios, σ_1, σ_2, that differ only in the choice of the winners (notice that the f_t's themselves might become different during the scenario run due to a previous choice of different winners). Let $v^*(\sigma_i) = min_{r(i) \leq t \leq d(i)}\{v_t^*(\sigma_i)\}$, and let t_i^* be the minimal time in which $v^*(\sigma_i)$ is obtained.

We first assume w.l.o.g. that $v_{t_2^*}^*(\sigma_1) > v_{t_2^*}^*(\sigma_2) = v^*(\sigma_2)$. Let us justify this. If $v^*(\sigma_1) \neq v^*(\sigma_2)$ then w.l.o.g. $v^*(\sigma_1) > v^*(\sigma_2)$ and therefore also $v_{t_2^*}^*(\sigma_1) > v^*(\sigma_2)$. If $v^*(\sigma_1) = v^*(\sigma_2)$ then, by the contradiction assumption, $t_1^* \neq t_2^*$, so w.l.o.g. $t_2^* < t_1^*$. Therefore $v_{t_2^*}^*(\sigma_1) > v^*(\sigma_1) = v^*(\sigma_2)$, as needed. Notice also that from this it follows that $t_2^* > r(i)$, as $A_{r(i)}(\sigma_1) = A_{r(i)}(\sigma_2)$.

Since $v^*_{t^*_2}(\sigma_1) > v^*_{t^*_2}(\sigma_2)$ then $f_{t^*_2}(\sigma_1) \neq f_{t^*_2}(\sigma_2)$, and therefore, by the prefix properties of section 4.7.1, $f_{t^*_2}(\sigma_1) \not\subseteq S_{t^*_2}(\sigma_2)$. Fix some $j \in f_{t^*_2}(\sigma_1) \setminus S_{t^*_2}(\sigma_2)$. Since $v(j) > v^*_{t^*_2}(\sigma_1)$ it follows that $j \notin A_{t^*_2}(\sigma_2)$. This implies that, in σ_2, j is the winner of some time $t' < t^*_2$, i.e. $j \in f_{t'}(\sigma_2) = P_{t'}(t', \sigma_2)$. As $d(j) \geq t^*_2$ then $P_{t'}(t', \sigma_2) = P_{t'}(t^*_2, \sigma_2)$. By claim 4.16 of section 4.7.1, it therefore follows that $v^*_{t'}(\sigma_2) \leq v^*_{t^*_2}(\sigma_2)$, contradicting the choice of t^*_2. ∎

Claim 4.4 *i's price in any strategy s_i is at least $v^*_{t^*}$ (where the other players play s_{-i}).*

proof: Recall that σ denotes the scenario in which i does not show up at all. Let σ' be the scenario in which i plays some strategy s_i and the others play s_{-i}. Denote by t_0 the minimal t with $i \in f_t(\sigma')$. We claim that there exists a scenario σ'', that differs from σ only in the choice of winners in f_t, such that $A_{t_0}(\sigma') = A_{t_0}(\sigma'') \cup i$. This follows by the an inductive argument: At time $t < t_0$, $A_t(\sigma') = A_t(\sigma'') \cup i$. Since $i \notin f_t(\sigma')$ then, by sub-claim 4.16, $f_t(\sigma') = f_t(\sigma'')$. Choose the winner in σ'' to be the winner of σ'. Therefore $A_{t+1}(\sigma') = A_{t+1}(\sigma'') \cup i$, and the inductive claim follows.

Now, at time t_0, since $i \in f_{t_0}(\sigma')$ then, by sub-claim 4.16, there exists some $j \in A_{t_0}(\sigma') \setminus S_{t_0}(\sigma')$ such that $j \in f_{t_0}(\sigma)$. Therefore i's price is at least $v(j) \geq v^*_{t_0}(\sigma'') \geq v^*_{t^*}(\sigma)$ (where the last inequality follows by claim 4.3, and the lemma follows. ∎

Claim 4.5 *The (recommended) strategy of arriving at time $r(i)$, declaring the true value and deadline and declaring a tentative deadline equals to t^* is a best response of i against s_{-i}.*

proof: If $v(i) \leq v^*_{t^*}$ then i cannot possibly gain positive utility, as claim 4.4 shows, and indeed any recommended strategy will not allocate any item to i.

If $v(i) > v^*_{t^*}$ then, if player i arrives at time t^* and declares tentative deadline t^* he will win item t^* for a price of $v^*_{t^*}$. Let σ be the scenario in which i does not show up at all and σ' be the scenario in which i arrives at $r(i)$ and declares tentative deadline t^*. We claim by induction that, for any $t < t^*$, the winners of σ and σ' are identical, and that i's tentative price is at most $v^*_{t^*}$. Therefore i will win item t^* for a price of $v^*_{t^*}$, and the claim follows. For any $t < t^*$, we have by sub-claim 4.16 and the construction of t^* that $min_{j \in P_t(t^*, \sigma)}\{v(j)\} \leq min_{j \in P_{t^*}(t^*, \sigma)}\{v(j)\} = v^*_{t^*}(\sigma) < v^*_t(\sigma)$. By the maximality of $S_t(\sigma')$ it follows that, in σ', i replaces the minimal player in $P_t(t^*, \sigma)$, therefore $f_t(\sigma) \subseteq S_t(\sigma')$, and so $f_t(\sigma) = f_t(\sigma')$. By claim 4.3 we can assume w.l.o.g. that the winner has not changed in the transition from σ to σ'. i's price at time t is (at most, as the mechanism has some freedom in setting this) $min_{j \in P_t(t^*, \sigma)}\{v(j)\} \leq v^*_{t^*}$, and therefore i's final price was not affected as well. ∎

And as an immediate result, we get:

Theorem 4.4 *The semi-myopic mechanism Set-Nash implements a 3-approximation of the welfare.*

4.6.2 Bad Examples

We would like to show, by an example, that the recommended strategies of the semi-myopic mechanism do not contain best responses to mixed strategies. We will only show it for correlated mixed strategies, i.e. it does not contain a best response against a distribution over all $R_{-i}(*)$. We will start with a basic problematic scenario, and then add to it a second scenario, together obtaining the counter example.

The basic problematic scenario demonstrates that a player might be tempted to arrive later, or to declare a deadline higher than his true one, although this is not his best response:

Example 4.2 *Consider the following scenario, where (v, d) denotes a player with value v and deadline d):*

- *At time 1 arrive players $(\epsilon, 1), (x_1, 4), (x_2, 4), (x_3, 4), (x_4, 4)$.*
- *At time 2 arrive players $(y_1, 2), (y_2, 3)$.*
- *At time 3 arrive players $(z_1, 5), (z_2, 5)$.*
- *At time 4 arrives a (very large) player $(z_3, 4)$.*

where the values satisfy: $\epsilon < x_2, x_3 < y_1 < x_1 < z_1, z_2 < y_2 < x_4 < z_3$.

If all players declare their true value and tentative deadline equals to their true deadline, a semi-myopic mechanism can choose the winners (first to last) x_1, y_1, y_2, z_3, z_1. So player x_4 looses. However, if he delays his arrival to time 2, or, equivalently, declares a deadline of 5, the winners will be $\epsilon, y_2, x_4, z_1, z_3$, so x_4 will win, with price x_1. Notice, however, that this is not his best response. His best response, to arrive at time 1 and declare tentative deadline 1, is still of-course recommended.

