Helsinki University of Technology
Department of Industrial Engineering and Management
Doctoral Dissertation Series 2009/12
Espoo 2009

BIDDER SUPPORT IN ITERATIVE, MULTIPLE-UNIT COMBINATORIAL AUCTIONS

Riikka-Leena Leskelä

Dissertation for the degree of Doctor of Science in Technology to be presented with due permission of the Department of Industrial Engineering and Management for public examination and debate in auditorium TU2 at Helsinki University of Technology, Espoo, Finland, November 14th, 2009 at 12 noon.

Helsinki University of Technology
Department of Industrial Engineering and Management

Helsinki University of Technology
Department of Industrial Engineering and Management
P.O. Box 5500
FI-02015 TKK, Finland
Tel. +358 9 451 2846
Fax +358 9 451 3665
Email: riikka-leena.leskela@tkk.fi
Internet: http://www.tuta.tkk.fi/en

ISBN 978-952-248-127-6 (print)
ISBN 978-952-248-128-3 (electronic)

ISSN 1797-2507 (print)
ISSN 1797-2515 (electronic)

http://lib.tkk.fi/Diss/2009/isbn9789522481283/

Yliopistopaino
Helsinki

ABSTRACT OF DOCTORAL DISSERTATION	HELSINKI UNIVERSITY OF TECHNOLOGY P.O. BOX 1000, FI-02015 TKK http://www.tkk.fi
Author Riikka-Leena Leskelä	
Name of the dissertation Bidder Support in Iterative, Multiple-Unit Combinatorial Auctions	
Manuscript submitted 15.5.2009	Manuscript revised 23.9.2009
Date of the defence 14.11.2009	

Faculty	Information and Natural Sciences
Department	Industrial Engineering and Management
Field of research	Decision and Negotiation Support
Opponent(s)	Prof. Wojtek Michalowski, University of Ottawa
Supervisor	Prof. Hannele Wallenius, Helsinki University of Technology
Instructors	Prof. Hannele Wallenius, Helsinki University of Technology
	Prof. Jyrki Wallenius, Helsinki School of Economics

Abstract

This thesis is about supporting the bidders' decision making in iterative combinatorial auctions. A combinatorial auction refers to an auction with multiple (heterogeneous) items, in which bidders can submit bids on packages. Combinatorial auctions are challenging decision making environments for bidders, which hinders the adoption of combinatorial mechanisms into practice. Bidding is especially challenging in sealed-bid auctions. Bidders do not know the contents of other bidders' bids and hence cannot place bids that would team up with existing bids to become winners. The objective of this study is to develop and test support tools for bidders in semi-sealed-bid, iterative combinatorial auctions. The tools are designed for reverse auctions, but can easily be applied to a forward setting.

The Quantity Support Mechanism (QSM) is a support tool, which provides the bidders with a list of bid suggestions. The bid suggestions are such that if submitted, they would become provisional winners. The QSM benefits both bidders and the buyer, because it chooses suggestions that are most profitable for the bidders while decreasing the total cost to the buyer. The QSM is based on a mixed integer programming problem.

The QSM was tested in two simulation studies. The results of the studies indicated that the QSM works well – it is much better to use the QSM than no support – but that it does not necessarily guide the auctions to the efficient allocation. The QSM was also integrated into an online auctions system, and tested with human subjects. The results of the laboratory experiment showed that the performance of the QSM is dependent on the bidders' behavior and the kind of bids they place in the auction. The user interface of the auction was good. I also observed bidders' strategies, and could identify different bidder types corresponding to those reported in earlier studies. The experiment also showed the importance of experience in complex bidding environments.

The simulation studies and the laboratory experiment showed that the QSM is too dependent on the existing bids in the bid stream, which causes the auctions to end in inefficient allocations. In order to overcome this problem we designed another support tool, the Group Support Mechanism (GSM). The main logic in the GSM is similar to the QSM. The main difference is that instead of solving for one bid that complements existing bids to become a winner, the GSM can suggest several bids for different bidders. Together this set of bids would then become provisionally winning. The preliminary tests show significant improvement in the efficiency of the auction outcomes when the GSM was used instead of the QSM.

Future research includes the further development of the GSM and its testing with simulations and human subjects. Also, bidder behavior, bidder strategies and the effect of learning and experience in combinatorial auctions should be further studied. This is important because bidders' behavior in the auctions affects the auction design and the requirements for the user interface.

Keywords Auctions, Combinatorial Auctions, Bidder Support, Online Auctions, E-auctions

VÄITÖSKIRJAN TIIVISTELMÄ	TEKNILLINEN KORKEAKOULU PL 1000, 02015 TKK http://www.tkk.fi
Tekijä Riikka-Leena Leskelä	
Väitöskirjan nimi Bidder Support in Iterative, Multiple-Unit Combinatorial Auctions	
Käsikirjoituksen päivämäärä 15.5.2009	Korjatun käsikirjoituksen päivämäärä 23.9.2009
Väitöstilaisuuden ajankohta 14.11.2009	

Tiedekunta	Informaatio- ja luonnontieteiden tiedekunta
Laitos	Tuotantotalouden laitos
Tutkimusala	Päätöksenteon ja neuvotteluiden tukeminen
Vastaväittäjä(t)	Prof. Wojtek Michalowski, University of Ottawa
Työn valvoja	Prof. Hannele Wallenius, TKK
Työn ohjaaja	Prof. Hannele Wallenius, TKK, Prof. Jyrki Wallenius, HSE

Tiivistelmä

Tämä tutkimus koskee tarjoajien päätöksenteon tukemista kombinatorisissa huutokaupoissa. Tarjousten tekeminen kombinatorisissa huutokaupoissa on haastavaa – etenkin suljetuissa huutokaupoissa. Voittavat tarjoukset täydentävät toisiaan; niiden summa on kaupan kohteena oleva hyödykekombinaatio. Suljetussa huutokaupassa tarjouksen tekijät eivät kuitenkaan tiedä toistensa tarjousten sisältöä, joten he eivät osaa tehdä tarjouksia, jotka täydentäisivät muita tarjouksia. Tutkimuksemme tarkoituksena on kehittää ja testata työkaluja tarjouksen tekijöille puolisuljettuihin, iteratiivisiin kombinatorisiin huutokauppoihin.

Kehittämämme työkalu, Quantity Support Mechanism (QSM), ehdottaa tarjoajille tarjouksia joista mikä tahansa olisi kyseisellä hetkellä voittajien joukossa. Tarjoajan tehtäväksi jää päättää, haluaako hän tehdä jonkin ehdotetuista tarjouksista. QSM hyödyttää molempia osapuolia, sillä sen tekemät ehdotukset ovat tarjoajille mahdollisimman voitollisia ja samalla vähentävät ostajan kokonaiskustannuksia (kun kyseessä on käänteinen huutokauppa). QSM pohjalla on kokonaislukuoptimointitehtävä.

QSM:ia testattiin simuloimalla. Simulointien tulokset osoittivat, että QSM toimii hyvin – on parempi käyttää QSM:ia kuin olla ilman tukea – mutta sen käyttö ei aina takaa, että huutokauppa päättyy tehokkaaseen allokaatioon. QSM myös integroitiin osaksi Internet-pohjaista huutokauppajärjestelmää. Tämä mahdollisti QSM:n testaamisen koehenkilöillä. Kokeen tulokset osoittivat, että QSM:n toimiminen riippuu siitä, millaisia tarjouksia tarjoajat ovat huutokaupassa tehneet. Huutokaupan käyttöliittymä todettiin toimivaksi. Tutkin myös koehenkilöiden käyttämiä strategioita ja tunnistin niiden joukosta samantyyppisiä strategioita kuin aikaisemmissa tutkimuksissa tarjoajien on havaittu käyttävän. Koe osoitti myös kokemuksen tärkeyden monimutkaisissa huutokaupoissa.

Simulaatiot ja koehenkilöillä tehty testaus osoittivat että QSM on liian riippuvainen olemassa olevista tarjouksista. Tästä seuraa mm. että huutokaupat eivät pääty tehokkaaseen allokaatioon. Ratkaistaksemme tämän ongelman kehitimme toisen tukityökalun, the Group Support Mechanismin (GSM). GSM toimii pääpiirteissään samalla tavalla kuin QSM. Suurin ero on, että GSM ehdottaa tarjouksia yhtä aikaa useammalle tarjoajalle. Yhdessä nämä kaikki tarjoukset pääsisivät voittajien joukkoon, mutta eivät yksinään kuten QSM:n ehdotukset. Alustavat testit osoittavat GSM:n parantavan huomattavasti huutokauppojen lopputulosten tehokkuutta QSM:iin nähden.

Jatkossa keskitymme kehittämään GSM:ia ja testaamme sitä simuloinneilla ja koehenkilöillä. Tarjoajien käyttäytymistä, strategioita ja erityisesti oppimisen ja kokemuksen karttumisen vaikutusta käyttäytymiseen tulisi myös tutkia lisää. Tämä on tärkeää, sillä tarjoajien käyttäytyminen vaikuttaa huutokaupan suunnitteluun ja käyttöliittymältä vaadittaviin ominaisuuksiin.

Asiasanat: huutokaupat, kombinatoriset huutokaupat, päätöksenteon tukeminen, Internet-huutokaupat

Acknowledgements

Just like no man is an island, no thesis is written without help and support from other people. Thus, now that I have completed my thesis, I want to express my gratitude to the people, who have either directly or indirectly influenced this process during the past years.

First and foremost I want to thank my advisors, Professors Hannele and Jyrki Wallenius. Besides introducing me to the interesting world of auction research, you have been a constant support throughout my PhD studies. You always had time for me, and you always, always had faith in me – even when I lost faith in myself. I also want to thank my co-authors, M.Sc. Valtteri Ervasti, whose programming skills made many of my ideas possible, Professor Murat Köksalan, with whom I have had several fruitful discussions over the years, and Professor Jeffrey Teich. It has been a privilege to work with you.

I am also grateful to the two pre-examiners of my thesis, Professor David Olson and Professor Karla Hoffman, for finding time in their busy schedules to get acquainted with my work and for giving comments that helped me improve my thesis. I also want to thank everybody in the Graduate School in Systems Analysis, Decision Making and Risk Management for creating an intellectually stimulating environment.

I want to thank the Emil Aaltonen foundation, the Kemira Oyj foundation, and the KAUTE foundation for financial support.

I have enjoyed the past five years of my life, and for that I want to thank my colleagues, my friends and my family. The faculty and staff of the Institute of Strategy have made this a great place to work at. Thanks to you and our 2:00 pm coffee (tea) breaks it is really nice to come to work every day. As for my friends and sisters, thank you for all the wonderful times we have shared together; all the laughs and tears, the singing and the travels, and all the great memories. You made sure I actually had a life outside work! I especially want to thank Johanna, without whom I would not have become excited about Economics, and without whom I would not be here today. I also want to thank my piano teacher Hanne for showing me that I too can be creative, and for giving me confidence to use my creativity.

Last but not least I want to thank my mother and my father. You taught me the value of education with your example, and you have been there for me every step of the way in all the ways that count.

Otaniemi, October 2009,

Riikka-Leena Leskelä

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

I INTRODUCTION

Almost everyone is familiar with the concept of an auction. When thinking of an auction, we immediately imagine a room full of people at an auction house such as Sotheby's or Christie's. The broker presents the item for sale – typically an expensive work of art – and announces the starting price. Bidders with numbered signs in their hands indicate their willingness to purchase the item as price rises. Today, however, the auction is not limited in time or space. The advent of the Internet allowed the transferring of the auction house online. Online auctions have become increasingly popular, and today eBay is at least as well known as Sotheby's or Christie's. The variety of products up for auction is incredible: you can buy pretty much anything from dinosaur eggs to real estate property, and from diamonds to used clothes. An increasing number of people have also participated in an auction themselves, either traditional or online.

The concept of an auction, however, is not as simple or narrowly defined as a typical bidder in an eBay auction may think. Imagine you are a philatelist, and that you are attending a stamp auction in hopes of adding to your collection of 19th century Finnish stamps. You are well aware of the fact that the stamps are more valuable as a complete series than individually. Thus, you would be willing to pay a lot more for the two stamps missing from one of your series than for them individually. Now, if in the auction all stamps were auctioned individually, one after the other, how much would you be willing to bid for the first one when you do not know whether you will win the second one? Would you not be happier, if you could indicate to the seller that you would be willing to pay more, if you were guaranteed both stamps? The question then becomes, why would the seller be selling the stamps individually and not as a series? Because not all bidders are interested in the whole series, but rather different subsets of the series, and the seller cannot know what kind of packages to build from the stamps. This simple example demonstrates that there are several cases in which the traditional Sotheby's style single-item auction is not optimal for the seller or the bidders. One is then tempted to ask, if the auction could somehow be modified to deal with the problem in this example. Fortunately, combinatorial auctions provide a solution to the problem, as we will see later in section 3.2.

The typical single-item forward auction prevalent in eBay, Sotheby's and Christies is only one of many different classes of auctions. McAfee and McMillan (1987) define an auction as "a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants." This definition is quite generic, and later on in this thesis we will see, how different kinds of mechanisms fall under the category of auctions. Auctions are often thought of as a special case of negotiations, with more strict structure and a set of rules than one-on-one negotiations. Some researchers even use the term "negotiation" to also include auctions, but in this thesis, the term negotiation will refer only to one-on-one negotiations. With auctions I refer to bidding processes between one seller and many buyers or one buyer and many sellers. Specifically, the case with one seller and many buyers is called a forward auction, and the case with one buyer and many sellers is called a reverse auction or a tender. Forward auctions are common in business-to-consumer (B2C) or consumer-to-consumer (C2C) transactions, whereas reverse auctions are often used in business-to-business (B2B) transactions. When there are multiple sellers and multiple buyers in the market, it is called a double auction (in which both sellers and buyers place bids, like in the stock market), or it is simply a regular market characterized by fixed prices.

The purpose of auctions and negotiations is to determine a price for the item(s) in question. They allow the seller to discover the buyers' valuation (and the correct price) during an auction process. Thus, auctions are good for selling and buying non-standard products, the market price of which is difficult to know beforehand.

Even though most people are familiar with the concept of an auction, the field of auction studies – the science of auctions – however, is not as well or broadly known. The question "Why study auctions?" is legitimate, and should be answered to justify any research in the area. McAfee and McMillan (1987) argue that the fact that they exist in practice is reason enough to study them. A rigorous theoretical treatment of auctions can help design better auctions in practice. Moreover, studying people's behavior in auctions empirically can help design better auctions. McAfee and McMillan (1987) also argue that auctions are a case of price setting under imperfect and possibly asymmetric information that is interesting in a game theoretic sense. Maasland and Onderstal (2006) mention that auctions are worth studying, because they can be applied to other situations as well, e.g. modeling incentive contracts. Auction-based approaches have also been used in scheduling and the design of manufacturing systems (Srivinas et al., 2004, and Kumar et al., 2006), and in distribution of computing resources over a grid of interconnected computers (Schnizler et al., 2008). Thus, advances in auction theory and design can have broad implications.

Systematic and scientific study of auctions began surprisingly late, considering how long auctions have been used in transactions. Auction research started essentially in 1961 with William Vickrey's seminal paper. At the beginning, auctions were of interest to only a small group of economists, who saw them as interesting applications of game theory. Since then, the field of auction studies has grown in both depth and breadth to the wide multidisciplinary field it is today. Economics and game theory still have a strong foothold in auction research, but there is a growing interest among computer scientists, mathematicians, and decision analysts towards auctions.

As the field of auction studies grew, also the concept of auctions was broadened. At first only traditional single-item, single-unit, price-only auctions were considered (see section 1.2 for definitions). Then researchers started to consider more complicated designs: multi-unit auctions, multi-attribute auctions, and finally combinatorial auctions (which are a case of multi-item auctions). The new designs allowed more and more goods to be sold and bought through auctions. The Internet has enabled the implementation of many of the more complicated designs, and it has brought up new research questions. Today, over forty years after Vickrey's article, both the use of auctions as a market mechanism, and the study of auctions are thriving.

This thesis contributes to the study of combinatorial auctions. In short, combinatorial auctions are multiple-item auctions, in which the bidders can place bids on packages of items. I discuss the complexities of combinatorial auctions, and the need for decision support for bidders in such auctions. In particular, the contribution of this thesis is the design of decision support tools for bidders in semi-sealed-bid, iterative combinatorial auctions (see section 6.3 for a more elaborate description of the objectives and methods of this study). The decision support tools are algorithms that provide the bidders with suggestions for bids to be placed at any given time in the auction. In this thesis I present the mathematical formulation for the decision support tools, the Quantity Support Mechanism (QSM), and its variation, the Group Support Mechanism (GSM). These tools are also easily extended to a multi-attribute combinatorial auction. I also present the design and results of two simulation studies we ran to test the QSM. Based on the results of the studies, some improvement ideas were developed, among them the GSM. I will present the improvement ideas, and mathematical formulations for them. The QSM was also implemented in an online auction system, CombiAuction. The user interface and the usability of the QSM were tested in a laboratory experiment with human subjects. Thus, this thesis also contributes to the design of auctions in practice. Also, I identified strategies bidders used in the combinatorial auctions in the experiment.

The structure of this thesis is the following. The Introduction contains this general introduction to auctions and auction research to be followed by a chapter reviewing the key concepts related to auction design (market types, auction mechanisms). The second part reviews relevant literature on auctions. The purpose of the literature review is to discuss combinatorial auction research, to position it in the field of auction studies, and to illustrate the links between other disciplines (economics, computer science, multicriteria decision making) and auction research. The third part discusses the need for support in combinatorial auctions, and shows the gaps in existing literature that this study attempts to fill. The third part also contains a more detailed description of the objectives and methods of the research. The fourth part presents the formulation of the QSM, and the designs and results of two simulation studies. The chapter discussing possible small improvements to the QSM concludes the fourth part, and demonstrates the need for more drastic changes in the QSM if significant improvement is to be achieved. The fifth part introduces the Group Support Mechanism (GSM), a support tool based on the QSM but with a few significant alterations, and presents an example to illustrate the mechanism. The sixth part describes the *CombiAuction*, an online auction system in which the QSM has been implemented, and an experiment I ran with human subjects. The seventh part concludes the thesis.

1 DESIGNING AUCTIONS

Understood as a traditional arts auction, auctions seem to be relatively simple market mechanisms. However, the auctions are actually a group of different transaction mechanisms, some more complicated than others. The common feature with all the transaction mechanisms categorized as auctions is that their purpose is to determine the price and allocation of the item(s) in question. Researchers have identified a number of desirable properties for auctions, and the purpose of auction design is to ensure that a desirable outcome is reached. Auction design takes place on many levels. On the macro level, the market setting, i.e. the type and number of goods for sale, affects the design. In addition, the auction owner must decide on the appropriate auction mechanism (e.g. ascending price or descending price, open or sealed-bid). Under any market setting, a number of different auction mechanisms can be used. On the micro level, auction design must also say something about the detailed design aspects (whether to set a reservation price, what kind of information to disclose, etc.) and the rules of the auction (who can participate, when will the auction end, how is defaulting by a bidder dealt with, etc.) to ensure fairness, maximal revenue for the auction owner, and impede cheating by the participants. In this chapter I will first review the desirable properties of auction designs. Thereafter I will review macro level concepts and definitions related to auction design including different market settings and auction types, and the most common auction mechanisms. The details of auction design and auction rules will be discussed more thoroughly in the context of Internet auctions (Chapter 5 of the literature review).

1.1 Desirable Properties of Auctions

There are several properties the auction organizer may want the auction to possess. From the economics point of view, an important property is allocative efficiency (or efficiency for short). Allocative efficiency is reached when the winner in the forward auction is the bidder with the highest valuation (or lowest production cost in reverse auctions). This also means that total welfare in the society is maximized. Several researchers (e.g. Bichler, 2001, Milgrom and Weber, 1982, and McAfee and McMillan, 1987) use the term “Pareto optimality” in a similar sense as allocative efficiency, but in my opinion it is not the best term to use here, and also Koppius and van Heck (2002) make a distinction between Pareto optimality and allocative efficiency. Thus, in this thesis I will use the term allocative efficiency in the context of auction design.

A property closely related to allocative efficiency is revenue maximization (or cost minimization). Often the seller would like to design the auction in such a way that her revenue is maximized. An auction mechanism that maximizes the seller's revenue is also called an optimal mechanism. Usually a revenue maximizing auction is also efficient, because the bidder with the highest valuation is the one who is willing to pay the most. However, an efficient auction need not be revenue maximizing, as we shall see later on in section 2.5.

The objective of incentive compatibility is also closely related to allocative efficiency and revenue maximization. Incentive compatibility means that the bidders have no incentive to shade their bids, and that they would be willing to report their true valuations, because lying would not increase their pay-off. When an auction is incentive compatible the bids reflect the magnitude of the bidders' valuations, and the ranking of the bids reflects the ranking of the bidders' valuations. Thus, the winner is the one with the highest valuation, and depending on the payment rule, the revenue to the bid taker can be maximized.

Other desirable properties include fairness, failure-freeness (robustness), resistance to cheating and manipulation, and low transaction costs. These properties are discussed in more detail in the literature review (Chapter 5).

1.2 Market Settings

The market setting under consideration sets the stage for the auction design. The different market settings can be classified based on 1) the number of different items (products or services) to be sold[1] in one auction, 2) the number of homogenous units of each item, and 3) the number of units each bidder wishes to acquire. In the simplest

[1] I will present the auction concepts in the context of forward auctions, which is the convention in auction literature. All the concepts apply in the reverse setting, with the only difference that items in the auction are to be bought (not sold), and thus bids represent supply rather than demand.

case there is only a single unit of one item for sale. Naturally, in this setting the bidders can only demand one unit. This auction type is often called a **single-item auction** for short. Single-item auctions were the focus of most early auction studies conducted by economists and game theorists. The properties of single-item auctions will be discussed in Chapter 2 of the literature review.

A logical extension of the single-item auction is the (single-item) **multiple-unit auction**. In such an auction at least two identical units of a particular item are to be sold. Depending on the nature of the item, each bidder can now demand either one unit (for example, if several identical licenses are auctioned, the bidders would want at most one license), or several units. See section 2.8 in the literature review for more discussion on multiple-unit auctions.

The most complex auctions are the **multiple-item auctions**, which are the focus of this thesis. There can be either one unit or multiple units of each item for sale. The bidders may wish to acquire only one of the items, or several. However, it is common to assume that at least one of the bidders wish to acquire more than one item in the auction. If each bidder had a demand for only one item, it is difficult to imagine, why hold one auction for a collection of items, unless the items were close substitutes. Thus, the case of multiple items but single-item demand is usually neglected in literature. Theoretically, in multiple-unit, multiple-item auctions bidders can either demand one unit or multiple units of the different items. Usually, however, in multiple-item auctions it is realistic to assume that bidders demand more than one unit. Thus, the case of multiple items and multiple units, but only single-unit demand has not been considered in auction literature.

In the multiple-item auctions there is also a big difference in the design of the auction depending on how the bidders are allowed to express their demand. Basically the choice is between allowing bids on combinations of items (also called **package bidding** and **combinatorial bidding**), or not allowing them. A combinatorial bid is a vector containing the desired quantities of each item, and a single price for the combination. The efficiency of the auction is improved, if bids on combinations are allowed, but the complexity of the auction increases a lot. I will return to combinatorial auctions and this trade-off in section 3.2. The different market characteristics described above result in eight different market settings (see Table 1).

Thus far I have implicitly assumed that bids are evaluated only based on the price attached to the bid. However, just as in many negotiations, the price of the items is not necessarily the only attribute of interest to the buyer in an auction. In such cases, the ability to make **multi-attribute**[2] bids can increase the efficiency of the auction, and it can even make it rational to sell through auctions some items, which have previously been sold through one-on-one negotiations. In these auctions, the bidders' bids are multidimensional vectors with one component for each attribute. Examples of non-price attributes often used in multi-attribute auctions are quality, terms of payment, delivery times etc. Multi-attribute auctions can be either single-unit or multiple-unit auctions. Guttman and Maes (1998) refer to multi-attribute auctions as win-win situations since the auction is no longer a zero-sum game and it is possible for both sellers and buyers to be better off. Thus, the nature of the multi-attribute auction is somewhat different compared to the price-only auctions. The addition of multiple attributes to the auction causes new kinds of complications in the auction process. For instance, the comparison of bids against each other becomes conceptually difficult. The characteristics of multi-attribute auctions, the arising problems, and the attempts to overcome the problems are briefly discussed in Chapter 4.

The possibility to consider other attributes besides price adds another eight different variants to the original eight market settings. Thus, every row in Table 1 corresponds to two market settings: one with price as the only bid attribute, and the other identical otherwise but with multiple attributes considered.

[2] Some researchers use the term 'multi-dimensional auction', but I will use the term 'multi-attribute auction' to avoid confusion, because former term can also refer to a combinatorial auction.

Table 1 Summary of market settings

# of Items	# of Units	Demand	Package Bidding	# of Attributes
single	single	single-unit	-	price-only/ multiple attributes
single	multiple	single-unit	-	price-only/ multiple attributes
single	multiple	multiple-unit	-	price-only/ multiple attributes
multiple	single	single-item	not allowed	price-only/ multiple attributes
multiple	single	multiple-item	not allowed	price-only/ multiple attributes
multiple	multiple	multiple-unit, multiple-item	not allowed	price-only/ multiple attributes
multiple	single	multiple-item	allowed	price-only/ multiple attributes
multiple	multiple	multiple-unit, multiple-item	allowed	price-only/ multiple attributes

1.3 Auction Mechanisms

An auction mechanism is defined as a set of rules telling how the winner is determined, how the payments of each bidder are determined, and how the bid information is collected from the bidder. According to Krishna (2002), a generic mechanism consists of three elements: the set of possible bids, the allocation rule, and the payment rule. The allocation rule determines the probability with which a bidder will win the object. The payment rule determines the payment the bidder with the winning bid must make. In every market setting described above there can be different auction mechanisms. Auction literature recognizes four basic mechanisms, which are most commonly used: the **English auction**, the **Dutch auction**, the **first-price sealed-bid auction**, and the **second-price sealed-bid auction** (also known as the **Vickrey auction** after its inventor William Vickrey). These basic mechanisms are special cases of the generic mechanism.

In a traditional ascending price English auction, the auctioneer[3] starts the bidding at the reservation price, if the auctioneer has set one. If no reservation price is specified, the starting price is set to the starting price specified by the seller. If no such price is determined, the starting price is zero. The bidders can then call out bids. A new bid has to exceed the currently highest bid to be acceptable. Depending on the specific rules of the auction, the bidders can either freely call out any acceptable bid, or the auctioneer calls out new prices which the bidders can accept. The latter version is referred to as a clock auction. The auction ends when no bidder is willing to increase her bid. The bidder i with the highest bid wins the item and pays the price equivalent to her bid. English auctions are most common in practice. They are especially popular in art and antiquities auctions, and consumer-to-consumer (C2C) auctions, such as eBay.

The Dutch auction reverses the logic of the English auction: the price descends in the Dutch auction. Thus, in the Dutch auction the auction clock is set at a very high price at the beginning. In fact, the price is set so high that no bidder would be willing to pay the price. When the auction begins, the price indicated by the auction clock is gradually decreased until one bidder indicates her willingness to pay the current price. The auction ends, and the bidder receives the item at the price indicated by the auction clock. If the auction clock reaches the seller's reservation price, and no bids have been made, the item is left unsold. The allocation rule in the Dutch auction is the same as is the English auction: the bidder with the highest bid wins the item. However, in the Dutch auction the allocation rule is trivial, since by definition there will only be one bid in the auction. Dutch auctions have been used a long time in flower auctions in the Netherlands (hence the name). Also other perishable items (e.g. fish) are auctioned through Dutch auctions, because they are fast to conduct.

In the first-price sealed-bid auction[4], all bidders simultaneously submit their bids to the auctioneer. This means that the bidders are unaware of the content of all bids except

[3] In this text the "auctioneer" or the "auction owner" refers to the bid taker, i.e. the seller of the good in a forward auction and the buyer in a reverse auction. It is also possible that the auctioneer is a neutral third party (e.g. an auction house) but that case is omitted from this discussion.

[4] In this text I will use "first-price auction" as shorthand for "first-price sealed-bid auction."

their own. The auctioneer goes through the bids and the bidder with the highest bid wins the item and pays the price equal to her bid (if the price exceeded the seller's reservation price). The allocation and payment rules are the same as in the Dutch auction. The second-price sealed-bid auction is identical to the first-price auction except for the fact that the winner (the bidder with the highest bid) pays the amount equal to the second highest bid. Sealed-bid auctions are a common practice in procurement situations, both public and private sector. One reason for their popularity is the fact that the bids may contain critical information of the bidders' cost structure or their competitive advantage. Thus, bidders prefer to keep their bids secret from their competitors.

The English auction and the Dutch auction are *open-cry* auctions as opposed to *sealed-bid* auctions. In open-cry auctions the bidders call out their bids. In other words, bidders are aware of the actions of their competitors. *Dynamic* (*progressive*, *iterative*) auctions have multiple rounds and bidders can revise their bids (place several bids). The English auction is a dynamic auction. The Dutch, first-price and Vickrey auctions, on the other hand, are static auctions. In all three auctions the bidders have only one chance to place the bid and no revision is allowed.

All the four mechanisms presented above were considered in the *forward* setting. However, all mechanisms can be used in the *reverse* setting as well. If in the forward English auction the price was ascending, it is descending in the reverse English auction. Similarly, in the reverse sealed-bid auctions the winner is the bidder with the lowest bid. In the reverse Dutch auction, the clock starts at a very low price and is then increased until some bidder agrees to take the item (e.g. a contract) at the current price. In the reverse first-price and Vickrey auctions the winner is the bidder with the lowest bid, and the price in the Vickrey auction is that of the second lowest bid.

Even though I have presented the four basic auction mechanisms only in the simple single-item, single-unit case, the same mechanisms can be extended to more complex settings. Also, most auction mechanisms considered in literature or used in practice contain elements of one or more of the four basic mechanisms, as will become apparent in the literature review.

II LITERATURE

2 THE ORIGINS OF AUCTION THEORY: SINGLE-ITEM AUCTIONS

Traditional auction theory discusses mainly the single-item, single-unit auction. The setting is simple enough so that equilibrium strategies can be solved analytically and comparisons between different auction mechanisms can be made. A good reference for a pure game theoretic discussion of auctions is Wilson (1992). However, the game theoretic approach requires that some specific assumptions have to be made about the bidders and the item on sale. Literature focuses on two main models that differ significantly in their assumptions about the bidders' valuation for the auctioned item: the Independent Private Values (IPV) model and the Common Value (CV) model. In the IPV model each bidder's valuation is assumed independent of other bidders' valuations. The bidder, however, knows her valuation with certainty. In the CV model the item's value is the same for all bidders, but the bidders only have an estimate of the item's "true" value. The affiliated values model presented in section 2.3 combines elements from the two extreme models. Studies are concerned with the expected revenue from the four basic mechanisms, the effects of relaxing assumptions behind the model, and the design of optimal auction mechanisms (i.e. revenue maximizing mechanisms) given the set of assumptions in the models. Both IPV model and the CV model, as well as the affiliated values model will be presented briefly in the following sections.

