UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 656

Bidding in Combinatorial Auctions

JIM WILENIUS

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2009

ISSN 1651-6214
ISBN 978-91-554-7554-3
urn:nbn:se:uu:diva-102960

Dissertation presented at Uppsala University to be publicly examined in Polacksbacken 2446, Lägerhyddsvägen 2, 752 37, Uppsala, Friday, September 18, 2009 at 13:15 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract
Wilenius, J. 2009. Bidding in Combinatorial Auctions. Acta Universitatis Upsaliensis. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology* 656. 59 pp. Uppsala. ISBN 978-91-554-7554-3.

This thesis concerns the interdisciplinary field of combinatorial auctions, combining the fields of computer science, optimization and economics. A combinatorial auction is an auction where many items are sold simultaneously and where bidders may submit indivisible *combinatorial bids* on groups of items. It is commonly believed that good solutions to the allocation problem can be achieved by allowing combinatorial bidding.

Determining who wins in a combinatorial auction is fundamentally different from a traditional single-item auction because we are faced with a hard and potentially intractable optimization problem. Also, unless we are limited to truthful mechanisms, game theoretic analysis of the strategic behavior of bidders is still an open problem.

We have chosen primarily to study the first-price combinatorial auction, a natural auction widely used in practice. Theoretical analysis of this type of auction is difficult and little has been done previously. In this thesis we investigate and discuss three fundamental questions with significant practical implications for combinatorial auctions.

First, because the number of possible bids grows exponentially with the number of items, limitations on the number of bids are typically required. This gives rise to a problem since bidders are unlikely to choose the "correct" bids that make up the globally optimal solution. We provide evidence that an expressive and compact bidding language can be more important than finding the optimal solution. Second, given a first-price (sealed-bid) combinatorial auction, the question of equilibrium bidding strategies remains an open problem. We propose a heuristic for finding such strategies and also present feasible strategies. And finally, is a first-price combinatorial auction worth pursuing compared to the simpler simultaneous (single-item) auction? We prove, through a model capturing many fundamental properties of multiple items scenarios with synergies, that the first-price combinatorial auction produces higher revenue than simultaneous single-item auctions. We provide bounds on revenue, given a significantly more general model, in contrast to previous work.

Keywords: combinatorial auction, multiple-object, first-price, sealed-bid, game theory, multiple items, simultaneous auction, integer programming, equilibrium, strategy, reveue

Jim Wilenius, Department of Information Technology, Division of Computing Science, Box 337, Uppsala University, SE-75105 Uppsala, Sweden

ISSN 1651-6214
ISBN 978-91-554-7554-3
urn:nbn:se:uu:diva-102960 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-102960)

To my wonderful wife and son.

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I Arne Andersson, Jim Holmström[1], Mattias Willman, *"An Auction Mechanism for Polynomial-time Execution with Combinatorial Constraints"*, Seventh IEEE International Conference on E-Commerce Technology, CEC. July 2005

II Jim Wilenius, Arne Andersson, *"Discovering Equilibrium Strategies for a Combinatorial First Price Auction"*, The 9th IEEE International Conference on E-Commerce Technology and The 4th IEEE International Conference on Enterprise Computing, E-Commerce and E-Services, CEC-EEE. July 2007

III Arne Andersson, Jim Wilenius, *"A New Analysis of Revenue in the Combinatorial and Simultaneous Auction"*, Technical Report 2009-001, Department of Information Technology, Uppsala University, January 2009. (Submitted to journal)

IV Jim Wilenius, *"Combinatorial and Simultaneous Auction: A Pragmatic Approach to Tighter Bounds on Expected Revenue"*, Technical Report 2009-013, Department of Information Technology, Uppsala University, May 2009.

Comments on my Participation

- Paper-I: This was a joint work with Arne Andersson and Mattias Willman. Mattias Willman and I worked together on designing the polynomial protocol and the rank remove auction, as well as the bidding agents. Arne Andersson took part in the discussions. I am responsible for most of the implementation. I executed all the experiments and compiled the results, as well as did nearly all of the writing. Arne Andersson took part in refining the text.

[1] My last name was changed to Wilenius when I got married.

- Paper-II: I am the principal author of this paper. Arne Andersson had the initial idea of sampling the probability of winning. I had the idea of using regression of specific models, and using best response iteration. I am responsible for all the writing, modeling and implementing, as well as all the experiments. Arne Andersson took part in refining the text.
- Paper-III: This was a joint work with Arne Andersson. I wrote most of the text. All proofs, lemmas and theorems were joint work, except for Theorem 1 to which my contribution was pointing out the need for Lemma 4.1. The proof of Corollary 5.1 is mostly my work, as well as the parameterization of the lemmas for the combinatorial auction. In a previous version of this paper, these lemmas were based on a fixed problem instance.
- Paper-IV: I am the sole author of this work.

Acknowledgments

I am sincerely grateful to my supervisor Professor Arne Andersson for giving me the opportunity to study for my PhD degree at Uppsala University, for valuable input and interesting discussions, and finally for always being available.

I wish to express my sincere thanks to my co-supervisor Justin Pearson for interesting discussions and valuable help as a source of great mathematical skills, as well as for interesting games of Go.

I would also very much like to acknowledge that this thesis would not have been completed without the support of my wife Maria. Thank you for your considerable support and encouragement. Thank you also for listening to my explanations of various proofs and problems, I know it must have required infinite patience.

I would also like to thank everyone at the department of Information Technology and the Computing Science Division, especially Magnus Ågren and Mattias Wiggberg for providing essential and stimulating coffee-break and after work discussions. And finally, a huge thanks to everyone that has provided valuable comments on various versions of this thesis.

Svensk sammanfattning

Auktioner har funnits länge. De första auktionerna anses ha ägt rum omkring 500 år före Kristus men den kanske mest spektakulära var sannolikt den då praetoriangardet auktionerade ut tronen till det romerska kejsardömet år 193. I dag används auktioner nästan överallt för försäljning och upphandling av högst varierande typer av objekt. Auktioner är speciellt lämpliga när man inte vet eller är osäker på vilket pris man bör sätta på en vara.

Ofta uppstår situationer där ett flertal varor eller kontrakt ska säljas (eller köpas) vid samma tidpunkt. När budgivare värderar specifika grupper av varor högre än de enskilda varorna, uppstår ett intressant problem eftersom dessa beroenden mellan varorna påverkar en budgivares vilja att betala. Betrakta följande enkla exempel: givet ett par skor är det naturligt att båda skorna tillsammans är mer värda än bara varje sko för sig.

Följaktligen, när man vill sälja många varor vid ett och samma tillfälle till köpare som har synergier mellan specifika varor, så finns flera sätt att gå till väga. Även om vi begränsar oss till auktionen som metod finns det fortfarande ett antal varierande alternativ. Till exempel kan man välja att auktionera ut varorna en och en efter varandra i något som kallas en sekventiell auktion. Ett alternativ är att auktionera ut varorna samtidigt i flera parallella auktioner, en så kallad simultanauktion. Båda dessa alternativ faller under kategorin enkelbudsauktioner, eftersom vinnaren bestäms per vara. Ett tredje alternativ är att använda sig av en kombinatorisk auktion, där budgivare tillåts bjuda på paket av varor. Det är den typen av auktioner som behandlas i denna avhandling.

I en kombinatorisk auktion tillåts alltså budgivare lägga bud på grupper av varor, så kallade kombinationsbud. Ett kombinationsbud är odelbart, antingen vinner budgivaren alla varor som specificerats i budet, eller ingenting. Kombinatoriska auktioner anses ge bra lösningar med avseende på samhällsekonomisk effektivitet. Det är även den allmänna uppfattningen att säljaren kan förvänta sig högre intäkt jämfört med traditionella enkelbudsauktioner, eftersom budgivare med kombinationsbud kan uttrycka sina preferenser med större precision och säkerhet än de kan med enkelbud på individuella varor.

Den grundläggande skillnaden mellan traditionella enkelbudsauktioner och kombinatoriska auktioner finner vi förvisso i budens utseende men framförallt ligger den stora skillnaden i hur man går tillväga för att bestämma vilka som vinner. I enkelbudsauktioner är det normalt att den högstbjudande vinner. I en kombinatorisk auktion måste vi lösa ett svårt optimeringsproblem. Normalt vill vi hitta det pussel av ej kolliderande kombinationsbud som ger högst intäkt. Detta kan medföra att den som är högstbjudande (på en grupp varor) kanske ändå inte vinner eftersom det budet eventuellt inte ingår den kombination av bud med högst totalvärde.

I denna avhandling har vi valt att huvudsakligen studera den vanligt förekommande kombinatoriska förstaprisauktionen, en kombinatorisk auktion där budgivarna betalar värdet av det lagda budet ifall de vinner. Att teoretiskt analysera denna typ av auktion är svårt och väldigt lite har gjorts tidigare. I avhandlingen diskuteras och undersöks tre grundläggande frågor som har stor praktisk betydelse för kombinatoriska auktioner.

Till att börja med, eftersom antalet möjliga bud växer exponentiellt med antalet varor, så är det i praktiken nödvändigt att begränsa antalet bud som får läggas. Denna begränsning medför ett allvarligt problem eftersom budgivarna med stor sannolikhet inte kommer att skicka in de bud som teoretiskt sett skulle ge den optimala lösningen. Vi belyser hur allvarligt detta problem är och visar att det kan vara viktigare att använda sig av ett uttrycksfullt men kompakt budgivningsspråk, än att lösa optimeringsproblemet optimalt.

Vidare, givet en kombinatorisk förstaprisauktion, vilken budgivningsstrategi ska en budgivare följa? Detta är väl undersökt i många enkelbudsauktioner men det är fortfarande ett öppet problem för denna typ av kombinatoriska auktioner. Vi föreslår en heuristik för att hitta jämviktsstrategier och tillhandahåller även tänkbara och rimliga strategier för denna typ av auktion.

Slutligen, lönar det sig att använda en kombinatorisk förstaprisauktion givet den markant enklare simultanauktionen (med enbart enkelbud)? Vi bevisar, givet en modell som inkluderar många av de grundläggande egenskaperna i scenarier med många varor och synergier, att den kombinatoriska förstaprisauktionen ger högre förväntad intäkt än simultanauktionen. Vi tillhandahåller övre och undre gränser på den förväntade intäkten i en modell som jämfört med tidigare arbeten är avsevärt mer generell.

Contents

1. Introduction

> This thesis is mainly centered around one simple observation: *in multiple-item auctions where bidders are uncertain about each other's preferences, knowing how much and which items to bid on is hard.*

Auctions have been around for a long time. It is generally accepted that the first auctions took place as early as 500 B.C. Perhaps one of the most spectacular auctions in history occurred in 193 A.D. when the throne to the Roman Empire was auctioned to the highest bidder by the Praetorian Guard, after having killed the emperor Pertinax (Edward Gibbon [15]).

Today auctions are ubiquitous and used for selling and procuring items of highly different nature. For example, auctions are used when trading oil, gas, timber, mineral rights, radio frequency rights, services contracts, collectibles of all sorts and much more. There are numerous on-line auction websites, and auctions have even made their way into the on-line gaming scene.

Considering what defines an auction, a general and open view could be that an auction is a method of commerce where the auctioneer elicits price information from bidders through the submission of bids. A winner is selected based on an allocation rule which takes into account only the submitted bids, and the winner pays some amount specified by a payment rule. This view is general and allows for a broad variety of formats. Normally, at least in traditional (standard) auctions, the allocation rule is to award the item to the highest bidder[1], but this does not have to be the case more generally.