Example 4.3 *Let scenario 1 be the scenario of example 4.2, where we consider the decisions faced by x_4, and scenario 2 be as follows:*

- *At time 1 arrive player $(x, 1)$ and our player $(x_4, 4)$.*
- *At time 2 arrive player $(x, 2)$.*
- *At time 3 arrive player $(x, 3)$.*

(where $x = x_4 - \epsilon$). The best response of x_4 to scenario 1 is to arrive at time 1 and declare deadline 1. The best response to scenario 2 is to arrive at time 1 and declare a deadline of 4 (thus winning item 0 with price 0). Now suppose that player x_4 knows/estimates that both scenarios have probability half. Then, a quick calculation shows that if x_4 plays some recommended strategy (and thus arrives at time 1) with tentative deadline lower than 4, then with probability half (for scenario

2) he will win of the items 1 *to* 3 *with a resulting utility (i.e. value minus price) of* ϵ*. If his tentative deadline will be* 4 *then with probability half (for scenario* 1*) he will lose. Therefore, any recommended strategy has resulting utility at most* $(x_4 + \epsilon)/2$*. However, if* x_4 *will arrive at time* 2 *and will declare deadline* 4*, a non-recommended strategy, his resulting utility will be half times* $x_4 - 0$ *(for scenario* 2*) plus half times* $x_4 - x_1$*, better than* $(x_4 + \epsilon)/2$ *for small enough* ϵ*.*

4.6.3 The Online Iterative Auction has a Set-Nash Equilibrium

We now show that our Online Iterative Auction is an ignorable extension of a semi-myopic mechanism, thus having a Set-Nash equilibrium which approximates the welfare, according to theorem 4.4. For this, we need to refine our intuitive definition:

Definition 4.15 (The Online Iterative Auction) *We apply the following modifications to Def. 4.1:*

Prices: *The auction maintains a tentative price* $p_t(i)$ *for each player* i *at time* t*, as follows: if* i *is a tentative winner at the end of the iterations of time* t *then* $p_t(i)$ *equals to the tentative price of* i*'s item, otherwise* $p_t(i) = 0$*. The winner* i *of time* t *pays* $max_{r(i) \le t' \le t}\{p_{t'}(i)\}$*.*

Recommended strategies: i*'s strategy is recommended if* i *chooses a tentative deadline* $d \le d(i)$*, plays myopically (as in Def 4.2) with value* $v(i)$ *and deadline* d *in all times* $r(i) \le t \le d$*, and plays myopically with value* $v(i)$ *and deadline* $d(i)$ *in all times* $t > d$*.*

It is not hard to verify that these recommended strategies are semi-myopic.

Theorem 4.5 *The Online Iterative Auction is an ignorable extension of a semi-myopic mechanism.*

Corollary 4.2 *The Online Iterative Auction Set-Nash implements a* 3*-approximation of the welfare.*

Proof (of theorem): We first prove that the iterative auction is an extension of a semi-myopic mechanism. We will then show that this extension is ignorable.

Claim 4.6 *If all players* i *play strategies in* $R_i(*)$ *then the iterative auction is a semi-myopic mechanism.*

proof: We need to map every recommended strategy of the iterative auction to a strategy of the semi-myopic mechanism, such that the result of the iterative auction (winners plus payments) will match the criteria of a semi-myopic mechanism. This is done as follows. At time t, map every players that plays myopically with (v, d) to a type (v, d), and denote this set of types as A_t. Let S_t be the optimal allocation of items $t, ..., M$ to the players of A_t. All we need to show is that the iterative auction selects a winner from f_t and sets correct payments. In what follows, we use the notion of a prefix and the claims of section 4.7.1. Let $Y[t, ..., M]$ and $p_t[t, ..., M]$ be the tentative allocation and prices of the iterative auction with the myopic strategies, at the end of time t. For any $d \geq t$, let $P_Y(d)$ be the appropriate prefix of Y, according to definition 4.19. Define $l(d) = min\{d' \geq t \mid P_Y(d') = P_Y(d)\}$, and

$$c_t(d) = max\{v(j) \mid j \in A_t \setminus Y \text{ and } d(j) \geq l(d) \}.$$

(notice that, by abuse of notation, we have defined both $c_t(d)$ for an item $d \in \{t, ..., M\}$, and $c_t(i)$ for a player i. Those are two differently defined terms, although we will see below that they are equal, for $d = Y[i]$).

Sub-Claim 4.5 *$p_t(d) \geq c_t(d)$.*

proof: Fix any $j \in A_t \setminus Y$ with $d(j) \geq l(d)$. If $d(j) \geq d$ then j has positive value for receiving d. Since j is myopic, it therefore follows that $p_t(d) \geq v(j)$. If $d(j) < d$, then since $P_Y(l(d)) = P_Y(d)$, By the construction of $P_Y(l(d))$, since it is equal to $P_Y(d)$, then there exist players $i_1, ..., i_k$ and items $t_1, ..., t_k$ such that, for any index $x \in \{1, ..., k\}$, $i_x = Y[t_x]$, $d(i_x) \geq t_{x+1}$, $t_1 \leq l(d)$, and $t_k = d$. Since $d(i_x) \geq t_{x+1}$ it follows that $p_t(t_x) \leq p_t(t_{x+1})$, otherwise i_x would have placed his name on item t_{x+1}. Therefore $p_t(d) = p_t(t_k) \geq p_t(t_1)$. Since $t_1 \leq l(d) \leq d(j)$ it follows that $p_t(t_1) \geq v(j)$, and the claim follows. ∎

Sub-Claim 4.6 *If $p_t(d) > p_{t-1}(d)$ then $p_t(d) \leq c_t(d)$.*

proof: Suppose by contradiction that d is the maximal one with $p_t(d) > c_t(d) + \epsilon$, for some small $\epsilon > 0$. Thus, at some point in the iterative process of time t, the price of item d was $c_t(d) + \epsilon/2$, and then some player, j, placed his name on item d, further increasing its price. Let $X[t, ..., M]$ be the tentative allocation at this point, just before j's action. Let us examine the identity of this player j. Notice that any item $d' < d$ has price at most $c_t(d)$, as any player that placed his name on d could have placed his name on d'. We first claim that $Y[l(d), ..., d] \subseteq X[l(d), ..., d]$. Otherwise, fix some $i \in X[l(d), ..., d] \setminus Y[l(d), ..., d]$. If $i \in Y[d+1, M]$ then i placed his name on an item with price strictly larger than $c_t(d) \geq c_t(d+1) \geq p_t(d+1)$ which is larger or equal to the current price of item $d+1$, a contradiction to the myopic behavior of i. If $i \in Y[t, l(d)-1]$ then, by the prefix properties, $d(i) < l(d)$, a contradiction. And if $i \in A_t \setminus Y$ with $d(i) \geq l(d)$ then $v(i) \leq c_t(d)$ by definition, therefore i placed his name on an item with price higher than his value, again a contradiction. Therefore $Y[l(d), ..., d] \subseteq X[l(d), ..., d]$. Now, j places his name on item d. But $j \notin Y[l(d), ..., d]$ as these players are already tentative winners, and $j \notin A_t \setminus Y[l(d), ..., d]$, by repeating exactly the same arguments from above, thus reaching a contradiction. ∎

Sub-Claim 4.7 *$Y = S_t$, and, for any $d \geq t$ and $i = Y[d]$, $c_t(d) = c_t(i)$ (as defined in eq. 4.2).*

proof: We first show that, for any $j \in A_t \setminus Y$, $Y \setminus i \cup j$ is independent w.r.t items $t, ..., M$ if and only if $d(j) \geq l(d)$. Since $Y[t, ..., l(d) - 1]$ is a prefix, any allocation X that contains it cannot allocate an item $\leq l(d) - 1$ to player $j \notin Y[t, ..., l(d) - 1]$. Therefore $d(j) \leq d$. In the other direction, if $d(j) \leq d$ then we can simply allocate d to player j instead of to i, thus having an allocation to $Y \setminus i \cup j$. Otherwise, $l(d) \leq d(j) \leq d$, and we can use the exact same chain argument of subclaim 4.5 to obtain an allocation, when replacing i with j.