A logical extension of the single-item, single-unit auction is the single-item, multiple-unit auction. I will call it the multiple-unit auction for short. In such auctions, a set of identical objects is sold (bought). The models for multiple-unit auctions are usually created under the IPV assumptions. Multiple-unit auctions are discussed in section 2.8.

All the auctions in this section are presented in the forward setting, because it is common practice among economists. I will try to keep the presentation non-technical, so the reader interested in the exact proofs of propositions is advised to look them up in the original references.

2.1 The Independent Private Values Model

The Independent Private Values (IPV) model was the first framework in which auctions were systematically studied. The earliest studies in the spirit of the IPV model were conducted already in the 1950s. Friedman (1956) analyses a situation in which each bidder can estimate the other bidders' behavior based on prior experience. The groundbreaking work was done by William Vickrey (1961). His main discovery was that under specific assumptions the second price auction he designed would generate the same expected revenues as the first-price and Dutch auctions.

The key assumptions behind the IPV models are (collected from Vickrey, 1961, McAfee and McMillan, 1987, Rothkopf and Harstad, 1994):

1) Each bidder knows the true value of the item for her, but she does not know the valuations of the other bidders.
2) The valuation of one bidder is statistically independent of any other bidder's valuation.
3) The bidder perceives the other bidders' valuations as drawn from some known probability distribution, and she knows that other bidders regard her valuation as being drawn from some distribution.

In addition, some other assumptions are either explicitly or implicitly presented in the context of the IPV model (see e.g. McAfee and McMillan, 1987, Rothkopf and Harstad, 1994, Maasland and Onderstal, 2006):

4) The bidders (and the seller) are risk-neutral.
5) The bidders are symmetric (i.e. they draw their valuations form the same distribution).
6) There is a single, isolated auction (not a multiple-stage game), and the number of bidders participating is fixed.
7) There is no collusion among bidders.
8) There are no externalities from the allocation of the item to the bidder, or from the payment made by the winner.

2.1.1 Strategic Equivalence of Auction Mechanisms in the IPV Model

In the IPV model, the Dutch auction and the first-price auction are strategically equivalent, as are the English and second-price auctions. The strategic equivalence of the Dutch and first-price auctions actually extends to other situations as well, but the strategic equivalence of the English and second-price auctions breaks down if bidders' valuations are not independent (see sections 2.2 and 2.3 for more discussion). However, for now let us return to the IPV model.

Several review papers provide a thorough explanation of the strategic equivalence. The following presentation is adapted from Paul Milgrom (1989), Martin Bichler (2001), and the somewhat more mathematical presentation of McAfee and McMillan (1987).

Consider first the English auction. The auction starts at a relatively low price and the price rises as the auction proceeds. Since the winning bidder has to pay her own bid, and every bidder knows her own valuation, it does not make sense for them to bid a price above their valuation (they would be better off not participating at all than paying more than their valuation). One by one, bidders drop out of the competition until only two bidders are left. I will denote them with B_1 (bidder with the highest valuation) and B_2 (bidder with the second highest valuation). Once the current bid price reaches the valuation of B_2 she will not place any more bids. Bidder B_1 could still bid higher, but it does not make sense to do so since she can win the auction by bidding only marginally higher than what B_2 is willing to bid. Thus, if B_1 bids rationally (and we will assume that she does), she wins the item and ends up paying the price equal (or almost equal) to the valuation of B_2. The dominant strategy for each bidder is to bid until the price reaches her valuation – regardless of what other bidders do.

The Vickrey auction seems different to the bidders since they cannot observe each others' bids and therefore obtain no information on other bidders' valuations. However, it does not matter, because as was explained in the previous paragraph, the opportunity to observe competitors' behavior in the English auction did not affect the bidders' optimal strategy. As it happens, in the Vickrey auction it is also a dominant strategy for each bidder to bid their own valuation. Remember that the bidder who wins only has to pay the price indicated by the second highest bid, i.e. the price is determined independent of the winning bid. If the bidder bids less than her valuation, she gains nothing with that move: if she is still the winner, she pays the price of the second highest bid, which would have been the same even if she had bid her valuation. Thus, by bidding less than her valuation the bidder only risks losing the item. If, on the other hand, the bidder bids higher than her valuation, she increases her chances of winning the auction. This strategy becomes effective only when she is not the bidder with the highest valuation, and by bidding higher than her valuation she manages to outbid the bidder B_1. She manages to win the auction, but the price she has to pay (i.e. the second highest bid) is now equivalent to the valuation of bidder B_1. This price, by definition, is higher than the winning bidder's valuation of the item, and she ends up with a negative pay-off. Thus it is clear that rational bidders bid a price equivalent to their valuation. The bidder with the highest valuation wins, and she pays the price equivalent to the valuation of the bidder with the second highest valuation. The strategies and the outcome are clearly the same as in the English auction.

The Dutch and first-price, sealed-bid auctions seem very different from one another at first glance. In the former, the auctioneer cries out prices whereas in the latter each bidder submits a sealed bid. However, assuming that the bidders plan their actions prior to the auction, their decision problem is exactly the same in either case. Let us assume that the bidder is bidding for an item in a Dutch auction. Assume that her strategy is the following. First she waits for the price to drop to p_1 and if no one has claimed the item, she will either place a bid, or wait. Assume that p_1 is greater than what she is willing to pay for the item, so she decides to wait. She now chooses p_2 as her new point of evaluation, and the same process repeats itself. As long as the price is higher than her valuation for the item, the choice is trivial. Once the price goes below her valuation, she needs to weigh the added utility from letting the price fall further against the risk of losing the item to another bidder. Finally, at price p the bidder places a bid and claims the item. Notice, however, that she has made her choices always under the assumption that no one else claimed the item. There is no point in considering the case when some one places a bid before her, because then the game is over. In the beginning of the auction she could have been asked to directly indicate the highest price at which she is willing to claim the item, and the answer would have been the same. The only additional information the bidder gets during the auction is either that somebody was willing to pay more, or that other bidders were not willing to pay as much as she was. Either type of information is useless to a bidder, when valuations are assumed independent and private. In the first-price auction the bidder is faced with the same trade off between larger utility and smaller probability of winning, because the Dutch auction offers no added information to the bidders in the IPV case. Thus she ends up bidding the same price p for the same item. This can be verified mathematically.

I have now shown that the price paid by the winner in the English and second-price auctions is equal to the second highest valuation, and in the Dutch and first-price, sealed-bid auction equal to the expected value of the second highest valuation. The auction owner organizing the auction does not know the valuations of the bidders, so her expected revenue from all four auction mechanisms is always equal to the expected value of the second highest valuation! Note that this does not imply that the actual outcomes of all auction mechanisms would always be the same. In the English and Vickrey auctions the price paid by the winner is always equal to the valuation of bidder B_2. In a first-price and a Dutch auction the price of the winning bid is the expected value of the second highest valuation. These two prices are identical only by accident, but on the average they are the same.

2.1.2 The Revenue Equivalence Theorem

Vickrey's result in 1961 regarding the revenue equivalence of the four basic auction mechanisms (as described in the previous section) was the preliminary version of the much celebrated revenue equivalence theorem. The revenue equivalence, however, can be extended to a much broader class of auction mechanisms. The exact formulation of the theorem was proposed by Myerson (1981). According to Myerson, the seller's revenue from the auction is completely determined by the allocation rule, and the utility gained by the bidder with the lowest possible valuation. As long as the auction mechanisms allocate the item to the same bidder, and the utility for a bidder with the lowest possible valuation is the same, the expected revenue for the seller is also the same.

Vickrey's discovery of the equivalence of the four basic mechanisms is thus a special case of the revenue equivalence theorem. In all four auctions the allocation mechanism is the same (highest bid wins), and the expected utility of the bidder with the lowest possible valuation is zero (she would never be the winner, and losing bidders do not have to pay anything). Only the payment rules are different, but according to

Myerson's theorem, the payment rule does not matter. One should keep in mind, though, that Myerson's theorem relies on the assumptions of independent private values, bidder symmetry and risk neutrality. When the IPV assumptions are relaxed, the equivalence of the four basic mechanisms breaks down. The practical implication of this is that in reality it does matter, which mechanism the seller chooses.

2.2 The Common Value Model

Not satisfied with the assumption of independent and private values (assumptions 1 and 2), researchers such as Wilson (1969) and Capen, Clapp and Campbell (1971) replaced it with an opposite assumption. The underlying assumption in the Common Value model is that the true value of the item is the same to all bidders, but at the time of the auction this value is unknown to all participants. This assumption is valid, for example, in the bidding for oil drilling rights, where the value of the asset (the value of the oil extracted) is the same to all participants, but no one knows beforehand the amount of oil in the area. The bidders make estimates of the "true" value of the item based on the information they have.

The interesting characteristic of a common value auction is that it is by definition always efficient. The downside of a common value auction is that the task of the bidders is much more difficult than in private value auctions. It is hard to determine a bid, when you do not know the value of the item to you. If the bidders are not careful they can easily fall prey to what is called the winner's curse. The following example presented by Milgrom (1989) illustrates the pitfall embedded in the common value auction. The example is presented in the reverse setting contrary to the previous sections. Let's assume that contractors are invited to bid on a job, which each of them can complete at a cost C. The contractors make unbiased estimates $c_i = C + \varepsilon_i$ of the cost, where ε_i is the estimation error of contractor i. The estimation errors are statistically independent, and the expected value of ε_i is zero (because estimates were assumed unbiased). Even though on average the estimates are correct, in every auction there will always be bidders whose estimates are higher or lower than the true value C. Because the expected value of the error term is zero, the expected value of the smallest estimate error must be less than zero. Each bidder is unaware of the other bidders' estimates, so she cannot know whether her estimate is above or below average. The bidder with the lowest cost estimate will bid for the lowest price and will win the contract at that price. However, as she was the bidder with the smallest c_i, her bid price is likely to be less than the true cost of the contract C. The winning bidder incurs losses and thus suffers from the winner's curse. In other words, the bidder's estimate of the production cost increases when she learns that she is the winner (and therefore the one with the lowest estimate). Unfortunately, at that point the auction has closed and she cannot revise her bid.

A wise bidder acknowledges the fact that if she wins, she has underestimated the production costs. Thus, she bids a higher price than bidders who do not acknowledge it. The optimal bidding strategy is based on the assumption that the bidder is the one with the lowest estimate, and the task is to, given this assumption, figure out the expected value of the second lowest estimate. Equilibrium bidding strategies can be derived just as for the IPV model, but they are a lot more complex (see, for example, Wilson, 1977).

In the common value case the English auction differs from the sealed-bid auctions, because the information available during the auction is now valuable to the bidders (contrary to the IPV case). The bidders can observe others' value estimates, and revise their own estimates accordingly. Thus there is not as much reason to correct the bids as the estimate becomes more accurate (Milgrom and Weber, 1982). Due to less uncertainty, bidders can also bid more aggressively and the English auction leads to higher expected revenues for the auction owner.

2.3 The Affiliated Values Model

In most cases it is not realistic to assume the bidders' valuations to be strictly common or independent. Items for sale have both a private and a common value element (Bichler, 2001). Unique pieces of art are usually used as examples in the IPV models, but more often than not the bidders are also interested in the resale value of the art work in addition to their private valuation. Thus there is a common value element in most IPV cases. Similarly, one can argue that firms differ in their resources and capabilities, and thus there can be a private value element in common value auctions.

Milgrom and Weber (1982) introduce a more general model – the Affiliated Values model – that combines elements from both the IPV and Common Value models. Affiliation between valuations means that if some bidders value the asset highly, it is more likely that the other bidders' valuations are high as well. The exact mathematical definition of affiliation is adopted from Milgrom and Weber (1982). Let v represent the vector of the valuations of the N bidders of the auction. Let f: $R^N \rightarrow R$ denote the joint probability distribution of the valuations. Finally, let $v \vee v'$ denote the vector containing the component-wise maximum, and $v \wedge v'$ the component-wise minimum of two valuation vectors v and v'. Then

$$f(v \vee v')f(v \wedge v') \geq f(v)f(v') \qquad (1)$$

This equation is roughly saying that it is more likely to have either a vector with relatively large valuations for all bidders ($v \vee v'$) or relatively low valuations ($v \wedge v'$) than a mixture of high and low values (either v or v').

The Affiliated Values model allows for statistical dependence between the bidders' value estimates as well as for differences in individual tastes. There can also be different degrees of affiliation. The IPV and the common value cases are included in the model as two extreme cases alongside numerous intermediate models, which are perhaps more realistic.

The revenue equivalence of the four basic mechanisms does not hold under the affiliated values model. In fact, Milgrom and Weber (1982) show that the four different mechanisms can be rank-ordered based on the expected revenue collected in the auction. It can be shown that the English auction generates more revenue than the other auction mechanisms. The situation is analogous to the pure common value auction: in an open-cry auction the bidders obtain information about each other's valuations and can then revise their own estimates. It can also be shown that the Vickrey auction yields a higher expected revenue than the Dutch and first-price auctions, which remain revenue equivalent. Although the revenue ranking of affiliated values auctions is similar to the pure common value auction, the observation that a common value auction is always efficient does not carry over to the affiliated values case (Goeree and Offerman, 2002). It is interesting that also the IPV auctions are theoretically efficient, yet the affiliated values auction, which is a mixture of both extremes, does not have this property. The intuition behind this result is quite straightforward. When both private and common elements are present, a bidder with a small private value but an overly optimistic estimate of the common value element may outbid a bidder with a higher private value.

2.4 The Popular English Auction

In light of the perhaps surprising results of the revenue equivalence theorem, it is interesting to observe that in practice auction owners are not indifferent between the auction mechanisms. For instance, the English auction is by far the most popular, and the Vickrey auction is hardly ever used. Milgrom (1989) offers an explanation to this discrepancy by saying that revenue is only one criterion to evaluate auction mechanisms. Others are robustness, efficiency, transaction costs, fairness, and immunity to cheating.

Robustness here refers to the vulnerability of the mechanism to changes in the IPV model assumptions. The English auction and the Vickrey auctions are more robust than the Dutch and first-price auctions. This is because bidders have a dominant strategy, which is independent of the distribution of other bidders' valuations, and the number of bidders in the auction (McAfee and McMillan, 1987). On top of that, the English auction is better than the Vickrey auction when there is a common value element in the auction. This is because the English auction is the only mechanism of the four basic mechanisms in which the bidders can observe each other's bids, and learn about other bidders' valuations. However, introducing risk averse bidders in an auction makes the first-price and Dutch auctions better in revenue terms (McAfee and McMillan, 1987). And since risk aversion among bidders is not an unreasonable assumption, this is a valid argument against the English auction. Another such argument is that the English auction is more vulnerable to collusion in the form of bidding rings (Robinson, 1985), and also cheating in the form of signaling simply because it is the only mechanism in which the other bidders can see the content of the bids.

Efficiency usually refers to the allocative efficiency of the auction outcome, which is always achieved under the IPV model. However, relaxing some of the assumptions breaks this result for the Dutch and first-price auctions. Milgrom (1989) broadens the concept of efficiency to include bid preparation costs. The more information gathering is required, and the more complicated the calculation of the bidding strategies, the higher the preparation costs. This argument also favors the English and second-price auctions, as the bidding strategies are simple, and in English auctions all the information available can be collected during the auction. According to Engelbrecht-Wiggans (2001) the lower participation cost of English auctions attracts more bidders, and on average the more there are bidders, the higher the expected revenue for the bid taker.

On the other hand, the preparation cost of bids in the English auction may be low, but it requires the bidders to actively participate in the auction for the whole duration of the bidding process. In 1989, when Milgrom wrote his article, participation also usually required physical presence. Nowadays with the Internet, physical presence is not that critical, but the English auction still requires the bidder to be alert and present online. Thus there is a time cost involved in the English auction, which increases its otherwise low transaction costs. A quick remedy for this would be to use bidding agents. The bidder would indicate the highest acceptable price, and the agent would bid on her behalf up to that price. In fact, bidding agents of one sort or another are becoming more and more common. For example, eBay uses bidding agents. Another solution would be to organize a Vickrey auction – this is the only solution, if we are talking of a traditional "off-line" auction. This is not usually done, however, because the Vickrey auction is highly susceptible to manipulation. Nothing stops the bid taker from inserting extra bids that increase the price charged from the winner. The possibility of such manipulation decreases the trustworthiness and attractiveness of the second-price auction. The Vickrey auction also requires adequate competition (as does the English auction, though); otherwise the price paid by the winning bidder can be much too low from the perspective of the seller. McMillan (1994) illustrates the problem of too little competition with an example form a spectrum auction in New Zealand. In that auction the winner with a bid of NZD 7 million only paid NZD 5000, which was the second-highest bid.

The strongest argument for the English auction, however, is that made by Milgrom and Weber (1982). They show that whenever there is a common value element in the auction, the English auction provides the highest expected revenue. It is reasonable to assume that in most auctions there is a common value element – the resale value of the item, if nothing else – and thus it is only natural that the English auction is so common.

Finally, an additional explanation for the popularity of the English auction, which is often omitted in auction literature, is the fact that the general public – the potential participants – are familiar with it. Throughout the centuries the English auction has been all but synonymous with the term auction, and it has become something of an industry standard. Bidders are more prone to participate in an auction which they are familiar with – especially the risk averse ones. Also, learning the rules of an auction always takes time and constitutes a transaction cost. Therefore, bidders tend to choose auctions for which they already know the rules. Thus it is in the interest of the auctioneer to organize such an auction, because the more there are participants the higher the expected revenue.

2.5 Optimal Auctions

The term "optimal auction" can either refer to an efficient auction (as understood by Vickrey 1961), or to a revenue maximizing auction (as understood by Myerson, 1981, and Riley and Samuelson, 1981). An efficient auction maximizes the welfare of the society as a whole, where as a revenue-maximizing auction maximizes the payoff to the seller. The four basic auction mechanisms described above are efficient, but not necessarily revenue maximizing. In this section I will discuss the design of a revenue maximizing auction. Also, in this thesis, the term "optimal auction" refers to a revenue maximizing auction.

Conceptually, the derivation of an optimal auction mechanism is straightforward: the task is to simply maximize the seller's expected revenue subject to individual rationality and incentive compatibility constraints. In practice, however the problem becomes easily intractable unless we resort to restrictive assumptions such as the IPV framework. In the IPV framework it can be shown that all the four basic auction mechanisms are optimal, if an appropriate reservation price is added (Myerson, 1981, Riley and Samuelson, 1981). The optimal reservation price is set to mimic the expected bid from the bidder with the second highest valuation, and it is strictly greater than the seller's own valuation for the item. The reasoning behind this is that if the reservation price is higher than the valuation of the second highest bidder, the seller's revenue is increased, because the winner now has to pay the reservation price, and not the second highest bid. However, it is also possible that the reservation price exceeds the valuation of the highest bidder as well, and no sale takes place even though the valuation of the seller was lower than that of the highest bidder. Thus, the optimal auction is not always efficient. If designing a real auction, one should consider the results of these theoretical models with caution, though. For instance, the number of bidders is treated as fixed and exogenous in the models. In reality, bidder entry is endogenous, that is dependent on the auction rules, and opening prices among others. Bajari and Hortaçu (2003) find out that when a secret reservation price is determined by the seller, fewer bidders entered the auction, which on average resulted in lower revenues for the seller.

Relaxing the IPV assumptions results in very complex optimal mechanisms. For instance, if the bidders are asymmetric, then it is optimal to favor the low-valuation bidders, because it forces the high-valuation bidders to bid higher than what they normally would. On the other hand, in the case of risk averse bidders it is optimal to subsidize high bidders who lose and penalize low bidders. The optimal auctions are very complex, and almost impossible to implement in practice, because the design of the optimal mechanism requires a lot of information about the bidders' valuations, risk attitudes etc. Efficient auction mechanisms are not as complex as optimal mechanisms, and they are better for the society as a whole. Hence, in the remainder of this thesis, efficiency will be used as a primary measure of the goodness of any auction mechanism.

2.6 Empirical Studies of Auctions

Theoretical studies of single-item auctions abound, as is evident from the previous sections. Another strand of economics is interested in the practical side of auctions. A key question is how realistic the theoretical models are (or how realistic the assumptions behind the models are). On the one hand researchers have studied real auctions, and on the other hand they have conducted controlled experiments to study human behavior in different auction settings.

2.6.1 Field Studies

Field studies of auctions observe bidding behavior in real auctions and try to test the predictions of auction theory. The problem with empirical analysis is that any tests should be able to first determine underlying risk preferences of the bidders, the independence (or interdependence) of bidders' valuations, and the symmetry of the bidders. The studies usually indicate that the theory does not hold in practice, but the results are contestable more often than not. Some field studies conducted in online auctions study the behavior of the bidders (see e.g. Bapna, Goes and Gupta, 2000, 2003). Among the most studied and documented real-life auctions are the radio spectrum license auctions held by the Federal Communications Commission (FCC) in the United States. The FCC auctions will be discussed further in section 3.1.2.2.

2.6.2 Experimental Studies

It is a bit difficult to analyze empirical data on auctions in the light of auction theory, because so many parameters (e.g. valuations, risk attitudes) are difficult to observe and to control. The school of experimental economics founded by Vernon Smith, Charles Plott and John Ledyard has attempted to test the theories in practice with human subjects, but under more controlled circumstances than what is possible in field studies. The experimental studies try to cut a balance between the realistic nature of the bidding situation and full control of the parameters.

IPV model

Numerous experimental tests of the IPV model have been conducted. Kagel (1995) reviews in detail experiments on all aspects of the traditional auction models. Bichler

(2001) gives a more general overview. The most commonly tested issues are the revenue equivalence of the four basic auction mechanisms and the efficiency of the mechanisms. Most research has been limited to the pairwise comparison of the Dutch and first-price auctions, and the English and Vickrey auctions. The revenue equivalence has not held even between the allegedly strategically equivalent auction pairs, so there has been no need to test the equivalence of all four mechanisms.

The experiments reviewed by Kagel (1995) conclude that subjects do not behave in strategically equivalent ways in first-price and Dutch auctions or in Vickrey and English auctions. The revenue equivalence between the auctions did not hold either. Equilibrium prices were consistently higher in first-price sealed-bid auctions than in Dutch auctions (Coppinger, Smith and Titus, 1980, Cox, Roberson and Smith, 1982). Similarly, prices in Vickrey auctions exceeded those of English auctions (Kagel, Harstad and Levin, 1987). Interestingly, the bids in Vickrey auctions consistently exceed the dominant strategy. This is probably due to the fact that the dominant strategy in Vickrey auctions is far from obvious.

Lucking-Reiley (1999) presents a more recent test of the revenue equivalence, which was conducted in the WWW-environment. He auctioned off collectable Magic cards worth around $2,000 by posting advertisements to news groups and using e-mail as a communication tool. His experiment differed from previous experiments in the sense that the bidders did not know they were participating in an experiment. On the one hand, this made the setting more authentic but on the other hand, it made it impossible to control for the assumptions underlying the IPV model. It was also impossible to control the number of bidders in each auction. Lucking-Reiley concludes that the equilibrium prices in the Dutch auctions were significantly higher than in the first-price auctions. The prices in the Vickrey and English auctions were about the same, although bid-level data indicated some tendency for bidders to bid higher in the English auctions. Interestingly, these results conflict with those reviewed by Kagel (1995). Lucking-Reiley explains the higher prices in the Dutch auction with the fact that the number of bidders in the Dutch auctions was on average higher than in the first-price auctions. A possible explanation for the higher than expected revenues in the English auctions is affiliation of bidders' valuations.

The results of the experiments indicate clearly that the revenue equivalence between auction mechanisms does not hold outside the ivory tower. Explanations for the experimental results have been sought in the restrictive and unrealistic nature of the IPV assumptions. The experimental tests of the efficiency of different auction mechanisms give further reason to doubt the realistic nature of the assumptions behind the IPV model.

Common Value Model

Most experiments in the context of the Common Value Model focus on the existence of the winner's curse. Kagel (1995) reviews numerous studies on the winner's curse. Experiments show that especially inexperienced bidders suffer from the curse (Bazerman and Samuelson, 1983, Kagel and Levin, 1986). With enough experience, though, bidders learn to bid below their estimates and are able to obtain profits from auctions (Garvin and Kagel, 1994). The experiments conducted by Levin, Kagel and Richard (1996) show that the English auction increases the expected revenue for the bid taker, as predicted by theory. However, if bidders make a mistake and do not take the winner's curse into consideration in sealed-bid auctions, they might bid higher than in the English auction (and end with a negative profit).

2.7 Criticism of the Traditional Single-Unit Models

Rothkopf and Harstad (1994) criticize the single-unit models presented in literature. Most of these models apply a game theoretic approach to auction design. Rothkopf and Harstad claim that this approach is too simplistic to accurately reflect real world auction situations. For instance, the models assume that the auction occurs in isolation of all previous and future auctions. This is not true in reality. Auction participants cannot optimize their behavior with respect to one single auction; they have to think about their reputation that affects the outcomes of future auctions. Also, the assumptions made about the bidders are too restrictive and do not reflect reality. The bidders are not symmetric, and they may have different risk attitudes. Bapna, Goes and Gupta (2000) studied online auctions and identified three different bidder types, each having a different risk attitude and bidding strategy. The results of numerous experiments reviewed in section 2.6 have demonstrated that the traditional single unit models do not reflect reality. The intricate optimal auction mechanisms presented in section 2.5 seem to be too complex to be implemented in practice. Rothkopf and Harstad's critique lies heavily on their observation that no bidder seems to be using a game-theoretic model to decide how much to bid. Rothkopf and Harstad's critique indicated a new direction for auction research, and since the 1990s the focus of auction research has been shifting from purely theoretical treatments of auction models toward more practice oriented studies.

2.8 Multiple-Unit Auctions

The first attempt to make auction models more realistic was the inclusion of multiple units of the same item in the auction. These units could be auctioned either sequentially (one at a time) or all at the same time. Using a sequential auction is also the simplest way to model interdependencies among auctions, and the effect of reputation. Because of the multiple units, the auctioneer has to decide whether the pricing scheme is discriminatory or competitive (non-discriminatory). Discriminatory pricing means that each winning bidder pays the price equal to her bid (i.e. all winners end up paying a different price for the units). In competitive pricing, all the winners pay the price equal to the lowest winning bid – or the price equal to the highest losing bid.

In the simplest version of the multiple-unit auction, each bidder only has use for one unit. This reduces the auction to a price-only situation, which is only slightly different from the single-unit case. All the four basic mechanisms extend to this setting easily. In the more complicated case, the bidders can bid for any number of units for sale. Here the bids $b_i = (p_i, q_i)$ are vectors with two components, one indicating the per-unit price (p_i) and the other indicating the desired quantity (q_i). The first-price, second-price and English auctions have their extensions here too, but the Dutch auction is not suitable for the general multiple-unit auction.

The general multiple-unit extension of the first-price auction is the pay-your-bid auction, where bidders provide a "demand schedule", that is, a price for each unit they are interested in. Each bidder receives the number of items she demands for the clearing price, and pays according to her bids. The general multiple-unit extension of the Vickrey auction is quite complicated. Bidders submit demand schedules, and the winners are the ones with the highest bids. However, the payment made by the bidder on the j^{th} unit she wins is equal to the j^{th} highest rejected bid of her opponents. A simple example adopted from Maasland and Onderstal (2006) will explain it clearly. Assume there are three bidders and three units for sale. The bids are presented in the table below:

Table 2 Example of a multiple-unit Vickrey auction: bid prices for the bidders

	Bidder 1	**Bidder 2**	**Bidder 3**
1st unit	10	8	6
2nd unit	9	4	3
3rd unit	7	3	3

Bidder 1 wins two units and Bidder 2 one unit. The highest losing bid from Bidder 1's competitors is p = 6 from Bidder 3, so Bidder 1 pays 6 for the first unit. The second highest losing bid from her competitors is p = 4 from Bidder 2, so the total payments from Bidder 1 are 10. Bidder 2's payment is 7, which is the highest losing bid made by her competitors.

It is sometimes mistakenly thought that the uniform-price auction is the multiple-unit extension of the Vickrey auction (Maasland and Onderstal, 2006). However, this is not the case, and actually the uniform auction is not even efficient, as the Vickrey auctions are reputed to be. The uniform-price auction is similar to the pay-your-bid auction, except that each bidder pays the same price for each unit, and the price is equal to the highest losing bid. Because it is assumed that bidders bid for more than one unit, it is possible that the highest losing bid is made by one of the winning bidders. Thus, there is the possibility that a bidder can affect the price she has to pay, and bidding your valuation no longer is a dominant strategy.

The multiple-unit extension of the English auction is the Ausubel auction (Ausubel, 2004). In the Ausubel auction the price starts from zero, and increases continuously. At each price level, the bidders announce the quantity they would be willing to purchase. Then, each bidder's demand in turn is taken out of the total demand. If the demand from all the other bidders exceeds supply, nothing happens, and the price is increased to the next level. However, if there exists a bidder, without whom the total demand would be smaller than the supply, this bidder is awarded the "shortage" for the current price. The awarded units are removed from the supply, and price is increased to the next level, and the auction continues until all the units have been sold.

Multiple-unit auctions have received a lot of attention (see the above-mentioned references and e.g. Vickrey, 1961, Wilson, 1979, Engelbrecht-Wiggans, 1988, Maskin and Riley, 1989 and Tenorio, 1999), but the research in the area has followed the tradition of the single-unit models. The auction models are usually based on traditional IPV assumptions and the goal is to obtain tractability and equilibrium strategies. However, despite the restrictive nature of the studies, the multiple-unit auction models with bid vectors, albeit consisting of only two components, paved the road for the study of multiple-item (see Chapter 3) and multi-attribute models (see Chapter 4).

3 MULTIPLE-ITEM AUCTIONS – AN INTERDISCIPLINARY FIELD

Where the traditional auction research is dominated by economists, the study of multiple-item auctions has been of interest to computer scientists, operations researchers and decision analysts as well as economists and game theorists. Most attention has been given to combinatorial auctions, which are multiple-item auctions in which package bidding is allowed. Multiple-item auctions, especially combinatorial auctions, are so complex to organize that it is almost impossible to do it without the aid of computers. Thus the research in combinatorial auctions has advanced along with the development of computers. The article by Rassenti, Smith and Bulfin (1982) is reputed to be among the first in the field of combinatorial auctions. Since then computers have developed greatly, and the invention of the Internet has lowered the threshold to organize all kinds of electronic auctions. Since combinatorial auctions are very complex they also provide a fruitful ground for a lot of different kinds of research. Thus, in the past decade or so combinatorial auctions have become a hot topic in auction research. Researchers with different backgrounds have all found a perspective on combinatorial auctions they can contribute to. Economists and game theorists construct mechanisms with theoretically desirable properties, such as efficiency. Operations researchers study the integer programming aspects of winner determination and feedback mechanisms, and computer scientists create algorithms for the winner determination problem. Decision analysts have been interested in the creation of bids, and developing tools to help bidders evaluate their preferences over bundles. In what follows I will present research from different fields, and show how they complement each other.