Many auctions concern the selling or procuring of a single item. Throughout this thesis, auctions that allow only individual bids on single items will be referred to as *single-item auctions*. Some of well known single-item auctions are (Krishna [22]):

- First-price sealed-bid auction – bidders submit their bids in a "sealed envelope", highest bidder wins and pays the amount bid.

[1] In a reverse (or procurement) auction, with one buyer and many sellers, the lowest bid presented by the sellers wins and is payed by the buyer. The analysis of selling auctions and procurement auctions is similar.

- Second-price sealed-bid (Vickrey) auction – bidders submit bids in a "sealed envelope", highest bidder wins and pays the value of the second highest bid.
- Open ascending price auction (English auction) – the price is raised with every new bid and bidders drop out of the auction when they are not willing to bid above the current price. The winner is the last remaining bidder, and he pays the price of his bid.

All of the above auctions have been thoroughly analyzed, and bidder behavior has been mapped for many settings. However, there are cases when a seller has many items that he wishes to sell simultaneously and where buyers have synergies on certain combinations of items. One such example is the shipping industry, and another is the spectrum license auctions held world wide.

There are several ways to approach the sale of multiple items, one could be to auction the items one after the other in a sequential auction. There is also the option of auctioning the items simultaneously in several parallel auctions, a simultaneous auction. Yet another way to auction multiple items is by allowing bidders to submit bids on indivisible bundles of items, which brings us to combinatorial auctions, our central topic.

This thesis concerns the interdisciplinary field of combinatorial auctions, combining the fields of computer science, operations research and economics. As such this thesis is limited to specific questions located in the intersection between these areas.

Analyzing combinatorial auctions is hard and compared to single-item auctions relatively little has been done. Since combinatorial auctions are being used more and more, there is a need to study them further. In this thesis a number of fundamental questions with great practical importance are investigated and discussed. Particularly and in a practical sense we deal with the following three specific questions.

(i) Combinatorial auctions are computationally hard, but is it always necessary to focus all energy on solving the auction optimally? We provide evidence that an expressive and compact bidding language can be more important than finding the optimal solution.

(ii) What strategy for bidding should a bidder use in a first-price combinatorial auction? We propose a heuristic for finding such strategies and also present feasible strategies.

(iii) Is a first-price combinatorial auction really worth pursuing compared to the much simpler simultaneous single-item auction? We prove, through a model capturing many fundamental properties of multiple-item scenarios with synergies, that the first-price combinatorial auction produces higher revenue than simultaneous single-item auctions. We provide bounds on revenue, given a significantly more general model, in contrast to previous work.

We approach these problems from a computer science background, which allows us to adopt a different mindset to that of traditional economics and auction theory. The work is partly to be considered practical in the sense that we investigate the widely used first-price combinatorial auctions, where the bidders pay the actual bids, and not a price determined by some abstract pricing function.

The thesis is organized as follows. After the introduction, in Chapter 2, combinatorial auctions are introduced, then in Chapter 3 a basic game theory background is covered. In Chapter 4 we discuss the auction game, give examples and discuss why analysis is hard. Chapter 5 contains summaries of the included papers.

2. Combinatorial Auctions

In this chapter we will cover the basics of combinatorial auctions, such as basic terminology, valuations and bids, and how to determine the winners. We conclude by shortly discussing some extensions and present three different types of combinatorial auctions.

2.1 Basic Terminology

In some cases the auctioneer has several items that he wants to put to auction simultaneously. It may be that bidders have *synergies* on combinations of items, that is, they value specific combinations of items higher than the items separately. When this is the case, combinatorial auctions allow bidders to express their specific preferences more precisely and with less risk than what is possible using separate bids on single items. This is achieved by allowing bidders to submit so called *combinatorial bids*, that is, bids on indivisible combinations. With a combinatorial bid you are guaranteed to either win all items specified in the bid on none at all. Table 2.1 illustrates a selling auction and a possible bidding scenario with three bidders and six items. For example, Bidder-A is willing to pay 33 if he gets the items 1, 3 and 4.

	Bid	Item-1	Item-2	Item-3	Item-4	Item-5	Item-6
Bidder-A	33	$\bullet_{A1}$		$\bullet_{A1}$	$\bullet_{A1}$		
Bidder-A	23	$\bullet_{A2}$	$\bullet_{A2}$				
Bidder-A	55				$\bullet_{A3}$	$\bullet_{A3}$	$\bullet_{A3}$
Bidder-B	32			$\bullet_{B1}$	$\bullet_{B1}$		
Bidder-B	27	$\bullet_{B2}$		$\bullet_{B2}$			
Bidder-C	29					$\bullet_{C1}$	$\bullet_{C1}$

Table 2.1: *Example of Combinatorial Bids*

An example: consider a company that wants to renegotiate all their shipping contracts. By using a combinatorial procurement auction (buying auction), logistics companies can submit bids on combinations of contracts where shipping lanes are close geographically, or perhaps combinations with both directions on certain or groups of lanes. Since obtaining such lanes will reduce overhead costs, better prices are possible.

In the field of combinatorial auctions, some of the key areas of study are:

- Computational complexity – the optimization problem of determining who wins belongs to a particular class of problems that are very hard, known as NP-hard problems.
- Bidding languages and communication – communicating and representing a possibly exponential[1] number of bids.
- Strategic behavior – how to bid given the private information and assumptions about the competitors?
- Mechanism design – designing the rules to achieve a specific outcome or strategic behavior.

Combinatorial auctions can be used for both selling and buying (procurement). However, we will restrict our discussions to selling auctions, although some examples may be procurement auctions[2].

The literature has yet to converge on a completely standard notation, we therefore choose a notation that we hope achieves readability and simplicity, and at the same time remains complete enough to avoid misunderstanding.

A combinatorial auction is an auction where M heterogeneous indivisible items are auctioned off simultaneously to N competing bidders. Bidders are allowed to submit bids on indivisible subsets of items, and the auctioneer decides who wins by selecting the non-conflicting bids that gives him the highest revenue. The winning bidders' payments may or may not be trivially calculated depending on the type of combinatorial auction being used. In Section 2.7 three main types of combinatorial auctions will be discussed. It is also conceivable that multiple units of each item are being sold in a multi-unit combinatorial auction, however this is outside the scope of this thesis. The description above is a somewhat simplified compared to some real world implementations, but it does capture the essence of what a combinatorial auction is.

In the literature it is sometimes assumed that bidders submit bids for all possible (non-empty) combinations of items, and that at most one of each bidder's bids can win. However, in this thesis bidders are allowed to submit bids on subsets of their choosing. No constraint is put on the number of bids

[1]Given items M there are $2^{|M|} - 1$ possible combinations.

[2]The analysis of procurement auctions is similar, however some details differ. For example unless negative bids are allowed, in a procurement auction there is a definite minimum, 0, but in a selling auction the limit is infinity.

that can win, and if needed, exclusive OR (XOR)[3] constraints may be expressed through the use of phantom items (phantom items are discussed in Section 2.4).

For a more complete description, the more common concepts are defined in the following sections.

2.2 Valuations and Bids

Bidders have *valuations*, which represent their willingness to pay. In the case of a single item for sale bidders have only one valuation. In a combinatorial auction, bidders typically have many valuations corresponding to the various combinations of interest. It is common to define the valuation as a function as in the following definition.

Definition 2.2.1 (Valuation)**.** A valuation function v is a function from the set of (non-empty) subsets of items to the positive reals (including 0), corresponding to the value a bidder associates with obtaining a specific set of items.

Although a bidder's valuations are defined as a function, they could equivalently be defined simply as a set of values, indeed this will be the preferred notation in the special case when bidders are interested in a single subset of items. Definition 2.2.1 above is unspecific when it comes to possible interdependencies between valuations. If, given two disjoint subsets of items S_1 and S_2, $v(S_1 \cup S_2) \geq v(S_1) + v(S_2)$, the valuation is said to be *super-additive* and S_1 and S_2 are *complements* which simply means that winning the items of both sets has at least the same if not a higher value. On the other hand if $v(S_1 \cup S_2) \leq v(S_1) + v(S_2)$ the valuation is said to be *sub-additive* and S_1 and S_2 are *substitutes*.

Free disposal is typically assumed either for the auctioneer or the bidders. When free disposal is assumed on behalf of the bidders this means that given some set of items S then for all T such that $S \subseteq T$ we have $v(T) \geq v(S)$, that is winning additional items to the ones specified will not decrease the value. However, since the main idea often is to allow for the possibility that some items will not be allocated to any bidder, an equivalent approach is to assume free disposal on behalf of the auctioneer, we have chosen the latter view. Without free disposal finding a feasible allocation of items may not always be possible. Without free disposal every item must be sold and bidders will not accept additional items to the ones they explicitly bid on. This creates a possi-

[3] Given two bids A and B, an exclusive OR (XOR) of the two bids means either A or B but not both, whereas an OR of the bids means A or B or both.

bly problematic situation where finding an allocation that includes all items is necessary. Also, if such an allocation is found it could possibly have a smaller total value than if some items are allowed to be unallocated. Quite often, the matter of free disposal is not mentioned at all, but implicitly assumed in the definition of the allocation problem.

Definition 2.2.1 assumes no inherent structure of the valuation function. In much of our work valuations have a structure that is dependent on the number of items and a synergy constant α. For example, the valuation of a set S of size $|S|$ could be $|S| \cdot (x + \alpha)$ where x is a random value in the interval $[0, 1]$ and $\alpha > 0$. This is one simple way to describe valuations where larger combinations have a greater value, but where both sub-additive and super-additive valuations can occur. Another (natural) way to describe valuations for combinations is to use single-item valuations as a basis and then add or subtract values to get the combination valuation.

Given that a bidder knows his valuations, he must then decide which combinations to bid on and what the values of his bids should be. A bid in a combinatorial auction constitutes in the simplest case a value associated with a set of items. In Section 4.2 we consider the strategic aspects of combinatorial auctions, i.e. how a bidder determines what to bid given a set of items and a valuation, for now however we will be satisfied with simply accepting that bidders submit bids. The following straightforward definition of a bid is chosen.

Definition 2.2.2 (Bid)**.** A bid is a tuple $\langle S, b \rangle$, where $S \subseteq M$ is a set of items and $b \in \mathbb{R}^+$ is the value of the bid.

Two bids $\langle S_i, b_i \rangle$ and $\langle S_j, b_j \rangle$ are disjoint (or non-conflicting) if and only if $S_i \cap S_j = \emptyset$, that is if they are bids on completely different items.

2.3 Winner Determination

An auctioneer in a combinatorial auction is thus faced with several bidders each bidding potentially an exponential number of bids. Given these bids, the auctioneer typically wants to find the revenue maximizing combination of non-conflicting bids, this problem is known as the *winner determination problem* and is generally a hard problem [10].

Definition 2.3.1 (Winner Determination Problem)**.** Given a set of bids $B = \{\langle S_0, b_0 \rangle, \ldots, \langle S_k, b_k \rangle\}$, the winner determination problem (WDP) is to find the subset of disjoint bids such that the sum of values b_i is maximized.

Compared to single-item auctions, where winner determination is trivial, determining winners a combinatorial auction requires solving a hard combinatorial puzzle. This is one of the fundamental differences between single-item auctions and combinatorial auctions. The winner determination problem is an NP-hard[4] optimization problem in the general case. However, Rothkopf, Pekeč and Harstad [48] provide several special cases for which winner determination is tractable. A few examples include: (i) when bids form a tree structure the optimal solution can be found in less than $o(|M|^3)$ time, and (ii) when bids are limited to geometric structures where items are totally ordered, then allowing bids only on consecutive items makes the winner determination problem tractable.