From this and subclaim 4.5 we have that, for any $i \in Y$ and $j \in A_t \setminus Y$ such that $Y \setminus i \cup j$ is independent w.r.t items $t, ..., M$, $v(i) \geq p_t(d) \geq c_t(d) \geq v(j)$. This property immediately implies, by the matroid basic properties, that Y is the optimal allocation. By using the above claim again we now get that $c_t(d) = c_t(i)$. ∎

From this last claim it follows that the winner of time t belongs to f_t, as $f_t \subseteq S_t = Y$, and therefore all first $|f_t|$ items of S_t must be sold to the players of f_t. It remains to show that the prices charged by the auction match the criteria of the semi-myopic mechanism.

Sub-Claim 4.8 *In the Online Iterative Auction, the winner i of time t pays $max_{r(i) \leq t' \leq t}\{c_{t'}(i)\}$.*

Let $p_t(i)$ be i's tentative price at time t. Let t' be such that $i = Y[t']$. By the above sub-claims, $p_t(i) = p_t(t') \geq c_t(t') = c_t(i)$. We additionally show that either $p_t(i) = p_{t-1}(i)$ or $p_t(i) = c_t(i)$, and the claim will follow. Assume $p_t(i) \neq p_{t-1}(i)$. Therefore i must have placed his name on item t' during the iterative process of time t. Thus $p_t(t') > p_{t-1}(t')$, and, by the above sub-claims, it follows that $p_t(i) = p_t(t') = c_t(t') = c_t(i)$. ∎

This concludes the proof of claim 4.6. ∎

We now continue with the proof of the theorem. By claim 4.6 it follows that the set $R(*)$ of the Online Iterative Auction forms a semi-myopic mechanism, and so the Online Iterative Auction is an extension of the semi-myopic mechanism. Fix any player i and and combination of recommended strategies of the other players, $s_{-i} \in R_{-i}(*)$. We need to show that i has best response to s_{-i} in $R_i(*)$. Since all players beside i are myopic with tentative deadline and then with final deadline, we can map them to types (v, d) as in claim 4.6. Let σ be this scenario, where i does not show up at all, and define t^* as in equation 4.3 of the proof of lemma 4.6 of section 4.6.1. Now suppose i plays some strategy $\bar{s}_i$, and denote this scenario by σ'. Let $Y_t(\sigma), Y_t(\sigma')$ be the tentative winners at time t in scenarios σ, σ', respectively. Let t_0 be the first time t such that $f_t(\sigma) \nsubseteq Y_t(\sigma')$. Therefore, for every $t < t_0$, σ' chooses a winner from $f_t(\sigma)$, and by sub-claim 4.3 we can assume w.l.o.g. that σ and σ' choose the same winner. Therefore $A_{t_0}(\sigma') = A_{t_0}(\sigma) \cup i$. Now suppose that $i = Y_{t_0}(\sigma')[d]$ for some d (if $i \notin Y_{t_0}(\sigma')$ then $d = M + 1$).

Sub-Claim 4.9 *Then $Y_{t_0}(\sigma')[t_0, ..., d - 1]$ is independent with respect to items $t_0 + 1, ..., d$.*

proof: By contradiction, let $f \subseteq Y_{t_0}(\sigma')[t_0, ..., d - 1]$ be its minimal prefix. Fix any $j \in f_{t_0}(\sigma) \setminus Y_{t_0}(\sigma')$. Since $j \in A_{t_0} \setminus Y_{t_0}(\sigma')$ we have that the tentative price of item t_0 in σ' is at least $v(j)$. By a chain argument as in sub-claim 4.6 if follows that every $j' \in f$ has value at least than $v(j)$. Since $j \in f_{t_0}(\sigma)$ it then follows that $j' \in Y_{t_0}(\sigma)$. Thus $f \subseteq Y_{t_0}(\sigma)$. Therefore $f_{t_0}(\sigma) = f \subseteq Y_{t_0}(\sigma')$, a contradiction. ∎

From this, and by using again the chain argument of sub-claim 4.6 we get that d's price is at least $v^*_{t_0}(\sigma) \geq v^*_{t^*}(\sigma)$. As i can win and pay $v^*_{t^*}(\sigma)$ by a strategy in $R_i(*)$ (e.g. arriving at time t^* and bidding only on item t^*), the claim follows. ∎

This concludes the proof of the Theorem. ∎

4.6.4 The Sequential Japanese Auction has a Set-Nash Equilibrium

To show that our Sequential Japanese Auction is an ignorable extension of a semi-myopic mechanism, we need to modify payments similarly to the modification of the Online Iterative Auction. For this, we need to handle simultaneous "drop" announcements more carefully: At any price level p, several players may want to drop. Furthermore, this may be an on-going process, as after one player drops, another one now wants to drop as well. We need to determine more accurately the order among them. This information is used in order to determine f_t (interestingly, we are not able to compute S_t entirely, only f_t, which is enough).

Definition 4.16 (The Sequential Japanese Auction) *The basic auction structure remains the same as in Def 4.5. Two additional points should be handled:*

Simultaneous "drop" announcements: *Define $D(p, n)$ as the set of players (among those who did not drop yet), that wish to drop when the price level is p and the number of remaining players is n. At every price level p, the auction solicits drop announcements by repeatedly accepting only one drop announcement out of $D(p, n)$, and decreasing n by 1.*[15] *When $D(p, n) = \emptyset$, the price increases. The winner is, as before, the last remaining player.*

Prices: *Prices $p_t(i)$ for every player i at every time t are maintained as follows: Let k be the number of non-drop-outs just before the price ended its time-t ascend, at a level of p^*. Let $D(p^*, k), D(p^*, k-1), ..., D(p^*, 1)$ be the order of drop-outs at this level. Define the critical number $x^* = \min\{0 < x < k \ : \ |D(p^*, x+1)| = 1\}$, and $D^* = \cup_{x \leq x^*} D(p^*, x)$. For any player i, if $i \in D^*$ set $p_t(i) = p^*$, otherwise $p_t(i) = 0$. The winner i of time t pays $max_{r(i) \leq t' \leq t}\{p_{t'}(i)\}$.*

Recommended strategies: *i's strategy is recommended if he arrives at $r(i)$, choose a tentative deadline $d \leq d(i)$, plays myopically with parameters $v(i), d$ until time d, and plays myopically with parameters $v(i), d(i)$ thereafter.*

Again, these recommended strategies are semi-myopic.