In this thesis I will not consider single-unit and multi-unit combinatorial auctions separately, because their properties are essentially the same. Our research assumes multiple units, but the single-unit case can be dealt with as a special case. Most research considers the single-unit case, and hence part of the following discussion is from the viewpoint of single-unit multiple-item auctions.

At first sight, the step from single-item, multiple-unit auctions to multiple-item auctions does not seem that big. However, the difference in complexity can be enormous. The complexity arises from the underlying assumption that the reason for holding a multiple-item auction is that bidders have nonlinear preferences over bundles of items. In other words, the items are either substitutes or complements, so that the value from a bundle is not the sum of the values of its components. For complements it holds that $v(A, B) \geq v(A) + v(B)$, and for substitutes $v(A, B) \leq v(A) + v(B)$. Examples of such preferences abound. In radio license auctions there may be synergies in obtaining licenses for adjacent areas, or two licenses of different frequency for the same area can be substitutes (Pekeč and Rothkopf, 2003). Also many reverse auctions exhibit nonlinear preferences. E.g. it is easy to imagine that transportation services (trucking) have complementarities: the cost per haul decreases, if the truck is full on all routes. In the reverse setting, complementarities between items translate to a *subadditive* cost function $c(A, B) \leq c(A) + c(B)$. Nonlinear preferences, such as the ones described above, make bidding a complex task, as the value of one item to the bidder is dependent on what other items she wins. Also, the task of determining the winners of the auction – which so far has been relatively straightforward – can become difficult.

The auction mechanism the auctioneer chooses to use has a major effect on the complexity of bidding and winner determination. On a macro level, the choice is essentially between allowing bids on combinations or not. Not allowing bids on combinations makes the auction a lot easier for the auctioneer to handle, and there are no problems with winner determination. The downside is that bidding is very difficult and the outcome can easily be inefficient. Allowing combinatorial bidding enables the bidders to better express their preferences, but winner determination becomes cumbersome. Bidding is still difficult, but for different reasons. In the following I will briefly discuss some non-combinatorial multiple-item mechanisms before going into a detailed discussion on combinatorial auction research.

3.1 Non-Combinatorial Auction Mechanisms

If bids can be made on single items only, the multiple-item auction is essentially a collection of several single-item auctions. These auctions can then be held sequentially or in parallel.

3.1.1 Sequential Auctions

In a sequential auction, each item in the bundle is auctioned one at a time, one after another. This design is simple for the auction owner, since it is easy to define the winner in each auction. The winner is simply the bidder with the highest bid. For the bidders, however, the sequential design imposes grave difficulties. A bidder's valuation of each item depends on what other items she wins in the ensuing auctions. Therefore, in order to establish her optimal strategy in one auction, she will have to try to guess the outcomes of the future auctions. This involves speculation of the possible strategies of the competitors, which in turn depend on the outcome of the auction at hand. The computational costs are high, and in auctions with relatively large numbers of items and bidders, the calculation of the optimal strategy becomes intractable. Hence, the outcomes of the auctions are easily inefficient: the bidders do not obtain combinations they wanted, or pay more than they would have wanted for the combinations they do get. This problem is referred to as the *exposure problem* (Rothkopf, Pekeč and Harstad, 1998, Pekeč and Rothkopf, 2003).

3.1.2 Simultaneous Auctions

An alternative to a sequential auction is a simultaneous auction. In this auction, the items are auctioned in separate auctions that run at the same time. Here I will first discuss the general properties of simultaneous auctions, and some improvement suggestions to the design. After that I will present the Federal Communications Commission's (FCC) radio spectrum license auctions as an example of simultaneous auctions. The FCC auctions have received a lot of attention in literature due to their large size, and the fact that there are clearly complementarities and substitutabilities between the items.

3.1.2.1 General Properties of Simultaneous Auctions

The exposure problem still plagues the bidders in simultaneous auctions, but not as badly as in sequential auctions. The determination of the winner in each auction is still as simple as in sequential auctions, but the bidders' task has become a little easier. The bidders can observe each other's bidding behavior in all the auctions[5], which reduces the need for speculation. Ledyard, Porter and Rangel (1997) compared simultaneous and sequential auctions, and concluded that simultaneous auctions are more efficient. However, some problems remain. The bidders still do not know which items they will receive when all the auctions are closed. Hence, they cannot determine their valuations for the items *a priori* and it is impossible to establish the optimal bidding strategy. Moreover, in simultaneous auctions each bidder would like to wait until the end to see what the going prices for the items will be, and then optimize her own bids taking the final prices into consideration. Because all bidders would prefer waiting, no bidding would begin in the first place. So called activity rules could be established to guarantee bidding (McAfee and McMillan, 1996). This means that each bidder must bid at least a certain volume by predetermined points in time, or her future bidding rights are reduced. The activity rules are sometimes referred to as Milgrom – Wilson activity rules (see Milgrom, 1998) after their developers.

Sandholm (2000) proposes some methods to improve the efficiency of sequential and simultaneous auctions. One approach would be to establish an after market where the bidders can exchange items once the auction has closed. This reduces the inefficiency of the auction outcome, but may require an impractically large number of exchanges. Another, more practical approach would be to allow bidders retract their bids. In this case it is important to guarantee that retractions do not diminish the auctioneer's payoff. There are many ways to take care of that. For example, if the closing price is less than the retracted bid, the bidder who retracted her bid has to pay the difference (McAfee and McMillan, 1996). Sandholm (2000) suggests a *leveled commitment protocol* be used. In this protocol, the penalty from retracting a bid is set up front, and

[5] This is true only if the auctions are held in the open-cry format. Simultaneous sealed-bid auctions are equivalent to sequential auctions, because no additional information can be obtained.

also the auctioneer is allowed to decommit from the auction outcome. This reduces the risk for the bidders, as the penalty from retracting is known in advance.

3.1.2.2 The FCC Spectrum License Auctions

The multiple-item auction that has received the most attention in the past decade is the Federal Communications Commission's (FCC) spectrum license auctions held since the mid-1990s. McMillan (1994) provides a thorough description of the auction process. The items for sale were regional radio spectrum licenses covering the wavelengths used for personal communications services (PSC), such as cellular phones and wireless computer networks. What made the auctions so unique were their size and complexity. Thousands of licenses were for sale, bidders were many and diverse, ranging from large national telecommunications companies to small local firms, and the estimated value of the 1994 auction was over $10 billion. Dozens of economists were hired by the telecommunications companies and the FCC to help design the auctions.

The complexity of the auctions was increased by the realization that there were potentially complementarities between the licenses. The potential efficiencies derived from the aggregation of licenses have both engineering and economics aspects. First, the fixed costs of technology acquisition and building up a customer base can be spread over several licenses. Second, there are often problems of interference at the boundaries of license areas so it is cost-efficient to operate in adjacent areas. Third, consumers will value the ability to use the same phone when traveling all over the country. The main question in the auction design became then how to best take the complementarities into consideration without compromising the functionality of the mechanism. Also, because the seller was the government, revenue maximization was not the primary goal, but rather the efficient allocation of licenses.

The FCC decided to run simultaneous multiple-round auctions (SMA or SMR) developed by Paul Milgrom, Paul Wilson and Preston McAfee (as documented by Milgrom, 1998 and 2000). Because the bidders were informed of the competing bids after each round, this format was close to an open-cry auction. The bidders were also allowed to withdraw bids, but if the equilibrium price ended up being lower than the retracted bid, the bidder was obliged to pay the difference. The parallel auctioning of the licenses combined with the option to retract bids gave flexibility for bidders to aggregate licenses. Moreover, this way the bidders were able to switch to their "back-up" combination, if their most preferred one turned out to be too expensive. The FCC contemplated allowing combinatorial bids, i.e. bids for a combination of licenses. This would have allowed the bidders to express their synergies over aggregation of licenses directly in monetary terms. Theoretically, this could have produced more efficient results than a parallel auction. However, the FCC was afraid that administrative or computer breakdowns would occur due to the computational complexity imbedded in combinatorial bidding. Allowing combinatorial bids could make the auction too complicated for the bidders causing the complexity costs to outweigh the potential efficiency gains. Also, the threshold problem creates incentives to free ride (see section 3.2.1.3 for a full explanation), which was seen as relevant problem impeding efficient outcomes. Thus, the FCC decided against combinatorial bidding.

Very recently, the FCC experimented with allowing package bidding in one of the spectrum license auctions (Auction #73 of the 700 MHz band). Combinatorial bidding was allowed in one license block containing 12 licenses (FCC, 2007). However, the allowed combinations were restricted to three packages: "50 states" (licenses 1-8), "Atlantic" (licenses 10 and 12) and "Pacific" (licenses 9 and 11). At the end of the auction, only the bid on the "Pacific" package was among the winners; all other licenses were sold individually (FCC, 2008).

3.2 Combinatorial Auctions

Combinatorial auctions are defined as auctions in which multiple but different items are sold, and bidders are allowed to make indivisible bids on packages (Pekeč and Rothkopf, 2003). Bids are vectors ($q_1, \ldots, q_K, p$), where the first K elements indicate the quantities for the items, and the last element indicates the price for the whole package. In single-unit combinatorial auctions the q_i's simply indicate whether a particular item is in the package or not. Indivisibility refers to the restriction that all bids have to be accepted as a whole or not at all; no partial bids can be accepted. Already Rassenti et al. (1982) acknowledged that allowing combinatorial bids alleviated many of the problems in sequential and simultaneous auctions. For instance, there is no need to estimate the opponents' strategies in other auctions (possibly later in time) when all items are sold in one auction. Researchers also argue that combinatorial bidding allows the bidders to better express their preferences over bundles, and therefore the auction outcome should be more efficient than in non-combinatorial auctions. Ledyard, Porter and Rangel (1997) ran a series of laboratory experiments to compare combinatorial auctions to sequential and simultaneous auctions. They concluded that combinatorial auctions are more efficient (they produce outcomes closer to the efficient allocation) than sequential or simultaneous auctions. Banks et al. (2003) compared the simultaneous multiple-round auction (SMA) used by the FCC to a combinatorial auction, and reached similar results.

3.2.1 Challenges with Combinatorial Auctions

Despite all the theoretical benefits accruing from combinatorial bidding, combinatorial auctions have not been used that much in practice. The reasons for this arise from three properties that distinguish combinatorial auctions from other auction types: complexity of winner determination, complexity of bid formulation, and the strategic gaming element, which leads to what is known in the literature as the *threshold problem* (Pekeč and Rothkopf, 2003).

3.2.1.1 Complexity of Winner Determination

The winners of a combinatorial auction are the bids that maximize the bid taker's revenue (or minimize the cost) and allocate each item to only one bidder. If there are multiple units of each item, the number of units allocated cannot exceed the number of units available. All bids are assumed indivisible, also called all-or-nothing bids. Thus, the solution to the auction is found from a set of disjoint bids, which maximizes the seller's revenue.

The winner determination problem (WDP) can be formulated as an integer programming (IP) problem. Because most applications of combinatorial auctions are in procurement situations, I will now present the WDP in the reverse auction setting.

Consider K items and assume that d_k ($k = 1, \dots, K$) units (nonnegative integers) of each of K items are requested by the buyer, defining demand. Now each bid j by bidder i is a (K+1)-dimensional vector: (q_{ij1}, q_{ij2}, … , q_{ijK}, p_{ij}), where $0 \le q_{ijk} \le d_k$ are nonnegative integers (quantities of item k) and p_{ij} (price of the bundle) is also a real positive number. In other words, bidder i's j^{th} bid is an offer to deliver q_{ijk} units of each item k for a total price of p_{ij}.

The WDP determining the status of each bid by each bidder at any given moment in the auction is formulated as an IP problem as follows:

$$\begin{aligned}
\min \quad & \sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij}\, p_{ij} \\
s.t. \quad & \sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij}\, q_{ijk} \ge d_k \qquad \forall\ k = 1,\dots,K \\
& x_{ij} \in \{0,1\}
\end{aligned} \tag{2}$$

The variable x_{ij} indicates whether bidder i's j^{th} bid is among the winners ($x_{ij} = 1$) or not ($x_{ij} = 0$), n_i is the number of bids placed by bidder i, where $i = 1, \dots, N$. In case there is only one unit of each item, $d_k = 1$ for all k, and $q_{ijk} \in \{0,1\}$ depending on whether item k is in the bid x_{ij} or not. This formulation allows any number of bids per bidder to be among the winners. In some auctions the auction owner may want to limit the number of winning bids per bidder to one. In that case, a constraint $\sum_{j=1}^{n_i} x_{ij} \le 1,\ \forall\ i$ should be added.

The problem of finding the optimal outcome in a reverse combinatorial auction is equivalent to a set-covering problem (SCP). The SCP is a close relative of the set-packing problem (SPP) of the forward auction (de Vries and Vohra, 2003). The set-packing problem is known to be NP-complete (Rothkopf, Pekeč and Harstad, 1998), which means that as such there is no revenue maximizing algorithm that can solve the problem in polynomial time[6]. The properties of the SCP and SPP are a little different, e.g. the SCP can be a little easier to approximate, but the most important results

[6] In the analysis of computational complexity, time is measured in the number of computations required for solving a problem.

regarding the computational manageability of the problems are very similar (de Vries and Vohra, 2003).

It is customary to study the "worst-case scenarios" for the time needed by any algorithm to compute a solution. Mathematicians have focused on how the computation time grows as a function of the input. A class of computational problems is labeled "computationally manageable", if an upper bound on computation time for all problems can be expressed as a polynomial function of the inputs. The notation $f(n) = O(g(n))$ means that for a function $f(n)$, which is the number of calculations required to solve a problem whose input size is n, there exists a limiting function $g(n)$ so that $f(n) \leq cg(n)$, where c is some constant, when n grows large (see Papadimitriou and Steiglitz, 1982, p.159). For example, in non-combinatorial auctions, where the winner determination can be solved by picking the highest bidders for each item separately, the auction can be solved in $O(NK)$ time, N represents the number of bidders and K the number of items. So clearly, the non-combinatorial auction can be solved in polynomial time no matter how large the input. The number of possible combinations in a combinatorial auction is 2^K-1 and the winner determination is solved in $O(K^K)$ time, so there is no polynomial function that would express the number of computations required to solve the problem. This also means that the exhaustive enumeration of all possible outcomes is not a viable method for searching the optimal allocation, unless the number of items is very small. The general WDP can thus be declared computationally unmanageable.

3.2.1.2 Complexity of Bid Formulation and Communication of Preferences

Combinatorial bidding allows bidders to express their preferences over the different items in the auction. However, this can be a complex task, because there are 2^K-1 combinations (or more, if there are multiple units) over which the bidder should be able to express her preferences. As K grows, it is impractical, not to mention time consuming, to evaluate all conceivable combinations. Somehow bidders should be able to identify, which combinations are interesting to them, and concentrate on evaluating those combinations. Hoffman, Menon and van den Heever (2004) and Jones and Koehler (2002) argue that bidders do not even think in terms of the items they want in the bundle, but have other objectives which can be achieved through different combinations. For instance, in the FCC radio license auctions bidders desired a certain level of population coverage and bandwidth (Hoffman et al., 2004), and not necessarily a particular license. Similarly, in auctions for airtime for TV advertisements the bidders are essentially interested in acquiring a large exposure among a particular demographic group, and obtaining particular slots is simply a means to achieving the objective (Jones and Koehler, 2002). Thus, the bidders need to translate their objectives and constraints (e.g. budget) into bid combinations. They also need to estimate the value of each bundle to them in order to attach prices to the bids, which is not necessarily a trivial task. In fact, An, Elmaghraby and Keskinocak (2005) report that in a combinatorial auction for transportation services most bidders abstained from placing combinatorial bids. A plausible explanation for this phenomenon is that bidders found the construction of combinatorial bids too difficult.

When the bidders have managed to construct a set of combinations they would like to bid on, they need to communicate their preferences to the bid taker. Usually auction owners have defined a specific bidding language that has to be used to encode bids and preferences. A bidding language both defines the exact syntax to be used in submitting bids, and defines what kind of interdependencies can be expressed between bundles. For instance, some bidding languages allow logical operators like "and", "or" and "not", but some do not. Thus, it depends on the language, how well the bidders can express their preferences. If the language is not fully expressive, an *exposure problem* similar to the one in sequential and simultaneous auctions can still occur (Pekeč and Rothkopf, 2003). Bidders may want to place bids on many different combinations in hope of winning at least something, but they might not want to win all of them (e.g. they may not have enough capacity to produce everything). A bidder would then like to express her bid in the form "either combination A or B, but not both". More primitive bidding languages (OR, XOR) do not allow bidders to make complicated bids that would allow bidders to express different preferences over a set of bundles. OR bidding language does not allow bidders to restrict the number of bids that might become winners. Basically, in OR language, any number of the disjoint bids placed by the bidder could become winners. XOR language is the other extreme. There at most one of the bidder's bids can be among the winners, so it requires the bidder to communicate every single combination she might be interested in. More advanced languages, such as combinations of OR and XOR languages (OR-of-XORs and XOR-of-ORs), enable more expressive bidding, and OR* is a compact and expressive language. Nisan (2006) provides an extensive review on different bidding languages and their properties.

The catch is that usually, the more expressive the bidding language is, the more difficult it is to use, and the more complicated the computations usually get. Thus the solution of the WDP slows down even further (Nisan 2000). However, there are exceptions to this rule; e.g. XOR is easier to compute than OR. One way to circumvent the use of complicated bidding languages is to use dummy items in bids. A dummy item is an item that costs nothing, so adding it to two otherwise disjoint bids makes them overlapping without any added costs. The artificial overlap created by the dummy item ensures that the two bids cannot both be among the winners (Fujishima et al. 1999). Even if there were no dummy items available, researchers have observed that bidders used the cheapest items in the auction as dummies.

3.2.1.3 Strategic Gaming in Combinatorial Auctions: the Threshold Problem

Even if the bidder has managed to sort out her preferences, placing bids in the actual auction is not simple. This is because in combinatorial auctions there can be multiple winners. In order to become a winner with other bidders, the bidder's bid has to complement the other bidders' bids. A phenomenon called the *threshold problem* is identified in literature (see e.g. Pekeč and Rothkopf, 2003). The threshold problem refers to the situation when small, "local" bidders bidding on single items cannot beat alone a currently winning bid on the whole bundle made by a "global" bidder. This leads to gaming between the bidders, as they all try to maximize their profit but still be among the winners. Consider the following simple example of a reverse auction. There are four bidders (*a*, *b*, *c* and *d*) bidding on three items (x_1, x_2 and x_3), and the demand is one unit for each item. The bids submitted by the bidders are $b_a(x_1) = b_b(x_2) = b_c(x_3) = 5$ and $b_d(x_1, x_2, x_3) = 13.5$. Assume that the costs for each bidder for each individual item $c_i(x_j) = 4$. Bidder *d* is currently the winner, because her bid price 13.5 is lower than the combined price of 15 offered by the other bidders. Bidders *a*, *b* and *c* could each afford to lower their bid price by one unit making the total cost 12, which would allow them to become provisional winners. In fact, it would be enough that two of them lowered their bid. However, none of them alone could afford to lower the price so much that the combined total cost would go below the bid of bidder *d*. The bidders would have to somehow come to a mutual agreement to lower their bids in order to oust the current winner. In this example, there is also the problem of potential free riding by one of the bidders. Since it is enough that only two bidders lower their bids from 5 to 4, every bidder hopes to be the one who can become a winner without having to reduce the price. Thus, it is very likely that none of them reduce their price, and the auction outcome is not efficient.

The threshold problem extends to any situation in which a number of bidders together are trying to beat a bid on a larger combination. Chances are that more than one bidder needs to adjust their bid price, and that the bidders have the incentive to free ride. A partial solution to the threshold problem would be to accept bids that do not become winners at that point in time in the auction (Pekeč and Rothkopf, 2003). This could make the design of activity rules more difficult, because it will become harder to distinguish between serious bids and attempts to merely fulfill activity rules without having to compromise on profits (see section 5.2 on auction rules). The threshold problem becomes even more difficult to overcome, if the auction is a sealed-bid auction. In that case, the bidders do not even know the prices in the other bids, which they are trying to coordinate with.

3.2.2 How to alleviate the problems?

A big part of combinatorial auction literature concerns alleviating one or several of the above-mentioned problems. Computer scientists and operations researchers have tackled the computational issues, and decision analysts and operations researchers have studied ways to support bidders in valuing bundles and placing bids.

3.2.2.1 Alleviating Computational Problems

There have been two different approaches offered as solutions to computational problems. The first approach aims at developing fast algorithms, which would enable the solution of larger auctions quickly. Also some approximation algorithms have been suggested. The second approach to the WDP is to restrict the bid space (the combinations that can be bid on or the number of combinatorial bids) so that computational manageability can be assured.

Exact and Approximation Algorithms for WDP

An exact algorithm is an algorithm which guarantees an optimal solution. Several exact algorithms, which use a variety of techniques, have been developed. Some algorithms utilize integer programming, others prune the search tree, and some are based on dynamic programming.

Rothkopf, Pekeč and Harstad (1998) present a dynamic programming algorithm, which makes it possible to solve the winner determination problem in $O(3^K)$ time, where K is the number of items. The algorithm uses the observation that for each possible combination S of the items, the maximal revenue comes either from a single bid $b(S)$ or from the maximal revenues of two disjoint exhaustive subsets of S. The algorithm starts from singletons and proceeds systematically to larger sets until it reaches M, the combination containing all items. The advantage of the algorithm is that it calculates the revenue maximizing solutions for the subsets only once each time the winner determination problem is solved. The weakness of the dynamic algorithm is that it makes the same number of calculations independent of the number of actual bids. This is because it goes through every possible combination S even if there is no bid for it. The algorithm enables the auction owner to determine the calculation time already prior to the auction, but it cannot take advantage of the potentially small number of actual bids. Due to this fact, the dynamic algorithm functions in the worst case only if the number of items for sale is 20-30.

Branch-and-bound algorithm is a general algorithm that finds an optimal solution to any integer programming problem. However, since essentially the branch-and-bound algorithm is based on enumerating all feasible solutions (although organized as a search tree), it can become slow when the size of the auction grows. Fujishima et al. (1999) were among the first to consider an algorithm based on an intelligent pruning of the search tree. Since then, several studies have been written on more efficient search algorithms.

Gonen and Lehmann (2000) have developed a branch-and-bound type of algorithm to solve the integer programming problem (the WDP). Their algorithm is a depth-first search, which calculates the values for each branch in turn, all the while updating the current best solution. To speed up the search, Gonen and Lehmann suggest that the algorithm estimates an upper bound[7] for the objective function that can be achieved from each branch. The upper bound is then compared with the current best solution. If the value of the objective function in the best solution so far is higher than the value indicated by the upper bound, that particular branch can be "pruned", i.e. excluded from further consideration.

Also Sandholm (2000) utilizes the intelligent pruning of the search tree in his algorithm. The key observation of Sandholm (2000) is that in larger auctions the bid space is necessarily very sparsely populated. For example, if the number of items for sale is 100, it would take longer than the life of the universe to bid on all the 2^{100}-1 combinations, even if every person on earth placed a bid every second. Even in smaller auctions there hardly ever is a bid for every conceivable combination. Sandholm (2000) proposes an algorithm that takes advantage of this sparseness in the bid space. The algorithm generates a tree where each path consists of a sequence of bids organized based on the items in the bids. The path ends when all items have been used. Each path represents a feasible allocation, the revenue of which can then be compared with other allocations. The algorithm is implemented as depth-first search. This enables the auction owner to find feasible allocations quickly. Also, the algorithm keeps track of the best solution so far, so in case the algorithm has to be terminated before all the paths have been generated, the best solution so far can be obtained. The most significant difference between Sandholm's algorithm and the dynamic

[7] Gonen and Lehmann consider forward auctions, hence it is useful to define an upper bound for the value of the objective function in some subset of bids. Correspondingly, in reverse auctions it would be useful to obtain a lower bound.

programming algorithm is that Sandholm's algorithm generates only the paths for which there are actual bids. In the worst case, Sandholm's algorithm takes $O(m^K)$ time, where K is the number of items and m the number of bids ($m = \sum_{i=1}^{N} n_i$ using notation from (2)), to find the optimal allocation.

Sandholm and Suri (2003) and Sandholm et al. (2005) improve the algorithm proposed in Sandholm (2000 and 2002). Their major revelation is that it is more efficient to branch on bids rather than on items (as was done in the earlier algorithm). They develop a new branching method, BOB, and an algorithm, CABOB, to be used in a combinatorial auction.

Also other improvements to a basic branch-and-bound algorithm have been suggested. Ono, Nishiyama and Horiuchi (2003) have developed a method for iterative combinatorial auctions that utilizes the previous solution of the WDP to increase the speed of the search algorithm. Their method can be combined either with different search algorithms. Mito and Fujita (2004) suggest a way to order bids so that once a branch-and-bound algorithm is applied, it finds an optimal solution faster. Günlük et al. (2005) on develop a solution algorithm based on "branch-and-price." Instead of operating with all variables (bids), branch-and-prices starts with a small subset of variables, and through the dual of the WDP it formulates and solves a pricing problem, which helps identify good variables to include in the solution. Günlük et al. (2005) test four different branching rules, and conclude that branching on items is better than branching on bids. Yang et al. (2009) suggest that regardless of the branching rule used, the search process would become faster, if "nagging" were used. Nagging refers to the parallelization of the search space, where portions of the search tree are distributed to individual processors operating simultaneously. Thus, instead of working though the search tree one branch at a time, the algorithm would work on several branches simultaneously.

Andersson, Tenhunen and Ygge (2000) note that the winner determination problem can also be formulated as a mixed integer programming problem. According to them, this formulation can utilize standard algorithms and the problem can thus be solved using commercial software. Andersson et al. (2000) test a software package called

CPLEX and they conclude that in most instances it performs very well and the computation times are smaller than achieved with Sandholm's (2000) intelligent algorithm, and comparable to those obtained with Fujishima et al.'s algorithm. According to Sandholm et al. (2005) their most recent algorithm, CABOB, is often drastically faster than CPLEX, and rarely drastically slower.

There have also been efforts to find heuristics (Mito and Fujita, 2004, Jones and Koehler, 2005, Guo et al., 2006, and Özer and Özturan, 2009) and approximation algorithms (see Crescenzi and Kann, 2006 for a review) that would be polynomial and produce a "reasonably good" result instead of an optimal one. However, for some instances even the approximation can be difficult, and a good solution may not be found (de Vries and Vohra, 2003). Another problem related to the use of approximation algorithms in auctions is that it can compromise the perceived fairness of the auction mechanism (Pekeč and Rothkopf, 2003). Even though the approximated optimum is close to the true optimum revenue-wise, it can be comprised of a totally different set of winning bids than the true optimum. Thus the bidders cannot be sure of the fairness of such auctions, and auctioneers are hesitant to implement them in practice.

Restricting Combinations

A different philosophy on the WDP is to not try to "bang one's head against the wall", but to constrain the bid space so that computational manageability is guaranteed. What this means is that the combinations bidders are allowed to bid on are decided prior to the auction. This way the auction owner can limit the number of combinations to a level which assures computational manageability. According to Rothkopf, Pekeč and Harstad (1998), there are three instances when structure of the bid space is such that it guarantees computational manageability: nested structures, cardinality-based structures and geometry-based structures. These structures are useful, because it is possible to identify situations in real life where these structures could be natural, and not limit the bidders from expressing their valuations.

Nested structures take advantage of the situation where there are disjoint groups of items with synergies within groups but none between groups. If this were the case, all bids could be restricted to contain items from only one group, the exception being a bid on all items in the auction. An example of this kind of a situation could be an auction in which the assets are on the East and West Coasts of the USA. Assume that there are no synergies to be obtained from mixing assets from both coasts. The auctioneer could then limit the permitted combinations to include assets from only one coast, the exception being a bid on all the grand combination, i.e. all assets in the auction. The optimal outcome of the whole auction is then the union of the optimal allocations of the subauctions of East and West Coast assets. According to Rothkopf et al., the optimal outcome of any such nested auction can be determined in $O(K^2)$ time, where K is the number of items in the auction.

On some occasions the auction owner might have an idea on how large combinations the bidders would like to submit bids on. In these situations it would be advantageous to use *cardinality-based structures*. For example, an apartment building is converted into condominiums which are then auctioned off. It is unlikely that there would be any synergies in buying two or three condominiums, but purchasing large enough a block to gain voting control might be in the interests of some real estate company. The auction owner could then restrict the allowable combinations to include singleton bids and large combinations. If a large combination is defined as $|S| > K/2$, i.e. it has to include more than half of all the items, then there can be only one such bid in the optimal outcome. It is easy to see that this kind of an auction is computationally manageable regardless of the number of items for sale. Cardinality-based structures can also be used when synergies can be obtained from small sets. For example, if bids are allowed only for pairs, the winner of the auction can always be found in polynomial time.

Geometry-based structures are applicable in situations where there are synergies between “neighboring” assets. The simplest of these structures is a line structure in which all assets can be ordered and placed on a line. Rothkopf et al. (1998) give an auction selling radio frequencies in the cities on the East Coast as a hypothetical example of such a situation. The cities could be numbered from north to south and ordered on a line. Allowing bids on only consecutive licenses ensures the computational manageability of the auction. The two-dimensional extension of the line structure to a plane of squares is computationally manageable only, if bidding is allowed on combinations including only rows or columns. An example of a situation in which the auctioneer could use a two-dimensional structure is a coin auction. Coin collectors are usually interested in either coins of different values (pennies, nickels, dimes) from the same year or the same coin (say nickel) from different years. If the coins to be auctioned are organized in a matrix according to their value and year, then permitting bids on single coins, rows and columns not only ensures the computational manageability, but also enables the bidders to express their valuations. The two-dimensional structure can be extended to *k* dimensions using the same logic.

It is appealing for the auction owner to predetermine allowable combinations to ensure the computational manageability of the WDP. After all, to be successful, the auction must be trustworthy. If the potential bidders cannot be sure that a winner can be found or that the winners are those who really made the best offers, they might choose not to enter the auction at all. However, there is a trade-off in restricting the bids. The efficiency of the auction is compromised, i.e. the auction owner's revenue and the bidders' utilities are not as large as they could be, if the bidders were allowed to bid according to their true valuations. This is due to the risk that the auction owner may lack the knowledge to be able to recognize all combinations that would be important for the bidders. Therefore, Park and Rothkopf (2005) suggest the bidders be able to determine the combinations that are to be bid on. The bidders are asked to submit a list of bids for all single items and also a prioritized list of combinatorial bids. The winner determination problem is solved iteratively. First only the bids on single items are considered. As this auction is always computationally manageable, it guarantees a lower limit for the solution. After the initial solution is obtained, the first combinatorial bids on each bidder's list will be considered in addition to the single bids, and the winner determination problem is solved. If a solution is obtained in reasonable time, the algorithm proceeds to the second bid on the lists, and so on. The auction ends when all bids on all lists have been considered, or if at some point the solution to the winner determination problem is not obtained in reasonable time.