The proof of NP-hardness of the general winner determination problem is based on a reduction from weighted set packing, a known NP-hard optimization problem. Set packing is one of the 21 original NP-complete decision problems presented by Karp [21]. The optimization version is the maximum set packing, and weighted set packing is the weighted version of the optimization problem.

Moreover, Håstad [20] proves that the maximum clique problem cannot be approximated within a constant factor[5] less than $n^{1-\varepsilon}$ for any $\varepsilon > 0$ (where n is the number of nodes in the graph). Based on Håstad's result Sandholm [50] provides the relevant analog for the winner determination problem. Thus, for any fixed $\varepsilon > 0$ there is no polynomial time algorithm that approximates the winner determination problem within $min(l^{1-\varepsilon}, |M|^{1/2-\varepsilon})$, where l is the number of bids and $|M|$ is the number of items.

Several authors have noted that the winner determination problem can be formulated as an integer program [45, 39, 58]. An alternative method to solving the winner determination problem is then to solve the corresponding integer program. Andersson, Tenhunen and Ygge [1] show that standard commercial optimizers for solving the integer program performs excellently on many problems. They also note that extending the model and solving more general problems demands less work than with specialized custom algorithms.

The following useful observation is made. Subsets of items M can be represented as binary vectors $\left[q_1, ..., q_{|M|}\right]$ where $q_j = 1$ if and only if item j is part of the subset, and $q_j = 0$ otherwise. Thus, bid i can be represented as a vector of $\left[q_{i1}, ..., q_{i|M|}\right]$, and a value b_i. This observation allows us to clearly and straightforwardly define an integer program for solving the winner determination problem.

[4]For more about the theory of NP-completeness, see [13]. A simplified explanation: we could say that for problems that are NP-complete (or NP-hard, depending on the formulation) there are no known *efficient* algorithms that solve all instances of a that problem. By efficient we mean that the running time of the algorithm is a polynomial function of the size of the input.
[5]Unless ZPP=NP. ZPP=Zero error Probabilistic Polynomial time.

Definition 2.3.2 (Integer Program for Winner Determination)**.**

$$\underset{x_i}{\text{maximize}} \sum_{i=1}^{k} b_i x_i \qquad \text{such that for each } 1 \leq j \leq |M|, \quad \sum_{i=1}^{k} q_{ij} x_i \leq 1$$

where:

k	is the number of bids.
b_i	is the constant value of bid i, as in Definition 2.2.2.
x_i	binary variable such that $x_i = 1$ if and only if bid i wins.
q_{ij}	is a binary constant denoting membership of item j in bid i.

The optimization problem is to find the optimal allocation of values $\{0,1\}$ to variables x_i such that the sum of bid-values is maximized while at the same time maintaining the constraint that each item j can be allocated to at most one bid. This definition implicitly assumes some form of free disposal since there is a possibility that some items may be left unassigned.

Closely related to the winner determination formulation is the concept of social welfare and efficiency. An auction that allocates items to the bidders that value them the most is said to be *efficient*. Assuming that bidders bid their true values for all subsets where they have a value greater than zero, then the optimal solution to the winner determination problem is the efficient solution. Economic efficiency and revenue are the two main goals which an auctioneer may want to maximize, typically there is a trade-off between the two. Companies or individuals normally want to maximize revenue, whereas in government held procurement auctions or the sale of public goods, efficiency may be the primary objective.

2.4 Bidding Language

It may not be feasible for bidders to submit bids for every combination of items, therefore a *bidding language* that describe how bids may be expressed is of practical importance.

The definition of the winner determination problem implicitly describes a bidding language, in Definition 2.3.1 the OR bidding language is implicit. The same definition could also handle XOR constraints if phantom items (see below) are introduced.

As an example of why efficient bidding languages are needed, consider a perfectly feasible auction with 30 items and 20 bidders. Assume that a bid can somehow be represented completely with 4 bytes and that bidders have valuations for all combinations. To express every possible bid, each bidder must communicate approximately $4 \cdot 2^{30} = 4GB$ of data, although somewhat large, it is not inconceivable. However, if some form of decision must be made when calculating valuations and bids, for example performing some relatively time consuming database lookups and/or performing non-trivial calculations then the time required to create all bids will be a severe problem. Assuming an average time of $1/4$ of a second for each combination, the time required for generating all the bids will be at least 8.5 years. Even if the average time required per combination is 10 milliseconds, it will still take more than four months. Leaving the bidder side problem, there is (currently) no efficient way to handle more than $21 \cdot 10^9$ bids in any solver. It is clear that restricting bids and employing more efficient representation through an appropriate bidding language is crucial.

Nissan [39] collects, formalizes and analyzes several bidding languages from the literature, for example the OR language, which corresponds to the one discussed so far, but also the XOR language, and combinations thereof. The most notable language is the OR*, which is the OR bidding language with *phantom items*. By introducing phantom items that hold no intrinsic value to anyone, it is possible to express XOR constraints on general collections of bids. The bids that contain the same phantom item are necessarily mutually exclusive. For example, to express that the bids $\langle\{1,2,3\},23\rangle$ and $\langle\{4,5\},13\rangle$ are mutually exclusive simply add a phantom item p_0 to both bids, $\langle\{1,2,3,p_0\},23\rangle$ and $\langle\{4,5,p_0\},13\rangle$. Since the phantom item p_0, like any ordinary item, can be allocated at most once, only one of the bids can win.

Many other languages are conceivable, for example min-bundles which are discussed in Paper-I. Min-bundles compactly expresses bids on all combinations of the specified items as long as the size is at least equal to the provided minimum constraint. For example, the min-bundle $\langle 2,\{a,b,c\}\rangle$ expresses bids on the following sets: $\{\{a,b,c\},\{a,b\},\{a,c\},\{b,c\}\}$. Submitting several min-bundles means that more than one min-bundle can win, which represents an OR of min-bundles.

2.5 Single-minded Bidders

We know that solving the winner determination problem is generally NP-hard, but what happens if we introduce the limitation that a bidder is interested in only one combination of items? Unfortunately the winner determination problem remains NP-hard even under such a drastic limitation. Also, approximating the optimally efficient allocation within a factor better than $|M|^{1/2-\varepsilon}$ is NP-hard (proposition 11.6 in [40]).

Definition 2.5.1 (Single-minded)**.** Bidder i, with valuation function v_i, is single-minded if for one specific subset $S^* \subseteq M$ there exists a value $v_i^* \in \mathbb{R}$ such that $v_i(S^*) = v_i^*$ and $v_i(S') = 0$ for all other subsets $S' \neq S$.

A *single-minded* bidder submits one combinatorial bid $\langle S_i^*, v_i^* \rangle$. If we relax the definition of single-mindedness to account for free disposal on the behalf of bidders, then $v_i(S) = v_i(S^*) = v_i^*$ for all sets S such that $S^* \subseteq S$.

Definition 2.5.2 (Weakly single-minded)**.** Bidder i is weakly single-minded if for one specific subset $S \subseteq M$ he has:

(i) For each item $s \in S$, a valuation $v_i(\{s\}) \geq 0$.

(ii) For the entire set S, a valuation $v_i(S) > \sum_{s \in S} v_i(\{s\})$.

(iii) For every subset $T \subset S$, $v_i(T) = \sum_{t \in T} v_i(\{t\})$.

(iv) For all other subsets $S' \subseteq M$, $v_i(S') = 0$.

A *weakly single-minded* bidder is single-minded in the sense that he is interested only in a specific group of items and that he receives a synergy when winning the entire combination. The difference from strict single-mindedness is that he is allowed to have a valuation for each individual item in the combination. This form of single-mindedness represents a fairly natural setting where each item in the set has a value, however where attaining the entire set is worth something extra. The results in Paper-III are based on this more relaxed version of single-minded bidders.

Given single-minded bidders it is possible to construct a tractable greedy algorithm that achieves a $\sqrt{|M|}$ approximation of the optimal efficiency. Lehman, O'Callaghan and Shoam [25] propose the following algorithm which is *incentive compatible*, that is, it is optimal for each bidder to bid his true valuation.

Each bid $\langle S_i^*, v_i^* \rangle$ is given a rank $v_i^* / \sqrt{|S_i^*|}$, and all bids are ordered according to rank. Let $B_1, \ldots, B_N$ be the ordering such that B_1 corresponds to the bid with highest rank and so on. The solution is constructed by greedily selecting bids that are disjoint starting with the highest ranking bid B_1. Losing bidders pay zero and each winning bidder pays the minimum amount that was required for him to be selected into the solution. Let $j > i$ be the index of the first bid after B_i such that for each index $k < i$, $B_k \cap B_j = \emptyset$ and $B_i \cap B_j \neq \emptyset$, then the payment for bid i is

$$p_i = \begin{cases} 0 & \text{if no such } B_j \text{ exists} \\ \frac{v_j^*}{\sqrt{|S_j^*|}} \cdot \sqrt{|S_i^*|} & \text{otherwise} \end{cases}$$

It is easy to see that with only one item for sale, this corresponds to the standard second-price single-item auction.

Mu'alem and Nisan [31] consider a model with single-minded bidders where the desired subset of each bidder (but not the valuation) is known by the mechanism. They provide a variety of polynomial mechanisms based on this assumption and also provide a truthful mechanism which guarantees an $\varepsilon\sqrt{|M|}$

approximation. Ledyard [24] constructs the optimal combinatorial auction under a similar assumption where the auctioneer knows the desired subsets.

2.6 Real World Application

Up till now winner determination has been mentioned only in the simple setting where the auctioneer wants to find the revenue maximizing collection of non-conflicting bids. However, both the auctioneer and the bidders may have additional constraints and business rules that must be satisfied. Therefore, the concept of winner determination sometimes needs to be expanded to handle the specific business rules and constraints that the auctioneer specifies.

These extensions may be needed in many practical settings, but for the remainder of this thesis, the original definitions will be enough, in fact in some cases even simpler models will be used.

There are many examples of combinatorial auctions for practical use. Rassenti, Smith and Bulfin [45] present a sealed-bid combinatorial auction for the allocation of airport time slots. In estate auctions, simple combinatorial auctions have been used for decades. Combinatorial auctions for radio spectrum rights were conducted in both the United States and Nigeria. Recently in the United Kingdom, combinatorial auctions based largely on the clock-proxy auction (Ausubel, Cramton and Milgrom [4]) have been proposed.

Some examples of first-price combinatorial auctions.

- Government construction projects (Bernheim and Whinston [6])
- School bus contracts around Manchester (Letchford [26])
- School meal contracts in Chile (Epstein et al. [12])
- Bus routes in London (Cantillon and Pesendorfer [7])

2.7 Types of Combinatorial Auctions

There are several types of combinatorial auctions. Although they all share the winner determination complexity they possess different properties that distinguishes them from each other in important ways. We will take a quick look at three different yet representative auctions, starting with the most straightforward.

• In the *first-price sealed-bid* combinatorial auction bidders submit their bids in a one-shot fashion (in a "sealed envelope"), without knowledge of each other's bids. The price payed by the winners is the amount of their winning bids. This is the most straightforward of all the combinatorial auctions. However, there are two problems that stand out: (i) the optimal bidder behavior is an open problem; (ii) because of the one-shot nature of the auction, the auctioneers faces a potential problem with the sheer number of bids that the bidders might submit. When bidders have expensive (computationally or monetarily) valuation problems, submitting a large number of bids is a problem also on this end.

On the other hand, some desirable properties carry over from the single-item case, such as resistance to both collusive behavior (Robinson [47]) and bid-signaling, i.e. using bids to signal agreements not to bid on each others combinations. Also, sealed-bid auctions in general, including combinatorial auctions, encourage participation (Pekeč and Rothkopf [44]). The first-price combinatorial auction is one of the more transparent combinatorial auctions. The rules are fairly simple and each bidder can trivially verify their payment in the event of winning, a property that may be considered important by bidders.