[15]E.g. if $D(p, n) = X$ then some $i \in X$ is chosen to be dropped, $X \setminus i \subseteq D(p, n-1)$ and $i \notin D(p, n-1)$.

Theorem 4.6 *The Sequential Japanese Auction is an ignorable extension of a semi-myopic mechanism.*

Corollary 4.3 *The Sequential Japanese Auction Set-Nash implements a 3-approximation of the welfare.*

Proof (of theorem): The theorem will follow by the following claims.

Claim 4.7 *If all players i play strategies in $R_i(*)$ then the Sequential Japanese Auction forms a semi-myopic mechanism.*

proof: Let p^* be the last price reached by the auction of time t, and suppose there are k players that did not drop out just before p^* was reached.

Sub-Claim 4.10 *Fix any $j \in A_t \setminus S_t$. As long as j does not drop, then every $i \in P_t(d(j))$ does not drop.*

proof: By contradiction, let $i \in P_t(d(j))$ be the first to drop, say at price p. Since j did not drop, $v(j) \geq p$. By the maximality of S_t, $v(i) > v(j)$. Thus i did not drop because of the price. But the number of non-dropped players is at least $|P_t(d(j))| + 1 > d(i)$. Therefore i could not have dropped at this point, a contradiction. ∎

Sub-Claim 4.11 $p^* = max\{\ v(j) \mid j \in A_t \setminus S_t\ \}$.

proof: Let j^* be the player with maximal value among those in $A_t \setminus S_t$. By the previous sub-claim, j^* will drop because the price will reach his value, as $|P_t(d(j^*))| \geq d(j^*)$. Thus $p^* \geq v(j^*)$. Suppose by contradiction that $p > v(j^*)$, and choose some p in between. Thus, when the price reaches p, all the non-drop-outs belong to S_t. Consider the one that receives, according to S_t, the latest item. The number of non-drop-outs is smaller than his deadline, so he will drop. The one that receives the item before last will next drop, by the same argument, and so on. Therefore the price will not increase beyond p, a contradiction. ∎

Sub-Claim 4.12 *For any $i \in f_t$, $p^* = c_t(i)$.*

proof: For any $j \in A_t \setminus f_t$, $S_t \setminus i \cup j$ is independent: choose an allocation in which i receives item t, and then remove i and allocate t to j. Therefore the claim follows from the previous sub-claim, and from the definition of $c_t(i)$. ∎

Sub-Claim 4.13 *For any l', $|D(p^*, l') \cup \cdots \cup D(p^*, 1)| = l'$.*

proof: Since $D(p^*, l'), ..., D(p^*, 1)$ includes only players that did not actually drop before phase (p^*, l'), and there are exactly l' of those, then $l' \geq |D(p^*, l') \cup \cdots \cup D(p^*, 1)|$. On the other hand, every player among the l' players that did not drop yet will drop in some phase $D(p^*, l'), ..., D(p^*, 1)$, so $l' \geq |D(p^*, l') \cup \cdots \cup D(p^*, 1)|$. ∎

Sub-Claim 4.14 *If* $|D(p^*, l' + 1)| = 1$ *then* $D(p^*, l') \cup \cdots \cup D(p^*, 1)$ *is a prefix.*

proof: Since $|D(p^*, l' + 1)| = 1$ then any $j \in D(p^*, l') \cup \cdots \cup D(p^*, 1)$ has deadline $d(j) < t + (l' + 1) - 1$, i.e. $d(j) \leq t + l' - 1$. Since $|D(p^*, l') \cup \cdots \cup D(p^*, 1)| = l'$ it follows from claim 4.12 that $D(p^*, l') \cup \cdots \cup D(p^*, 1)$ is a prefix. ∎

Sub-Claim 4.15 *Let* x^* *be the critical number of drop-outs, and* $D^* = \cup_{x \leq x^*} D(p^*, x)$*, as in def. 4.16. Then* $D^* = f_t$*.*

proof: Let $l = |f_t|$. Notice that, for any $l' > l$, $f_t \cap D(p^*, l') = \emptyset$: If $i \in f_t$ then $v(i) > p^*$ and $d(i) \leq |f_t| + t - 1 < l' + t$, so i will not drop. This, in turn, implies that a player in f_t will drop in one the of phases $(p^*, l), ..., (p^*, 1)$, so $f_t \subseteq D(p^*, l) \cup \cdots \cup D(p^*, 1)$. Since $|D(p^*, l) \cup \cdots \cup D(p^*, 1)| = l$, we conclude that $f_t = D(p^*, l) \cup \cdots \cup D(p^*, 1)$. It is left to show that $x^* = l$. As $D(p^*, l) \subseteq f_t$ and $f_t \cap D(p^*, l + 1) = \emptyset$ then $D(p^*, l + 1) \cap D(p^*, l) = \emptyset$. This implies that $|D(p^*, l + 1)| = 1$, so $x^* \leq l$. But if $x^* < l$ then $D(p^*, x^*) \cup \cdots \cup D(p^*, 1) \subsetneq f_t$ is a prefix, contradicting the minimality of f_t (by claims 4.13, 4.15). Therefore $x^* = l$ and $D^* = f_t$. ∎

From all the above, the proof of the claim immediately follows: First, the winner belongs to $D^* = f_t$. Second, all time t prices for players not in f_t equal 0, and for players in f_t, time t prices equal $p^* = c_t(i)$, i.e. as required by the price rules of the semi-myopic mechanism. ∎

We now continue with the proof of the theorem. Using the above claim, it only remains to show that, fixing some player i and some strategies $s_{-i} \in R_{-i}(*)$ of the other players, i has a best response in $R_i(*)$. Consider some strategy s_i of i. Let t_0 be the first time in which i enters D^*. We first notice that, in every time prior to t, i can wave participation without affecting the winner: If the price when i participates reached a level p^*, then clearly, when i does not participate the price cannot rise above p^*. By definition, a player in D^* will not drop before there will be at most $|D^*| - 1$ other non-drop-outs (as the price does not reach his value). Therefore the last non-drop-outs will be exactly all players in D^*, and so the winner will be the same.

Now suppose the price level at time t, in which i entered D^*, is p^*. Therefore i's price will be at least p^*. We claim that, by arriving at time t and playing the fixed confidence strategy $(p^*, 1)$, i can win item t for a price p^*. Since this strategy is in $R_i(*)$, the claim will follow. To see this, observe that $|D(p^*, x)| > 1$ for any $1 < x < x^*$ (since $D(p^*, x) \cap D(p^*, x - 1) \neq \emptyset$). Therefore, even if i will not be willing to drop out until being the last non drop out, all others will drop out at price p^*, and so i will win t and will pay p^*. ∎

4.7 Useful Properties of Offline Allocations and Matroids

This section summarizes useful properties that we have used throughout our proof. Most of the properties here are new, and are interesting in our context. For completeness, we begin with a short introductory summary of Matroids and their relevant properties.