Park and Rothkopf (2005) argue that their approach has two advantages. First, the auction will be regarded as fair by the bidders since they can choose the combinations they want to bid on. Second, most of the efficiency is likely to be captured as the bidders can freely express any synergies they might have and the bidders can be expected to bid for the most important combinations first. The problem with Park and Rothkopf's method is that it is basically a sealed-bid auction. Once the lists are submitted, the bidders cannot go back and improve their offers. The bidders face the same decision problem as in single item price-only situations: they have to weigh the extra profit against the probability of winning without knowing the decisions of other bidders. If the auction is a multiple-unit auction there may easily be inefficiencies due to a mismatch in the item quantities. Also, the bidder can never know when comprising the list of bids, how many of her bids will be considered before the auction ends. All this, I think, makes bidding in the auction perhaps overwhelming for inexperienced bidders who have not studied auction theory. It is thus possible, that the outcome of the auction is not as efficient as Park and Rothkopf assume it to be.

3.2.2.2 Preference Elicitation and Bidder Support

The problems combinatorial bidding poses for the bidders are receiving increased attention from researchers. Most research is focused on preference elicitation. Researchers have produced tools to help bidders translate bidder' goals and constraints into bids, and attach appropriate prices to the interesting bundles. These ideas are similar to those presented in the fields of decision making and decision aiding, and they have close links to the multi-attribute utility theory (MAUT). The bids in a combinatorial auction can be seen as expressing a multi-attribute utility function in which each item is an attribute (Sandholm and Boutlier, 2006). For some reason, the gaming elements in the auction are disregarded. For instance, the threshold problem is rarely addressed directly.

Because most auction mechanisms considered in the combinatorial setting are iterative, also most preference elicitation schemes are designed for iterative auctions. An iterative auction alone with feedback on prices and provisional allocations can be understood as a kind of bidder support. The iterative format can help the bidders because bidders are no longer required to supply bids on all 2^K-1 combinations. Instead, they can only supply a few bids every round, and adjust their strategies according to information feedback coming from the auction (Pekeč and Rothkopf, 2003). However, the iterative auctions as such do not offer any explicit support for preference elicitation, and they are often framed as auction mechanisms rather than support tools, hence I will discuss them later in section 3.2.3.2 on different iterative mechanisms.

Conen and Sandholm (2001) propose a selective preference elicitation approach. They argue that in an iterative auction the auctioneer does not have to elicit each bidder's preferences over all combinations, but ask for preference information only on relevant combinations. To see this, consider the following example adopted from Sandholm and Boutlier (2006). For instance, assume that bidder *i* has indicated she prefers bundle X over bundle Y, and the lowest cost she is offering for bundle X is 100 €. If the auctioneer has a better offer for bundle Y from someone else, there is no point for asking bidder *i* to express her valuation of bundle Y. This method reduces the number of packages that need to be valued, but offers no help in the actual valuation process. It also requires the bidders to submit preference information to the auctioneer.

In some cases firms may prefer not to reveal any preference information to outside parties as it could reveal the source of their competitive advantage. Therefore, Hoffman, Menon and van den Heever (2004) have developed a support tool for the bidders' private use. The tool is created specifically for the FCC license auctions, but it could be adjusted for other environments as well. The tool has an interface through which the bidders insert their preferences in the form of constraints (minimum population required, overall budget constraint, lowest level of profit acceptable, minimum bandwidth required etc.). Price information from the current round is then used to optimize the combinations the bidder should bid on. This second step relies on the design of the FCC auction, according to which minimum acceptable prices for each bandwidth are announced after each round.

A somewhat similar, but even more straightforward approach was suggested by Jones and Koehler (2002). They designed an auction in which bidders only submit rules they want their bids to follow (i.e. a set of constraints that must be fulfilled). The specific bids are then constructed by the auction mechanisms when it calculates the optimum allocation. The auction is iterative so the bidders can revise the restrictions they place on the bids.

Adomavicius and Gupta (2005) offer different kind of bidder support. They do not try to elicit bidders' preferences. Instead, they offer metrics by which bidders can assess the potential of their bids being among the winners. That way the bidders get an idea of which bids to improve on. However, the metrics do not help in constructing new bids, because they do not give indication on how to improve the bids.

Teich et al. (2001, 2006) propose a "price support" tool for bidders to use in a multi-attribute auction, but the tool is also directly applicable in a combinatorial auction. Based on linear programming (integer programming in combinatorial auctions), the "suggested price" tool calculates the maximum price for a given combination of items (or attributes) that brings the bid among the provisional winners. The quantities and items need to be predetermined by the bidder. Gallien and Wein (2005) present a similar system and the underlying theory for an optimization-based multi-item auction mechanism to minimize the buyer's cost under the suppliers' (known) capacity constraints. They assist suppliers in finding a winning bid price. However, the underlying assumption is that the suppliers are willing to disclose their cost functions to a supposedly neutral third party auction organizer.

3.2.3 Designing Combinatorial Auction Mechanisms

Mechanism design for combinatorial auctions is not concentrated on the equilibrium strategies or revenue comparisons between mechanisms, which were the focus of interest in single-item auctions. The combinatorial environment is so complex that equilibrium strategies are hard to analyze (Pekeč and Rothkopf, 2003). Some attempts to construct optimal mechanisms in the Myerson (1981) sense exist, but they are very limited (de Vries and Vohra, 2003). Instead, combinatorial auction design focuses on the usability of the auction mechanisms. This involves dealing with the potential problems identified in the previous section. Also the allocative efficiency of the mechanisms is considered desirable. The key dilemma in the design is the trade-off between efficiency and complexity.

The elements of mechanism design are the same as in simpler auctions: the way of communicating bids, determination of winners, and payment rule (see section 1.3). The first choice is between a one-shot (single-round) and an iterative auction. If an iterative auction is chosen, it must also be decided whether it is continuous (bids are allowed to enter at any given time, and WDP is solved after every bid), or round-based (WDP solved only after each predetermined round). In addition, the designer needs to decide what information is revealed to the bidders in between the rounds (or bids), that is, whether bids are open or sealed. The determination of winners is based on revenue maximization (or cost minimization): the winning bids are the ones that are the most favorable for the auction owner. There are several possible payment rules, e.g. uniform pricing, pay-your-bid (first-price), and Vickrey pricing.

The design parameter that affects the mechanism the most is the choice between one-shot and iterative auction mechanisms. Thus, in the following I will discuss these two instances separately. Mechanism design literature in combinatorial auctions is traditionally presented in the forward auction setting, so I will adhere to that. Also, it is customary to only consider the single-unit case, so that is what I will do as well, unless mentioned otherwise.

3.2.3.1 One-shot auctions

Multiple-item extensions of single-item, multiple-unit auctions (first-price sealed-bid auction and uniform-price sealed-bid auction) are not very suitable for a combinatorial setting, and difficult to implement. In a first-price, sealed-bid auction all combinatorial bids are submitted before an announced deadline, after which the WDP is solved once to determine the winners. The winners then pay the amount indicated by their bids. According to Pekeč and Rothkopf (2003) the benefits of these auctions are that they are resistant to collusion, and they are transparent as everyone pays the price they bid for. The main problems are potential computational unmanageability and the complexity of the bidding task due to strategic complexities and the large number of bids bidders may wish to construct. Some of the combinatorial auctions implemented in practice, however, have been one-shot, sealed-bid auctions (e.g. Epstein et al., 2002).

When selling multiple units of a single item, market-clearing (uniform) prices have a certain appeal. However, the idea of market-clearing prices is difficult to translate into a multiple-item setting (Pekeč and Rothkopf, 2003). The uniform, linear[8] market-clearing prices for the items may not exist, due to the fact that the bidders' valuations are superadditive (or subadditive). A simple example adopted from Wurman and Wellman (2000) illustrates this. Assume that there are two bidders in an auction and two items for sale. Bidder 1's valuations are superadditive: she values the individual items at 0, but has a value of 3 for the pair. Bidder 2's valuations are subadditive: she values both items at 2 individually, but gets no extra benefit from obtaining both. The efficient outcome would be to allocate both items to Bidder 1. However, there are no linear prices that support this allocation. The prices for the individual items should be at least 2 in order for Bidder 2 to not be upset over losing, but Bidder 1 would not be willing to pay 4 for the pair. Due to the superadditive valuations, the linear relaxation of the WDP may not have an integer solution, and shadow prices for the items do not exist. In that case, no linear prices exist that would separate winning bids from the losing ones. Also, in single-item auctions it is customary to use the highest losing bid as the uniform price, but in a combinatorial auction the concept of a highest losing bid is not well defined, because bids contain different items.

The Vickrey (second-price) auction has been generalized to combinatorial setting by Clarke (1971) and Groves (1973). The VCG mechanism is an efficient mechanism under fairly general conditions (Maasland and Onderstal, 2006), and it can be used in other frameworks than the combinatorial framework as well. The main restrictions are that utility must be additively separable in money, and bidders' valuations must be independent. In the VCG mechanism bidders announce their valuations over all bundles (= their type) and the mechanism calculates the optimum allocation and determines payments.

The payments are determined so that it is a weakly dominating strategy for bidders to announce their valuations truthfully. This leads to each bidder paying a different price. The idea behind Vickrey pricing is that bidder *i*'s payment is the difference in

[8] Having linear prices in combinatorial auctions means that the package price is the sum of the item prices in the package.

"welfare" of the other bidders without her, and with her in the auction. De Vries and Vohra (2003) present this formally. Let M denote the set of all items, and S any subset of M. Thus, $v^i(S)$ denotes the value that bidder i attaches to subset S. Additionally, let y(S, i) = 1 if subset S is allocated to bidder i. Assume that V is the aggregate value from the auction to the bidders in the optimum allocation y^*. Let V^{-i} and y^{-i} denote the maximum aggregate value and optimum allocation from an auction in which bidder i is not present. The payment to bidder i is determined by

$$p^i = V^{-i} - \left[V - \sum_{S \subset M} v^i(S) y^*(S,i) \right] \tag{3}$$

where $\sum_{S \subset M} v^i(S) y^*(S,i)$ is the value for bidder i from the winning allocation. Thus, the term in brackets describes the aggregate value of all the other bidders in the auction in which bidder i participates in. Another interpretation for the payment is then that it is the reduction in other bidders' welfare due to the fact that bidder i by entering the auction takes a piece of the cake. Notice, that if bidder i is not among the winners (i.e. y^*(S, i) = 0 for all S), $V = V^{-i}$ and her payment is zero. And, if bidder i is the only winner, her payment equals V^{-i}. Payment is always nonnegative, since V^{-i} (value from all items M) must be greater than the aggregate value of the subset $M \backslash S^i$ (S^i = subset allocated to bidder i) of items to the same set of bidders. A simple example in Pekeč and Rothkopf (2003) illustrates the VCG payments. Assume that there are two items, a and b, for sale, and two bidders. The first bidder is offering 10 for {a}, 5 for {b}, and 15 for {a, b}, and the second bidder 1, 6 and 12 respectively. The auctioneer's revenue is maximized when item a is sold to the first bidder, and item b to the second bidder (sum of bids = 16). Without the first bidder the total revenue of the auction would be 12, and her reported valuation for item a is 10, so according to Equation (3) her payment is 12-[16-10] = 6. Similarly the payment for the second bidder is 15-[16-6] = 5. The total revenue for the auctioneer is 11.

It can be proved that the VCG payments make it a (weakly) dominant strategy (see e.g. Ausubel and Milgrom, 2006 for a compact proof) to truthfully reveal one's preferences. This is the main benefit of the VCG mechanism aside from the fact that it is efficient.

It eliminates all gaming elements from the bidding process, so it presumably reduces the costs of participation.

Even though on a conceptual level the VCG mechanism is very appealing, it has several disadvantages which make it impractical to implement in practice (Isaac and James, 2000, Pekeč and Rothkopf, 2003, Ausubel and Milgrom, 2002, Ausubel and Milgrom, 2006, Maasland and Onderstal, 2006). First of all, the dominant strategy is far from obvious, especially to inexperienced bidders. In the laboratory experiment of Isaac and James (2000), only 13.6% of bidders bid their exact valuation, and 49.4% bid close to their valuation. And even if the bidders knew the dominant strategy, they might still be unwilling to reveal their valuations to the bid taker. They fear that the bid taker can use the information in later auctions, and harm the bidders. In larger auctions, the communication of valuations becomes complicated, as the VCG mechanism requires bidders to announce their valuation for every conceivable combination. A valuation should be stated even if the bidder is sure she cannot win a particular combination, because the payments of the auction depend on losing bids. Thus, the omission of one losing bid can potentially change the final payments.

Actually, the fact that the final payments are not based on the bidders' own bids – nor are they easily identifiable from other bidders' bids – creates a potential problem. This is because the bidders may not appreciate the lack of transparency in the pricing. In fact, the determination of the bidders' payments requires the solution of an IP problem for each winner (Porter at al., 2003). The bidders may not trust an auction in which they cannot verify the mechanism through which their payments were calculated.

Also, even though the Vickrey auction is efficient, it is not necessarily revenue maximizing. In fact, it can result in low revenues for the auctioneer. This is because the Vickrey prices are not necessarily in the core of the auction game. A core is the set of allocations and prices in which the auction owner cannot negotiate a better deal with the losers (Ausubel and Milgrom, 2002). Consider the following example of Ausubel and Milgrom (2002). There are two items for sale (A and B), and three bidders (B1, B2 and B3). B1 only wants the whole package and v_1(A,B) = \$2 billion. B2 and B3 bid for the single licenses, and v_2(A) = v_2(B) = v_3(A) = v_3(B) = \$2 billion. In the winning allocation, the items are then allocated to B2 and B3 (one item each). The VCG payment of both bidders is 2-(4-2) = 0. These prices are not in the core, because the auction owner would like to go to the losing bidder B1 and offer to sell the items for her for $2 billion. Bidder B1 would accept this offer, and the revenue for the seller would increase from 0 to $2 billion. An interesting twist to this example is to consider what happens if bidder B3 does not enter the auction at all. Now the winner is either B1 or B2 (they are both tied with a bid of $2 billion), and the price the winner has to pay is now 2-(2-2) = 2. Such sensitivity of revenue to the number of bidders and the kind of bids they place is clearly unacceptable. The examples above are of course extreme examples, but in fact the same phenomena are present whenever the items are not substitutes for even one of the bidders. *Substitutes preferences* means that the bidder's demand for one item does not decrease when the price of another item increases. In these examples the items were perfect complements for bidder B1 (a single item is of no value to her), which violates the *substitutes preferences* assumption – with drastic consequences. Recently, researchers have developed auction mechanisms which would always choose core solutions and prices (Day and Raghavan, 2007, and Day and Milgrom, 2008), but which would choose the VCG prices whenever they are in the core. This would solve the problem of low revenues of the VCG mechanism but still preserve the benefits of the VCG mechanism (allocative efficiency and incentive compatibility).

When the *substitutes preferences* assumption is violated, the bidders can try to take advantage of the loopholes making the VCG mechanism susceptible to collusion and shill bidding by bidders or cheating by the bid taker. Shill bidding refers to the incentive to use multiple identities or hire someone to pose as a new bidder, and then buy the items from her after the auction. For instance, assume that there are two items for sale, and two bidders have made bids of 1000 and 900 for the combination of the two items. A third bidder, who values the pair at 800, cannot place a competitive bid. However, if she hires a shill bidder, and they both bid anything above 500 for one of the goods, they will become winners together. Because there are no other bids for single items, their payments would be zero. Shill bidding is even easier in internet auctions, because a bidder can easily enter the auction with multiple identities (Yokoo, Sakurai and Matsubara, 2004). Bidder identification is usually based on information such as email address, and it is very simple and cheap to acquire multiple email addresses. The bid taker has an incentive to cheat as well. Once she observes all the bids in the auction she can increase her revenue by inserting false bids just below the winning bid prices, and it will be hard for the bidders to detect this.

Even if no cheating occurred, the outcome of the VCG auction might be politically unacceptable (e.g. if it is an auction collecting revenue for the government). The public may be outraged when they see that the bidders were willing to pay more (as indicated by their bid prices), but were in fact charged the VCG prices, which are less (McMillan, 1994).

Finally, the whole VCG mechanism is designed under the assumption of independent private values. Common value multiple-item auctions have not been studied theoretically, but studies of single-item auctions with common value elements (e.g. Klemperer, 1998) show that Vickrey auctions lead easily to very low revenues to the bid taker. It is quite reasonable to assume that in many auctions there is either a common value element to the items or at least one bidder who has complementarities between the items. Thus, it comes as no real surprise that VCG auctions are not common in practice.

Rassenti et al. (1982) designed a one-shot, sealed-bid combinatorial auction for allocating airport time slots. The difference to the Vickrey auction is that Rassenti et al. use uniform per-item pricing (although bidders announce only package prices). The uniform prices are more transparent and much simpler to compute than the Vickrey prices. Because market clearing prices may not exist, it is possible that a package with a price above the final prices is not accepted causing frustration among the bidders. Rassenti et al. solve this dilemma by defining two sets of prices: bid rejection prices, and bid acceptance prices to be announced to the bidders. The acceptance prices cannot sum up to more than the prices in the winning bids. In case market clearing prices exist, the two sets converge; otherwise, the bid acceptance prices are lower than the rejection prices. The abandonment of Vickrey prices means that the auction is no longer incentive compatible. Rassenti et al. argue, though, that strategic behavior is very risky for the bidders in a sealed-bid, one-shot auction. Nevertheless, it is not easy for bidders in the auction to determine what kind of bids to place.

3.2.3.2 Iterative Mechanisms[9]

Iterative mechanisms have distinct advantages over single-round auctions in the combinatorial setting. The most important advantages are that bidders do not have to bid for every possible combination in advance, and information can be obtained during the bidding process (de Vries and Vohra, 2003). The bidders' task is easier, because they can place bids when needed, and they can revise them based on feedback obtained during the auction. Also, if we assume affiliated values, iterative auctions are more efficient than single-round, sealed-bid auctions due to information revelation (Parkes, 2006). Most combinatorial mechanisms presented in literature are iterative. There are already many mechanisms, even though the research field is quite young. The mechanisms could be classified in many ways, but I will use the classification of Parkes (2006), who divides iterative mechanisms to price-based and non-price-based mechanisms. The essential difference between these two groups is that in the price-based mechanisms bidders are provided with information on what prices to bid for.

Non-Price-Based Mechanisms

Among the first iterative mechanisms is the Adaptive User Selection Mechanism (AUSM) introduced by Banks, Ledyard and Porter (1989). Bidding in AUSM is continuous, and bids on all combinations are allowed. The provisional winning bids at any given time are announced to all bidders. In AUSM, the computational burden is delegated to the bidders. Anyone willing to submit a new bid must suggest a combination of bids that complements her own bid, and demonstrate that they together provide more revenue than the current winning combination. To help identify good bids, a standby queue is maintained. Bidders can "advertise" their willingness to make certain bids, and bidders can use these bids as complements to their own bids. When a new bid comes in, the bid taker has to verify that the

[9] Some authors use term *iterative auctions* as synonymous with round-based auctions (e.g. Kwasnica et al., 2005). However, in this text iterative auctions refer to any kind of auction in which bidders have the opportunity to improve upon their old bids.

complementing bids are either in the standby queue or among provisional winners, and that the new combination in fact produces more revenue. The use of a standby queue partially alleviates the threshold problem (Kelly and Steinberg, 2000), but the incentive to try to "free ride" still remains. According to the results of simulations conducted by Ledyard et al. (1997), auctions using AUSM increased the efficiency of the final allocations compared to simultaneous and sequential, single-item auctions. However, AUSM only supports additive-OR bids (Parkes, 2006). This means that bidders cannot restrict the number of disjoint bids that can become winners. The only way to indicate e.g. substitutabilities among bids is to make them overlapping. A version of AUSM was implemented by Sears Logistics to procure trucking services (Ledyard et al., 2002).

In *proxy auctions*, automated agents bid on the behalf of the bidders (Ausubel and Milgrom, 2002). Prior to the auction, bidders express their preferences to the proxy agents, who then bid to maximize the bidders' profit. Proxy agents bid until there is no room for improvement. Winners pay the price bid by the proxy agents (i.e. the lowest price on a particular combination that allowed the bid to become a winner). Proxy auctions are efficient provided that bidders can (and are willing to) express their preferences to the proxy agent. However, communicating preferences to the proxy agent is every bit as complicated as communicating them to a bid taker in a VCG auction. Also, no learning can take place in the auction process. The only major improvement is that they are more resistant to collusion than VCG auctions, and the failure of the *substitutes preferences* assumption is not as devastating. In fact, from the bidders view point, a proxy auction is very similar to a single-round auction, but with the added option of revising the information given to the proxy agent.

Another problem with the proxy auction is that it can become computationally infeasible. Whereas in the VCG mechanism, the WDP is solved only once, the proxy auction advances progressively in increments, and the WDP is solved after every round. In order to end in an efficient allocation, the increment has to be small enough. This slows down the convergence of the auction by increasing the number of proxy bids that have to be placed. The number of rounds could be thousands even for an auction with six items and ten bidders (Hoffman et al., 2006). Several researchers have suggested methods to speed up the convergence of the proxy auctions (see Hoffman et al., 2006 for a review and comparison).

Another non-price-based mechanism is the direct mechanism suggested by Conen and Sandholm (2001), where bidders do not have to place bids. Instead, the auctioneer asks for preference information iteratively from the bidders, but only as little as needed to determine the optimal allocation.

A *clock-proxy auction* designed by Ausubel, Cramton and Milgrom (2006) is a hybrid auction combining a clock auction with a proxy round. In the first stage an ascending clock auction is organized. The auctioneer announces prices for items, and bidders report the quantities they demand for those prices. Auctioneer increases prices for goods with excess demand, and bidders report new quantities. The process continues until there is no excess demand. The prices established in the clock auction act as a lower bound on the prices for the proxy round. The clock auction does not allow bids on combinations, so potential synergies are not realized until in the proxy phase, and therefore the prices are expected to increase. A clock-proxy auction is actually a hybrid between price-based and non-price-based auctions. The proxy phase is not price-based, but bidders receive price information in between the clock auction and the proxy phase.

Another such hybrid auction with a price-based first stage is the Progressive Adaptive User Selection Environment (PAUSE) developed by Kelly and Steinberg (2000). PAUSE is a combination of a simultaneous ascending auction and an AUSM-like second phase. The simultaneous ascending auction provides information on market prices, which can then be used in the ensuing AUSM auction, where combinatorial bidding is allowed. The prices in the package bids in the AUSM phase must be at least as high as the sum of the prices determined in the simultaneous ascending auction.

Day and Raghavan (2008) on the other hand have designed a three-stage auction, the purpose of which is to combine the benefits of AUSM and clock-proxy auctions while avoiding some of their problems (mainly the incentives to free ride and the distortion caused by linear prices). In the first stage the bidders can submit bid tables in which they can express substitutabilities between items. In the second stage the can "probe" the auction, that is, ask for prices that would allow specific combinations to become provisional winners. This is very similar to the "suggested price" tool of Teich et al. (2001, 2006). The third stage is a proxy auction, after which the winners are determined. The prices the winners pay are can be either VCG payments or Pareto-efficient core prices, which ensure incentive compatibility.

Price-Based Mechanisms

As noted earlier, linear market-clearing prices may not exist for a combinatorial auction. However, different kinds of approximations are available. The approximated prices have been used in iterative auctions to guide bidders' bids towards an efficient allocation. These price approximations can be linear (per-item) approximations, in which the bundle prices are simply the sums of the per-item prices, or nonlinear approximations where there are separate approximations for each package. The linear prices can be used either as per-item "ask prices", which act as lower bounds on new bids the bidders create, or "clock prices" at which the bidders announce their most preferred combinations (= the ones that maximize their payoffs). The non-linear prices are used only as clock prices. Moreover, the ask prices and clock prices can be either anonymous, which means that the same prices are announced to all bidders, or non-anonymous (personalized). In the following I will briefly describe seven combinatorial mechanisms based on ask prices or clock prices.

Kwasnica et al. (2005) describe a Resource Allocation Design (RAD), which uses linear, anonymous ask prices. According to the authors, RAD combines in one auction elements from a simultaneous multi-round auction and AUSM. However, the only resemblance to these auctions is that RAD allows package bidding (as does AUSM), and it quotes per-item prices that bidders must beat (as does the simultaneous multi-round auction of the FCC). Thus, there is only one auction and not two separate stages as in PAUSE (which also is a combination of a simultaneous auction and AUSM). The RAD auction proceeds in rounds. After each round, a set of linear prices (one for each item) is calculated, and new bids must beat these price. Minimum prices for packages are obtained as sums of the minimum prices for the individual items in the package. The calculation of the prices are based on three principles: 1) in order to keep "pay-your bid"-feature, the ask prices should be such that the winning bidders end up paying what they bid for, 2) Prices should be higher than what the losing bidders bid for, 3) whenever possible the prices should be such that if the losing bidders bid according to them, they would become winners. When principle 3) holds, ask prices convey information about opportunities in the auction for the next round, which is desirable. Kwasnica et al. formulate an LP problem to solve for the prices. The objective of the problem is to minimize the deviation from the two latter principles (the first principle must always hold). The prices can be called "pseudo dual prices", as they are prices that minimize the deviation from the dual prices of the linearized WDP. One problem with these pseudo dual prices is that they can oscillate a lot from round to round (Dunford et al., 2004). This means they can decrease as well, which is counterintuitive. Thus, Dunford et al. (2004) introduce a smoothed anchoring method to solve for pseudo dual prices that would deviate as little as possible from the prices quoted in the previous round.

The Combinatorial (CC) auction of Porter et al. (2003) uses linear prices like RAD. However, in the CC auction the prices are presented as clock prices. In each round there is a set of prices at which the bidders are asked to announce their demand for each item. If there is excess demand for some items, the auction continues to the next round. Prices for items with excess demand are increased before bidders are asked to announce their demand. The good aspect about the CC auction is that it is directly extendable to the multiple-unit setting. However, Porter et al. give a very vague description on how the final prices are determined in the case when there are no linear prices to support the winning allocation. Also, Porter et al. (2003) do not tell how the price increases are determined, so it is impossible to compare the CC auction and RAD.

Other price-based mechanisms use nonlinear prices, where prices are determined for each combination and each item separately. Thus the price of a combination does not have to be the sum of the prices of its components. The benefit of nonlinear prices is that they do not compromise the efficiency of the final allocation like linear prices do. However, they are more tedious to solve. Four such mechanisms are A*k*BA auction by Wurman and Wellman (2000), *i*Bundle by Parkes (1999) and Parkes and Ungar

(2000), dVSV by deVries, Schummer and Vohra (2007), and the Vickrey-Dutch Auction (VDA) by Mishra and Veeramani (2007).

In the A*k*BA auction Wurman and Wellman formulate an assignment subproblem that solves for the minimal (anonymous) prices that support the solution of the WDP. Supporting prices are defined as prices which, if announced as posted prices for the combinations, would not cause any changes in the bidders' behavior. The winners from the WDP would be willing to purchase the combinations at the posted prices, and the losers would not. Because there can be a range of such prices, the assignment subproblem is designed to choose the smallest one to maximize the bidders' payoff. The prices are treated as ask prices.

In both *i*Bundle (Parkes, 1999, and Parkes and Ungar, 2000) and dVSV (deVries et al., 2007) the prices are treated as clock prices, and the bidders are asked to announce their "demand set", i.e. the combination(s) which maximize her profit at the current prices. In addition to being nonlinear, the clock prices in both auction mechanisms are non-anonymous; that is, each bidder can be announced a different price on the same combination (although Parkes and Ungar also propose a version of *i*Bundle that uses anonymous prices). After each round, the prices are only increased for those bidders and those combinations, which are in the bidder's demand set, but are not part of the provisional allocation. In both, *i*Bundle and dVSV, the prices to be increased and the size of the increase are determined based on an LP solution algorithm. The difference is that *i*Bundle uses the subgradient algorithm, whereas dVSV uses the primal-dual algorithm. In both auction mechanisms the winning bidders pay the price indicated in the bids. However, Mishra and Parkes (2007) develop extensions of both auctions (extended dVSV and *i*BEA), in which the winners pay a price lower than in their bid. The benefit of this minor change is that the new mechanisms lead to efficient allocations also when the *bidders are substitutes* condition does not hold.

The Vickrey-Dutch Auction (Mishra and Veeramani, 2007) also uses nonlinear and non-anonymous clock prices. However, they follow in the Dutch auction logic of decreasing prices in forward auctions, and increasing prices in reverse auctions[10]. In the other mechanisms above, prices start low, and increased when there is excess demand. In the VDA, prices start high, and they are dropped by ε for each bidder for each combination, which is not in her demand set. The buyer's demand set contains the combinations which maximizes her payoff at current prices. The auction ends, when all the combinations are in the buyers' demand sets. In essence, the auction reveals the buyers' valuations for each combination. The prices in the final price vector are then adjusted so that they correspond to Vickrey prices. This can be done, because the auction owner has complete information on the bidders' valuations. And because the payments are Vickrey prices, the bidders should not have incentives to misrepresent their valuations during the auction. Mishra and Veeramani (2007) admit that the VDA mechanism is not scalable to large auctions, because the number of combinations – and thereby the size of the price vectors – increases exponentially with the number of items. However, they do not discuss the fact that bidder may be hesitant to participate in an auction in which their valuation for all possible combinations – even the ones they do not win – is revealed to the auction owner. Also, the auction process, in which announcing your demand for any price facing you gives no indication of whether you will win or not, may cause frustration among the bidders.

A potential problem with nonlinear clock auctions such as *i*Bundle, dVSV, and VDA is that clock prices in each round are given only to predetermined bundles. Thus, unless the auction owner wants to quote prices for each conceivable combination (and have bidders evaluate all the combinations), the auction can be inefficient. Both mechanisms lead the auction to an efficient outcome, but only with respect to the combinations included in the auction. The only way to ensure true efficiency is to quote prices for each possible combination, which is infeasible in large auctions. However, even if the auction owner quoted prices for each combination (which is not possible in multiple-unit combinatorial auctions), the good news for the bidders is that

[10] Mishra and Veeeramani (2007) present the VDA mechanism in a reverse setting. Since all the other mechanisms described in this section have been presented in the forward setting, I will transform the mechanism to the forward setting. This should make comparisons to other mechanisms easier.

in each round they only have to evaluate the combinations for which the price changed.