• An *iterative* combinatorial auction is a combinatorial auction that executes in multiple rounds. This class of combinatorial auctions is designed to address settings when bidders have hard or costly valuation problems. If a bidder must solve for example a potentially hard scheduling problem when determining each valuation, then calculating all or perhaps just some of the $2^{|M|} - 1$ valuations even for a small number of items may not be feasible. The purpose of the multiple round construction is to guide bidders so that they evaluate valuations only for presumably meaningful combinations. Bidders are guided by feedback in each round with regards to prices and allocations. A natural drawback with the iterative nature is that winner determination must be solved in each round. Another issue of a strategic nature is that bidders may choose not to participate until the last round of the auction, in this case the purpose is lost completely. Iterative auctions are also sensitive to both bid signaling and collusive behavior. Iterative combinatorial auctions include the AUSM [5], iBundle [42, 43] and AkBA[59], and others [3, 2].

• The most famous combinatorial auction is probably the *Vickrey-Clarke-Groves* (VCG) mechanism which corresponds to a second-price version of a combinatorial auction. This auction is widely used as a baseline for comparisons throughout the literature. The VCG stems from the work of Vickrey [54], Clarke [9] and Groves [16] and has a number of desirable properties, but also some very undesirable properties. The VCG for the sale of M heterogeneous items to N bidders requires that each bidder i bids (truthfully) for every subset of items, even for subsets that he believes unlikely to win. Omitting bids may affect the payments by the other bidders.

Determining winners is done by selecting the value maximizing allocation of disjoint bids, that is, solving the winner determination problem where bidders bid their true valuations. The VCG payments are calculated as follows. Let V^* be the total value of the optimal allocation, and for each bidder i let V^*_{-i} be the total value of the optimal allocation when bidder i does not participate. Let V_i be the value of bidder i's winning subset. Losing bidders pay zero and the payment made by each winning bidder is $p_i = V_i - (V^* - V^*_{-i})$. The payment calculations are not trivial, the winner determination problem must be solved anew for each payment calculation.

An important advantage of the VCG is that the optimal bidding strategy for each bidder is to bid the true values regardless of the behavior of the other bidders. This means bidders do not have to spend resources and time devising strategies. Despite this nice property, the VCG suffers from some serious problems (unless valuations are limited to substitues, see Cramton, Shoham and Steinberg [10] for more details).

- Low or even zero revenues can occur.
- Sensitive to collusion and multiple bidding identities.
- Non-monotone: increasing bids may lead to lower revenues.

Much work has been done regarding the Vickrey-Clarke-Groves mechanism, however this is outside the scope of this thesis, see Milgrom [29], Krishna [22] for more details on the VCG, and Rothkopf [49] for more details on the problems of the VCG. Attempts have been made to boost revenue in the VCG, see for example Likhodedov and Sandholm [27].

The following is an example of the VCG with 2 bidders and two items.

	item 1	item 2	item 1 & 2
Bidder 1	10	5	15
Bidder 2	1	6	12

Clearly the optimal allocation has a value of 16, assigning item 1 to bidder 1 and item 2 to bidder 2. The optimal solution without bidder 1 is 12 and the optimal solution without bidder 2 is 15. We thus have:

$$
\begin{array}{c|c|c}
V^* = 16 & V^*_{-1} = 12 & V_1 = 10 \\
 & V^*_{-2} = 15 & V_2 = 6
\end{array}
$$

The payments by bidder 1 and 2 are:

$$
\begin{aligned}
p_1 &= V_1 - (V^* - V^*_{-1}) = 10 - (16 - 12) = 6 \ , \\
p_2 &= V_2 - (V^* - V^*_{-2}) = 6 - (16 - 15) = 5
\end{aligned}
$$

and the auctioneer receives $p_1 + p_2 = 11$. Now, consider the same example as above, but where bidder 2 instead of bidding 6 for item 2, bids 10. Now the optimal allocation has a total value of 20. The new payments will be $10 - (10 + 10 - 12) = 2$ for bidder 1, and for bidder 2, $10 - (10 + 10 - 15) = 5$. The result of bidder 2 submitting a higher bid illustrates the non-monotonicity problem, the new revenue, 7, is lower compared to before.

3. Game Theory

3.1 Introduction

So far we have not mentioned anything regarding the strategic problems that bidders face when bidding in an auction. Game theory is is an important tool in the theory of auctions, it is used to analyze the strategic behavior of bidders, and to analyze equilibrium outcomes in different auctions to obtain a measure of efficiency and revenue.

Game theory has been heavily influenced by the works of von Neumann [55], von Neumann and Morgenstern [56], Nash [36, 35, 34, 37, 38], and given the type of games that we are concerned with here, perhaps most notably by Harsanyi [17, 18, 19] for his contribution to the theory of Bayesian games. In this section a brief introduction will be provided, for a more complete introduction see for example Myerson [33] and Osborne and Rubinstein [41]. We begin by looking at a well known yet informative example.

> ***Prisoner's Dilemma*** *(Osborne and Rubinstein [41])*
> Two suspects in a crime are put into separate cells. If they both confess, each will be sentenced to three years in prison. If only one of them confesses, he will be freed and used as a witness against the other, who will receive a sentence of four years. If neither confesses, they will both be convicted of a minor offense and spend one year in prison.

This situation can be modeled as a *strategic game* where each *player* (suspect) has two *(pure) strategies* $\{\text{Silent}, \text{Confess}\}$. Let a pair (s_1, s_2) correspond to the strategies chosen by player 1 and player 2 respectively. Each player has a preference for all pairs of strategies. Specifically for player 1, the strategy pairs ordered from best to worst are:

$$(\text{Confess}, \text{Silent}) \succ_1 (\text{Silent}, \text{Silent}) \succ_1 (\text{Confess}, \text{Confess}) \succ_1 (\text{Silent}, \text{Confess}) ,$$

that is, player 1 prefers the scenario where he chooses Confess and player 2 chooses Silent over when they both choose Silent, and so on. The corresponding preference ordering for player 2 is:

$$(\text{Silent}, \text{Confess}) \succ_2 (\text{Silent}, \text{Silent}) \succ_2 (\text{Confess}, \text{Confess}) \succ_2 (\text{Confess}, \text{Silent}) .$$

Given the preferences of the players, we can now assign values to each preference such that the order is maintained. A function that does this is called a *payoff* (or *utility*) function. Let u_1 be the payoff function of player 1 and u_2 that of player 2. The following simple assignments maintain the preference order for player 1, $u_1(\text{Confess}, \text{Silent}) = 3$, $u_1(\text{Silent}, \text{Silent}) = 2$, $u_1(\text{Confess}, \text{Confess}) = 1$, $u_1(\text{Silent}, \text{Confess}) = 0$, and for player 2, $u_2(\text{Silent}, \text{Confess}) = 3$, $u_2(\text{Silent}, \text{Silent}) = 2$, $u_2(\text{Confess}, \text{Confess}) = 1$, $u_2(\text{Confess}, \text{Silent}) = 0$. We can favorably represent this in a table as in Table 3.1, with payoffs as pairs (u_1, u_2).

		Player 2	
		Silent	Confess
Player 1	Silent	(2,2)	(0,3)
	Confess	(3,0)	(1,1)

Table 3.1: *A representation of a strategic game in Normal-form*

Writing the set $\{Y_i\}_N$ for short when meaning $\{Y_1, ..., Y_N\}$, then a strategic game can be defined as follows.

Definition 3.1.1 (Strategic game)**.** The tuple $\left(N, \{C_i\}_N, \{u_i\}_N\right)$ is a strategic game with N players. Each player i has a set of strategies C_i, and a payoff function $u_i : C \to \mathbb{R}$, where $C = C_1 \times \cdots \times C_N$.

The game is actually defined on a preference relation, however when a payoff function is used to represent the relation then Definition 3.1.1 will be used. A strategic game is said to be *finite* when, as in the case with the Prisoner's dilemma, the set C is finite and otherwise *infinite*.

Now, what should the suspects do? Game theory assumes that players are *rational*. This means that players are aware of their alternatives and chooses the best option at hand, that is, each player wants to choose a strategy that gives the maximal payoff given the strategies of the other players.

Consider the rational player 1, if player 2 chooses Silent then player 1 should choose Confess, this will give him a payoff of 3, the best possible outcome. If player 2 chooses Confess, then player 1 should still choose Confess, because if player 1 were to chose Silent, then he will get a payoff of 0 which is less than the payoff from choosing Confess. The same reasoning also applies to player 2, which means that both players should choose Confess, landing them solidly in prison for three dreadful years each. This may seem strange considering that if they both choose Silent then they will spend only one year in prison. The reason for this is the rational behavior, it is not rational of player 1 to assume that player 2 will be benevolent and choose Silent when he can gain by choosing Confess.

Thus, each player will chose the pure strategy Confess, and there is no incentive for any of the players to deviate from this choice. We call this concept an *equilibrium*, a steady state in the game. In some sense this is the solution to the Prisoner's dilemma. Game theory offers different solution concepts that are sensible and stable in some way. The solution to the Prisoner's dilemma constitutes a *dominant strategy equilibrium*, where dominant refers to the fact that the optimal strategy is to choose Confess, no matter what the strategy of the other player may be. Solutions, if they exist, for different games do not always have this dominant property. Furthermore, strategies do not have to be pure, the alternative to a pure strategy is a *mixed* strategy consisting of a set of pure strategies each associated with a probability of being chosen.

3.2 Bayesian Games

The Prisoner's dilemma is a useful example as it introduces many of the basic concepts in a concise way. However, there are certainly scenarios when players would have private information not known to the other players. For example, it is certainly conceivable in auctions. When this is the case we are faced with a game of *incomplete* information. We distinguish between the following, *incomplete* information is when some of the parameters necessary for game theoretic analysis are missing, the opposite is a *complete* information game where every player knows all parameters, this precludes the existence any private information. A game of *imperfect* information is when some players may be uninformed about some or all of the moves (actions taken) made by other players, and *perfect* information means that all moves have been observed by all players.

Harsanyi [17, 18, 19] introduced the notion of Bayesian games, an approach that deals with the case when players have private information (i.e. a game of incomplete information) by considering that players have beliefs about other players' private information. In a Bayesian game, the uncertainty that some players have about the private information of other players is modeled as stochastic variables, taking into account players' beliefs about the distribution of these variables. This changes the incomplete structure to a completely specified structure with stochastic variables and known probability distributions. Adopting this point of view, the game can be seen as a game with complete but imperfect information.

A generalization of the strategic game is used to represent games of incomplete information. Writing the set $\{Y_i\}_N$ for short when meaning $\{Y_1, ..., Y_N\}$ and borrowing from Myerson's definition (Myerson [33]). A game of incomplete information (*Bayesian* game) $\left(N, \{C_i\}_N, \{T_i\}_N, \{p_i\}_N, \{u_i\}_N\right)$ with N players, has for each player $1 \leq i \leq N$ a set C_i of *actions*, a set T_i of *types* representing a players private information, a probability function p_i, and a payoff function u_i. Let the set of action profiles be $C = C_1 \times \cdots \times C_N$, with elements $c = (c_1, \ldots, c_N)$. Let the set of type profiles be $T = T_1 \times \cdots \times T_N$, with elements $t = (t_1, \ldots, t_N)$. We write T_{-i} for short when we mean the product set $T_1 \times \cdots \times T_{i-1} \times T_{i+1} \times \cdots \times T_N$ and similarly we write t_{-i} when we mean a type profile without player i. A *pure strategy* in a Bayesian game is a function $s_i : T_i \rightarrow C_i$, that is, action as a function of type. Although mixed strategies are needed in general, we will limit ourselves to only pure strategies.