Definition 4.17 (A Matroid) *A Matroid is a finite set S and a collection $I \subseteq 2^S$ of independent sets, such that:*

1. $\emptyset \in I$
2. *If $X \in I$ and $Y \subseteq X$ then $Y \in I$.*
3. *If $X, Y \in I$ and $|X| = |Y| + 1$ then there exists $j \in X \setminus Y$ such that $Y \cup j \in I$. If $X \subseteq S$ but $X \notin I$ then it is a dependent set. A base of a matroid is a maximal independent set, and a cycle is a minimal dependent set.*

Claim 4.8 *The offline allocation of M items among a set A of players is a matroid, where S is the set of players, and a subset X of players is independent if there exists an allocation of (part of) the items to all the players in X.*

proof: The first two conditions of the matroid are trivially satisfied. Let us verify the third one. Let X, Y are be two independent sets with $|X| > |Y|$. We first claim that there exists allocations for X, Y such that, for any $j \in X \cap Y$, j receives the same item in both allocations. To see this, start from arbitrary two allocations, and choose some $j \in X \cap Y$. Suppose j receives items t_1, t_2 in the allocations $1, 2$, where $t_1 < t_2$. Suppose that player j' receives item t_2 in allocation 1. Then we can swap players j and j' in allocation 1, so that j will receive item t_2 (this is valid as we know he can receive this item) and j' will receive item t_1 (this is valid as $t_1 < t_2$). Notice that we have strictly decreased the number of players in $X \cap Y$ that receive different items, and so repeating this implies the result. Now, choose some item t which is being allocated for X but not allocated to any player of Y. Suppose that t is allocated to j in the allocation of X. By our assumption, $j \notin Y$, and so $Y \cup j$ is independent: use the previous allocation of Y, and allocate item t (that beforehand was not allocated) to j. ∎

The following claim lists some useful matroid properties. For extensive discussion and proofs, see e.g. the textbook [100].

Claim 4.9 *Let $M = (S, I)$ be a matroid. Then:*

1. *If $X, Y \in I$ and $|X| < |Y|$ then there exists $Z \subseteq X \setminus Y$ such that $|X \cup Z| = |Y|$ and $X \cup Z \in I$.*
2. *If B_1, B_2 are bases then $|B_1| = |B_2|$.*

3. *If B_1, B_2 are bases, then, for any $j \in B_1 \setminus B_2$ there exists $j' \in B_2 \setminus B_1$ such that $B_1 \setminus j \cup j' \in I$ and $B_2 \setminus j' \cup j \in I$.*

The following claims are slight alterations of classical properties:

Claim 4.10 *Let $X, Y \in I$, and $X \not\subseteq Y$. Then, for any $j \in Y \setminus X$ such that $X \cup j \notin I$ there exists $j' \in X \setminus Y$ such that $X \setminus j' \cup j \in I$ and $Y \setminus j \cup j' \in I$.*

proof: If $|X| = |Y|$ then we can assume w.l.o.g. that both are bases (as $I' = \{ Z \in I \mid |Z| \leq |X| \}$ are also the independent sets of a matroid), and the claim immediately follows.

If $|X| > |Y|$ then assume, as before, that X is a base. There exists $Z \subseteq X \setminus Y$ such that $B = Y \cup Z$ is a base. Since $j \in Y \setminus X$ then $j \in B \setminus X$ and so there exists $j' \in X \setminus B$ such that $X \setminus j' \cup j \in I$ and $B \setminus j \cup j' \in I$. Since $Y \subseteq B$ and $j' \in Y \setminus X$ as well, the claim follows.

If $|X| < |Y|$ then assume that Y is a base, take some $Z \subseteq Y \setminus X$ such that $B = X \cup Z$ is a base, and notice that $j \notin Z$ as $X \cup j \notin I$. Thus we can essentially repeat the above logic: j is also in $Y \setminus B$ so there exists $j' \in B \setminus Y$ such that $B \setminus j' \cup j \in I$ and $Y \setminus j \cup j' \in I$. Since $B \setminus Y = X \setminus Y$, and $X \subset B$, then the claim follows. ∎

Claim 4.11 *Let B be a base of the matroid, and $Y \in I$ such that $|B \setminus Y| = 1$. Then $|Y \setminus B| \leq 1$.*

proof: $|B| \geq |Y| = |B \cap Y| + |Y \setminus B| = |B| - |B \setminus Y| + |Y \setminus B| = |B| - 1 + |Y \setminus B|$. Therefore $|Y \setminus B| \leq 1$, as claimed. ∎

4.7.1 Some Useful Properties of Offline Allocations

For the following discussion, it will be convenient to assume the following ϵ-assumption: There exists many small valued players in A_t that desire any one of the items $t, ..., M$.

Definition 4.18 (A prefix) *A subset $X \subseteq S_t$ is called a prefix if it is a prefix of any allocation $S_t[1, M]$ of S_t.*

Claim 4.12 *$X \subseteq S_t$ is a prefix if and only if for all $j \in X$, $d(j) \leq t + |X| - 1$.*

proof: Suppose first that X is a prefix, and, by contradiction, that there exists some $j \in X$ with $d(j) > t + |X| - 1$. Let $S_t[t, M]$ be some allocation of S_t. Since $j \in X$ and X is a prefix then j is allocated some item $\leq t + |X| - 1$. Suppose player j' is allocated item $d(j)$. Then we can switch between j and j' and have an allocation in which X is not a prefix, a contradiction. In the other direction, if $X \subseteq S_t$ and $d(j) \leq t + |X| - 1$ for any $j \in X$ then, in any allocation, $j \in S_t[t, t + |X| - 1]$. Therefore $X \subseteq S_t[t, t + |X| - 1]$, and since $|S_t[t, t + |X| - 1]| = |X|$ then it follows that $S_t[t, t + |X| - 1] = X$, i.e. it is a prefix. ∎

Definition 4.19 *For any $t \leq d \leq M$, we build the set of players $P_t(d)$ using the following process (fix any allocation of S_t):*

1. *Let $x_0 = d$.*
2. *For $i > 0$, define inductively $x_i = max\{\ d(j) \mid j \in S_t[t, x_{i-1}]\}$.*
3. *Let k be some index such that $x_{k+1} = x_k$, and fix $P_t(d) = S_t[t, x_k]$.*

Claim 4.13 *$P_t(d)$ is the prefix with minimal length among all prefixes with length $\geq d - t + 1$.*

proof: First notice that, from the ϵ-assumption it immediately follows that $|P_t(d)| = x_k - t + 1$. Also notice that, by our construction, any $j \in P_t(d)$ has $d(j) \leq x_k = t + |P_t(d)| - 1$. Therefore, by claim 4.12, $P_t(d)$ is a prefix. Suppose by contradiction that there exists a prefix P' with $d \leq t + |P'| - 1 < x_k$. Choose index i such that $x_i \leq t + |P'| - 1 < x_{i+1}$. But then, by the construction process of $P_t(d)$, we must have a player in P' with deadline at least x_{i+1}, contradicting claim 4.12. ∎