In order to remedy the potential problem, Kwon et al. (2005) propose a mechanism, the Endogenous Bidding Mechanism, which combines aspects of linear pricing mechanisms (such as RAD) with *i*Bundle. Their mechanism provides (nonlinear) prices for combinations, just as in *i*Bundle. In addition to that, they offer a vector of single-item prices. Bidders can use these single-item prices when constructing new combinations after the first round. The ask price for any new combination can be derived from the single-item prices linearly; for the old combinations, the nonlinear ask prices apply. The mechanism of Kwon et al. (2005) improves the efficiency of *i*Bundle, because it removes the problem of predetermined bundles discussed in the previous paragraph. However, the authors are not clear about the performance of their mechanism relative to RAD.

Table 3 summarizes the key characteristics of the iterative combinatorial auction mechanisms presented above. I chose to include the one-shot mechanism of Rassenti et al. (1982) as well, because the way they have calculated the winning and losing prices resembles that of the iterative mechanisms developed later.

Table 3 Summary of combinatorial auction mechanisms

Mechanism	Type	Price-based	Type of Prices	Determination of Payments	Bidding	Deriving and Updating Prices
Rassenti et al. (1982)	one-shot	no	N/A	linear per-unit prices (sum ≤ bid prices)	package bidding	two sets of prices (for accepted and rejected bids) from the restricted dual
AUSM Banks et al. (1989)	iterative (continuous)	no	N/A	pay-your-bid	package bidding; bidder must show her bid increases seller revenue	N/A
PAUSE Kelly and Steinberg (2000)	iterative (1st stage round-based, 2nd continuous)	1st stage yes, 2nd stage no	N/A	pay-your-bid	1st stage: bids on single items, 2nd stage: as in AUSM	result of 1st stage is lower limit for 2nd stage
RAD Kwasnica et al. (2005)	iterative (round-based)	yes	linear, anonymous ask prices	pay-your-bid	package bidding; bids must satisfy ask prices	minimizes infeasibility in the restricted dual
Endogenous Bidding Mechanism Kwon et al. (2005)	iterative (round-based)	yes	linear and nonlinear, anonymous ask prices	pay-your-bid	package bidding; bids must satisfy ask prices	minimizes infeasibility in the restricted dual
Combinatorial Clock Auction Porter et al. (2003)	iterative (round-based)	yes	linear, anonymous clock prices	pay-your-bid	bidders announce their demand at current prices	prices are increased for items with excess demand
A*k*BA Wurman and Wellman (2000)	iterative (round-based)	yes	nonlinear, anonymous ask prices	pay-your-bid	package bidding; bids must satisfy ask prices	minimal prices that support the solution of the WDP
*i*Bundle Parkes and Ungar (2000)	iterative (round-based)	yes	nonlinear, non-anonymous clock prices	pay-your-bid	bidders announce profit maximizing packages at current prices	determined through subgradient algorithm
dVSV de Vries et al. (2007)	iterative (round-based)	yes	nonlinear, non-anonymous clock prices	pay-your-bid	bidders announce profit maximizing packages at current prices	determined through primal-dual algorithm
*i*BEA Mishra and Parkes (2007)	iterative (round-based)	yes	nonlinear, non-anonymous clock prices	discount on bidprice	bidders announce profit maximizing packages at current prices	determined through subgradient algorithm
DVA Mishra and Veeramani (2007)	iterative (round-based)	yes	nonlinear, non-anonymous ask prices	Vickrey prices	bidders announce profit maximizing packages at current prices	prices decreased by ε for combinations not in the demand sets

3.2.4 Combinatorial Auctions in Practice

There are several reports of combinatorial auctions taking place in practice. At least Sears, Roebuck, and Co. (Ledyard et al., 2002), Mars Inc. (Hohner et al., 2003), Motorola (Metty et al., 2005) and Procter & Gamble (Sandholm et al., 2006) have successfully implemented combinatorial auctions in some form in their procurement process. Sears organized combinatorial auctions to acquire transportation services. The usefulness of combinatorial bidding in transportation auctions has been discussed in other articles as well (Sheffi, 2004, Caplice and Sheffi, 2006, and Caplice, 2007). Caplice (2007) reports that since 1997, hundreds of companies have used combinatorial, electronic auctions to purchase truckload transportation.

Mars Inc., Motorola and Procter & Gamble have used combinatorial auctions to find suppliers. The reasoning behind adopting auctions as a part of their procurement process is the same for every firm. The standard practice used to be to negotiate with each potential supplier individually. These negotiations were lengthy, and it was difficult for the negotiators to indicate to one supplier what they wanted, since they did not necessarily know yet, what other suppliers had to offer. Auctions were thought of as a way to cut down the time and money spent on negotiations, and hopefully to even find a better set of suppliers to work with. All the firms also had similar concerns about switching to auctions. They feared that it would ruin the relationships they had with their suppliers. Indeed, if the reliable suppliers viewed the new auctions as hostile action attempting to squeeze out all profits, they might choose not to enter the auction. The auction might not have enough good participants, and the firms would be left without the required supplies.

The auctions organized by the four firms were all customized to their specific needs, and therefore they were all different in many ways. However, there were significant similarities between them too. First of all, they all embraced the complexity of the environment rather than trying to simplify matters. The auctions were designed so that they could capture all the benefits from economies of scale and scope the supplier might have – while making sure that the bidders do not feel being ripped-off. This leads to complicated designs, in which bidders are able to make very expressive bids on bundles, but also things like quantity discounts can be expressed. Procter & Gamble even allow bidders to announce other kinds of conditional discounts. Bidders can also announce capacity constraints. The buyers can restrict the number of winning bidders (because dealing with a large number of suppliers is more costly), and they can favor trusted suppliers, if they wish. All firms report positive results from their preliminary experiences with auctions. Costs have gone down, and the suppliers are still happy doing business with them.

Combinatorial auctions have also been utilized by the public sector. Epstein et al. (2002) describe how using combinatorial auctions in procuring meals for schools saved the Chilean government $40 million annually. More importantly, the reduction in cost did not come at the expense of the equality of the meals. Another application of combinatorial auctions in the public sector is the spectrum license auction of the FCC (Auction #73), which was described already in section 3.1.2.2. Because there has only been this one combinatorial spectrum license auction so far, it is no possible to estimate the increase in government revenue resulting from combinatorial bidding.

Auctioning bus routes has become popular in cities and metropolitan areas. Recently, there have also been bus route auctions allowing combinatorial bidding (Cantillon and Pesendorfer, 2006, Tukiainen, 2008). Allowing combinatorial bidding makes sense since there are synergies between bus routes originating at the same place (and possibly near the garage of the firm). However, both Tukiainen (2008) and Cantillon and Pesendorfer (2006) report that rather few combinatorial bids were submitted. One reason for this could be that only a small subset of all bus routes are up for auction each year, and at least in the case described by Cantillon and Pesendorfer (2006), the routes were divided into smaller auctions with 4 routes in one auction on average.

4 MULTI-ATTRIBUTE AUCTIONS

So far all the auctions I have discussed have been based on price alone (and quantity in multiple-unit auctions). A completely different take on auctions is to include more attributes into the bids. Using price as the only bidding attribute is sufficient when selling existing, clearly defined products, such as agricultural products or works of art. All other attributes related to the products are predetermined and the information is available for the bidders. However, often in procurement situations when the item auctioned does not exist yet and can have many varieties, negotiating merely a price is not sufficient. The buyer will want to agree on other issues (quality, and terms of payment and delivery to name a few) before agreeing to sign a contract. Multi-attribute auctions are designed to take these issues into consideration already in the auction process. Auctions are often considered a special case of negotiations, and including multiple attributes into bids brings auctions a step closer to negotiations. Also, as combinatorial auctions are used in procurement situations, it would be important to be able to combine these two auction types in one auction mechanism.

Traditionally, companies and governments have sent out requests for quotes (RFQs) in procurement situations. Based on the quotes, the company chooses the most promising candidates and begins one-on-one negotiations with them. The contracts are signed as a result of the negotiations, which can sometimes be long and tedious. Multi-attribute auctions allow for a more structured process, which should save both the buyer's and the seller's time – and money. In auctions, the rules of the game are clearly defined and the bidders are aware of the attributes according to which their bids are evaluated. The process becomes more automated in a way, especially if the Internet is used as the medium for the auction. Multi-attribute auctions are flexible in the sense that they can be single-unit, multiple-unit[11] or even multiple-item auctions.

Multi-attribute auctions resemble combinatorial auctions in some ways. First of all, bids in both auctions are vectors. Only now the vector components indicate the levels of attributes, where as in combinatorial auctions they indicated desired quantities of

[11] Teich et al. (2004) use the term "multiple issue auction" when referring to a multi-unit, multi-attribute auction.

different items. Secondly, winner determination is complicated in both auctions, although for different reasons. In combinatorial auctions the difficulties were computational. In multi-attribute auctions the difficulty is in comparing the bids with each other – it is like trying to compare apples and oranges. The bid taker's preferences over the multiple attributes need to be elicited prior to the auction, and the mechanism designed according to these preferences. Again there is a certain resemblance to combinatorial auctions, where the bidders needed to elicit their preferences over the different items in the auction. The problem is essentially the same, even though this time it is the bid taker, who has to solve it. Naturally, the bidders also need to make similar evaluations when placing bids. The fields of multiple criteria decision making and decision support specialize in developing tools to help decision makers express their preferences over multiple attributes. In the following sections I will review different approaches to multi-attribute auction design found in auction literature. The review will follow along the lines of Teich, Wallenius, Wallenius and Koppius (2004). First I will present the scoring function method, which is based on value function theory and commonly found in theoretical articles. Then I will proceed to less rigorous and perhaps more user-friendly auction designs.

4.1 The Scoring Function Approach

One way to evaluate multi-attribute bids is to formulate a scoring function $S(\mathbf{x})$: $R^n \rightarrow R$, where $\mathbf{x}$ is the vector of n attributes. The scoring function assigns each bid a score based on which the bids can then be ranked. The winner of the auction is simply the bidder whose bid produces the highest score. The concept of the scoring function is similar to that of the value function. The construction of value functions has been extensively studied within the multi-attribute utility theory (MAUT). The value function assigns a value for the level of each attribute and then combines the values from all attributes. A commonly used value function is the additive value function with weights.

The focus of auction studies implementing the scoring function differs from that of the MAUT. Where the MAUT studies different approaches to the elicitation of preferences and construction of value functions, most theoretical auction studies take the scoring functions as given. The point of interest in auction papers is either the optimal (i.e. utility maximizing) scoring rule under different auction mechanisms (Che, 1993, Branco, 1997, and Beil and Wein, 2003) or the study of the economic implications of the multi-attribute setting compared to the single-attribute auctions (Bichler, 2000, and Chen-Ritzo, Harrison, Kwasnica and Thomas, 2005).

The scoring function approach follows the research tradition established by the game-theoretical price-only auction studies. The scoring function reduces the multi-attribute auction to a single attribute auction, which allows the development of mathematically beautiful models similar to auction models such as the IPV and Common Value models. In other words, the scoring function auction models are an extension of the traditional models. This is apparent in the objectives of the studies, which include testing the efficiency of different mechanisms and the revenue generated by the auctions, and in the formulations of the game-theoretic models.

Bichler's (2000) approach is more practice-oriented even though the objectives of his study are related to the efficiency of the outcome and the comparison of the payoffs of different auction mechanisms. The major difference is that Bichler has not only designed an auction, but also implemented it. He tests the payoff equivalence hypothesis and the efficiency of a multi-attribute auction (using a scoring function) in the WWW-environment using MBA students as test cases. Bichler arrives at the conclusion that multi-attribute auctions produced higher payoffs than single-attribute auctions. However, he does not mention, how the values for the non-price attributes were chosen in the single-attribute auctions. Hence the basis of such a comparison is left vague. More convincing are Bichler's results indicating that the auction mechanisms were not payoff equivalent (first-score auction produced a higher revenue than English or second-score auctions). Bichler offers the heterogeneity (asymmetry) of bidders as a possible explanation for the discrepancy between theory and the results of the experiment.

Following Bichler's example, Chen-Ritzo et al. (2005) also compare multi-attribute auctions to price-only auctions. Their experiment is more complicated than Bichler's (they use three attributes instead of two), but their results are the same: multi-attribute auctions are more efficient and result in higher payoffs. Chen-Ritzo et al. (2005) use the utility maximizing levels of the non-price attributes in the price-only auctions, which makes their comparisons more reliable than Bichler's (2000).

4.2 Other Approaches

The explicit assessment of value functions, which the scoring function approach requires, has been criticized by several researchers over the years (e.g. Simon 1955, Larichev, 1984, and Korhonen and Wallenius, 1996). It requires practice and expertise to be able to express one's preferences in the form of a value function. In the auction setting, the assessment of the auction owner's preferences should be easy in order to entice managers to resort to auctions in the procurement process. The applicability of the scoring function method is thus questionable. Bichler (2000) added a decision-aiding tool to help auction owners construct their value functions. However, he does not describe the tool. If the process is not transparent and understandable to the user, it will not evoke trust in the procurement manager and she might decide not to use the auction system.

The focus of recent studies in multi-attribute auctions has diverted from the focus of traditional auction research. Now the focus is not as much the efficiency or the revenue (utility) equivalency of auction mechanisms, but the functionality of the designed auction. Functionality refers here to the ease of use for both the seller and the buyers. Even though the primary goal is not to design an auction that maximizes the auctioneer's revenue (utility), it is, of course, important that the auctioneer receives a satisfactory utility. Multi-attribute auctions can be seen as cooperative negotiations (see e.g. Guttman and Maes, 1998), as there can be opportunities for joint gains. In a way, multi-attribute auctions resemble traditional one-on-one negotiations. Therefore, some multi-attribute auction design builds upon negotiation theory, e.g. the The Leap Frog Method and the Auction Owner Controlled Bid Mechanism suggested by Teich, Wallenius and Wallenius (1999).

Cripps and Ireland (1994) propose a method, which uses quality thresholds. This would eliminate all the other attributes besides price, and render the auction to the price-only situation. They consider three specific designs. In (a) a price-only auction is held only after bidders have submitted their quality plans (which contain information on all non-price attributes) and the plans have been accepted. In (b) the auction is held first, and quality plans are requested in the order indicated by the auction results, starting with the winner. The first bidder, whose quality plan is approved, obtains the contract. In (c) price and quality plans are submitted simultaneously, and the contract is awarded to the best priced (i.e. the cheapest deal in the reverse auction setting) plan that satisfies the predetermined quality requirements.

Teich, Wallenius, Wallenius and Zaitsev (2001 and 2006) have implemented the "pricing out" method in their Internet-based hybrid auction called *NegotiAuction*™. Pricing out can be used without having to explicitly formulate the auctioneer's value function. Instead, it probes into the implicit preferences of decision makers. This makes it a popular approach among decision analysts. It is, however, possible to construct a value function with pricing out, but Teich et al. (2001, 2006) choose not to do so.

Pricing out can be used in all situations, where there is a natural monetary attribute (price or cost) attached to the auctioned asset. Borcherding, Eppel, and von Winterfeldt (1991) compared pricing out with three other value elicitation techniques. They concluded that weights generated with pricing out corresponded closest with weights generated externally by a group of experts.

Keeney and Raiffa (1976) provide a general description of the pricing out method. The underlying idea is to express the decision maker's (in this case the auction owner's) preferences over multiple attributes in monetary terms. Assume that there is a monetary attribute M and n different non-monetary attributes $X_1, X_2, \dots, X_n$ related to a product. Lower case letters m and $x_1, \dots, x_n$ denote the values the attributes can assume. In pricing out the auction owner is asked to identify the monetary value m^* for a bundle $\mathbf{x}^*=(x_1^*, \dots, x_n^*)$ which makes her indifferent between the bundle $(m^*, \mathbf{x}^*)$ and a predetermined reference bundle $(m_0, \mathbf{x}_0)$. That is, $(m^*, \mathbf{x}^*) \sim (m^0, \mathbf{x}^0)$. The difference m^* - m_0 depicts the auctioneer's willingness to pay for the possibility of transforming the bundle $\mathbf{x}^0$ into $\mathbf{x}^*$.

The approach explained above becomes tedious and time-consuming when the number of vectors that need to be evaluated is large. Then, it pays off to make simplifying assumptions. Keeney and Raiffa (1976) state that pricing out can be made easier when

1) The difference between m^* and m^0 (i.e. the willingness to pay) does not functionally depend on the base value of m^0
2) The monetary attribute M and attribute X_i as a pair are preferentially independent of the complementary set of attributes

When these assumptions apply, the pricing out can be done individually for each attribute. The method is the same as in the general case described above, but here the value of only one attribute will be changed at a time. For each attribute X_i the auctioneer is asked to state a monetary value m^* that will satisfy the indifference equation

$$(m^*, x_1^0, \ldots, x_{i-1}^0, x_i^*, x_{i+1}^0, \ldots, x_n^0) \sim (m^0, \mathbf{x}^0) \tag{4}$$

Thus the auctioneer only has to make n such assessments instead of pricing out all possible combinations of the attributes X. This simplifies the preference elicitation process and formulates it in such a way that it is easy for the auctioneer to make the assessments.

Pricing out also provides computational advantages, because it reduces the bids to two-dimensional vectors containing only price and quantity components. The bidders still submit multi-dimensional bids, but all other attributes are "priced out" before the winner determination problem is solved.

The various preference elicitation methods described above try to achieve two goals: the realistic and truthful description of the preferences of buyers and sellers and the ease of use of the method for all participants. Unfortunately, most of the time these goals conflict. The more elaborate the preference elicitation scheme is, the more difficult it is for a novice to use. The scoring method clearly emphasizes the first goal and is therefore more suitable for theoretical studies. The rest of the above mentioned approaches prioritize the usability of the methods attempting to generate mechanisms that could be implemented in practice. Regardless of the method in question, it is clear that in the multiple-issue case, bidding becomes more difficult for the bidders, especially inexperienced ones. Also, it is more difficult for the auction owner to set up an auction that would produce the kind of results that would match her true preferences. Hence, all sorts of decision aid tools become important in auctions (Bichler, 2000), and decision making theory becomes relevant for auction theory. The introduction of Internet auctions provides an excellent medium to include decision support with auctions, as will be discussed later in Chapter 5.

4.3 Multiple Attributes in Combinatorial Auctions

Since both multi-attribute auctions and combinatorial auctions are very suitable for procurement situations, it would make sense to combine aspects of both auction types into one auction. However, as both auction types alone are already quite complicated, their combination cannot be expected to be any simpler. There has not been much work on multi-attribute combinatorial auctions, but researchers have identified the potential benefits of such combinatorial auctions.

Sandholm and Suri (2006) propose the use of a weighting function to translate the multiple attributes into monetary terms. The weighting function $f(p_i, a_i)$ takes a bid b_i and weights its price p_i with the values of the attribute vector a_i and returns a new price, which is then used to compare bids with each other. The approach suggested by Sandholm and Suri bears great resemblance to the pricing out method used by Teich et al. (2001, 2006). Teich et al. present the pricing out method in a single-item setting, but the extension to a multi-item auction is straightforward. In a multi-attribute combinatorial auction the incoming bid would actually be a matrix, where the columns indicate the items, the last column being the price, and the rows the values of different attributes (first row is for the item quantities and subsequent ones for non-price attributes).

Epstein et al. (2002) describe a practical application of a multi-attribute combinatorial auction. In the auction the Chilean government procured school meals. The auction was designed to be as simple as possible in such a complex setting. Thus, the multiple attributes were taken care of by pre-auction approval. The government set several criteria for the food to be delivered, and only firms able to fulfill the criteria were allowed to place bids. This approach is similar to the quality threshold method proposed by Cripps and Ireland (1994).

The "expressive bidding" procedure of Sandholm (2007) is an ambitious attempt to combine combinatorial and multi-attribute auctions for procurement purposes. His auction system, CombineNet, also allows the bidders and the auction owner to include a variety of side constraints (capacity constraints, minimum or maximum number of winning bidders, etc.). The system also allows the bidders to express discount schedules. The underlying goal in the development of CombineNet is create as much flexibility as possible so that the bidders and the buyer could find win-win solutions (Pareto improvements).

5 ONLINE AUCTIONS AND AUCTIONS IN PRACTICE

A big part of auction literature concentrates on auctions in theory, and forget about the practice. According to Klemperer (2002) most of the traditional studies (also reviewed in this thesis) are of little help when designing auctions in practice.

The development of computers and the introduction of the Internet have had a tremendous impact on the practical side of organizing auctions, and thereby also on auction research. First of all, computers have enabled the organization of new kinds of auctions, like combinatorial and multi-attribute auctions, and they are helpful in multiple-item auctions (Pinker et al., 2003). Secondly, besides enabling the use of more complicated auction mechanisms, the Internet has affected the traditional and simple single-item auctions creating new potential problems for the design.

The Internet has introduced new elements into auctions, which were not present in traditional auction settings. First of all, Internet has reduced transaction costs from organizing auctions and participating in them. This has broadened the spectrum of products that can be sold through auctions. Many standard products, which earlier were always sold with posted prices, are nowadays also sold through auctions. Because the items sold in online auctions are not unique, it is likely that similar products are sold in several separate auctions. Some of these auctions can be ongoing at the same time, and some occur later in time. In any case, it is not reasonable to treat such auctions as isolated and independent incidents.

Other new elements include an increased number of potential participants, and the endogenous entry of new bidders (i.e. bidders can enter auctions even after they have begun). Also, the most common auction type in the online environment seems to be a uniform price multiple-unit auction (Bapna, Goes and Gupta, 2000). This is in clear contrast to the pre-Internet era, when the single-unit auction was the prevalent auction type. Also, the duration of online auctions can be much longer than that of traditional auctions. All of the new elements mentioned above affect the analysis of all auctions, and the theoretical results of traditional auction theory focused on single-unit auctions (e.g. Vickrey, 1961, Myerson, 1981, and Milgrom and Weber, 1982) may not apply anymore. According to Pinker et al. (2003), there has been very little research so far that would consider these new elements.

Auctions, even the more complex ones, allegedly have lower transaction costs than negotiations. There are some concerns, though (Pinker et al., 2003). First of all, procurement auctions have been accused of squeezing out all surplus from the bidders. The second concern is that people advocating the low transaction costs do not take into consideration all costs related to procurement by auction. Mainly these ignored costs refer to switching costs incurred whenever changing to a new supplier. Also, time costs from participation can be significant. The third argument is that the trend in supply chain management has favored vertical integration and partnerships, and auctions do not support this development. Tight partnerships cannot be established, if the supplier base is renewed annually based on results of auctions. All of these concerns were voiced in the cases of Mars, Inc, Motorola and Procter & Gamble (see section 3.2.4). The potential problems were considered when designing the auction, and the results have been positive on all accounts.

The fact that auctions are used more often in different kinds of market transactions, and the demand for different kinds of auction designs has grown, there is also need for more research on auction design in practice. The fact is that designing a successful auction is by no means trivial. Already the choice of the appropriate type (single- or multiple-item, single- or multiple-unit, multi-attribute or price-only) and mechanism (English, Dutch, first-price sealed-bid, Vickrey, and their extensions) is difficult, and abundant theoretical research is not necessarily of much help, because it relies on a set of restrictive assumptions. Once the auction type and mechanism are chosen, they already determine some of the major design issues such as bid type, pricing rule and winner determination. However, this is not sufficient. One must also take into consideration minor issues, which are not inbuilt in the auction mechanism. In the following I will discuss additional design issues that have not been presented yet, because they have not been considered in the theoretical models. I will also discuss the need for additional rules for the auction, the purpose of which is to minimize the risk for the auctioneer from bidders cheating or defaulting.

5.1 Additional Design Issues

Additional design issues are issues that need to be determined when designing an auction, but which are not tied to any particular auction mechanism. Such issues are the bid increment (or decrement in descending price auctions), possible reservation price and the complete information architecture of the auction.

The bid increment is the minimum amount by which the new bid must exceed the current highest bid in order to become the new provisional winner. Similarly, the bid decrement is the minimum amount by which the price must decrease (in a reverse auction) before the bid can become a provisional winner. In theoretical models the bids are assumed continuous, that is, an increase of the size of ε is always possible, but in practice in almost every auction a fixed increment is defined. The increment can be of either a fixed dollar value or a percentage. In high value auctions, an increment of 1% could be millions of dollars, hence there usually is a maximum dollar value for the increment as well. The size of the increment/decrement naturally affects the convergence speed of the auction. The larger the increment, the fewer bids are needed to reach the final outcome. However, the larger the increment the larger the expected gap between the auctioneer's actual revenue and potential revenue. For example, assume that the bidder with the highest valuation has the valuation of 100, and the bid decrement is 10. If the current highest bid in the auction is 91, the bidder cannot bid anymore. If the increment were 9 or less, the bidder could place a bid, and the auctioneer's revenue would increase. The choice of bid increment/decrement has not been studied much. Teich et al. (2001, 2006) use a bid decrement in their formulation of the *NegotiAuction* system, but they do not study the effect of different decrements. Bapna, Goes and Gupta (2000, 2003) also have a bid increment in their auction mechanism. Bapna et al. (2003) conclude based on theoretical and empirical considerations that the bid increment has a significant effect the auctioneer's revenue, but their result is tied to the multiple-unit, uniform-price auction. Also, the bid increment in the studied actual auctions is devised as a fixed dollar amount and not proportional to the bid amount. However, clearly the role of the bid increment should be studied more.

The reservation price means the price below which the auctioneer will not sell the item (in a forward auction) or the price above which the auctioneer will not buy the item (in a reverse auction). The reservation price was mentioned already in the context of optimal auctions in section 2.5. In the case of optimal auctions the reservation price is set to equal the expected valuation of the second highest bidder. However, such a decision rule is of little help, if the distribution of the bidders' valuations in unknown – provided that there even exists such a distribution. The main idea behind setting a reservation price is to shield the auctioneer from an undesirable outcome. For example, little competition can lead to very low prices in a forward auction. A good example of this is the spectrum auction in New Zealand (see section 2.4) in which the winner paid a very low price because the difference between the highest and second highest bids was huge. A reservation price anywhere in the gap would have guaranteed the government a higher revenue. The downside of the reservation price is that it can discourage bidders from entering, and the item might not get sold at all. An additional question related to the reservation price is whether to disclose it to the bidders or keep it as a secret. Bajari and Hortaçsu (2003) have discovered that in eBay auctions the number of participants decreases, revenues are lower, and items are left unsold more often, if a secret reservation price is applied.

In online auctions, the duration of the auction also becomes a design issue. In traditional auctions, which take place in auction houses, the duration of the auction is measured in minutes, and there is no need to predetermine the duration. However, in online auctions where the bidders are not present when the auction opens, the auction duration becomes an issue. Potential bidders need to have time to find the auction. The longer the duration, the more bidders find the auction and participate in it (Vakrat and Seidmann, 2000). Hence, in principle the number of participants should increase as the duration increases, which should be better for the auction owner. However, if the duration is very long, the bidders do not have an incentive to bid at the beginning. By bidding early, they only run the risk of increasing the price. Bidders prefer waiting until the auction is about to close to observe the level of competition. Thus, it is possible that the number of bids remains low. Also, the longer the auction, the larger the transaction costs for the participants. Bidders need to monitor the auction, plan their strategies and place bids. In high stakes auctions firms usually have teams of experts involved in the bidding process throughout the auction, which can be a big cost to the firms. Take for example the FCC auctions in which all major telecom companies had hired their own team of auction researchers to advise them in the auctions.

Information architecture as understood by Koppius and van Heck (2002) encompasses all information flowing between the bidders and the auctioneer. Part of the information architecture is thus determined by the auction mechanism, for example the type of bids (multiple-item, multi-attribute or price-only) and whether the bids are openly visible to every participant (open-cry auctions) or kept a secret (sealed-bid auctions). There can be various degrees of bid openness, though. In some auctions only the provisionally winning bids are announced and in others the bidders only know the status of their own bids (provisional winner or not), but not the content of any other bids. Information architecture also contains all the information the bid taker wishes to disclose to the bidders. This can be the number and identity of all the bidders, the bid taker's reservation price, or the bid taker's preferences over the multiple attributes in multi-attribute auctions. The bid taker can also choose to misrepresent some information, or reveal only partial information about her preferences. Koppius and van Heck study the impact of information architecture in multi-attribute auctions. According to them, information architecture is even more important in multi-attribute auctions, because it enables the bidders and bid taker to identify possibilities for Pareto-improvements. Thus, they hypothesize that more information would improve the efficiency of the final allocation. Their experiments support this hypothesis, but they also discover a saturation point after which more information did not affect the auction outcome significantly.

5.2 Auction Rules

When organizing a real auction, choosing the mechanism and designing the details is not enough; some fine tuning is required. The need for more detailed rules is especially pronounced in electronic auctions, because it becomes more difficult to observe and control bidder behavior and prevent cheating. In this section, I will discuss rules, which concern issues and behavior that are ignored in theoretical models (e.g. defaulting on bids, cheating and collusion), but which are important in practical applications. There has been a lot of discussion on auction rules in the context of the FCC Auction #31 (FCC, 2000, Pekeč and Rothkopf, 2000, Vohra and Weber, 2000), which ironically after all the planning and discussions never took place. Also articles by Pekeč and Rothkopf (2003), Kelly and Steinberg (2000) discuss different auction rules, and Klemperer (2002) describes several real auctions where design flaws have led to undesirable auction outcomes.

The reasoning behind rule design is that bidders will try to take advantage of the auction design in any way they can think of. For instance, in iterative auctions with a predetermined ending time bidders may want to wait until the last minute before making a bid. This type of behavior is known as *sniping*. Snipers hope to be able to surprise the competition with a last minute bid and leave no time for the competitors to respond (Roth and Ockenfels, 2002). If they are successful, the revenue from the auction remains low. To combat this kind of behavior, *activity rules* have been introduced in iterative auctions which involve substantial amounts of money. *Activity rules* state how often and what kind of bids a bidder must enter in order to be allowed to continue bidding. Activity rules are often linked to minimum bid increments. *Closing rules* affect activity rules. Activity rules become more important, if the auction has a predetermined closing time. However, if the auction closes only after a certain period of inactivity, activity rules are not as critical, but it does not mean they would be unimportant. In an extreme case there would be no bidding activity until right before the closing time. There would still be competition as the ending would be pushed back, but the process of price discovery would not be as efficient undermining one of he benefits of iterative mechanisms. Auction owners may want to use *eligibility rules* as well. Eligibility rules require the bidders to prove their solvency prior to entering the auction. For instance, they might have to make a deposit, and the size of the deposit determines the size of bids they are allowed to place.