The payoff $u_i(t,c)$ is a function $u_i : T \times C \rightarrow \mathbb{R}$, that is the payoff that player i receives when players types are t and players choose actions c. The conditional probability function $p_i(t_{-i}|t_i)$ represents player i's belief of the distribution of types t_{-i} given that player i's type is t_i. In the case when for every player i, $p_i(t_{-i}|t_i) = p(t)/p_i(t_i)$, where $p(t)$ is the probability that players' types are t and $p_i(t_i)$ is the probability that player i's type is t_i, players beliefs are said to be consistent with a common prior, this is a very common assumption in many models. In this case a Bayesian game can be represented by an equivalent complete but imperfect information game where a special player "nature" or "chance" is added to make an initial move to decide the type profile t based on the distribution over T which is considered to be common knowledge amongst the players. Nature's move is only partly observed by each player which makes the game one of imperfect information. This interpretation implies that we imagine the game beginning at an earlier point in time, before the players know their private information.

As an example, consider the following simple game:

> ***Two player card game*** *(Myerson [33])*
> Player 1 and player 2 each put one dollar in the pot. Next player 1 draws a card from a randomized deck in which half the cards are black and half the cards are red. After having looked at the card without showing it to player 2, player 1 decides if he want to fold or raise. If player 1 folds, the game ends and he shows the card to player 2. If the card is black player 2 wins the pot, but if it is red then player 1 wins the pot. If player 1 instead raises then he adds one dollar to the pot and now player 2 must decide if he wants to meet or pass. If player 2 passes then the game ends and player 1 wins the pot. If player 2 meets then he must add another dollar to the pot and the game ends with player 1 showing the card to player 2. As before, if the card is black then player 2 wins the pot and if it is red then player 1 wins the pot.

This game can be represented in its extensive form as in Figure 3.1, where we have added a root node with two edges representing a chance move that determines the type of player 1, it is the state of the game before player 1 draws his card from the deck. From the root node, two states are reachable each with probability 0.5. The labels 1.*red* and 1.*black* represent player 1's information sets. Player 1 draws either a red or black card from the deck and this determines his information set. By drawing the card player 1 observes the move that was made by the player *chance*.

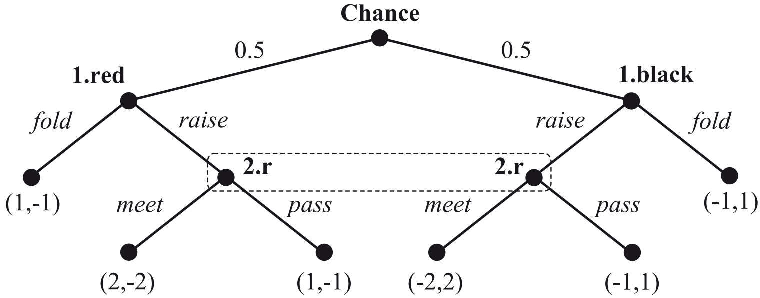

Figure 3.1: Extensive form representation of a simple card game (Myerson [33])

The label 2.r means that it is player 2's turn and that he has observed player 1 making the move *raise*, he does not know if player 1 has a red or black card (he did not observe the move made by *chance*), he can therefore not distinguish from his two states in the game, this is emphasized by the dashed box. The choice of *red*, *black* and r as labels for the different states is has no inherent meaning and can be chosen arbitrarily. The numbers in parenthesis (u_1, u_2) represents the payoffs for each player from that particular outcome.

Changing our perspective slightly by assuming that player 1 draws the card from the deck and looks at it before the game begins, the game then becomes a Bayesian game (of incomplete information) with representation $N = 2$, $T_1 = \{red, black\}$, $T_2 = \{r\}$, $C_1 = \{raise, fold\}$, and $C_2 = \{meet, pass\}$. Player 1's conditional probability function is $p_1(r|red) = p_1(r|black) = 1$ since player 1 has no uncertainty regarding the type of player 2. Player 2 assigns equal probability that player 1 is of either of his type's, $p_2(red|r) = p_2(black|r) = 0.5$. The payoff functions depend only on c_1, c_2 and t_1 and are defined in Table 3.2, where (for completeness) all combinations of actions are displayed.

		Player 2			Player 2	
	$t_1 = red$	meet	pass	$t_1 = black$	meet	pass
Player 1	raise	(2,-2)	(1,-1)	raise	(-2,2)	(-1,1)
	fold	(1,-1)	*(1,-1)*	fold	(-1,1)	*(-1,1)*

Table 3.2: *The payoffs in the two states of the game. Player 1 knows which state the game is in but player 2 does not.*

3.3 Equilibrium

So far we have mentioned little about steady states or equilibria in games. An equilibrium in a game, either with mixed or pure strategies, is a desirable property. It allows us to say something about the behavior of players and to analyze outcomes in a more or less predictable manner. In the Prisoner's dilemma both criminals will choose the Inform strategy which constitues a dominant strategy equilibrium[1]. John Nash [35] introduced an equilibrium concept that is now known as Nash equilibrium which is the most widely used in game theory. A pure strategy Nash equilibrium is when no single player can find an alternative strategy that gives a higher payoff given the strategies played by the other players, that is, every player's strategy is the best response to the strategies of the other players. Using the same short notation as before, that is, writing $\{Y_i\}_N$ instead of $\{Y_1, ..., Y_N\}$ then formally,

Definition 3.3.1 (Nash equilibrium)**.** A Nash equilibrium of a strategic game $\big(N, \{C_i\}_N, \{u_i\}_N\big)$, is a strategy profile $c^* = (c^*_{-i}, c^*_i)$ such that for each player i it is true that for all strategies $c_i \in C_i$,

$$u_i(c^*) \;\; \geq \;\; u_i((c^*_{-i}, c_i)) \; .$$

In a Bayesian game, a Bayes-Nash equilibrium states that each player chooses the expected payoff maximizing strategy given the strategies of the other players and his beliefs of the other players types. The definition is very similar to the Nash equilibrium, in fact a Bayes-Nash equilibrium is a Nash equilibrium in the type centric view of a Bayesian game. Instead of players each with several possible types, the type centric view simply means that each type of each player is itself regarded as a player. For example, in a Bayesian game with two players each with two types then in the type centric view there are four players. In the Nash equilibrium of the type centric Bayesian game, each (type centric) player is maximizing the expected payoff given the other players strategies and his own beliefs about their types. Let $s(t)$ be the vector of strategies $(s_1(t_1), s_2(t_2), \ldots, s_N(t_N))$ and $s_{-i}(t_{-i})$ the vector without the element $s_i(t_i)$. As before we write$\{Y_i\}_N$ instead of $\{Y_1, ..., Y_N\}$.

Definition 3.3.2 (Pure Bayesian Nash equilibrium)**.** A pure strategy Bayesian Nash equilibrium of a Bayesian game $\big(N, \{C_i\}_N, \{T_i\}_N, \{p_i\}_N, \{u_i\}_N\big)$, is a pure strategy profile $s^* = (s^*_{-i}, s^*_i)$ such that for each player i and types $t_i \in T_i$ it is true that for all pure strategies $s_i : T_i \to C_i$,

$$\sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \cdot u_i(t, s^*(t)) \;\; \geq \;\; \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) \cdot u_i(t, \; (s^*_{-i}(t_{-i}), s_i(t_i)) \;)$$

[1] Although, this is not the global payoff-maximum in the game.

If all players use the same strategy in equilibrium, we call it a *symmetric* equilibrium.

With infinite type sets the theory of Bayesian Nash equilibrium can be quite involved (see for example Myerson [33]). However, in the auction games considered here when the type sets are compact subsets of the reals, for example when types are in the interval $[0, 1]$, in most cases, the summations in Definition 3.3.2 can be replaced by integrals.

3.4 Mechanisms

In this section we will briefly and informally introduce the concept of mechanisms. Auctions fit into a more general concept called *mechanisms*. In an auction the price is determined by some form of competition amongst bidders and the winner is selected based completely on the submitted bids. A mechanism also consists of an allocation rule and a payment rule but is a more general concept in the sense that no restriction is put on the allocation rule. For example, a perfectly feasible allocation rule could be to randomly choose a winner (this would not be an auction). In mechanism design theory the purpose is to design the rules of the system (or game) to achieve a specific outcome or behavior. A mechanism typically constitutes a Bayesian game.

We will not go further into mechanism design theory, however there are two useful concepts that frequently occur and are worth stating.

Definition 3.4.1 (Incentive Compatible)**.** A mechanism is incentive compatible if the equilibrium strategy is to reveal ones true type.

Similarly, an auction is incentive compatible if bidders achieve the highest expected payoff in equilibrium by bidding their true valuations. The second-price sealed-bid single-item auction and its generalization to the multiple items case are both dominant strategy incentive compatible (see e.g. the books [22, 29, 10]).

A well known theorem called the Revelation Principle (Myerson [32]) states the following.

Theorem 3.4.1 (Revelation Principle)**.** *Given any mechanism with a Bayesian Nash equilibrium, there exists an incentive compatible mechanism with the same allocations and payments.*

Proof sketch. Given a mechanism $M' = (a', p')$ with allocation rule a' and payment rule p', and where $(s_1, \ldots, s_N)$ constitutes a Bayesian Nash equilibrium. Consider the direct revelation mechanism $M = (a, p)$ such that

$$p(t_1, \ldots, t_n) = p'(s_1(t_1), \ldots, s_n(t_n)) \ .$$

$$x(t_1, \ldots, t_n) = x'(s_1(t_1), \ldots, s_n(t_n)) \ .$$

It is easy to see that truthfully reporting one's type is a must, since reporting a different value $z \neq t$ would be the same as the original mechanism M' reporting $s(z) \neq s(t)$, but since reporting $s(t)$ is the equilibrium then reporting $s(z)$ can not be better. The allocation and payments are clearly the same in both mechanisms. □

There are variants of the revelation principle for other types of equilibria, the proofs are essentially the same. But since we will be concerned mostly with Bayesian games as implied by many auctions and Bayesian Nash equilibrium, it seems appropriate to choose the corresponding version.

From the auction perspective, the revelation principle simply means that given any auction with an equilibrium, there exists an incentive compatible auction with the same outcome. This theorem is mainly used to motivate studying direct truthful mechanisms.

It is common that combinatorial auctions are discussed in the context of mechanisms, especially when the goal is to design a combinatorial auction with specific strategic properties such as incentive compatibility. However the focus in this thesis is mostly on the first-price combinatorial auction, where the mechanism is already completely determined. Since we are more interested in considering the strategic options of the bidders, we will keep the auction perspective instead of the mechanism perspective.

4. The Auction Game

In this chapter we will look at the single-item auction and go through the equilibrium analysis of the second-price and first-price single-item auctions. In Section 4.2 we define the combinatorial auction game and note the important differences to the single-item case that make the standard tools ineffective.

4.1 Single-Item Auctions

From a game theoretic perspective auctions are typically modeled as *infinite* games of incomplete information, that is, Bayesian games with infinite action sets and types. The players of the game are the bidders, the types correspond to a bidder's valuation of the item being auctioned, the actions are the bids and the strategies determine the bid given a valuation.

We will consider only a specific information environment, the standard *independent private values* model originally introduced by Vickrey [54]. In this model, bidders' valuations are independent of other bidders' valuations and a valuation is the private information of each bidder. However, bidders have beliefs about the distributions of other bidders values and these distributions are common knowledge.