Claim 4.14 *$j \in P_t(d)$ if and only if there exists an allocation of S_t in which $j \in S_t[t, d]$.*

proof: If $j \in S_t[t, d]$ then by definition $j \in P_t(d)$. Let us verify the other direction. Fix any allocation of S_t, and compute $P_t(d)$ by that allocation. Assume $j = S_t[d']$ for some $d' > d$ (otherwise the claim immediately follows). Let j_i be the player that determined x_i. Then we have $j_1 \in S_t[t, d]$. Consider the following allocation replacements: allocate item x_1 to player j_1 (this is his deadline, so this is valid), j_2 will get item x_2, ... , j_k will get item x_k. Finally, allocate j's item to j_{k+1} (that received x_k), and allocate j_1's item to j. Therefore we have an allocation in which j receives some item $\leq d$, as claimed. ∎

Claim 4.15 *$f_t = P_t(t) = P_t(|first_t| + t - 1)$.*

proof: If $j \in first_t$ then there exists an allocation of S_t such that $j = S_t[t]$. Since $P_t(t)$ is a prefix of $S_t[1, M]$ then $j \in P_t(t)$. On the other hand, claim 4.14 tells us that for any $j \in P_t(t)$ there exists an allocation such that $j = S_t[t]$, and therefore $j \in f_t$. We conclude that $f_t = P_t(t)$. From claim 4.13 we now get also that $P_t(t) = P_t(|first_t| + t - 1)$, as $P_t(t)$ is a prefix with length $|first_t|$. ∎

Claim 4.16 *For any t, d with $t < d$, $min_{j \in P_{t+1}(d)}\{v(j)\} \geq min_{j \in P_t(d)}\{v(j)\}$.*

proof: We will actually show that $min_{j \in P_{t+1}(d)}\{v(j)\} \geq min_{j \in P_t(d) \setminus ON[t]}\{v(j)\}$. Let $x = |P_t(d)| + t - 1$, the last item allocated to a player in $P_t(d)$. By the above claims, for any $j \in P_t(d)$, $d_j \leq x$, and $x \geq d$. Let j be the player with minimal value in $P_{t+1}(d)$, and assume by contradiction that $v(j) < min_{j \in P_t(d) \setminus ON[t]}\{v(j)\}$. Therefore $j \notin P_t(d) \setminus ON[t]$. Consider some allocation of S_{t+1} such that j receives item $\leq d$. Now consider $P_t(d) \setminus ON[t]$ and $S_{t+1}[t+1, x]$. These are two bases of the matroid over items $t+1, ..., x$. Since $j \in S_{t+1}[t+1, x] \setminus (P_t(d) \setminus ON[t])$ then there exists $j' \in P_t(d) \setminus ON[t] \setminus S_{t+1}[t+1, x]$ such that $S_{t+1}[t+1, x] \setminus j \cup j'$ is independent (w.r.t items $t+1, ..., x$). As $d_{j'} \leq x$, it follows that $j' \notin S_{t+1}$, and therefore $S_{t+1} \setminus j \cup j'$ is independent as well. As $j' \in A_{t+1}$, and by the maximality of S_{t+1}, we must have $v(j) > v(j') \geq min_{j \in P_t(d)}\{v(j)\}$, a contradiction. ∎

Claim 4.17 *f_t is independent w.r.t items $t+1, ..., M$ if and only if S_t is independent w.r.t items $t+1, ..., M$.*

proof: Since $f_t \subseteq S_t$ then the right to left direction is immediate. Let us verify the other direction, i.e. that if f_t is independent w.r.t items $t+1, ..., M$ then so is S_t. Let $\tilde{A}_t, \tilde{f}_t, \tilde{S}_t$ be the variables after adding many ϵ players, as in the ϵ-assumption. By the optimality of S_t it follows that $S_t \subseteq \tilde{S}_t$ (when ϵ is small enough). As $\tilde{f}_t$ is a prefix, it cannot be independent w.r.t items $t+1, ..., M$. Thus there exists $j \in \tilde{f}_t \setminus f_t$. By definition, $\tilde{S}_t \setminus j$ is independent w.r.t items $t+1, ..., M$, and therefore $S_t \setminus j$ is independent w.r.t items $t+1, ..., M$. If $j \in S_t$ this will therefore imply $j \in f_t$, a contradiction. Thus j is an ϵ player, and $S_t \subseteq \tilde{S}_t \setminus j$. Since $\tilde{S}_t \setminus j$ is independent w.r.t items $t+1, ..., M$, then this implies that so is S_t, as needed. ∎

Sub-Claim 4.16 *Let $A'_t = A_t \cup j'$. Let S_t, S'_t and f_t, f'_t be derived from A_t, A'_t, respectively. Then:*

1. *If $j' \in f'_t$ then $f_t \setminus S'_t \neq \emptyset$.*
2. *$f_t \neq f'_t$ if and only if $j' \in f'_t$.*

proof: From the prefix properties detailed in sub-section 4.7.1 it immediately follows that, if $f_t \subseteq S'_t$ then $f'_t = f_t$, and thus the first claim follows. This also implies the right to left direction of the second claim. We are left to show that, if $f_t \neq f'_t$ then $j' \in f'_t$. By the maximality of S_t, S'_t it follows that either $S_t = S'_t$, or $S'_t = S_t \setminus j \cup j'$ for some $j \in S_t \setminus S'_t$. Since $f_t \neq f'_t$, the latter alternative must hold. If $j \notin f_t$ then $f_t \subseteq S'_t$, implying that $f'_t = f_t$, a contradiction. Thus $j \in f_t$. Therefore there exists an allocation with $j = S_t[t]$. Since $S'_t = S_t \setminus j \cup j'$ then there exists an allocation with $j' = S'_t[t]$ (simply use the previous allocation, changing only the player who receives item t from j to j'). By definition, this implies that $j' \in f'_t$. ∎

Chapter 5

Conclusions

This thesis studies auction theory in computational contexts. The guidelines of classic algorithmic theory differ from those of the classic auction theory in many ways: worst-case analysis vs. average case analysis, the tradeoff between near optimal outcomes and computational efficiency, and different models with new algorithmic goals. On the other hand, algorithmic theory classically ignores incentive issues, necessary in large distributed systems like the Internet. This thesis examines the integration of the two theories, and the new possibilities and impossibilities that emerge.

Three groups of results are presented. In the first, we suggest the new notion of *online auctions*, in which buyers arrive over time, and the auction is required to make decisions before knowing the entire sequence of bidders. We present incentive compatible (truthful) auctions, i.e. where bidders are motivated to reveal their true types, that achieve optimal worst case guarantees with respect to both the optimal social welfare, and the seller's revenue. This demonstrates the above mentioned integration between the algorithmic theory and auction theory in two ways. First, the "online" setting from CS, in which the input is revealed over time and the algorithm must make decisions before knowing it entirely, is examined in the context of auction theory. Second, we harness worst case analysis tools to analyze the suggested auctions, instead of average case analysis through the assumption of a specific, known probability distribution over the input.

The second set of results studies the limits of the truthfulness notion. This question gains new meanings due to the computational limitations that we require from our mechanisms to satisfy, and because of the new algorithmic goals we wish to consider. We ask what types of algorithmic goals can be achieved by a computationally efficient incentive compatible auction. We study this question in the context of combinatorial auctions. This is a general auction model that captures many classic computational settings, as well as new "real world" settings. The classic possibility result from auction theory, the VCG mechanism, is computationally infeasible here, by this emphasizing the critical role that the additional computational assumptions can take. We show that under few additional assumptions, of which the most significant one resembles Arrow's IIA condition, the only algorithmic goal that can be truthfully implemented is the weighted maximization of the social welfare. This also implies that no computationally efficient truthful combinatorial auction that satisfy our additional requirements can approximate the optimal welfare within "reasonable" bounds.