Collusion in the form of collusion rings agreed upon prior to the auction was discussed already in section 2.4. Collusion does not have to be explicit, though. During the auction bidders can try to signal their intentions to competitors. Maasland and

Onderstal (2006) and Klemperer (2002) report a case of signaling in the German GSM auction. In 1999, the German government put 10 licenses up for auction, and required a 10% bid increment. One of the competing firms, Mannesmann, made a bid of 20 million DM on five licenses, and a bid of 18,8 million DM on the other five licenses. Mannesmann's main competitor T-Mobile was able to calculate that topping the current high bid of 18,8 million with the required bid increment of 10% would bring the price of these licenses to about 20 million as well. Actually, this was Mannesmann's way of signaling to T-Mobile its willingness to share the licenses equally; it did not necessarily have to bid for the other five licenses at all. T-Mobile understood the signal, and the auction ended very shortly with the bidders sharing the licenses at the price of 20 million DM. The efficiency of the outcome is hard to determine, but clearly the final price obtained by the seller (German government) was artificially low. This example shows that no explicit collusion among the bidders is necessarily needed for them to reach a silent agreement to not drive up the prices.

In some auctions bidders have used the lower digits of bid prices to signal the items they are interested in (Kelly and Steinberg, 2000). To combat this, the auction owner can require that bid prices follow certain increments. Of course, sealed-bid auctions do not suffer from this type of signaling. Another form of signaling is called jump bidding, which was observed in earlier FCC auctions (McAfee and McMillan, 1996). Jump bidding refers to aggressive bidding behavior, where a bidder places a bid way above the required increment in order to warn other bidders. One way to prevent jump bidding is to use a clock auction, in which the bidders simply announce their demand at the price indicated by the auction clock instead of calling out their own prices (Banks et al., 2003).

Cheating and fraud are also a concern in auctions, and especially in online auctions. It is easier for bidders to remain anonymous, and thereby use multiple bidder identities in the same auction (Pinker et al., 2003, Yokoo et al., 2004). In a similar vein, it is also possible for the bid taker to place bids under a false identity in order to force the auction to a more favorable outcome. These cases of cheating were discussed already in the context of the VCG mechanism in section 3.2.3.1. Bidders could also choose not to pay, or sellers refuse to provide the product after receiving a payment. Auction designers have tried to design mechanisms to detect and prevent fraud. Also reputation of both buyers and sellers has become a factor in the electronic market places.

Pekeč and Rothkopf (2003) point out that auction rules should also define how ties are to be broken. Theoretically ties are not interesting, because with continuous distributions their probability is zero. In practice, however, it is quite possible that a tie arises, because bidders tend to round prices, or the auctioneer requires bidding at certain bid increments (decrements). In combinatorial auctions it is even more likely because ties can occur in different ways, because the same total revenue can result from many different combinations. In the name of fairness, the ties should be broken based on predetermined rules. One way to break ties is to pick the winning combination randomly. A more sophisticated way is to use time stamps. Each incoming bid receives a time stamp when it enters the auction. The winning combination is then either the one that was completed first (highest time stamp value is lower than the highest stamp value of other combinations), or the one with the lowest average time stamp. Using the time stamps may not be entirely fair though, because due to differences in traffic loads on the Internet, some bids might be at a disadvantage.

Additional issues to be decided on are how to deal with bid withdrawals or defaulting on winning bids. Withdrawals are usually allowed in simultaneous and sequential auctions to alleviate the exposure problem. However, in combinatorial auctions exposure problem is not as crucial, so allowing withdrawals may only make collusion or signaling easier. Penalizing withdrawals and defaulting on winning bids can thus help reduce cheating and “gaming” behavior. Setting eligibility rules and requiring the bidders to make a deposit also ensures that the bidders are capable of paying the penalties.

The important thing to keep in mind in auction design is that each auction is unique and therefore requires a unique combination of design parameters. This fact is emphasized by Binmore and Klemperer (2002), who report their experiences on telecom license auctions in the UK and other European countries. According to them it is not enough that the items for sale are identical to justify using an identical auction design, because market conditions (number of potential participants, attractiveness of market) vary. Sometimes it is enough to adjust minor rules, but in some cases the whole auction mechanisms needs to be modified. In Binmore and Klemperer's study, what worked well in the British telecom auctions, did not work in the Netherlands or Switzerland, which clearly demonstrated that auction design should not be copied.

III MOTIVATION AND OBJECTIVES

6 MOTIVATION AND OBJECTIVES

As discussed in the literature review, combinatorial auctions are characterized by complexity. The winner determination is computationally complex, and the construction of bids requires the elicitation of complex preferences over a set of different items. Good bidding strategies are not easy to calculate, as good bids for one bidder depend not only on her preferences and cost structure, but also on bids made by other bidders. There is a lot of literature concerning the efficient solution of the WDP (see section 3.2.2.1), and recently there has also been some research into the elicitation of bidder preferences (section 3.2.2.2) and the design of auction mechanisms (section 3.2.3). The addition of multiple units of each item makes the mechanism design more complicated – and complicates further the bidders' task of submitting bids. Also, only a few of the mechanisms are easily extended to the multi-unit setting. Bidder support has been neglected in existing literature. Thus, there is a need for a multi-unit auction combinatorial, which would be easy for bidders to participate in.

6.1 Need for Mechanisms for Multiple-Unit Combinatorial Auctions

Thus far, the extension of the combinatorial auction mechanisms for multiple-unit cases has not been discussed in literature much. Out of the mechanisms reviewed in section 3.2.3.2 five are easily extendable to the multi-unit setting: AUSM, PAUSE, RAD, the Endogenous Bidding Mechanism, and the combinatorial clock auction. However, all of them have their shortcomings, one of them being that they compromise efficiency. In multiple unit auctions, the effect of linear ask prices on efficiency would be even worse because they remove the possibility of expressing economies of scale (i.e. quantity discounts). Using non-linear pricing improves efficiency, but it cannot really be used in multi-unit auctions, or large single-unit auctions for that matter. This is because auctions with multiple units, or a large number of items have too many possible combinations to be evaluated in reasonable time. Of course, not all combinations need to be considered – e.g. iBundle only considers combinations for which there are bids from previous rounds – but that compromises efficiency. Thus, there is room for a different approach in combinatorial auction mechanism design.

6.2 Identification of the Puzzle Problem and Need for Quantity Support

Combinatorial auction literature recognizes the threshold problem – as noted in section 3.2.1.3 – but little has been done to try to alleviate it. In addition, I believe the threshold problem does not adequately describe all the problems facing the bidders bidding in a combinatorial auction. Firstly, the threshold problem refers to a situation in which a large number of "local" bidders – bidders bidding for single items or small packages – are trying to coordinate their bid prices to outbid a "global" bidder – a bidder bidding for the whole bundle (or a few big bidders). However, the bid price is only one dimension in the bid vector in a combinatorial auction. The bidder can also choose to vary the item combination associated with the price, and this, I believe, creates a whole new problem.

In combinatorial auctions, a successful bid complements existing bids, placing all of them among the winners (unless the winning bid is for the entire bundle). This brings a cooperative flavor into the auction even though bidders are still in competition with each other and are not allowed to collude. The threshold problem is one phenomenon arising from this cooperative nature of bidding. The threshold problem – the way it is presented in literature – is confined to price adjustments. However, in combinatorial auctions, the item combination in a bid plays as large a role in determining whether the bid is among the winners or not. A combinatorial auction is like a puzzle: in addition to the prices being right, the bids need to fit together to form the whole bundle like puzzle pieces fit together to form a complete puzzle. However, in a combinatorial auction, the size and shape of the puzzle pieces are not predetermined; it is the task of the auction mechanism to endogenously determine them. The WDP is then analogous to the process of choosing which pieces to use to compile the puzzle. Due to this very fitting analogy, I call the problem of finding and placing bids that complement other bidders' bids (i.e. bids that will be chosen by the WDP), the *puzzle problem*. An implication of the puzzle problem is that even if a bidder has managed to identify her most preferred combinations, it may not make sense to bid on them, if there are no complementing bids coming from other bidders.

In open-cry auctions, the puzzle problem is not more difficult to overcome than the threshold problem. But in sealed-bid auctions the puzzle problem becomes almost impossible to overcome, and it is a serious threat to allocative efficiency in such auctions. In sealed-bid auctions the bidders do not know the contents of the competing bids, and thus it is impossible for them to deliberately place complementing bids. The puzzle problem can thus arise without the coordination issues, which are at the heart of the threshold problem. A bidder could be able to place a bid that would make her (and a group of existing bids) winners, but does not know which bid it would be.

Most procurement auctions use a sealed-bid format, because they are preferred by the bidders (Jap, 2003). In sealed-bid auctions bidders do not have to worry what information their bids could reveal to their competitors. Also, for the auction owner a sealed-bid auction has the advantage that it removes the possibility of signaling, and jump bidding will not be as effective because competitors cannot observe it. The fact that combinatorial procurement auctions are often held in a sealed-bid format means that the concerns of auction outcomes being inefficient due to the puzzle problem are very relevant. The problems are aggravated in multiple-unit combinatorial auctions, because not only do the items in the bids complement each other, but the quantities also need to add up to the total demand. Thus, it can easily happen that a bidder with a low cost structure (in a reverse auction) loses, because she did not bid for the "right" combination.

Even though winning in a combinatorial auction depends on other bidders' bids, the support mechanisms presented in section 3.2.2.2 or the iterative auction mechanisms in section 3.2.3.2 do not attempt to find bids to form coalitions with other bidders. Offering price information in an iterative auction guides the bidders to bid for items with a relatively high price, but it does not help in determining the quantities for each item (this is of course relevant only in multi-unit cases). I feel that this aspect of bidding in combinatorial auctions has been neglected in existing literature, and it is important that the puzzle problem be addressed.

6.3 Objectives and Methods of the Study

The objective of this study is to overcome the puzzle problem present in multi-unit sealed-bid and semi-sealed-bid[12] combinatorial auctions. In this research project, we consider only iterative auctions, as one-shot auctions present very limited opportunities to a) gather any kind of information from the auction, and b) to support bidders. Also, our focus is on continuous, iterative auctions, and not on round-based auctions.

We try to reach the objective by developing support tools for bidders. The task of the support tools is to find bid suggestions that would complement the existing bids (that is, identify the size and shape of possible missing pieces of the puzzle). These bids should be beneficial for both the buyer (total cost should decrease), and the bidder (bidder should make a profit).

Our hypothesis is that providing this kind of "quantity support" the sealed-bid or semi-sealed-bid auction would reach a more efficient outcome. The support tools should also be considered fair by both the buyer and the sellers, because they try to maximize the bidders' profit, all the while decreasing the total cost to the buyer. Also, providing support for the bidders – and thereby making bidding easier and less costly – the auction would be more attractive, and more bidders would participate. More competition should improve the buyer's position, as she can expect to obtain the items for a lower total cost.

The main methods used in this study are simulations and laboratory experiments. The simulations were used to study the performance of the support tools. Through simulations we could observe the efficiency of the final allocations, as well as the total cost to the buyer. The laboratory experiment was used to study whether the simulated results could be reproduced with human users. The laboratory experiments were also used to observe bidders' behavior in combinatorial auctions. Based on bidders' behavior I identified different bidding strategies. In addition, I could draw conclusions on how difficult a bidding environment is for inexperienced bidders, and how good the usability of the user interface is.

[12] A semi-sealed-bid auction is a sealed-bid auction in which the bidders know which of their own bids are among the provisional winners (= active) and which are not (= inactive).

IV DESIGNING AND TESTING THE QUANTITY SUPPORT MECHANISM

7 THE QUANTITY SUPPORT MECHANISM[13]

The contribution of this chapter is to present the Quantity Support Mechanism (QSM), a bidder support tool we developed for continuous, semi-sealed-bid combinatorial auctions. I also provide an example auction to illustrate how the QSM works in an auction. The auction mechanism we consider is such that bidders are free to enter bids at any time in the auction, and the WDP is solved after every incoming bid. In this auction, the set of provisional winners changes if the new bid decreased the total cost to the buyer (in a reverse auction) by at least a predetermined decrement. The decrement is public knowledge. The bidders receive information on their bid status (whether they are provisional winners or not), but they do not know the contents of the competitors' bids. The auction has a predetermined closing time, but the time will be pushed back if there is bidding activity in the last minutes of the auction. The winning bidders receive the price indicated by their bid (or in forward auctions, pay their bid price). We chose to consider an iterative, pay-your-bid auction mechanism, because it resembles ideology of the English auction many bidders are familiar with (see discussion in section 2.4).

The QSM has been designed for reverse auctions – and therefore the following discussion will be from the reverse auction perspective – but it can easily be applied to forward auctions. The purpose of the QSM is to suggest bids (price-quantity combinations) to bidders who wish to become provisional winners. The QSM would use the existing bids to solve for good complements, and then suggest these complementing bids to the bidder without revealing the contents of the other bidders' bids. Also, the idea is that the bid suggestions would be the best possible for the bidder, i.e. bids which would maximize the bidder's profit while having low enough a price so that they would become winners at that time in the auction. Continuing the puzzle

[13] Material in this chapter and section 8.1 has been published in Leskelä, Teich, Wallenius and Wallenius (2007).

analogy, the QSM solves for the size and shape of the missing piece in the puzzle. And because by choosing different pieces (existing bids), different puzzles can be compiled, the QSM chooses the one it thinks is the most profitable for the bidder.

7.1 The Quantity Support Problem

If the bidder could express her costs in a functional form, and if she would be willing to disclose the cost function to the bid taker (or, if it could be arranged, to a neutral third party), it would be fairly straightforward for the bid taker to solve for the bid that maximizes the bidder's profit. The quantity support problem for bidder *m* (QSP_m) would reduce to a standard mixed integer programming problem. The objective (5) of the problem is to maximize the profit of bidder *m* by solving for the new price $p_{m,new}$, the vector $Q_{m,new}$ of item quantities $q_{m,new,k}$ and the values for the bid status variables x_{ij}. In order for the new bid to become active (provisional winner), the current total cost to the buyer C^* is required to decrease by a predetermined decrement δ as a result of the new bid (6), and the demand for each item d_k must be fulfilled (7). The item quantities $q_{m,new,k}$ should not exceed the bidder's corresponding capacities a_{mk} (10)[14]. It is also assumed that at most one bid per bidder can be active at a time (8), (9) to simplify the bidding language. The formulation of the QSP_m is thus:

[14] We have included only these simple, per item capacity constrains in our formulation. Allowing the bidders to announce capacity constraints for combinations of items would increase the number of constraints in the formulation, reducing its readability. If desired, such more complex capacity constraints can be incorporated in the design.

$$
\begin{aligned}
&(\mathrm{QSP_m}) \quad \max \quad p_{m,new} - \tilde{c}_m(Q_{m,new}) && (5)\\
&\text{s.t.} \quad \sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij} p_{ij} + p_{m,new} \le C^* - \delta && (6)\\
&\sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij} q_{ijk} + q_{m,new,k} \ge d_k \quad \forall\ k = 1,\ldots,K && (7)\\
&x_{mj} = 0 \quad \forall\ j = 1,\ldots,n_m && (8)\\
&\sum_{j=1}^{n_i} x_{ij} \le 1 \quad \forall\ i = 1,\ldots,N && (9)\\
&q_{m,new,k} \le a_{mk} \quad \forall\ k = 1,\ldots,K && (10)\\
&x_{ij} \in \{0,1\} \quad \forall\ i, j\\
&p_{m,new}, q_{m,new,k} \ge 0 \quad \forall\ k = 1,\ldots,K
\end{aligned}
$$

where $\tilde{c}_m(Q_{m,new})$ is the cost function of bidder *m*, and as in the WDP (2), x_{ij}'s indicate which bids are among the winners, and p_{ij} indicates the price and q_{ijk} the quantity of item *k* in bidder *i*'s j^{th} bid.

Note that even though we call our tool the "quantity support mechanism", it also suggests a price to attach to the quantities. The quantity support problem in ($\mathrm{QSP_m}$) is presented for a price-only auction, but if we use the "pricing out" –method as in Teich et al. (2001, 2006), also multiple attributes could be included in the bids. Also, if the bids in the bid stream were disclosed to all bidders (an open-cry auction), each bidder could formulate her own quantity support problem similar to ($\mathrm{QSP_m}$) replacing $\tilde{c}_m(.)$ with her own cost function (or an approximation of it).

However, it is not realistic that the $\mathrm{QSP_m}$ as such could be applied into practice. It is possible that the bidders are not able to express their costs in a functional form. However, we do assume that the bidders are still capable of comparing the profitability of different bids. Also, it is unlikely that they would be willing to disclose their cost functions to the auction owner or even a neutral third party, even if they could specify the functions. Therefore we need to find a way to approximate the bidders' cost functions, and preferably without having to ask for information from the bidders.

In this study we used a linear approximation of the bidders' cost using the dual prices of the demand constraints of the WDP as the variable cost parameters. The dual prices can be interpreted as market prices for the items (see discussion on combinatorial auction mechanisms in section 3.2.3.2 of the literature review), hence they can be expected to reflect the underlying costs as well. Because in the integer programming case there are no dual prices, we used the dual prices of the linear relaxation. This means that the WDP (2) is solved again, but the binary constraints $x_{ij} \in \{0,1\}$ are replaced with $0 \leq x_{ij} \leq 1$. The cost function $\tilde{c}_m(Q_{m,new})$ in (5) is replaced with the linear cost function

$$\tilde{c}_m(Q_{m,new}) = \sum_{k=1}^{K} \mu_k q_{m,new,k} \tag{11}$$

where μ_k is the dual price of the k^{th} quantity constraint in the linear relaxation of the WDP, and thereby the dual price for item *k*.

According to economic theory, firms have two kinds of costs: variable and fixed. We have considered only the variable costs in our linear approximation. A fixed cost term could easily be added, but as it is a constant, it would not affect the solution of the maximization problem. The lack of the fixed cost element affects the value of the objective function, but in this problem the approximated profit indicated by the objective function is not interesting – only the allocation and bundle price are.

The resulting linear cost function has several limitations that need to be taken into account. First of all, it cannot portray economies of scale. Thus, in case the bidders experience economies of scale, the bidders' costs for large bundles are systematically overestimated. This would mean that there would be a bias in the QSP towards smaller bids. However, since the objective is to maximize absolute profit (not relative), we expect the very small bids to be eliminated anyway. Secondly, a linear cost function is incapable of portraying economies of scope (= subadditive cost function). Because of that the QSP has no incentive to try to bundle many items into the same bid. It is the existence of economies of scope, which was the reason to organize a combinatorial auction in the first place, hence we should assume the bidders' true costs to exhibit economies of scope. Therefore, one purpose of this study is to determine, whether a linear approximation of the cost functions is good enough.

Thirdly, notice that the linear cost function estimate (11) is the same for all bidders, hence it does not account for individual differences in the bidders' cost functions. Thereby the bid solutions of QSP are anonymous (i.e. the QSP gives the same suggested bids regardless of the bidder). The only differences in the suggestions can arise from the requirement that maximum one bid per bidder can be active at a time. Thus, the bid suggestion for bidder m cannot be such that it would team up with bidder m's previous bids, whereas for any other bidder these bids can be teamed-up with. We also recognize that there are differences between the bidders' cost structures, so the "anonymous" suggestion offered by the QSP may not be the best for all bidders.

Due to these limitations, it is possible that the solution of the QSP is not an acceptable bid suggestion for any of the bidders. Therefore a shortlist of alternative solutions is generated, and the shortlist alternatives are presented to the bidders together with the solution from the QSP, and the bidder can choose the bid which is the most profitable one for her. This can be done in two ways. First, we can go through the neighboring pivots of the original quantity support problem. Pivoting in the integer case is interpreted as solving the QSP over and over again, but each time forcing one status variable x_{ij} with a zero value in the original QSP solution to assume the value "one". Hence, the number of pivots depends on the number of bids in the bid stream (as there is one status variable for each bid), and the number of bids in the optimal combination (those which assume the value "one" in the original QSP solution). The second alternative is to solve the QSP over and over again, but this time setting different Q_i's to zero to generate new combinations. If all the possible combinations are searched through, there are 2^K-2 combinations to go through, so in larger auctions it would not be feasible. Oftentimes, though, the different pivots, as well as the different combinations produce the same bid suggestion, so the number of non-identical items on the shortlist hardly ever reaches the theoretical maximum. The QSP together with the shortlist forms the core of the bidder support tool we call the QSM.

7.2 An Example of a Combinatorial Auction with the QSM

Consider a procurement situation, where a single buyer desires to buy a bundle of items: 100 units of item A, 100 units of item B, and 200 units of item C. Her reservation price for the entire package is \$100,000. If there were other attributes, they have been "priced out". Let us further assume that there are three bidders: X, Y, and Z. Finally, assume that the desired price decrement δ is equal to \$3,000 from one bid to the next.

Assume bidder X makes her first bid of 50 units of item A and 100 units of item B (but no units of C) for a total bid price of \$25,000. In vector notation, this bid is (50; 100; 0; \$25,000). Naturally, since this was the first and only bid so far, it is inactive (i.e. not among the provisional winners), but it is entered into the bid stream. Next, assume that bidder Y enters the following bid (100; 50; 200; \$79,000). The bidder is informed that her bid is not a provisional winner. Considered jointly with the bid of bidder X, they would meet the quantity demand of the buyer. In fact they would exceed the demand for items A and B, since bid X_1 + bid Y_1 = (150; 150; 200; \$104,000). That would be acceptable, except that the buyer's reservation price (\$100,000) is exceeded, making the combination of bid X_1 and Y_1 not feasible. However, as previously, we retain bid Y_1 in the bid stream. Next, bidder Z enters the following bid (100; 100; 0; \$32,000). This bid is also inactive. Bidder Z is informed. The other two earlier bids remain in the bid stream with inactive status.

Assume that bidder X decreases the price on her original bid to \$21,000. Now together the bids of bidders X and Y become provisional winners. They are notified of their changed status. Note that the new bid X_2 contains the same quantities as X_1 but for a lower price. Thus we can drop X_1 from the bid stream because each bidder can have only one bid active simultaneously.

Bidder Z decides to use the quantity support tool to find an "active" bid. First, in order to obtain the dual prices for the quantity constraints, we formulate and solve the LP relaxation of the WDP:

$$
\begin{aligned}
\min\ & 21{,}000x_2 + 79{,}000y_1 + 32{,}000z_1 \\
s.t.\ & 50x_2 + 100y_1 + 100z_1 \geq 100 \\
& 100x_2 + 50y_1 + 100z_1 \geq 100 \\
& 200y_1 \geq 200 \\
& 0 \leq x_2 \leq 1 \\
& 0 \leq y_1 \leq 1 \\
& 0 \leq z_1 \leq 1
\end{aligned} \tag{12}
$$

where x_2, y_1, and z_1 are the bid status variables. The solution of the problem is $x_2 = 0.5$, $y_1 = 1$, $z_1 = 0$. The dual prices are 0, 210, and 342.5 for the three quantity constraints. Using these dual prices as the coefficients for the linear cost function we can formulate the following quantity support problem. Denote the unknown price by p_{new} and the unknown quantities by $q_{new,A}$, $q_{new,B}$, and $q_{new,C}$. Here we assume that any bidder can have only one active bid, so bid Z_1 is deleted from the quantity support formulation. The QSP_z can then be formulated as

$$
\begin{aligned}
\max\ & p_{new} - 210q_{new,B} - 342.5q_{new,C} \\
s.t.\ & 21{,}000x_2 + 79{,}000y_1 + p_{new} \leq 97{,}000 \\
& 50x_2 + 100y_1 + q_{new,A} \geq 100 \\
& 100x_2 + 50y_1 + q_{new,B} \geq 100 \\
& 200y_1 + q_{new,C} \geq 200 \\
& x_2, y_1 \in \{0,1\}
\end{aligned} \tag{13}
$$

The suggested bid is (100; 100; 200; \$97,000) with $x_2 = 0$ and $y_1 = 0$. The shortlist is generated by forcing the bid status variables of inactive bids to assume the value "one" in turn. Problem (13) is solved twice, first with the additional constraint $x_2 = 1$, and then with $y_1 = 1$. The shortlist consists of the following three bids: (100; 100; 200; \$97,000), (50; 0; 200; \$76,000), (0; 50; 0; \$18,000).

Bidder Z decides to accept the second bid from the shortlist, which is added to the bid stream. The new bid Z_2 becomes a provisional winner together with X_2; Y_1 becomes inactive. The bidders X and Y are informed.

Next bidder Y requests a suggested bid. The linear relaxation of the WDP (similar to (12)) is solved to obtain the dual prices. They are: 0, 210, and 380. The quantity support problem is formulated as in (13), but this time Y_1 is left out, and X_2, Z_1, and Z_2

are included. The total cost to the buyer cannot exceed \$94,000. The suggested bid is (50; 0; 200; \$73,000) with $x_2 = 1$ and $z_1 = z_2 = 0$. The following shortlist is generated: (50; 0; 200; \$73,000), (0; 0; 200; \$62,000), (50; 100; 0; \$18,000).

Bidder Y decides to accept bid Y_2 = (0; 0; 200; \$62,000), which is then added to the bid stream. Bid Y_2 teams up with Z_1, and Bidder X becomes inactive. The bidders are informed about their new status. The auction continues along these lines until no bidder is willing to place a new bid.

8 TESTING THE QUANTITY SUPPORT MECHANISM: SIMULATION STUDIES

The contribution of this chapter is the analysis the results of the simulation studies conducted to test the QSM, and the insights and deepened understanding of combinatorial auctions obtained through the analysis. Especially the insights from the second simulation study were important and led to new ideas and significant improvements on the efficiency of the quantity support tools.

Simulations are a convenient way test solution algorithms and auction mechanisms. Many researchers have used them in their research (e.g. Kelly and Stenberg, 2000, Parkes, 2001, and Sandholm et al., 2005). The advantage of simulations over laboratory experiments – which are another common research method – is that they are faster and cheaper to set up, and it is much easer to test large-scale (say 30 bidders) auctions through simulations. Therefore, we chose to test the QSM first with simulations. The laboratory experiments with human subjects are presented in Chapter 12.

Two separate simulation studies were conducted to test the properties of the quantity support mechanism. The auctions in both studies are reverse auctions. The mechanism used is a first-price, semi-sealed auction, as used in the example auction in the previous chapter.

The first study compares the QSM using dual prices in the objective function with a quantity support mechanism, where the dual prices have been replaced with random coefficients. The purpose of this study is to obtain validation for the use of dual prices. The second simulation study studies the convergence properties of an auction where the QSM is used. The main point of interest is the final auction outcome, i.e. what is the total cost to the buyer, and how the items are allocated to the bidders.

8.1 First Simulation Study

The first simulation study consisted of three phases. In the first phase all the auction parameter values were chosen and cost functions were created for the bidders. Also a set of initial bids had to be generated into the bid stream because the QSM cannot be used before there are some bids to team up with. In the second phase, all bidders used the quantity support mechanism, and the most profitable bid from the shortlist for each bidder was recorded. In this simulation study we were only looking at one step in the auction, so no further bids were considered. In the third phase, the simulation results were compared against two benchmark cases. In the first benchmark case, the QSM was modified so that the dual prices in the objective function were replaced with random numbers. This represented the "no information" case. In the second benchmark case the approximate cost function is replaced by each bidder's true cost function in turn. This benchmark represented the ideal "perfect information" case, and constituted the largest possible profit obtainable for any bidder at that point in the auction. The profits obtained by bidders in phase two were then compared against the two benchmark cases.

8.1.1 First Phase: Setting Up the Auctions

In any simulation study, some initial values have to be assumed. In this study we needed to choose the number of bidders, number of items and the quantities demanded of each item. The bidders needed to be assigned cost functions so that it became possible to evaluate the profitability of the shortlist items offered by the QSM. Also, a set of initial bids from the bidders had to be generated, because the QSM cannot be used without some bids already in place in the bid stream.

8.1.1.1 The Cost Function

We wanted the cost function to be as simple as possible, but also we wanted it to portray both economies of scale and scope. The use of combinatorial auctions is justified in a situation in which there are synergies between the items. In reverse auctions, synergies between items can be understood as synergies in the production process of the items, i.e. economies of scope. Thus, it would not make sense to use a cost function that would not allow for synergies in production.

The simplest form for a cost function exhibiting economies of scale is a linear function with a fixed cost element:

$$C(q_k) = F_k + c_k q_k \tag{14}$$

where F_k is the fixed cost, and c_k the per-unit cost of producing item *k* alone.

For the two-item case the cost function would take the form of

$$C(q_1, q_2) = F_{12} + c_1 q_1 + c_2 q_2 \tag{15}$$

The existence of economies of scope in this framework implies simply that the inequality of Equation (16) between the fixed cost parameters holds:

$$F_{12} < F_1 + F_2 \tag{16}$$

The above presented multi-product cost function (Equation (15)) is very simplistic. It is theoretically very restrictive, as it implies constant marginal costs and monotonically decreasing average costs. The function is discontinuous at points when the level of one or more outputs is zero, which makes it difficult to use in optimization problems. The function is easy and convenient to use only in situations with relatively few products. It can easily be seen that the number of different fixed cost elements increases rapidly as the number of products increases. In the case of two products, there are only three parameters F_1, F_2, and F_{12}. With three products there are seven parameters: F_1, F_2, F_3, F_{12}, F_{13}, F_{23},, and F_{123}, where F_{ij} indicates the fixed cost of producing goods *i* and *j*. When the number of products is increased to five there are already 31 fixed cost parameters, with six products there are 63, and with seven products 127 parameters. However, the simplicity of the function makes it intuitive and it is very flexible as it can represent economies of scope of different magnitudes between different items – and even diseconomies of scope between so items, if necessary. Thereby it is very appealing in a theoretical framework such as ours.

Another reason for our choice of cost function was that there are not very good alternatives available. Cobb-Douglas and CES (constant elasticity of scale) forms can be used for multi-product cost functions, but they would have to be linearized before they could be used in linear or integer programming problems. A commonly used form for the cost function is the translog cost function (see Equation (17)), which is a function of the output quantities (y_i) and input prices (p_j).

$$\ln C = \alpha_0 + \sum_1^n \alpha_i \ln y_i + \sum_1^m \beta_j \ln p_j + \frac{1}{2}\sum_1^n \sum_k^n \sigma_{ik} \ln y_i \ln y_k$$
$$+ \frac{1}{2}\sum_j^m \sum_h^m \gamma_{jh} \ln p_j \ln p_h + \sum_i^n \sum_j^m \delta_{ij} \ln y_i \ln p_j \quad (17)$$

The translog function is more versatile, as it does not force homogeneity or constant elasticity on the cost structure. Due to its flexibility, the translog cost function has been popular in empirical studies which aim at estimating real world cost functions for some firm or industry (see e.g. Murray and White, 1983, and Cho, 2003).

The translog cost function, however, is too complex for the purposes of this study. In a simulation study it would be good to minimize the number of parameters to choose. Thus, the inclusion of the input prices in the cost function is an unnecessary complication. We chose to use the simple form of the cost function presented in Equation (15). It is intuitive, easy to use, and sufficient for the purposes of this study, where we only want to find out whether the QSM works under some circumstances.