The model considered here will be further limited by assuming that bidders are *symmetric* and *risk neutral*. Here, this means that bidders' valuations are identically (but independently) distributed and that bidders want to maximize their expected profit (see e.g. Krishna [22]). We also assume that bidders are able and willing to pay up to their valuations. Finally, we only consider *sealed-bid* auctions, meaning that bids are privately submitted in a one-shot fashion.

Definition 4.1.1 (Single-item auction game)**.** A sealed-bid single-item auction constitutes an infinite game of incomplete information with the following elements.

- N risk-neutral bidders.
- $v_i \in [0,1]$ is bidder i's privately known valuation, a realization of a stochastic variable X_i with independent distribution F and continuous density f.

- $\beta_i : [0,1] \to \mathbb{R}_+$ is the strategy of bidder i.
- $\Pi_i : [0,1] \times \mathbb{R} \to \mathbb{R}$ is the expected payoff of bidder i given his valuation and bid.
- A payment-rule.
- Bidders submit bids ($b_i = \beta_i(v_i)$) in a one-shot fashion.
- The *highest* bidder wins and pays an amount specified by the payment rule.

In a first-price auction the payment rule is straightforward, the winner pays the value of his bid. In a second-price (Vickrey) auction the winner pays the amount of the second highest bid. However, there are also other types of auctions, for example in an all-pay auction the highest bidder wins but every bidder pays the amount of their bids.

Relating to the definition of Bayesian game in Section 3.2 the auction can be written in Bayesian form where the sets of actions are $\mathbb{R}_+$, the sets of types are $[0,1]$, the probability distribution is represented by the cumulative distribution F with density f and the (expected) payoff functions are Π_i.

Given a particular auction, we are interested in finding strategies β_i in equilibrium, that is, strategies that represent some stable state in the sense that bidders have no incentive to deviate by choosing other strategies. When we know these equilibrium strategies we can examine the outcome of the auction in terms of revenue and efficiency. The following two sections illustrate of how equilibrium strategies can be derived.

4.1.1 Example – Second-Price Auction

In the second-price auction the optimal strategy is to bid one's true value. When an auction has this property, it is said to be *incentive compatible*. A winning bidder in the second-price auction pays the value of the second highest bid. The expected payment made by a winning bidder is simply the expected value of the second highest valuation. To see why this is so, consider a bidder with valuation v. His view of the other bidders is simply a set of random variables $X_2, \ldots, X_N$ corresponding to their valuations, unknown to bidder 1 but each independently distributed according to F with density f. Define Y to be the stochastic variable denoting the maximum of $X_2, \ldots, X_N$, i.e. the highest order statistic of the $N-1$ other bidders, and let G and g be the distribution function and density. The expected payment by the winning bidder is

$$G(v) \cdot E[Y | Y < v],$$

that is, the probability of v being the highest valuation times the expected value of the highest of the other bidders' valuations, conditional on v being highest. But since

$$E[Y|Y < v] = \frac{1}{G(v)} \int_0^v y g(y) dy$$

then the expected payment is simply $E[Y]$, that is, the expected value of the highest of $N-1$ valuations. In the event that v is not highest, the bidder pays nothing and the contribution to the expected payment is 0.

We mentioned earlier that the second-price auction is incentive compatible, this is stated more formally in the following theorem.

Theorem 4.1.1. *In a second-price sealed-bid single-item auction with the standard independent private values model, the pure symmetric (Bayesian Nash) equilibrium strategies are defined by*

$$\beta(v) = v \, .$$

Proof. A proof by contradiction follows. Assume we bid something other than v. There are two cases, either bidding less than or greater than v.

Case I: Assume we bid $v' > v$, and let a be the highest of the other bids.

i) If $a < v < v'$ the payoff is the same as bidding v.

ii) If $v < a < v'$ the payoff from bidding v' is less than 0.

Case II: Assume we bid $v' < v$, and let a be the highest of the other bids.

i) If $a < v' < v$ the payoff is the same as bidding v.

ii) If $v' < a < v$ the payoff from bidding v' is 0.

Therefore, bidding anything other than v may in some cases reduce the expected payoff but will never increase it. It is therefore optimal to bid v. □

As we can see, the second-price auction is incentive compatible and has a dominant strategy equilibrium.

4.1.2 Example – First-Price Auction

In the first-price auction bidding your true valuation is bidding too much. Informally this is easy to see, a bidder who submits a bid equal to his valuation will receive a payoff of 0, this means he could submit a slightly smaller bid $v - \Delta$ that would have almost the same chance of winning but which would give an expected payoff greater than 0.

As an example, we will derive the pure symmetric Bayesian Nash equilibrium strategy in the first-price sealed-bid auction with N bidders[1]. As before let b_{-i} be all elements except b_i. Consider bidder i with valuation v_i and bid b_i. Given the bids of the other bidders b_{-i}, he receives

$$\text{payoff} = \begin{cases} v_i - b_i & \text{if } b_i > \max(b_{-i}) \\ 0 & \text{if } b_i < \max(b_{-i}) \end{cases}$$

If two bidders bid the same amount the item is awarded to either of the bidders with equal probability. It is common to assume that the equilibrium strategy β is strictly increasing and differentiable, this allows us to take the inverse and state the first-order condition. Making this assumption also removes the separate treatment that would occur if bidders bid identically. Since there is zero probability that bidders have identically realized valuations their bids will also be different. In an auction with discrete bids this is not the case.

Some basic properties can be stated before we begin. A bidder with valuation $v_i = 0$ will never submit a bid $b_i > 0$ and negative bids are not allowed. Similarly a bidder will never submit a bid $b_i > v_i$ since whenever he wins he will have negative payoff and is thus better of bidding v_i or not bidding at all. We have now established that the bidding strategy must be to bid somewhere in the interval $0 \leq \beta(v_i) \leq v_i$.

Consider bidder 1, who wins the auction whenever he submits the highest bid. His view of the other bidders is simply a set of random variables $X_2, ..., X_N$ corresponding to their valuations, unknown to bidder 1 but each independently distributed according to F with density f. As in the previous section, we define Y to be the stochastic variable denoting the maximum of $X_2, ..., X_N$, i.e. the highest order statistic of the $N-1$ other bidders, and let G and g be the distribution function and density.

Assume that bidders 2 through N are following the symmetric, strictly increasing and differentiable equilibrium strategy β. Suppressing the bidder subscript 1 and writing b and v instead of b_1 and v_1, the condition of bidder 1 winning can thus be stated as $b > \max_{j \neq 1} \beta(X_j)$, but since β is strictly increasing this is the same as $b > \beta(\max_{j \neq 1} X_j) = \beta(Y)$, or $\beta^{-1}(b) > Y$. The expected payoff Π of bidder 1 with valuation v bidding b is the probability that his bid is the highest multiplied by the profit, that is

$$\Pi(v, b) \quad = \quad G(\beta^{-1}(b)) \cdot (v - b) \,. \tag{4.1}$$

[1]This is similar to the derivation in Krishna [22]

Now, bidder 1 is trying to maximize his expected payoff with respect to bid b given his valuation v. Since

$$\frac{d}{dx}\beta^{-1}(x) = \frac{1}{\beta'(\beta^{-1}(x))}$$

the derivative of the expected payoff is (and setting it equal to zero)

$$\frac{\partial}{\partial b}\Pi(v,b) \;=\; \frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(v-b) - G(\beta^{-1}(b)) \;=\; 0\,. \tag{4.2}$$

Given that we are in a symmetric equilibrium then it must be that $b = \beta(v)$ and we can rewrite Equation 4.2 as

$$\begin{aligned}
\frac{g(v)}{\beta'(v)}(v-\beta(v)) - G(v) &= 0 \\
&\Leftrightarrow \\
vg(v) &= G(v)\beta'(v) + g(v)\beta(v)) \\
&\Leftrightarrow \\
vg(v) &= \frac{d}{dv}\big(G(v)\beta(v)\big)\,.
\end{aligned} \tag{4.3}$$

Since $\beta(0) = 0$, Equation 4.3 can be re-written as

$$\beta(v) \;=\; \frac{1}{G(v)}\int_0^v yg(y)dy \;=\; E[Y|Y<v]\,. \tag{4.4}$$

That is, if β is a pure symmetric equilibrium strategy and given that all other bidders follow β then β is defined according to Equation 4.3. Although Equation 4.3 is a payoff maximizing strategy, it remains to be proven that β actually constitutes a Bayesian-Nash equilibrium, that is, it remains to be proven given that everyone else is using this strategy that there exists no profitable deviation for any bidder. The following theorem more formally states the equilibrium strategy.

Theorem 4.1.2. *In a first-price sealed-bid single-item auction with the standard independent private values model, the pure symmetric Bayesian Nash equilibrium strategy is*

$$\beta(v) \;=\; \frac{1}{G(v)}\int_0^v yg(y)dy \;=\; E[Y|Y<v]$$

where Y is the stochastic variable denoting the maximum of $N-1$ independently and identically drawn valuations, with distribution G and density g.

Proof. Assume that all bidders except one are following the continuous and strictly increasing strategy β as stated in Equation 4.4, this implies that among these bidders the bidder with the highest valuation will also make the highest bid, we want to show that a bidder cannot profit from deviating.

We have assumed that one bidder deviates from using the strategy β. This bidder has valuation v but he bids some value $b \neq \beta(v)$. Since β is strictly increasing and continuous there exists a value w such that $b = \beta(w)$. We proceed by showing that given a valuation v bidding anything other than $\beta(v)$ will not result in a higher expected payoff.

The expected payoff from bidding $b = \beta(w)$ is

$$\begin{aligned}
\Pi(v, \beta(w)) &= G(w)(v - \beta(w)) \\
&= vG(w) - G(w)\frac{1}{G(w)}\int_0^w xg(x)dx \\
&= vG(w) - \left(wG(w)\Big|_0^w - \int_0^w G(x)dx \right) \\
&= G(w)(v - w) + \int_0^w G(x)dx
\end{aligned}$$

The expected payoff from bidding $\beta(v)$ is derived similarly and is

$$\Pi(v, \beta(v)) = \int_0^v G(x)dx \,.$$

The difference in expected payoff is therefore

$$\begin{aligned}
\Pi(v, \beta(v)) - \Pi(v, \beta(w)) &= \int_0^v G(x)dx - \left(G(w)(v - w) + \int_0^w G(x)dx \right) \\
&= G(w)(w - v) + \int_w^v G(x)dx \; \geq \; 0 \,.
\end{aligned}$$

This is true both for $v > w$ and $w > v$. To see why, consider the cases.
Case I $(v > w)$:

$$G(w)(w - v) + \int_w^v G(x)dx \; \geq \; 0 \;\; \Leftrightarrow \;\; \int_w^v G(x)dx \geq G(w)(v - w) \,,$$

which is true since the r.h.s. is the under approximation of the integral.
Case II $(w > v)$:

$$G(w)(w - v) + \int_w^v G(x)dx \; \geq \; 0 \;\; \Leftrightarrow \;\; G(w)(w - v) \; \geq \; \int_v^w G(x)dx \,,$$

which is true since the l.h.s. is the over approximation of the integral. □

Theorem 4.1.2 applies to the general distribution, however, it might be illustrative go give a more concrete example. Consider the case when valuations are distributed uniformly over [0,1]. The symmetric equilibrium strategy in the first-price auction is then to bid a constant fraction of the private valuation. Specifically

$$\beta(v) = \frac{N-1}{N} \cdot v$$

where N is the number of bidders.

Auctions such as the first-price single-item auction and similar can, with variations to the model, be analyzed with the techniques used here.