As this impossibility result is widely believed to be even stronger than what we have shown, we next turn to investigate, in the third set of results, other solution concepts. We study an online model of "gradually expiring items", for which we prove that no (deterministic) truthful auction can obtain a good approximation ratio. We suggest natural ascending auctions that obtain near optimal performance for a wide family of selfish player behaviors. The conceptual shift is from auctions that motivate the player to take a single (dominant) strategy, to auctions for which many behaviors still lead to the approximate social optimum, and so we are not required to "convince" the players to follow a specific action. In our concrete setting, we first give an intuitive explanation why we expect the player to limit his attention to choosing one out of this set of behaviors. We then describe a general game-theoretic framework that captures this logic. We suggest the notion of "Set-Nash" equilibria, where the player has a set of best response actions to choose from, and deviating from this set cannot increase his utility (with a variety of increasing strength interpretations to the term "best response actions"). The overall argument is therefore two-fold: first, we cannot expect to predict which of these actions the player will choose, we can only expect him to choose one out of these actions, and, second, *any* combination of such actions will lead to an approximately optimal outcome.

Much more understanding and exploration of the possibilities-impossibilities border is needed, and it seems that this is one of the main tasks of the field of algorithmic mechanism design. This thesis took this trail in parallel from both ends. We have initiated a search of the impossibilities of the strong truthfulness notion, as well as a search for other solution concepts that are "distribution-free". The exact borders, however, remain unclear, and the need to exactly pin-point them is stronger and clearer than ever.

This thesis also suggests more technical-in-nature open questions. The most intriguing one seems to be the necessity of the IIA requirement for the impossibility result of chapter 3 to hold. A wide range of monotonicity conditions, that lie between the weak and strong monotonicity conditions we have defined, is waiting to be explored. Another interesting question is the usefulness of the hierarchy of alternative solution concepts posed in chapter 4. We provide a concrete example only for one of the suggested concepts. Can other interesting examples be found? This is left open for future research to come.

Bibliography

[1] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Eva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *Proc. of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*, 2003.

[2] Aaron Archer and Eva Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42st Annual Symposium on Foundations of Computer Science(FOCS'01)*, 2001.

[3] Aaron Archer and Eva Tardos. Frugal path mechanisms. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02)*, 2002.

[4] Kenneth Arrow. *Social Choice and Individual Values*. Wiley, 1951.

[5] Kenneth Arrow. The property rights doctrine and demand revelation under incomplete information. In M. Boskin, editor, *Economies and Human Welfare*. Academic Press : NY, 1979.

[6] Lawrence Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 2004. To appear.

[7] Lawrence Ausubel and Paul Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1), 2002.

[8] Baruch Awerbuch, Yossi Azar, and Adam Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proc. of the 35th ACM Symposium on Theory of Computing (STOC'03)*, 2003.

[9] Moshe Babaioff and Noam Nisan. Concurrent auctions across the supply chain. *Journal of Artificial Intelligence Research*, 21:595–629, 2004.

[10] Moshe Babaioff and William E. Walsh. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decision Support Systems*, 2004. To appear.

[11] Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *Proc. of the 13th Symposium on Discrete Algorithms (SODA'02)*, 2002.

[12] Yair Bartal, Francis Y. L. Chin, Marek Chrobak, Stanley P. Y. Fung, Wojciech Jawor, Ron Lavi, Jiri Sgall, and Tomas Tichy. Online competitive algorithms for maximizing weighted throughput of unit jobs. In *Proc. of the 21st Symposium on Theoretical Aspects of Computer Science (STACS'04)*, 2004.

[13] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi-unit combinatorial auctions. In *Proc. of the ninth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'03)*, 2003.

[14] Yair Bartal and Ron Lavi. Analysis of the β-opt algorithm, 2002. Unpublished manuscript.

[15] Kaushik Basu and Jorgen W. Weibull. Strategy subsets closed under rational behavior. *Economics Letters*, 36:141–146, 1991.

[16] B. Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 52:1007–1028, 1984.

[17] Sushil Bikhchandani, Shurojit Chatterji, and Arunava Sen. Incentive compatibility in multi-unit auctions, 2003. Working paper.

[18] Duncan Black. On the rationale of group decision-making. *The Journal of Political Economy*, 56:23–34, 1948.

[19] Avrim Blum, Jeffrey Jackson, Toumas Sandholm, and Martin Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research (JMLR)*, 5:649–667, 2004.

[20] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Proc. of the 14th Symposium on Discrete Algorithms (SODA'03)*, 2003.

[21] Avrim Blum, Toumas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. In *Proc. of the 13th Symposium on Discrete Algorithms (SODA'02)*, 2002.

[22] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communication. In *Proc. of the 43rd Annual Symposium on Foundations of Computer Science (FOCS'02)*, 2002.

[23] Allan Borodin and Ran El-Yaniv. *On-Line Computation and Competitive Analysis.* Cambridge University Press, 1998.

[24] Michel Le Breton and John Weymark. Arrovian social choice theory on economic domains. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare: Volume 2.* North-Holland: Amsterdam, Forthcoming.

[25] Michel Le Breton and John A. Weymark. Arrovian social choice theory on economic domains. In Kenneth Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare.* Elsevier Science, 2003.

[26] Kim-Sau Chung and Jeffrey C. Ely. Ex post incentive compatible mechanism design, 2004. Working paper.

[27] E. Clarke. Multipart pricing of public goods. *Public Choice*, 8:17–33, 1971.

[28] Wolfram Conen and Toumas Sandholm. Minimal preference elicitation in combinatorial auctions. In *Proc. of the 17th International Joint Conference on Artificial Intelligence (IJCAI'01)*, 2001.

[29] Claude d'Aspremont and Louis-Andre' Ge'rard-Varet. Incentives and incomplete information. *Journal of Public Economy*, 11:25–45, 1979.

[30] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

[31] Nikhil Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V.Vazirani. Market equilibrium via a primal-dual-type algorithm. In *Proc. of the 43rd Annual Symposium on Foundations of Computer Science (FOCS'02)*, 2002.

[32] Nikhil Devanur and Vijay Vazirani. The spending constraint model for market equilibrium: Algorithmic, existence and uniqueness results. In *Proc. of the 36th ACM Symposium on Theory of Computing (STOC'04)*, 2004.

[33] John Duggan and Michel Le Breton. Dutta's minimal covering set and shapley's saddles. *Journal of Economic Theory*, 70:257–265, 1996.

[34] John Duggan and Michel Le Breton. Dominance-based solutions for strategic form games, 1998. Working paper.

[35] John Duggan and Michel Le Breton. Mixed refinements of shapley's saddles and weak tournaments. *Social Choice and Welfare*, 8:65–78, 2001.

[36] R. El-Yaniv, A. Fiat, R.M. Karp, and G. Turpin. Optimal search and one-way trading online algorithms. *Algorithmica*, 30(1):101–139, 2001.