8.1.1.2 Parameters

In order to reduce the sensitivity of the results to the initial values, we decided to vary some of them. The number of bidders was fixed at 10, but the number of items to be auctioned was either 3 or 5. The quantity demanded was 1000 for each item. The variable cost parameters were drawn from the same uniform distributions in each design. In the three-item auctions the variable cost parameters were drawn from the range [30, 50] for c_1, [40, 60] for c_2, and within [60, 70] for c_3. For the five-item auction the variable costs fort items 1, 2 and 3 were drawn from the same range as in the three-item auction, and for the additional items from the range [15, 45] for c_4 and [20, 55] for c_5. The distributions for the variable cost parameters (and fixed cost parameters) were the same for all bidders, so I have omitted the index indicating the bidder from the notation.

The uniform distribution from which the fixed cost parameters (F_{ijk}) were drawn had two possible levels: "high" and "low". The ranges for the fixed cost parameters were chosen so that it was very likely that economies of scope would exist. This meant that the lower bound of F_{12} was less than or equal to the sum of the lower bounds of F_1 and

F_2. A similar logic was applied to the upper bounds. It was also kept in mind that total fixed cost should not decrease from the addition of a new product. Thus the lower bound of F_{12} was set higher than the upper bounds of F_1 and F_2.

When the costs were low, the lower bounds ranged from 700 (for F_3) to 2300 (F_{123}) and the upper bounds from 1000 (F_3) to 3200 (F_{123}) in the three-item auctions. Again, in the five-item auctions the parameter ranges were the same for the first three items. The lower bounds ranged from 700 (for F_3) to 5500 (F_{12345}) and the upper bounds from 1000 (F_3) to 7500 (F_{12345}). The exact ranges for all fixed cost parameters can be seen in Appendix 1. When the costs were high, the lower bounds for the fixed costs ranged from 5000 (F_3) to 42000 (F_{12345}) and upper bounds from 7000 (F_3) to 50000 (F_{12345}). The new ranges can also be seen in Appendix 1. The higher fixed cost values were used to test the effect of more pronounced economies of scope on the results. With the lower level of fixed costs, the proportion of fixed costs in the total cost was only about 10%. After the increase the proportion of fixed costs was almost 30%.

The initial bids in the bid stream were created so that each bidder was assumed to have placed one bid. Thus, the number of bids in the initial bid stream equaled the number of bidders. The quantities in the initial bids were drawn randomly from a uniform distribution [100, 500], and rounded to the nearest 50. Some of the bid quantities, however, were chosen to be zero in order to create some sparsity in the bid matrix. It is realistic to assume that all bidders would not place bids for all products but a subset of them. The level of sparsity was 20%. The constraint $q_{new,k} \leq 500$ was added to the quantity support problem to simulate the capacity constraints of the sellers. The capacity constraint was set to simulate the fact that no bidder alone would be capable of producing the whole demand. The bid price in the initial bid was the cost for the bidder of producing that specific bundle, to which an initial mark-up of 30% was added.

The simulation study consisted of four different experiment settings displayed in Table 4. Five replications of each setting were conducted.

Table 4 Design of the first simulation study

Experiment	Items	Bids	Fixed Cost
I	3	10	Low (~10%)
II	5	10	Low (~10%)
III	3	10	High (~30%)
IV	5	10	High (~30%)

8.1.2 Second Phase: Simulations

The simulations advanced as follows. First, the winner determination problem was solved for the initial bid stream. The solution of the WDP determined the first set of "active" bidders (or provisional winners), and the current lowest total cost for the buyer. The linear relaxation of the same WDP was also solved to obtain the dual prices to be used in the quantity support problem (QSP). Even though we knew the bidders' cost functions, we assumed that the auction owner solving the QSP would not know them. The decrement δ with which the total cost to the buyer must decrease in the new winning combination was set to 5% of the total cost in every simulation. The size of the decrement generally has an impact on the convergence of auctions (as discussed in section 5.1 on minor auction design issues). However, since we only study the first incoming bid – and not the convergence – the choice of the decrement does not have a major impact. Thus, different levels of the decrement were not tested.

With the decrement defined, the QSP could then be solved. The shortlist was compiled through solving the QSP again and again adding the constraint $x_i = 1$ for each original non-basic variable (all x_i for which $x_i = 0$ in the WDP) in turn. The profit for each bidder from each shortlist item was calculated as the difference between the suggested bid price and the production costs for the suggested quantities. The largest profit was recorded. In case all shortlist items produced a loss, the profit was set to zero indicating the fact that the bidder would choose not to bid anything. The shortlist items were evaluated also for the active bidders even though it would be more logical to evaluate them only for inactive bidders. The reasoning behind this decision was that active bidders may want to use the QSM to find out if they could obtain a bigger profit than the 30% mark-up in their initial bids. The largest absolute profit for each bidder was recorded[15].

8.1.3 Third Phase: Benchmarks

Using the same initial data (cost function parameters and initial bids), two benchmark cases for the QSM were solved. The first benchmark case repeated the procedure of phase two with the exception that instead of dual prices, randomly chosen cost parameters were used in the QSP. The parameters were chosen from a continuous uniform distribution with the range [0, 250]. This interval was chosen because the dual prices generally seemed to fall in the same range. The random parameters did not utilize any information available on the cost functions of the bidders, and therefore represented the extreme case of "no information".

In the second benchmark – the case of "perfect information" – the whole approximated cost function in the QSP was replaced by each bidder's true cost function in turn. The IP formulation of the profit maximizing problem is not trivial due to the discontinuous cost function. The formulation is presented in Appendix 2. In the perfect information case there is no shortlist to be created because the optimum is found directly for each bidder.

8.1.4 Results of the First Simulation Study

We were primarily interested in comparing the profits obtained by the bidders using quantity support to profits in the two benchmark cases. On the one hand we wanted to know, whether the QSM using dual prices performed better than the QSM with random cost parameters, and on the other we wanted to find out, how close to the best possible profit the QSM could get. Thus, when presenting the results I will mainly focus on these comparisons. However, as a sideline I have also looked at the bidders' mark-ups from the most profitable bids to see how they change from the 30% set in the initial bids. The mark-ups are considered only for the bids generated by the dual price QSM.

[15] Using the largest profit as bidders' decision rule is only one alternative. Other alternatives would be the largest mark-up (ratio of profit to total cost), or largest turnover (price) which is still profitable.

8.1.4.1 Profits

The most interesting aspect concerning the bidders' profits from the use of the QSM is to look at how close to the maximum profit (or perfect information case) we could get. Another point of interest is to compare the performance of random parameters (the no information case, or "random support") to the quantity support. The performance measure used to compare the outcomes was the percentage of maximum profit available that each bidder could get with both quantity support and random support:

$$performance_D = \frac{\pi_{iD}}{\pi_i^*} \tag{18}$$

where π_{iD} is the profit of the most profitable bid suggested to bidder i by the QSM using dual prices, and π_i^* the optimal profit of the perfect information case calculated using bidder i's true cost function. The performance of random support, $performance_R$, is calculated by substituting into the numerator of (18) the profit of the most profitable bid suggested to bidder i by the QSM using random cost parameters.

The performance indicators for QSM with dual prices (D) and "random support" (R) for the first three experiments (see Table 4 for the details of experiment designs) are presented in Table 5 (Experiment I), Table 6 (Experiment II) and Table 7 (Experiment III).

Table 5 Results of Experiment I: Performance indicators of the dual price quantity support (D) and random support (R)

Replication		Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
1	D	0,86	0,87	0,76	0,14	1,00	0,95	0,87	0,95	0,76	0,00
	R	0,59	0,57	0,39	0,33	0,72	0,62	0,66	0,36	0,61	0,33
2	D	0,88	1,00	0,87	0,88	0,92	1,00	0,86	0,00	0,90	0,91
	R	0,27	0,55	0,68	0,59	0,60	0,58	0,66	0,67	0,07	0,55
3	D	0,00	0,60	0,73	0,83	0,74	0,79	0,00	0,84	0,82	0,75
	R	0,82	0,72	0,74	0,73	0,75	0,74	0,93	0,77	0,73	0,75
4	D	0,92	0,88	0,00	0,79	0,91	0,86	0,42	0,88	0,95	0,91
	R	0,42	0,41	0,12	0,40	0,44	0,38	0,41	0,43	0,41	0,26
5	D	0,76	0,83	0,92	0,00	0,87	0,83	0,66	0,87	0,86	0,00
	R	0,65	0,76	0,82	0,85	0,75	0,16	0,72	0,80	0,83	0,00
# cases where D≥R		4	4	3	3	4	5	3	4	5	4

Table 6 **Results of Experiment II: Performance indicators of the dual price quantity support (D) and random support (R)**

Replication		Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
1	D	0,78	0,73	0,78	0,49	0,56	0,72	1,00	0,79	0,73	0,75
	R	0,38	0,26	0,15	0,00	0,33	0,21	0,22	0,24	0,40	0,29
2	D	0,90	0,74	0,80	0,78	0,74	0,00	0,77	0,80	0,82	0,80
	R	0,73	0,62	0,57	0,00	0,67	0,00	0,69	0,60	0,00	0,58
3	D	0,85	0,65	0,80	0,00	0,85	0,78	0,76	0,00	0,84	0,86
	R	0,24	0,20	0,26	0,00	0,00	0,23	0,30	0,00	0,38	0,35
4	D	1,00	0,86	1,00	1,00	1,00	0,00	0,94	1,00	1,00	0,00
	R	0,34	0,28	0,29	0,05	0,30	0,00	0,31	0,39	0,39	0,44
5	D	0,98	0,76	0,84	0,76	0,00	1,00	0,91	1,00	0,47	0,00
	R	0,34	0,34	0,42	0,00	0,37	0,17	0,36	0,37	0,33	0,11
# cases where D≥R		5	5	5	5	4	5	5	5	5	3

Table 7 **Results of Experiment III: Performance indicators of the dual price quantity support (D) and random support (R)**

Replication		Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
1	D	0.50	0.55	0.47	0.45	0.48	0.00	0.44	0.46	0.48	0.51
	R	0.26	0.23	0.23	0.24	0.25	0.00	0.31	0.30	0.32	0.51
2	D	0.98	0.53	0.68	0.55	0.00	0.73	0.62	0.65	0.66	0.73
	R	0.98	0.50	0.68	0.55	0.31	0.73	0.62	0.65	0.66	0.73
3	D	0.52	1.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	R	0.06	0.22	0.28	0.17	0.24	0.25	0.19	0.28	0.18	0.28
4	D	0.63	0.57	0.59	0.54	0.59	0.62	0.57	0.31	0.64	0.53
	R	0.67	0.55	0.66	0.61	0.40	0.58	0.49	0.65	0.57	0.64
5	D	0.83	0.84	0.73	0.89	0.85	0.86	0.68	1.00	0.87	0.83
	R	0.68	0.76	0.58	0.73	0.61	0.00	0.68	0.79	0.81	0.71
# cases where D≥R		4	5	3	4	4	5	5	4	5	4

The results were very promising in general. The QSM using dual prices categorically found profits that were above 70 % of the maximum, as can be seen from the tables. Also, in some cases the bidder could actually obtain the maximum profit with the help of the dual QSM (indicated by "1" in the tables). This means that the optimal bid for the bidder in question was on the shortlist. The profits obtained with the random approach were often lower, and varied more. In pairwise comparisons, the dual price approach (D) performed as well as or better than the random cost approach (R) 86% of the time. I also tested the statistical significance of the results using a pairwise t-test. The p-values (two-tailed) for the three respective tables were highly significant (0.007, 0.0000, 0.0001).

In Experiment IV (5 products, 10 bids and higher fixed costs), the quantity support tool using dual prices did not perform as well as the one with random cost parameters. The performance indicators for dual price quantity support and random support in Experiment IV can be seen in Table 8.

Table 8 **Results of Experiment IV: Performance indicators of the dual price quantity support (D) and random support (R)**

Replication		Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
1	D	0,70	0,50	0,67	0,69	0,69	0,69	0,60	0,70	0,74	0,67
	R	0,97	1,00	1,00	1,00	1,00	1,00	1,00	1,00	0,39	0,82
2	D	0,47	0,00	0,72	0,77	0,76	0,79	0,69	0,87	0,70	0,69
	R	0,90	0,74	0,81	0,88	0,88	0,84	0,65	0,00	0,81	0,74
3	D	0,00	1,00	0,85	0,00	1,00	0,91	*	1,00	0,92	0,72
	R	0,46	0,57	0,33	0,52	0,65	0,58	*	0,55	0,61	0,61
4	D	0,87	0,86	0,70	0,00	0,87	0,84	0,90	0,00	0,87	0,87
	R	0,81	0,96	0,89	1,00	0,95	0,88	0,98	0,00	0,99	0,96
5	D	0,33	0,14	0,00	0,34	0,34	0,47	0,08	0,34	0,00	0,32
	R	0,71	0,87	0,67	0,68	0,29	0,82	0,66	0,68	0,70	0,72
# cases where D≥R		1	1	1	0	2	1	1	3	2	1

* Lindo could not find a solution

The better performance of the random support is also statistically significant (p-value 0.003). However, as can be seen from Table 8, the dual price approach is still oftentimes reasonably good generating profits well above 60% of the maximum. The random parameters simply worked even better.

One must keep in mind here, though, that the random parameters were chosen from the interval within which the dual prices varied. In reality, only one or the other approach would be used. So, if we were using random parameters alone, we would not know the range within which the dual prices varied. We reproduced Experiment IV with random cost parameters, but this time chosen from the range [0, 2000]. The random parameters worked poorly producing very small profits categorically (see Table 9).

Table 9 **Results of Experiment IVb: Performance indicators of the random support (R) with parameters drawn from the range [0, 2000]**

Replication		Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
1	D	0.70	0.50	0.67	0.69	0.69	0.69	0.60	0.70	0.74	0.67
	R	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	D	0.47	0.00	0.72	0.77	0.76	0.79	0.69	0.87	0.70	0.69
	R	0.03	0.00	0.12	0.11	0.12	0.05	0.11	0.06	0.11	0.11
3	D	0.00	1.00	0.85	0.00	1.00	0.91	*	1.00	0.92	0.72
	R	0.00	0.02	0.21	0.00	0.25	0.04	*	0.25	0.22	0.23
4	D	0.87	0.86	0.70	0.00	0.87	0.84	0.90	0.00	0.87	0.87
	R	0.42	0.40	0.25	0.00	0.42	0.28	0.45	0.00	0.00	0.43
5	D	0.33	0.14	0.00	0.34	0.34	0.47	0.08	0.34	0.00	0.32
	R	0.36	0.00	0.36	0.01	0.02	0.44	0.35	0.35	0.01	0.35
# cases where D≥R		4	5	4	5	5	5	4	4	4	4

* Lindo could not find a solution

The performance of the random parameters approach understandably seems to depend on the chosen range of the cost parameters. The approach of generating random parameters is beautiful in its simplicity, but their use is complicated by the fact that an unsuitable range can significantly deteriorate the results. Thereby, we feel that using dual prices is a better and more robust approach. It is also fairly simple, even though an additional (linear) optimization problem has to be solved each time someone wishes to use the quantity support tool.

8.1.4.2 Mark-Ups

An interesting detail related to the mark-ups (profit divided by production cost) in the bidders' bids can be observed in the results. Remember that the mark-up percentage in the initial bids was set at 30%. One would assume that the mark-ups would have to decrease in order for the new bids to become active. Interestingly, the profits generated by the dual price QSM resulted in higher than 30% mark-ups in over 80% of the cases. The mark-ups in the four experiments are presented in Tables 10-13. The mark-ups correspond to the dual price QSM bids for which the profit ratios are presented in Tables 5-8. In some cases the profits were actually larger than costs, i.e. mark-up was over 100%. This demonstrates the benefits that can be obtained through searching for combinations that team up well together. When the auction progresses, the mark-ups will naturally have to start to decline.

Table 10 Mark-ups in Experiment I

	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
Replication 1	0.29	0.31	0.41	0.08	0.35	0.36	0.31	0.33	0.32	0.00
Replication 2	0.37	0.54	0.51	0.46	0.44	0.43	0.44	0.00	0.42	0.37
Replication 3	0.00	0.42	0.50	0.56	0.51	0.34	0.00	0.49	0.60	0.46
Replication 4	0.52	0.71	0.00	0.62	0.64	0.80	0.59	0.62	0.50	0.65
Replication 5	0.44	0.47	0.48	0.00	0.44	0.55	0.42	0.53	0.58	0.00

Table 11 Mark-ups in Experiment II

	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
Replication 1	0.36	0.36	0.31	0.46	0.55	0.45	0.39	0.42	0.29	0.31
Replication 2	0.42	0.68	0.38	0.43	1.11	0.00	0.50	0.64	0.35	0.49
Replication 3	0.51	0.53	0.77	0.00	0.99	0.65	0.62	0.00	0.35	0.93
Replication 4	1.05	0.89	0.98	1.02	1.14	0.00	0.86	1.10	1.11	0.00
Replication 5	0.82	0.59	0.58	0.59	0.00	0.85	0.63	0.70	0.24	0.00

Table 12 Mark-ups in Experiment III

	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
Replication 1	0.80	0.88	0.47	0.59	0.75	0.00	0.66	0.70	0.86	0.33
Replication 2	0.44	0.33	0.54	0.40	0.00	0.43	0.45	0.46	0.45	0.46
Replication 3	0.31	0.53	0.00	0.59	0.45	0.65	0.61	0.51	0.45	0.58
Replication 4	0.38	0.18	0.31	0.27	0.22	0.27	0.24	0.38	0.29	0.18
Replication 5	0.48	0.47	0.46	0.42	0.51	0.55	0.45	0.48	0.54	0.55

Table 13 Mark-ups in Experiment IV

	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
Replication 1	0.59	0.36	0.68	0.39	0.67	0.55	0.64	0.73	0.41	0.60
Replication 2	1.06	0.00	1.09	0.96	1.08	1.00	1.04	0.95	1.03	0.97
Replication 3	0.00	0.67	0.69	0.00	0.63	0.64	0.63	0.67	0.56	0.47
Replication 4	0.93	0.92	1.65	0.00	1.08	0.82	1.21	0.00	1.07	0.82
Replication 5	0.41	0.26	0.00	0.44	0.35	0.44	0.26	0.42	0.00	0.35

8.2 Second Simulation Study[16]

After concluding in the first simulation study that the dual price quantity support was helpful, and better than random support, we wanted to test the QSM further. In the first study we only considered the first bid added to the initial bid stream with the help of the QSM. While that is an indicator of the QSM's ability to find good complementing bids, it still does not tell us anything about the final outcome of the auction. Thus, in the second simulation study we wanted to look at the auction outcomes. Our primary interest was in finding out what the buyer's final cost ended up being, and how close to an efficient allocation we could get with the help of the QSM. In addition, we wanted to examine more thoroughly the sensitivity of the auction outcome to the chosen input parameters. Our secondary interest was in studying the speed of convergence, i.e. how long it took to reach the final allocation. The convergence speed tells us the time the participants have to spend in the auction, and is therefore an indication of the time cost involved. A mechanism that ultimately results in an efficient or nearly efficient outcome may not be popular, if it takes too long to reach to outcome. Therefore, the length of the auctions should not be ignored in the testing phase.

The auction setting and mechanism in the second auction study was essentially the same as in the first one. The second simulation study resembled the first one also in the sense that it also consisted of three phases: defining the input parameters, the actual simulation phase, and defining (and calculating) the benchmark cases. The biggest difference between the two studies was naturally the fact that in the second study the auctions are run all the way to the closing, but there were also some differences in the

[16] Material presented in this section is based on joint work with Valtteri Ervasti (Ervasti and Leskelä, 2009, forthcoming in European Journal of Operational Research).

values of the simulation parameters, choice of parameters to vary, and in the way the shortlist was compiled.

8.2.1 Phase One: Parameters

The logic by which the input parameters were created in the second study resembled the first simulation study in many respects. The demand was the same across items; $d_k = 600$ for all $k = 1, \ldots, 5$. The number of bidders was either 15 or 30. The bidders' maximum capacity was again 50% of total demand, i.e. 300. However, this time we also tested the effect of unequal capacities. In the unequal capacities case bidders had a maximum capacity a_{ik} of 300 (with 50% probability), 225 (with 25% probability) or 150 (with 25% probability). The bidders' capacities were defined separately for each item. Thus, the bidders' capacities varied from item to item. In the second simulation we tried to improve the robustness of our results by increasing the number of replications to 50. The buyer required that the total cost reduces by 2% each iteration. Thus, δ in the QSP is $0.02C^{*}$. The size of the decrement affects the auction in two main ways. The smaller the δ, the closer to the efficient production cost the total cost in the auction can potentially go. However, the smaller the decrement, the slower is the convergence of the auction. We chose 2% so that the auction would take long enough that there would be relatively many bids in the auction for the QSP to use, and still not have the auction take too long. Also, changing the decrement to 1% did not change the results significantly (basically, the auction only continued one step further), hence we decided not to vary the decrement in our final design.

Thus, the first two simulation design parameter values varied in the experiment were number of bidders (15 or 30), and the bidders' capacities (equal and unequal).

8.2.1.1 Cost Function

We used the same form for the cost function as in the first study (see Eq. (15)). Thus, the cost function limited the number of items that could be handled easily. We chose to use 5 items in every simulation. We made the choice based on pilot studies, which indicated that changing the number of items from 3 to 4 or 5 did not change the results significantly.

The cost function parameters were chosen again from uniform distributions, and the distributions were the same for all bidders (i.e. symmetric situation). Thus, in this section I have omitted the index *i* indicating the bidder from the notation. The variable cost parameters c_k were chosen from the uniform distribution [53.33, 66.67], with $E(c_k) = 60$ for all five items. The fixed cost parameters had again two different levels, but this time the difference was in the degree of economies of scope, and not in the proportion of the fixed costs of total cost. We denote these two cases as "normal" economies of scope, and "large" economies of scope. These labels should not be taken too literally, though. For simplicity, all items were treated identically, i.e. the distribution for the fixed costs was dependent only on the number of items in the combination. Normal economies of scope for a combination *L* was defined as

$$E(F_L) = E(F_{L\setminus k}) + 0.5E(F_k) \tag{19}$$

so the inclusion of an additional item increases the expected fixed cost by 50% (and not 100% as in a linear case). The "large" economies of scope were defined in a similar fashion as "normal" economies of scope, but the multiplier used was 0.4 instead of the 0.5 used in (19). The expected value for the single-item fixed costs was set so that the expected proportion of fixed costs over total cost in a single-item bid at full capacity (= 300) would be 50%. The expected value of variable costs is $60 \times 300 = 18{,}000$. Thus, $E(F_k) = 18{,}000$. The spread of the distribution around the mean was designed so that the upper limit would be 25% higher than the lower limit. This translates to a maximum of 25% cost advantage of one bidder over another. The ranges for the fixed cost parameters with "normal" and "large" economies of scope are presented in Table 14 and Table 15.

Table 14 Ranges for fixed cost parameters in "normal" economies of scope

L	Lower Limit	Mean	Upper Limit
1 item	16,000	18,000	20,000
2 items	24,000	27,000	30,000
3 items	32,000	36,000	40,000
4 items	40,000	45,000	50,000
5 items	48,000	54,000	60,000

Table 15 Ranges for fixed cost parameters in "large" economies of scope

L	Lower Limit	Mean	Upper Limit
1 item	16,000	18,000	20,000
2 items	22,400	25,200	28,000
3 items	28,800	32,400	36,000
4 items	35,200	39,600	44,000
5 items	41,600	46,800	52,000

Thus, the third simulation design parameter that was varied in the study is the size of the economies of scope (normal or large).

8.2.1.2 Initial Bids

In the second simulation study we wanted to generate the initial bids "intelligently". In the first study we simply drew the bid quantities randomly form a uniform distribution, and then randomly forced some quantities to zero. Intuitively, it would make more sense to have the bid quantities depend on the cost functions. Bidders in any auction should have some idea on their cost level or the general cost level of the industry even if they do not have exact information. We designed three different ways to generate the bid quantities: one based on the evaluation of variable costs ("Bid1"), one based on fixed costs ("Bid2"), and one based on the combination of fixed and variable costs ("Bid3"). In each case, the price for each bid was obtained by calculating the production cost for the bundle, and adding a 20% mark-up on top of the cost.

In Bid1, bidders' variable costs were compared to the expected value of the distribution of the variable costs. If $c_{ik} < E(c_{ik})$, the bid quantity q_{i1k} (quantity of item *k* in bidder *i*'s first bid) was set to the bidder's maximum capacity of item *k*. Otherwise $q_{i1k} = 0$. For example, assume that the variable costs of bidder B1 for the five items were (55.9, 56.9, 61.0, 63.2, 57.3) respectively. The expected value of the distribution, $E(c_{ik}) = 60$. Thus, bidder B1's variable costs are less than average for items 1, 2, and 5. If B1's capacities were (300, 150, 150, 225, 300), the bidder's initial bid quantities according to Bid1 would be (300, 150, 0, 0, 300).

In Bid2, a similar comparison was done with the fixed cost parameters. All the 31 fixed cost parameters were compared to their expected values, and the one that was proportionally most below the expected value was chosen. The bid quantities of the items corresponding to the most advantageous fixed cost combination were set to maximum capacity. Other q_{i1k} = 0. For example, bidder B1's fixed cost for items 2, 4 and 5 were 33,024 under normal economies of scope, the ratio to the expected fixed cost (= 36,000 for three items) would be 0.917. If this were the smallest ratio among all 31 comparisons, the bidder's initial bid quantities according to bid2 (assuming same capacities as in the previous example) would be (0, 150, 0, 225, 300).

In Bid3, the comparisons considered the total production costs (fixed cost + variable cost). Using the cost function parameters c_{ik} and F_{iL}, we calculated the production cost for each item combination (in total 31 combinations including the single items) at maximum capacity of bidder i. The reference point was the expected total cost of producing each of the 31 combinations at maximum capacity of bidder i. More formally, the reference point is:

$$E(C_{i,L}) = E(F_L) + \sum_{k \in L} E(c_{ik}) a_{ik} \tag{20}$$

where $E(F_L)$ is the expected fixed cost of combination L (see Table 14 and Table 15), a_{ik} is the maximum capacity of bidder i for item k, and $E(c_{ik}) = 60$ for all i and all k. The bidders' actual fixed costs and variable costs are substituted in (20) to obtain the actual cost, which is then compared to (20). The combination for which the ratio of actual total cost to the expected cost, $C_{i,L} / E(C_{i,L})$, was the lowest, was chosen. Again, the quantities in the most advantageous combination were set to maximum capacity, and other quantities were set to zero. Continuing the example from above, assume that B1's actual fixed cost $F_{1,245}$ = 34,687 and variable costs the same as earlier. Then B1's actual total cost for items 2, 4 and 5 would be 34,687 + 56.9×150 + 63.2×225 + 57.3×300 = 74,632. expected cost for the bidder B1's bid for items 2, 4 and 5 at full capacity would be 36,000 + 60 ×(150+225+300) = 76,500. The ratio of these two costs is 0.976, and if it were the smallest ratio among all 31 ratios, bidder B1's initial bid quantities would again be (0, 150, 0, 225, 300).

Thus, fourth simulation design parameter to be varied in the study was the way in which the initial bid stream was created (Bid1, Bid2, Bid3).

8.2.2 Phase Two: Simulations

For the simulation phase, the auction process is divided into four steps and presented as an algorithm:

Step1: Winner Determination. The WDP (Eq.(2)) is solved for the set of bids in the bid stream (in the beginning the set of bids in the bid stream are the initial bids). Based on the solution, a set of bidders become provisional winners, and the rest are inactive. Proceed to Step2.

Step2: Quantity Support Mechanism. The WDP of Step1 is solved again in the linearized form creating a vector of dual prices. The QSP presented in section 7.1 is solved using these dual prices as variable costs in the objective function.

Next, the shortlist is created. In this study, we used two different ways to create the shortlist. The one called "full" shortlist contains the original solution of the QSP, and the solutions of the QSP with one of the following additional constraints in turn:

$$\begin{aligned} & x_{ij} = 1 \text{ in turn for all } x_{ij} = 0 \text{ in } QSP_0 \\ & q_k = 0 \ \ \forall \ k \in L \text{, for each subcombination } L \text{ of the set of all items } M \end{aligned} \tag{21}$$

The "express" shortlist only contains constraints of the latter type. The shortlist type was the fifth (and last) simulation design parameter to be varied in the study. Proceed to Step3.

Step3: Selection of Bidder and Bid. From the set of inactive bidders determined in Step1 a bidder i is chosen randomly as the bidder "using" the quantity support. The cost function of bidder i is used to evaluate the bids on the shortlist. The bid for which the profit is the largest – i.e. $\max_{p_r, q_r} p_r - (F_{i,L} + \sum_{k \in L} c_{ik} q_{k,r})$, where L is any subset of items, p_r is the price of the r^th^, shortlist item, and $q_{k,r}$ is the quantity of item k in the r^th^ shortlist item – is selected. If profit is nonnegative[17], the bidder places the bid, and it is added into the bid stream. Move back to Step1. If all profits are negative, another bidder from

[17] We used zero profit as the limit for an acceptable bid. The lowest acceptable profit is only a matter of normalizing, and it is customary in economics to assume that "normal" profit is already included in the cost function (cf. "economic cost").

the set of inactive bidders is chosen, and she evaluates the short list items. If no bidder has accepted the suggested bids on the shortlist, and there are no inactive bidders left, move to Step4.

Step4: Auction Ends. No inactive bidder accepted a bid suggestion made by the QSM. Note that we assumed that after the initial bids are entered, no new bids are entered outside the suggestions of shortlist. Thus, no new bids enter the auction, and the auction ends. The provisional winners from Step1 become the actual winners, and the total cost to the buyer equals the value of the objective function in the solution of the WDP.

8.2.3 Benchmark Cases

In the second simulation study, we used three benchmark cases. As in the first study, we had one benchmark case to represent the situation with less information (and support), and one to represent the case of perfect information (the “first best” solution). Also, as in the first simulation study, the “less information” case is studied as an alternative to the QSM, and the “perfect information” case is what the two other cases are compared against. The third benchmark case we used was an auction in which the items were auctioned off individually, i.e. bids on combinations of items could not be placed (the “non-combinatorial” case).

The case of less information means that the QSM is not available for the bidders. Instead, in Step 2, they can only use the “suggested price” tool proposed by Teich et al. (2001 and 2006). Teich et al. have presented the suggested price tool in the context of multi-attribute auctions, but the idea is directly transferable to the multiple-item setting. The suggested price tool in a combinatorial auction returns the price that will make a given item combination a provisional winner. The bidder needs to specify the quantities b_k of each item beforehand. The suggested price problem for bidder m (SPP_m) can also be formulated as an integer programming problem:

$$
\begin{aligned}
&(\mathrm{SPP_m}) \quad \max \quad p_{new} && (22)\\
&s.t. \quad \sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij} p_{ij} + p_{new} \le C^* - \delta && (23)\\
&\sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij} q_{ijk} + b_k \ge d_k \qquad \forall\ k = 1,\dots,K && (24)\\
&x_{mj} = 0 \qquad \forall\ j = 1,\dots,n_m && (25)\\
&\sum_{j=1}^{n_i} x_{ij} \le 1 \qquad \forall\ i = 1,\dots,N && (26)\\
&x_{ij} \in \{0,1\}
\end{aligned}
$$

We are assuming that the auction mechanism allows at most one bid from each bidder to be active, and thus the bid status variables x_{mj} for the bidder's earlier bids are forced to zero. The suggested price tool does not convey any information the bidder could not obtain by herself with trial and error. However, it would be very impractical to have to place a large number of bids with decreasing prices to discover the price that would make the bid active. Thus, the suggested price tool is a welcomed convenience that expedites the auction.