The expected payment and subsequently the expected revenues of the first-price and second-price auctions are the same. More generally the *revenue equivalence theorem* (originally due to Myerson [32], and Riley and Samuelson [46]) states: given the standard private values model with independently and identically distributed valuations, and where bidders are risk neutral, any symmetric and increasing equilibrium of any single-item auction where the highest bidder wins, gives the same expected revenue to the seller, as long as the expected payment by a bidder with valuation 0, is 0. The proof of this can be found in any standard book on auction theory, see for example Krishna [22] or Milgrom [29].

Myerson [32] formulated and solved the optimal (single-item) auction design problem. He considered not a specific auction format but a general mechanism description, answering the question given all possible auction formats, which mechanism the auctioneer should choose to maximize his revenue. Concurrently Riley and Samuelson [46] looked at essentially the same problem.

4.2 Combinatorial Auction Game

Analysis of the first-price combinatorial auction is different from the single-item case because of the winner determination problem. In a single-item auction a simple parsing of all the bids will reveal the highest bidder. In a combinatorial auction the problem of determining winners is far from trivial, a bidder may be the highest bidder and still not be part of the globally optimal allocation. The hardness of the allocation problem has serious implications on what can be done in terms of game theoretic analysis in the first-price setting.

In this section we will extend Definition 4.1.1 of the single-item case to include valuations as functions as defined in Definition 2.2.1. In the most general case, a bidder has one valuation for each possible combination of items.

Definition 4.2.1 (Combinatorial auction game)**.** The sealed-bid first-price combinatorial auction constitutes an infinite game of incomplete information with the following elements.

- N risk-neutral bidders.
- For each bidder i, and for all non-empty subsets $S \subseteq M$.
  - $v_i(S) \in \mathbb{R}_+$, is the privately known valuation function that defines the private realization of the valuation of the items S.
  - $\beta_i(S, v_i) \in \mathbb{R}_+$ is the strategy that determines the bid value for a set of items.
  - $B_i = \{\langle S, b\rangle : \forall S \subseteq M, S \neq \emptyset, \text{such that } b = \beta_i(S, v_i)\}$ is the set of all bids by bidder i.
  - $\Pi_i(B_i, v_i) \in \mathbb{R}$ is the expected payoff given all bids and the valuation function.

The valuation function is intentionally left unspecified since valuations may be modeled in various ways. For example, the valuation could be defined as an independently drawn value for each subset of items, this is probably the most general model. Another model, as mentioned in Section 2.2, imposes a structure on the valuation; a subset S has value $|S| \cdot (x + \alpha)$ where x is independently drawn and α is a public synergy[2]. In the single-item case the model is straightforward, valuations are independently distributed according to some distribution F. In the combinatorial auction it is only natural that valuations can be modeled in many ways. Most of our work assumes some form of structure in the valuation function.

4.2.1 Analyzing the First-Price Combinatorial Auction

Before we can even attempt a similar analysis as we did in the single-item case we would need to know the probability of winning given our bids B_i. In the single-item auction, knowing that the strategy is strictly increasing and differentiable together with the distribution is enough to know the probability of winning given a valuation. This is enough since winning means being highest. However, the corresponding probability function is not trivially determined in the combinatorial auction. Knowing the distribution of valuations, assuming that strategies are increasing, is not enough to construct a useful expression for the probability of winning. The expression describing this probability would have to cover all possible combinations of outcomes, which means that even if we managed to write it down it would still be impractical to evaluate much less use it in our analysis.

[2] Krishna [23] also uses a public synergy α.

Unfortunately we are stuck before we can begin, we cannot even determine one of the most fundamental parts needed for the analysis.

When bidders are single-minded the game resembles that of the single-item case only in the sense that the valuation is a single random variable and a bidder's set of bids contains one element, however we are still faced with an NP-hard winner determination problem. In this model the strategy determines the bid value as before, and the expected payoff, $\Pi(b,v)$, is defined for the one bid b and valuation v. Unfortunately, as in the unrestricted case, we still don't have any useful expression for the probability of winning, except for very small problems with very few bidders and bids. On the other hand the model is sufficiently simple to allow some form of estimation of the probability function, this line of reasoning is pursued in Paper-II.

In summary, game theoretic analysis in the first-price sealed-bid combinatorial auction is a hard problem, and standard techniques for analysis simply are not practical. In Paper-III we approach this problem from a different angle, by deriving bounds on the achievable revenue in pure symmetric (Bayesian Nash) equilibrium, instead of focusing on the actual equilibrium strategies. Miltersen and Santillan [30] prove the existence of pure strategy Bayesian Nash equilibrium in the first-price sealed-bid combinatorial auction in the same model as we study in Paper-III, that is, with synergy-bidders bidding one bid on a combination of items and with single-bidders each bidding on a single item. In a *complete information* setting where bidders are completely informed, Bernheim and Whinston [6] provide an equilibrium analysis of a first-price menu-auction.

4.3 Finite Game – Discrete Strategies

When analyzing auctions the focus is mostly on the format and the informational structure with no restrictions on bids. In reality on the other hand, the monetary system implies a discrete action space. Quite often in iterative auctions there is also a minimum bid increment rule.

In the single-item setting, Chwe [8], Yu [60], David et al. [11], and Mathews and Sengupta [28], have studied several discrete auctions, we only mention this briefly as this thread of research is outside the scope of this thesis.

Studying the finite combinatorial auction game with discrete strategies, we are faced with a search problem where the evaluation of every strategy profile is a possibly intractable problem, since the winner determination problem must be solved to evaluate the payoff. Also, unless we severely limit the strategy space, representing the game in normal form becomes a problem since the size is exponential in the number of bidders. If the strategy-space is of size s

and there are N bidders then the normal form matrix is of size s^N. Some work in this area try to alleviate the representation issue by using heuristics such as best-response dynamics, local search and genetic algorithms for searching and evaluating only part of the payoff-matrix, see for example Sureka [51], and Sureka and Wurman [53, 52]. In the case of a discrete strategy space, even if equilibrium strategies are found, the resulting strategies may be far from the actually sought strategies in the infinite game case, depending on the coarseness of the discretization.

5. Summaries of Included Papers

In this chapter we present a short summary of each paper and some reflections about the work. Each paper in essence deals with the hard problem of bidding although the perspective is different in each work.

1. In Paper-I we look at the *bid limitation problem*. When there is a constraint on the number of bids a bidder may place, choosing which bids to submit is hard and important. We argue that expressiveness in bidding may in some cases be more important than finding optimal allocations of items.

2. In Paper-II we consider the question of what to bid given a valuation. We propose a heuristic for the hard problem of finding equilibrium strategies in the first-price combinatorial auction.

3. Paper-III concerns the aggregate perspective of bidding. We analyze the question of revenue in the first-price combinatorial auction compared to simple simultaneous single-item auctions and provide bounds on the achievable revenue.

4. Paper-IV empirically extends the results from Paper-III in the case when the number of bidders is small, and also tightening the upper bound on the expected revenue in one-shot simultaneous single-item auctions.

5.1 Paper-I: An Auction Mechanism for Polynomial-time Execution with Combinatorial Constraints

Since it is typically not feasible to submit all possible combinations of bids, the auctioneer will most likely need to limit the number of bids that a bidder may submit. This phenomenon we refer to as the *bid limitation problem*. This is an important problem since if bidders submits the "wrong" bids then the resulting allocation will be inefficient. The following example illustrates this problem clearly. Given two bidders and two items c_1 and c_2. The bidders valuations are as follows.

	c_1	c_2	c_1 & c_2
Bidder 1	100	40	140
Bidder 2	80	80	160

Now, assume that bidders are limited to submitting at most one bid each. Suppose that bidder 1 and bidder 2 both submit one bid each on the combination $\langle c_1, c_2 \rangle$. Given these two bids, the solution to the winner determination problem is to assign both items to one of the bidders, the highest bidder. However, the optimally efficient solution, when there are no restrictions on the number of bids that may be submitted, is to award item c_1 to bidder 1 and item c_2 to bidder 2. This is a fairly trivial example where submitting all possible bids could very well be allowed, for larger auctions this would not be possible and with many items, the probability that a bidder chooses the optimal subset of bids is very low.

With this work three main problems with combinatorial auctions are addressed, (i) bid limitation, (ii) protocol transparency and (iii) algorithmic efficiency. We present a polynomial time algorithm that accepts bids termed *min-bundles*. A min-bundle is a bundle of single-bids associated with a minimum constraint on the least number of bids that must win for the bundle to be valid. When we have a set of several min-bundles, they constitutes an OR of min-bundles. This language is especially useful when it is more important for a bidder to get a specific number of items than a specific set of items.

We show trough simulation that the impact of the bid limitation problem is severe, in fact, given a standard combinatorial auction where the number of bids are limited we can statistically achieve more efficient allocations even when using a polynomial time allocation algorithm in conjunction with a more expressive bidding language. More specifically we compare the economic efficiency of a standard combinatorial auction accepting the OR bidding language as described in Chapter 2, but with a limitation on the number of bids, to the economic efficiency of an auction using a polynomial time algorithm for determining winners together with a more compact bidding language. We show empirically that the there are instances when focusing on a more compact bidding language is more important than actually solving the winner determination problem optimally.

Another main issue is that of transparency. Solving the winner determination problem optimally is not trivial. In fact, the algorithms used to solve this problem are somewhat complicated and to any non-specialist the winner determination algorithms are truly a black box. It may be important for bidders to understand how a solution is arrived at in a transparent way, we address this issue by describing a fairly simple protocol for assigning winners.

The final problem is that of the general intractability of solving the winner determination problem optimally. The protocol we provide can be executed in polynomial time which may be significant from a predictability viewpoint.

The rest of the paper deals with the experimental evaluation of our winner determination protocol using software agents that act as bidders with natural behavior in the context of an iterative auction. The economic efficiency is measured as an average over several trials.

A comment: to highlight the severeness of the bid limitation problem, we compare our polynomial protocol to a standard one-shot first-price combinatorial auction. In fact to make the situation more favorable for the standard combinatorial auction we assume that bidders in that auction bid their true valuations, and we also provide a simple and natural iterative version of the standard combinatorial auction to sell unallocated items. This proposed iterative combinatorial auction should be viewed simply as a fair benchmark and not as a serious candidate for an iterative combinatorial auction.

5.2 Paper-II: Discovering Equilibrium Strategies for a Combinatorial First Price Auction

This work deals with the problem of finding pure (Bayesian Nash) equilibrium strategies in the first-price sealed-bid combinatorial auction, an open and hard problem. We consider the infinite game of incomplete information.

The strategy problem is studied in a symmetric model, where valuations are independently drawn from identical distributions, and specifically for two settings: (i) bidders submit bids on all combinations, and (ii) bidders have a synergy on one random combination and bid individual single-bids on the remaining items.

Since analytically determining equilibrium strategies is hard, we approach the problem from a less strict viewpoint by designing a heuristic which uses model fitting and best response dynamics. Vorobeychik, Wellman and Singh [57] employ learning of payoff functions in infinite games, our approach is different in the sense that instead of focusing on the payoff function, we instead sample the probability that a bid of a certain value will be part of the winning allocation given the bidders current strategy sets. The advantage of our approach is that we can calculate the best response strategy analytically which significantly reduces execution time.

Consider a particular bidder. In setting (i) he submits one bid for every combination of items and every bid he submits will affect the probability of his other bids winning. The interactions between bids are fundamentally hard to predict since he essentially needs to predict the solution of the NP-hard optimization problem. In setting (ii), he bids one combination bid and single-bids on the remaining items, in this setting the interactions are less significant since the values of the combinatorial bids are high compared to single-bids.