[37] Ran El-Yaniv. Competitive solutions for on-line financial problems. In A. Fiat and G. Woeginger, editors, *Online Algorithms: The State of Art.* Springer-Verlag (LNCS 1442), 1998.

[38] Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63(1), 2001.

[39] Amos Fiat, Andrew Goldberg, Jason Hartline, and Anna Karlin. Competitive generalized auctions. In *Proc. of the 34th ACM Symposium on Theory of Computing (STOC'02)*, 2002.

[40] Eric Friedman and David Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *Proc. of the 5th ACM Conference on Electronic Commerce (EC'03)*, 2003.

[41] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.

[42] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proc. of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'01)*, 2001.

[43] Rica Gonen and Daniel Lehmann. Multi unit combinatorial auctions: Branch and bound heuristics. In *Proc. of the 2nd ACM Conference on Electronic Commerce (EC'00)*, 2000.

[44] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[45] Faruk Gul and Ennio Stacchetti. The english auction with differentiated commodities. *Journal of Economic Theory*, 92:66–95, 2000.

[46] Bruce Hajek. On the competitiveness of on-line scheduling of unit-length packets with hard deadlines in slotted time. In *Proc. of the 2001 Conference on Information Sciences and Systems*, 2001.

[47] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David Parkes. Adaptive limited-supply online auctions. In *Proc. of the 6th ACM Conference on Electronic Commerce (EC'04)*, 2004.

[48] John Harsanyi. Games of incomplete information played by bayesian players. *Management Science*, 14:159–182, 320–334, 486–502, 1967.

[49] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.

[50] Bengt Holmstrom. Groves' scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.

[51] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004.

[52] Ellis Horowitz and Sartaj Sahni. *Fundamentals of Computer Algorithms*. Computer Science press, Rockville, Maryland, 1978.

[53] Ehud Kalai, Eitan Muller, and Mark Satterthwaite. Social welfare functions when preferences are convex, strictly monotonic, and continuous. *Public Choice*, 34:87–97, 1979.

[54] Alexander Kesselman, Zvi Lotker, Yishay Mansour, Boaz Patt-Shamir, Baruch Schieber, and Maxim Sviridenko. Buffer overflow management in qos switches. In *Proc. of the 33rd ACM Symposium on Theory of Computing (STOC'01)*, 2001.

[55] Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proc. of the 44rd Annual Symposium on Foundations of Computer Science (FOCS'03)*, 2003.

[56] Paul Klemperer. *Auctions: Theory and Practice*. Princeton University Press, 2004.

[57] Anshul Kothari, David Parkes, and Subhash Suri. Approximately-strategyproof and tractable multi-unit auctions. In *Proc. of the 5th ACM Conference on Electronic Commerce (EC'03)*, 2003.

[58] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *Proc. of the 44rd Annual Symposium on Foundations of Computer Science (FOCS'03)*, 2003.

[59] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science*, 310:159–180, 2004.

[60] Aurel A. Lazar and Nemo Semret. The progressive second price auction mechanism for network resource sharing. In *Proc. of the 8th International Symposium on Dynamic Games*, 1998.

[61] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proc. of the 3rd ACM Conference on Electronic Commerce (EC'01)*, 2001.

[62] Daniel Lehmann, Liadan O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[63] Herman B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

[64] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auctions. In *Proc. of the 2nd ACM Conference on Electronic Commerce (EC'00)*, 2000.

[65] Kevin Leyton-Brown and Yoav Shoham. Bidding clubs: Institutionalized collusion in auctions. In *Proc. of the 2nd ACM Conference on Electronic Commerce (EC'00)*, 2000.

[66] Jeffrey K. Mackie-Mason and Hal R. Varian. Pricing the internet. In B. Kahn and J. Keller, editors, *Public Access to the Internet*. Prentice Hall, 1994.

[67] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford university press, 1995.

[68] Eric Maskin. Nash equilibrium and welfare optimality. *Review of Economic Studies*, 66:23–38, 1999.

[69] Eric Maskin and John Riley. Optimal multi-unit auctions. In *The Economics of Missing Markets, Information, and Games*. Oxford University Press, Clarendon Press, 1989.

[70] Moritz Meyer-Ter-Vehn and Benny Moldovanu. Ex-post implementation with interdependent valuations. Working paper, 2002.

[71] H. Moulin. *The Startegy of Social Choice*. North-Holland, 1983.

[72] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *AAAI-02*, 2002.

[73] Rudolf Muller and Rakesh Vohra. On dominant strategy mechanisms, 2003. Working paper.

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

[75] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[76] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

[77] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proc. of the 2nd ACM Conference on Electronic Commerce (EC'00)*, 2000.

[78] Noam Nisan. The communication complexity of approximate set packing and covering. In *Proc. of the 29th International Colloquium on Automata,Languages, and Programming (ICALP'02)*, 2002.

[79] Noam Nisan and Amir Ronen. Computationally feasible vcg mechanisms. In *Proc. of the 2nd ACM Conference on Electronic Commerce (EC'00)*, 2000.

[80] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[81] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting lindahl prices. Working paper, 2002.

[82] Noam Nisan and Edo Zurel. An efficient approximate allocation algorithm for combinatorial auctions. In *Proc. of the 3rd ACM Conference on Electronic Commerce (EC'01)*, 2001.

[83] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory.* The MIT press, 1994.

[84] David Parkes. ibundle: An efficient ascending price bundle auction. In *Proc. of the 1st ACM Conference on Electronic Commerce (EC'99)*, 1999.

[85] Ryan Porter. Mechanism design for online real-time scheduling. In *Proc. of the 6th ACM Conference on Electronic Commerce (EC'04)*, 2004.

[86] Ryan Porter and Yoav Shoham. On cheating in sealed-bid auctions. *Decision Support Systems*, 2004. To appear.

[87] Kevin Roberts. The characterization of implementable choise rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences. Papers presented at the first European Summer Workshop of the Econometric Society*, pages 321–349. North-Holland, 1979.

[88] Amir Ronen. On approximating optimal auctions. In *Proc. of the 3rd ACM Conference on Electronic Commerce (EC'01)*, 2001.

[89] Amir Ronen and Amin Saberi. Optimal auctions are hard. In *Proc. of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)*, 2002.

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

[91] Irit Rozenshtrum. Dominant strategy implementation with quasi-linear preferences. Master's thesis, Dept. of Economics, The Hebrew University, Jerusalem, Israel, 1999.

[92] Toumas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. Cabob: A fast optimal algorithm for combinatorial auctions. In *Proc. of the 17th International Joint Conference on Artificial Intelligence (IJCAI'01)*, 2001.

[93] Tuomas W. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. In *Proc. of the 16th International Joint Conference on Artificial Intelligence (IJCAI'99)*, 1999.

[94] Mark Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[95] Lloyd Shapley. Some topics in two person games. *Annals of Mathematics Studies*, 52, 1964.

[96] Scott Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. In *Proc. of the ACM SIGCOMM*, 1994.

[97] William Vickrey. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[98] Mark Voorneveld. Preperation. *Games and Economic Behavior*, 42:403–414, 2004.

[99] Michael P. Wellman, William E. Walsh, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35:271–303, 2001.

[100] Dominic J. A. Welsh. *Matroid theory.* London : Academic Press, 1976.