In order to create a tractable search problem the simplifying assumption is made that each bid submitted by one and the same bidder is in fact independent of the other bids, thus disregarding the interactions. This allows us to separately find a strategy for each combination size. In essence this view is equivalent to considering one bidder to be a set of single-minded bidders.

At the core of our heuristic is the probability function $P(\cdot)$, the probability of a bid of a certain size and value winning. Given a combination size, (i) we sample $P(\cdot)$ for a finite set of evenly distributed bid values (given the current strategy set of all bidders), and (ii) use the sample data to regress a model function $\overline{P}(\cdot)$.

At our disposal are several model functions that can be adapted to the probability sample data. Each model has been chosen beforehand such that we can express the best-response strategy β in closed form by maximizing the expected payoff $\overline{P}(\beta(v))(v-\beta(v))$. Thus, given the parameter values provided by regression of the probability models, we have predefined closed form expressions for the best response strategies.

With this method of finding the best response strategy in place, we use standard best reply dynamics to iterate the procedure until the system reaches a (within tolerance) fixed point. The idea is quite simple, and when the search converges we are in pure symmetric equilibrium. Best reply dynamics is unfortunately not guaranteed to converge, however during all the experiments non-convergence was never a significant problem.

5.3 Paper-III: A New Analysis of Revenue in the Combinatorial and Simultaneous Auction

In this work we examine the fundamental relationship between the one-shot simultaneous single-item auctions and the one-shot first-price combinatorial auction with regards to revenue. It is commonly believed that using a combinatorial auction will generate higher revenue than selling items individually since bidders are able to more precisely and with no risk state their preferences. Although this is an important topic, until now no proofs has been presented that actually confirm the revenue superiority of combinatorial auctions. In fact, the only known theoretical analysis (by Krishna and Rosenthal [23], and Albano, Germano and Lovo [14]) has been done for the case of two items for sale and indicates the opposite; the combinatorial auction gives a lower revenue.

The case with only two items is however very far from most real cases. In real-world combinatorial auctions we may have many items and bidders with synergies on independently and seemingly randomly selected combinations of items. Bidders are also faced with an incomplete information setting and the winner determination problem is complicated, meaning more than one combination can win and that some items may be allocated to single bids.

Striving to analyze scenarios that depict reality as closely as possible, we study a model where there are two types of bidders, (i) single-bidders interested in one single item, and (ii) synergy-bidders that are interested in one specific though randomly selected combination of items. Miltersen and Santillan [30], has proven the existence of a pure Bayesian Nash equilibrium in the same model. In this work, we construct bounds on the achievable revenue in both auctions given a pure symmetric equilibrium. We show that asymptotically as the number of bidders approaches infinity, the combinatorial auction is revenue superior to the simultaneous auction. We also provide parameterized bounds on revenue and provide, for many cases, the required parameters such that the combinatorial auction achieves a higher revenue than the simultaneous auction.

Although the model we study is specific, it is still vastly more general than any other model for which similar results have been attempted. We have chosen a model with only uniform distributions since this is a common distribution in examples in the literature. This could be considered a limitation, however, it should not be viewed as particularly serious since the question of distribution is not our main point[1]. The most important property of our results is that they are achieved for a model with many bidders and many items, with random combinations of more than two items. The analysis of complicated situations such as these require approaches that differ from the common methods of analysis in auction theory. In this work we present the first step towards analyzing these complex settings.

5.4 Paper-IV: Combinatorial and Simultaneous Auction: A Pragmatic Approach to Tighter Bounds on Expected Revenue

This work examines the revenue bounds on combinatorial and simultaneous auctions from pragmatic viewpoint. In our previous work (Paper-III), due to

[1] A generalization to other continuous distributions is likely to be possible.

proof-technical reasons, the number of bidders are relatively high. In this work we look at the expected revenue in the first price combinatorial auction in a setting where the number of bidders is small and show empirically, based on experiments, that the expected revenue of the combinatorial auction exceeds that of the simultaneous auction also with a small number of bidders. Furthermore, given few bidders a tightening of the upper bound on the simultaneous auction is also possible.

We make the observation that Lemma 5.5 from Paper-III allows us to simplify the experimental model by replacing the single-bidders with a fixed low bid, thus allowing us to use the search heuristic from Paper-II. Since we are undercutting the single-bid, this is a reasonable compromise. The search heuristic is used for finding pure symmetric equilibrium strategies for the synergy-bidders in the first-price sealed-bid combinatorial auction. Given these strategies, we are then able to sample the expected revenue of the combinatorial auction. Our results are empirical, however they offer strong indications that the expected revenue superiority of the combinatorial auction, to the simultaneous auction, holds also in the case of a small number of bidders. This suggests that the results in Paper-III hold more generally than theoretically proven.

Bibliography

[1] A. Andersson, M. Tenhunen, and F. Ygge. Integer programming for combinatorial auction winner determination. *Fourth International Conference on Multiagent Systems, Boston, 2000.*, 2000.

[2] Arne Andersson, Jim Holmstrom, and Mattias Willman. An auction mechanism for polynomial-time execution with combinatorial constraints. *E-Commerce Technology, 2005. CEC 2005.*, pages 17–24, 2005.

[3] Lawrence Ausubel. An efficient ascending-bid auction for multiple objects. *The American Economic Review, AER*, 94:17–24, 2004.

[4] Lawrence M. Ausubel, Peter Cramton, and Paul Milgrom. The clock-proxy auction: A practical combinatorial auction design. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 5, pages 115–138. MIT Press, 2006.

[5] Jeffrey S. Banks, John O. Ledyard, and David P. Porter. Allocating uncertain and unresponsive resources: An experimental approach. *The RAND Journal of Economics*, 20(1):1–25, 1989.

[6] B. Douglas Bernheim and Michael D. Whinston. Menu auctions, resource allocation, and economic influence. *The Quarterly Journal of Economics*, 101(1):1–32, 1986.

[7] Estelle Cantillon and Martin Pesendorfer. Combination bidding in multi-unit auctions, 2007.

[8] Michael Suk-Young Chwe. The discrete bid first auction. *Economics Letters*, 31(4):303 – 306, 1989.

[9] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[10] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. The MIT Press, 2006.

[11] E. David, A. Rogers, J. Schiff, S. Kraus, and N. R. Jennings. Optimal design of english auctions with discrete bid levels. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 98–107, New York, NY, USA, 2005. ACM.

[12] Rafael Epstein, Lysette Henríquez, Jaime Catalán, Gabriel Y. Weintraub, and Cristián Martínez. A combinational auction improves school meals in chile. *Interfaces*, 32(6):1–14, 2002.

[13] Michael R. Gary and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[14] Fabrizio Germano Gian Luigi Albano and Stefano Lovo. A comparison of standard multi-unit auctions with synergies. *Economics Letters*, 71(1):55–60, 2001.

[15] Edward Gibbon. *The History of the Decline and Fall of the Roman Empire*, volume 1. Strahan & Cadell, 1776-1789.

[16] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[17] John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part i. the basic model. *Management Science*, 14(3):159–182, 1967.

[18] John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968.

[19] John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part iii. the basic probability distribution of the game. *Management Science*, 14(7):486–502, 1968.

[20] Johan Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182(1):105–142, 1997.

[21] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

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

[23] Vijay Krishna and Robert W. Rosenthal. Simultaneous auctions with synergies. *Games and Economic Behavior*, 17(1):1–31, November 1996.

[24] John O. Ledyard. Optimal combinatoric auctions with single-minded bidders. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 237–242, New York, NY, USA, 2007. ACM.

[25] Daniel Lehmann, Liadan Ita Oćallaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, 2002.

[26] Adam N. Letchford. Allocation of school bus contracts by integer programming. *Operational Research Society*, 47(3):369–372, 1996.

[27] Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, pages 237–242, July 2004.

[28] Timothy Mathews and Abhijit Sengupta. Sealed bid second price auctions with discrete bidding. *Applied Economics Research Bulletin*, 1(1):31–52, 2008.

[29] Paul Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[30] Peter Bro Miltersen and Rocio Santillan Rodriguez. Existence of equilibria in first-price combinatorial auctions with synergy. Work in progress: https://www.daimi.au.dk/ rocio/reports/synergy.pdf, 2009.

[31] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612 – 631, 2008. Special Issue in Honor of Michael B. Maschler.

[32] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1), 1981.

[33] Roger B. Myerson. *Game theory : analysis of conflict*. Harward University Press, 1991.

[34] John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[35] John F. Nash. Equilibrium points in n-person games. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 36, pages 48–49, January 1950.

[36] John F. Nash. *Non-cooperative games*. PhD thesis, Mathematics Department, Princeton University, May 1950.

[37] John F. Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, September 1951.

[38] John F. Nash. Two-person cooperative games. *Econometrica*, 21(1):128–140, 1953.

[39] Noam Nisan. Bidding and allocation in combinatorial auctions. In *EC '00: Proceedings of the 2nd ACM conference on Electronic commerce*, pages 1–12. ACM Press, 2000.

[40] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

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

[42] David C. Parkes. ibundle: an efficient ascending price bundle auction. In *EC '99: Proceedings of the 1st ACM conference on Electronic commerce*, pages 148–157, New York, NY, USA, 1999. ACM Press.

[43] David C. Parkes and Lyle H. Ungar. Iterative combinatorial auctions: Theory and practice. In *AAAI/IAAI*, pages 74–81, 2000.

[44] Aleksandar Pekeč and Michael H. Rothkopf. Combinatorial auction design. *Manage. Sci.*, 49(11):1485–1503, 2003.

[45] S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics*, 13(2):402–417, 1982.

[46] J G Riley and W F Samuelson. Optimal auctions. *American Economic Review*, pages 381–392, 1981.

[47] Marc S. Robinson. Collusion and the choice of auction. *The RAND Journal of Economics*, 16(1):141–145, 1985.

[48] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[49] Michael H. Rothkopf. Thirteen reasons why the vickrey-clarke-groves process is not practical. *Operations Research*, 55(2):191–197, 2007.

[50] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artif. Intell.*, 135(1-2):1–54, 2002.

[51] Ashish Sureka. *Techniques for finding Nash equilibria in combinatorial auctions*. PhD thesis, Department of Computer Science, North Carolina State University, 2004. (Available from www.lib.ncsu.edu/theses/available/etd-07072005-100355/unrestricted/etd.pdf).

[52] Ashish Sureka and Peter R. Wurman. Applying metaheuristic techniques to search the space of bidding strategies in combinatorial auctions. In *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 2097–2103. ACM Press, 2005.

[53] Ashish Sureka and Peter R. Wurman. Using tabu best-response search to find pure strategy nash equilibria in normal form games. In *AAMAS'05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 1023–1029. ACM Press, 2005.

[54] Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 41:8–37, 1961.

[55] John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.

[56] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2002.

[57] Yevgeniy Vorobeychik, Michael P. Wellman, and Satinder Singh. Learning payoff functions in infinite games. *Machine Learning*, 67(1-2):145–168, 2007.

[58] Peter Wurman. *Market Structure and Multidimensional Auction Design for Computational Economies*. PhD thesis, Department of Computer Science, University of Michigan, 1999.

[59] Peter R. Wurman and Michael P. Wellman. Akba: a progressive, anonymous-price combinatorial auction. In *EC '00: Proceedings of the 2nd ACM conference on Electronic commerce*, pages 21–29, New York, NY, USA, 2000. ACM.

[60] Jin Yu. *Discrete approximation of continuous allocation mechanisms*. PhD thesis, Division of the Humanities and Social Sciences, California Institute of Technology, 1998.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 656

Editor: The Dean of the Faculty of Science and Technology

A doctoral dissertation from the Faculty of Science and Technology, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology".)

Distribution: publications.uu.se
urn:nbn:se:uu:diva-102960

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2009