What are then the quantities b_k that the bidder enters into the suggested price tool in the simulation? In our simulations we assume that the bidder anchors on her initial bid. She first asks for the suggested price for the initial bid, and then adjusts the bid quantities and asks for new suggested prices. The additional combinations she considers are all the subsets of the original bid (on item level), and a bid that is 50% of the original bid. For example, assume that the bidder's initial bid had the quantities [0, 300, 150, 0, 225]. She would then ask for a suggested price for the following set of bids:

[0	300	150	0	225]
[0	0	150	0	225]
[0	300	0	0	225]
[0	300	150	0	0]
[0	0	0	0	225]
[0	0	150	0	0]
[0	300	0	0	0]
[0	150	75	0	112.5]

In Step3 the bidder evaluates the resulting price-quantity combinations against her cost function, and places a new bid, provided that at least one combination produces a nonnegative profit.

The "perfect information" case in the second study was the efficient allocation of the items between the bidders. By efficient I mean the allocation that minimizes the total production cost of the bundle demanded by the buyer. When solving for the efficient allocation, the bidders' cost functions were assumed to be known, so the true optimum could be discovered. The conceptually simple optimization problem is in fact quite difficult (and lengthy) to formulate due to the discontinuous cost function, so the formulation is presented in Appendix 3. The formulation is based on the same logic as the optimum quantity support bid calculation in the first simulation study (Appendix 2).

In the non-combinatorial case we solved for the efficient allocation in a situation in which each item was auctioned in a separate, multiple-unit auction. To simplify matters we assumed that the winners of each auction would be the bidders with the lowest production cost for that particular item that can fulfill the demand. Thus, no strategic bidding that would consider the possible outcomes of the auctions of the other items was taken into account. We also assumed that the outcome of the single-item auctions would be efficient, and that there would be enough competition to drive the prices close to the production costs (or normal profit). The comparison of the auctions was then done based on the final cost to the buyer and not based on the production costs as was done with the two other benchmark cases.

8.2.4 Results of the Second Simulation Study

In the simulation design there were five design parameters that we varied in the study: number of bidders (15 or 30), bidders' capacities (equal or unequal), the magnitude of economies of scope (normal or large), the way the initial bid stream was compiled (Bid1, Bid2 or Bid3), and the way the shortlist was compiled (full or express). Thus, there were 48 different auction designs, and with 50 replications of each design, the total number of simulated auctions was 2400. Each of the 2400 auctions was run though first with price support and then with quantity support, and finally the efficient allocation for each auction was solved. Note that the price support auctions do not use the shortlist, which is part of the QSM. Hence, there were actually only 24 different designs for the price support auctions, but still all the 2400 auctions were run through with price support as well. Conducting price support auctions were for all "full" and "express" shortlist auctions allowed for pairwise comparison with the corresponding quantity support auctions.

In this second simulation study, our primary interest was in examining the final outcome of the auctions. Thus, we studied the outcomes of the auctions from many perspectives. We studied both the total cost to the buyer as well as the efficiency of the final allocation. We contrasted the final outcomes of the quantity support auctions and the price support auctions with the efficient allocation. The final outcomes of the quantity support auctions were also compared against the outcomes of the individual, non-combinatorial auctions. In addition, we studied the effects of the initial parameters on the auction outcomes. Besides total cost to the buyer and allocative efficiency, we were interested in the speed of convergence in the auctions and the usefulness of the shortlist.

8.2.4.1 Total Cost to the Buyer

The total cost to the buyer at end of the auction is the primary concern of the buyer (auction owner); it is only us researchers and perhaps the government as a buyer who are interested in the efficiency of the final allocation. The buyer wants to minimize the price she has to pay, and she does not care whether the allocation is efficient or not. Naturally, the more efficient the allocation, the lower the total cost can potentially go. However, even if the final allocation were efficient, the cost to the buyer can be high if there is not much competition and the mark-ups remain high even at the end of the auction. Hence the buyer will want the auction to be relatively efficient, but only as a means to an end. If the QSM does not produce outcomes with prices acceptable to the buyer, she will not want to use the mechanism. Thus, in our simulation study we were interested in finding out how the use of the QSM affected the total price paid by the buyer.

Table 16 contains the ratios of the average total cost to the buyer from the winning allocations and the efficient production cost (theoretically the lowest possible cost to the buyer) in quantity and price support auctions. For example, a ratio of 1.070 means that the total cost to the buyer is 7% above the efficient production cost. The cost to the buyer is averaged over all the simulation settings in which the parameter considered has the same value. For example, there were 1200 auctions with 15 bidders, and 1200 auction with 30 bidders, and 800 auctions with Bid1, Bid2 and Bid3 respectively.

Table 16 Total cost to the buyer as ratio of efficient production cost

		Quantity Support		Price Support	
		Mean	St.dev	Mean	St.dev
# of bidders	15	1.070	0.054	1.239	0.120
	30	1.054	0.039	1.157	0.085
Econ. Of scope	Normal	1.062	0.048	1.196	0.111
	Large	1.062	0.048	1.200	0.113
Initial bids	Bid1	1.065	0.051	1.168	0.098
	Bid2	1.061	0.048	1.164	0.091
	Bid3	1.060	0.045	1.261	0.117
Capacities	Equal	1.025	0.014	1.128	0.069
	Unequal	1.099	0.042	1.267	0.103
Shortlist	Full	1.058	0.042	1.196	0.114
	Express	1.066	0.053	1.199	0.110
ALL		1.062	0.048	1.198	0.112

As can be read from Table 16, the total cost to the buyer is much lower when quantity support is used. The price the buyer has to pay is on average only 6.2 percent above the efficient production costs in the quantity support auctions, where as it was on average almost 20 percent above the efficient production costs in the price support auctions. One plausible explanation is that the QSM helps the bidders find profitable combinations so that the auction continues longer and either the efficiency of the bids improves, or the mark-ups on bids are driven down (the required decrease in the buyer's total cost is 2% with every new bid), or both.

Paired, one-tailed t-tests were conducted to test the statistical significance of the difference in the total cost to the buyer in quantity support and price support auctions. Results of the t-tests for all the 48 designs indicate that the difference in the total cost to the buyer (in favor of the quantity support auctions) is statistically significant (all p-values smaller than 10^{-12}).

There does not seem to be large variations in the total cost to the buyer in the quantity support auctions. The only larger differences are between auctions with 15 vs. 30 bidders and auctions with equal vs. unequal bidder capacities. It is natural, that the auctions with more bidders result in a lower total cost. The more there are bidders, the more likely it is that there are bidders with low costs, and the average cost of the winners should be lower. Also, the more there are bidders, the more there is competition and the total cost to the buyer is driven closer to the production costs of the bidders. One explanation for the better performance of auctions with equal bidder capacities is that it increases the competition among the bidders, when all the bidders can place the same bids. In that case, only cost matters. With unequal capacities a low cost producer can be at a disadvantage due to a smaller capacity compared to another bidder. Also, in the case of equal capacity case the bid quantities of each item (300 or 0) sum easily up to the total demand (600). Other reasons for the difference between auctions with equal and unequal bidder capacities are discussed in section 8.2.4.6.

8.2.4.2 Efficiency of the Final Allocation

In the previous section I concluded that the total cost to the buyer is much lower in the quantity support auctions than in the price support auctions. However, because the bid prices also have mark-ups above the bidders' production costs in them, we cannot estimate the improvement in the efficiency of the final allocation. Because of larger mark-ups, the final cost to the buyer in auction A could be higher than in auction B, even though the allocation in A is more efficient. Thus, we need to "clean" out the mark-ups in the bids. This is done by using the bidders' cost to produce the bids. We defined efficiency to be the ratio of the winning bidders' combined cost of producing the winning allocation and the total cost of the efficient (lowest total cost) allocation. Let I_Q^* denote the combination of winning bidders, and $Q_{i,Q}^*$ the combination of items in bidder *i*'s winning bid in a quantity support auction, and $q_{ik,Q}^*$ the quantity of item *k* in that bid (notice that $q_{ik,Q}^* = 0$ for all items not in $Q_{i,Q}^*$). The efficiency indicator can then be expressed as

$$efficiency_Q = \frac{\sum_{i \in I^*_Q} \left(F_{i,Q^*_{i\,Q}} + \sum_{k=1}^{K} c_{ik} q^*_{ik,Q} \right)}{C_e} \tag{27}$$

where C_e indicates the efficient allocation. Efficiency indicator for a price support auction (*efficiency*$_P$) is defined similarly. Only in that case the combination of winning bidders I^*_P and the combinations and quantities of items $Q^*_{i,P}$ and $q^*_{ik,P}$ are taken from the winning allocation of the price support auctions. If the final allocation does not coincide with the efficient allocation, the ratio will be grater than one. For example, an efficiency of 1.026 means that the cost of producing the winning allocation is 2.6% above the efficient production cost. When efficiency improves, the ratio approaches unity.

Table 17 presents the mean efficiency indicators of the auctions with quantity support, and the benchmark case (price support). Again, the data have been grouped according to the design parameters in order to study the effects of the chosen values.

Table 17 Efficiency ratios of final allocations in quantity and price support auctions

		Quantity Support		Price Support	
		Mean	st.dev	mean	st.dev
# of bidders	15	1.028	0.030	1.096	0.064
	30	1.025	0.025	1.065	0.039
Econ. Of scope	Normal	1.026	0.028	1.079	0.055
	Large	1.026	0.028	1.082	0.056
Initial bids	Bid1	1.027	0.028	1.065	0.048
	Bid2	1.025	0.026	1.067	0.042
	Bid3	1.027	0.029	1.109	0.062
Capacities	Equal	1.008	0.010	1.051	0.031
	unequal	1.045	0.028	1.110	0.059
Shortlist	Full	1.025	0.027	1.080	0.054
	express	1.027	0.029	1.080	0.056
ALL		1.026	0.028	1.080	0.055

As can be seen from the figures in Table 17, the auctions where quantity support was available for the bidders resulted in more efficient outcomes than the ones where only price support was available. The cost of producing the winning allocation was on average 2.6 percent higher than in the efficient allocation when quantity support was used, and 8 percent higher, when only price support was used. Thus, part of the difference in the final cost to the buyer in quantity support auctions vs. price support auctions depicted in Table 16 can, in fact, be attributed to the higher efficiency of the winning bidders in the quantity support auctions. Paired, one-tailed t-tests show that the difference in the efficiency of quantity support and price support auctions is statistically significant for all 48 designs (p-values all below 0.001, and vast majority even below 10^{-10}).

Notice also that the difference between the average efficiencies (Table 17) of quantity and price support auctions is much smaller than the difference in the total cost to the buyer (Table 16). This indicates larger profits for bidders (and a higher cost to the buyer) in price support auctions. Based on these two tables we could also estimate the size of the average mark-ups in the winning bids to be around 4% in the quantity support auctions and 11% in the price support auctions. The higher mark-ups together with the less efficient final outcomes caused the buyer to pay almost 20% extra over the efficient production costs when only price support is available. When quantity support is available, the buyer paid only a little over 6% extra in our experimental setting.

Looking at the efficiency ratios of the quantity support auctions, it can be seen that the different design parameter values have very little effect on the efficiency of the final allocation, except for the bidders' capacities. When the bidders are symmetric in their production capacities (and the capacities conveniently sum up to the total demand), quantity support helps the bidders get very close to the efficient allocation. T-tests suggest that the differences in the final efficiency between auctions with equal and unequal bidder capacities are statistically significant. The variance of the production costs of the winning allocation was also the smallest with equal production capacities. This means that the teaming-up happens more easily when the bids are more similar. The practical implication of this intuitive result is that it improves the efficiency of a multi-unit auction to restrict the quantities that can be bid on to only a few levels. The fact that there was very little difference in the efficiency of the final allocation between the two types of shortlists used, and also very little difference in the total cost to the buyer indicates that the "express" shortlist could be used. This would decrease the bidders' evaluation efforts when using the QSM and not worsen the buyer's situation significantly.

8.2.4.3 *Total Cost to the Buyer in Non-Combinatorial Auctions*

The total cost to the buyer in the quantity support auctions was also compared to the total cost to the buyer in non-combinatorial auctions in which the items are auctioned off individually (see section 8.2.3). Because of the assumptions we made about the non-combinatorial auctions, the total cost to the buyer in those auctions coincides with the lowest production costs for the single items. We wanted to compare the total costs to the buyer (presented in Table 16) because we wanted to demonstrate that even though allocative efficiency was not reached in the quantity support auctions, they were still more profitable to the buyer than auctioning the items individually – even if assuming that the efficient allocations of the individual item auctions could be reached with competitive (= no excess profit) prices. The ratios of the total cost to the buyer in quantity support auctions to the total cost in non-combinatorial auctions presented in Table 18 verify this argument. Again, the data have been grouped according to the design parameters.

Table 18 The ratio of total cost to the buyer in a quantity support auction and the lowest total cost in a non-combinatorial auction

		Mean	St.dev
# of bidders	15	0.910	0.069
	30	0.873	0.052
Econ. of scope	Normal	0.897	0.064
	Large	0.886	0.064
Initial bids	Bid1	0.905	0.067
	Bid2	0.885	0.063
	Bid3	0.884	0.059
Capacities	Equal	0.844	0.032
	Unequal	0.939	0.052
Shortlist	Full	0.889	0.059
	Express	0.894	0.068
ALL		0.892	0.064

The total cost to the buyer is consistently around 10 percent lower in combinatorial auctions than in single-item auctions. Naturally, the advantage of combinatorial auctions is increased when the economies of scope are larger. Bidders can express their synergies in combinatorial auctions, but not in single-item auctions. However, if the non-combinatorial auctions are held as simultaneous auctions similar to the FCC auctions described in section 3.1.2.2, and the bidders had better chances of winning favorable combinations, the outcome of the benchmark case might not be as inefficient as portrayed in this study.

8.2.4.4 Convergence Speed

In this simulation we also studied the convergence speed of the auction. The convergence speed is an indicator of how much effort the bidders need to put into the auction process. In our simulations an iteration is defined as a new bidder entering the set of provisional winners. And because after the initial bids, the bids placed in the auctions are all suggestions from the support mechanisms, moving from one iteration to the next implies that the bidder found at least one bid suggestion profitable. The total cost to the buyer was required to decrease by two percent from iteration to iteration. Thus, the lower the final total cost – and usually also the better the efficiency – the more there would be iterations in the auction. Thus, in our simulation setup it is preferable to have the auction go on for as many iterations as possible – although at the cost of bidder effort increasing.

We have already concluded in section 8.2.4.1 that the total cost to the buyer was on average much lower in the quantity support auctions. From that directly follows that on average the auctions with quantity support went on for more iterations than the price support auctions, regardless of the design. The average number of iterations in all quantity support auctions was 11.2 whereas in price support auctions the number of iterations was 5.4.

However, the number of iterations only records the times the use of the support mechanisms resulted in a new bid. The time and effort put into evaluating bid suggestions which turned out to be unprofitable is not measured at all. One iteration can take a long time if bidder after bidder uses the support mechanism to no avail before finally one bidder finds a profitable suggestion. Thus, a better measure for the convergence speed of the auctions would be the average number of times the support mechanisms (price or quantity depending on the auction) was used in each iteration. The number of bidders in the auction naturally affects the number of times a support mechanism is used, because there are more inactive bidders to go through. Thus, the number of times support was used per iteration is considered separately for the 15 and

30 bidder auctions. With 15 bidders, the average number of times support was used per iteration was 5.3 for price support auction and 2.9 for quantity support auction. The average total effort (number of times support was used during the entire auction) of the bidders was then 28.62 in price support auctions and 32.48 in quantity support auctions. With 30 bidders support was used per iteration 11.7 times in the price support auctions, and 5.7 times in quantity support auction. Total effort of bidders was 63.18 in price support auctions and 63.84 in quantity support auctions. Thus, even though the price support auctions last for fewer iterations, the amount of effort the bidders have to put in during the bidding process is almost as high as in the quantity support auctions, and each iteration there are more futile attempts to find bids.

8.2.4.5 Usefulness of the Shortlist

We included the shortlist in the QSM because we believed that the original solution of the QSP might not be the most profitable bid – at least not to all bidders. Thus it is also interesting to study, whether the additional shortlist items proved to be useful or not.

In the simulations, whenever a bidder placed a bid suggested by the QSM, we recorded whether the bid was the original (first) solution of the QSP, or whether it was one of the additional bid suggestions on the shortlist. In the auctions where the bidders' capacities were equal, the bidders chose the original shortlist item 94% of the time. Thus, in those auctions, having the additional suggestions did not improve the performance of the QSM that much. However, in the unequal capacities case the original shortlist item was chosen over the other items only 57% of the time. Thus, we concluded that having the additional shortlist items improved the performance of the QSM, especially in the unequal capacities case.

8.2.4.6 Further Observations from the Quantity Support Auctions

Taking a closer look at the progress of simulated quantity support auctions, some further observations can be made on how the auctions progress and what factors affect the final outcome. In sections 8.2.4.1 and 8.2.4.2 I already demonstrated that one simulation design parameter – bidders' capacities - clearly affected both the efficiency of the final allocation and the total cost to the buyer. However, looking into the quantity support auction more closely allowed me to make more detailed observations on what affects the progress and outcome of the auctions.

In addition, in the previous sections I have only considered the final outcome of the auctions – be it the efficiency of the final allocation or the total cost to the buyer. Another interesting viewpoint would be to study the change in the efficiency of the allocation from the initial bid stream to the final (winning) allocation.

Factors Affecting the Efficiency of the Auction Outcome

In Table 17 (section 8.2.4.2) I used average efficiencies to compare the outcomes of different simulation designs. However, Parkes (2001) actually argues that the frequency of exactly efficient outcomes is a better measure of the efficiency than e.g. average outcome. Thus, I decided to study the outcomes of the quantity support auctions more carefully. I observed that almost half of the auctions (1,136 out of 2,400) ended in an efficient allocation (417/2,400), or a "pseudo efficient" allocation (719/2,400). With "pseudo efficient" I mean two different situations in which the final allocation is not efficient (the winners are not the efficient bidders or the winning bids are not the efficient combinations), but should still be regarded as such in the analysis. Firstly, in a situation in which the total cost to the buyer is already within 2% (the required decrease in the total cost) of the efficient allocation production costs, it is not possible for the efficient bidder to place the efficient bid without incurring a loss. Secondly, in a situation in which the total cost is not within 2% of the efficient cost but in which existing bids have such high mark-ups in their bids, it would require the incoming efficient bidder to decrease her bid price below production costs to decrease the total cost to the buyer by 2% (this is the threshold problem). These both situations are considered "pseudo efficient" in the sense that the failure to reach the efficient allocation is not the fault of the QSM. In fact, in almost all the pseudo efficient cases, the QSM proposes the efficient bid to the efficient bidder, but the price required to make the bid active is too low.

Of course, the definition of pseudo efficiency is tied to the bid decrement δ, which was fixed at 2% in our simulations. However, because of that, as δ changes, the definition for pseudo efficiency changes. For example, choosing a smaller δ, say 1%, would narrow down the definition of pseudo efficiency leaving some of the current pseudo efficient outside the scope. However, a smaller δ could potentially allow the auctions to proceed one step further, making the final allocations again efficient or pseudo efficient. Thus I do not expect different values of δ to change our results, because these two effects may cancel each other out.

A common factor in the efficient and pseudo efficient cases is that the efficient allocation usually consists of two bids only (they split the lot). Two bids is the minimum number required, as the capacity constraints were adjusted so that it was not possible for anyone to bid for the whole lot. This is the case in 1,068 out of the 1,136 efficient and pseudo efficient auctions. In the rest of the efficient or pseudo efficient simulations, there are three bids in the efficient allocation, except for one simulation in which there were four bids. Interestingly, in many of the pseudo efficient cases in which the efficient allocation would have consisted of three bids, the actual winning allocation still consists of only two bids. The production cost of the winning allocation is so close to the efficient one that it was possible to bring the total cost to the buyer within 2% of the efficient production cost.

The vast majority of auctions, in which the efficient allocation consisted of two bids, were auctions in which the bidders' capacities were equal. In the unequal capacities case many bidders had capacities smaller than 300, so more than two bids were required to fulfill the demand. Thus, based on this it would appear that the reason why auctions with equal capacities led to significantly better efficiency than auctions with unequal capacities (see Table 17) is the fact that in those cases the efficient allocation consisted of only two bids. Also, when the capacities are equal (and sum up to total demand), the solution space is in fact rather small. Thus, the efficient allocation is easier to find, and it was found in 1,064 of the 1,200 simulations, while in the unequal capacities case the efficient allocation was found only in 72 cases out of 1,200. Figure 1 below contrasts the cumulative distribution of the efficiency of auctions in which capacities were equal to the efficiency of auctions with unequal capacities.

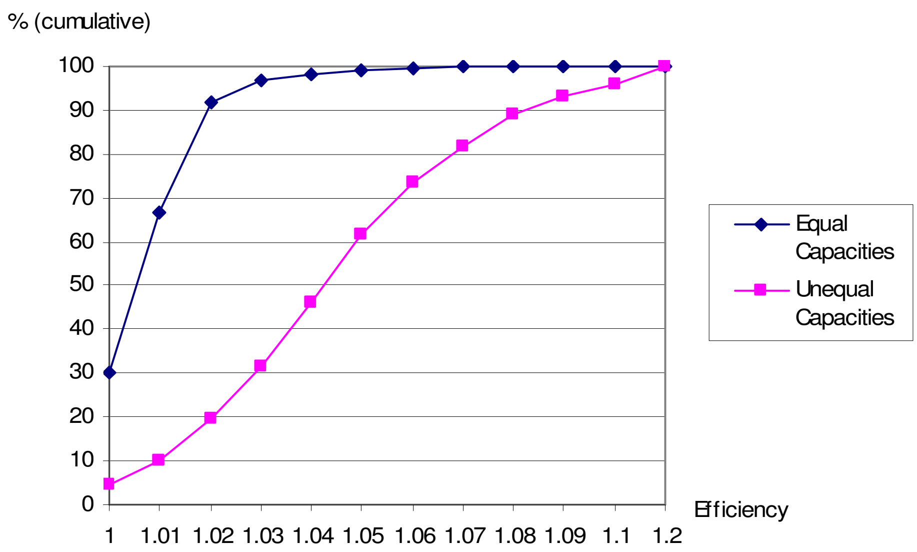

Figure 1 Cumulative distributions of the efficiency (compared to the production cost of the efficient allocation) of auctions with equal capacities and unequal capacities

By looking at the distributions, it is clear that the efficiency of auctions with equal capacities is generally better than in auctions with unequal capacities. About 90% of auctions with equal capacities have efficiency ratios below 1.02, whereas the same is true for only 20% of auctions with unequal capacities.

The assumption of equal capacities is analogous to a combinatorial auction in which the demand is two units for each item, and bidders can only bid for one unit. This comparison is straightforward, because all quantities in bids were either 0 or 300, and total demand was 600. The single-unit combinatorial auction considered widely in literature is even simpler, and thus I anticipate that the quantity support mechanism would work even better in such auctions. The addition of multiple levels for capacities increases the solution space tremendously. The initial bid stream (which has the same number of bids as there are bidders in the auction) now covers only a very small portion of the solution space. The QSP always searches the complements for the new bid from the bid stream, and hence the content of the initial bid stream becomes very important. In our simulations, when the efficient allocation consisted of more than two bids, it was necessary that at least one of the efficient bids be in the initial bid stream. However, having one or more efficient bids in the initial bid stream was not sufficient; the final allocation of the auction might still not be efficient. When the efficient allocation consisted of two bids, however, the efficient bids needed not be present in the initial bid stream; the efficient allocation was still found almost every time.

The total cost to the buyer behaved similarly as the efficiency of the final allocation. Naturally, the efficiency of the final allocation is correlated with the final cost to the buyer, because efficient bidders can afford to drive the price lower. However, an efficient allocation does not guarantee a low total cost to the buyer (12% of efficient final allocations resulted in a total cost more than 5% above the efficient production cost). It would appear that the best indicator is again the capacities of the bidders (or the number of bids in the efficient allocation, which is closely related to the capacities of the bidders). Figure 2 depicts the cumulative distributions of the total cost to the buyer in auctions with equal and unequal capacities.

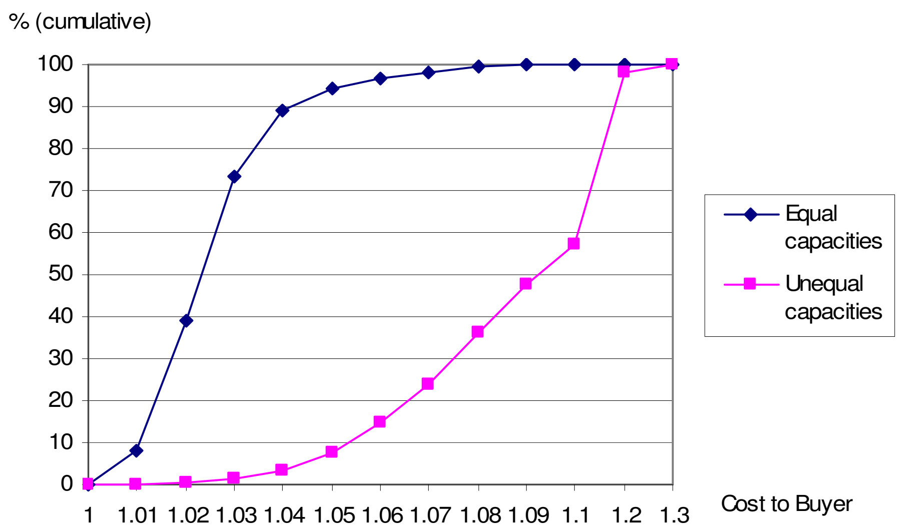

Figure 2 Cumulative distributions of buyer's total cost (compared to the production cost of the efficient allocation) in auctions with equal capacities and unequal capacities

From Figure 2 it is clear that the auctions with equal capacities led to lower total cost to the buyer than auctions with unequal capacities. Almost 90% of auctions with equal capacities ended up in a total cost to the buyer with 4% of the efficient production cost, whereas less than 5% of auctions with unequal capacities achieved the same level.

Final Efficiency vs. Initial Efficiency

Thus far only the efficiency of the final allocation has been studied. However, a valid question is, what the efficiency of the auction was before the bidders began using the QSM. This is especially valid since the set of initial bids was determined in the simulation design. Therefore, it is a possibility that certain auctions end in a more efficient outcome simply because they started off with a more efficient allocation. Also, it would be interesting to know, if the QSM improves the efficiency of certain types of auctions more than some others. According to Table 17, the way the initial bid stream was constructed did not have a very big effect on the efficiency of the final allocation. However, from Table 17 one cannot study possible effects two or more design parameters could jointly have.

Thus, I have calculated the difference between the initial efficiency (= the production cost of the provisional winners in the initial bid stream) and final efficiency (= the production cost of the winning bidders at the end of the auction). The average improvement in efficiency, i.e. the decrease in the production costs, is reported in Table 19 as a percentage of the efficient production cost. The data in Table 19 is aggregated over all simulations according to the simulation design parameters.

Table 19 Average improvement in efficiency in the quantity support auctions

Design variable		**Average improvement** (as % of eff. prod.cost)
Capacities	equal	3.6
	unequal	9.0
Bidders	15	8.5
	30	4.1
Econ. of scope	normal	6.2
	large	6.4
Shortlist	express	6.2
	full	6.3
Initial bid	bid1	5.5
	bid2	4.7
	bid3	8.6

The results in Table 19 indicate that the improvement in efficiency during the auction is larger if the bidders' capacities are unequal, there are only 15 bidders, or if the logic of "Bid3" is used to create the initial bid stream. Studying the interactions of these three design parameters, Table 20 was constructed. In the table, improvements in the efficiency were averaged over the three interesting design parameters. The size of the economies of scope and the compilation of the shortlist do not seem to have an effect, so they were not studied further.

Table 20 **Average improvement in efficiency (as % of efficient cost) broken down according to three design parameters**

		15 bidders		30 bidders	
		Equal	Unequal	Equal	Unequal
Initial bids	Bid1	3.3	12.1	1.6	4.8
	Bid2	2.9	9.5	1.9	4.5
	Bid3	7.5	15.3	4.1	7.5

It is clear from Table 20 that the efficiency improved the most in the auctions in which there were 15 bidders and the bidders' capacities were unequal. Also if the logic Bid3 was used to construct the initial bid stream, the improvement in efficiency was larger than if other logics were used.

In order to understand what is behind the figures in Table 20, I constructed a table from the average final efficiency ratios in Table 17 , but regrouped them according to the grouping in Table 20.

Table 21 **Average final efficiency ratios broken down according to three design parameters**

		15 bidders		30 bidders	
		Equal	Unequal	Equal	Unequal
Initial bids	Bid1	1.009	1.049	1.009	1.042
	Bid2	1.009	1.046	1.008	1.039
	Bid3	1.006	1.050	1.008	1.043

As can be seen from Table 21, the differences in final efficiency are not dependent on the number of bidders or the logic by which the initial bids are constructed. The only design parameter explaining the differences is the bidders' capacities – the conclusion to which we have already come in the previous sections.

Notice that there are large differences in how much the efficiency improved during the auctions, but that there is much less variation in the final efficiencies of the auctions. This means that there must be large differences in the initial efficiencies of the auctions. Indeed, the average initial efficiency is worse (i.e. efficiency ratio is higher) for auctions in which there are 15 bidders and the bidders' capacities are unequal, or when the initial bids are created with Bid3. What is very interesting is that the QSM seems powerful enough to smooth out the initial differences caused by the smaller number of bidders, and the bid creation logic Bid3. However, it would appear that the QSM it cannot quite tackle the inefficiency caused by the unequal bidders' capacities.

9 IDEAS TO IMPROVE THE QSM[18]

The results of the second simulation study indicated that there is still room for improvement in the QSM. Whenever there were more than two bids in the efficient allocation (which corresponds to the case of unequal bidder capacities in our simulation setting), it became difficult for the QSM to lead the auction to the efficient allocation. The contribution of this chapter is to discuss improvement ideas to the QSM and their applicability.

9.1 Varying the Bid Decrement

The bid decrement δ refers to the amount by which the total cost to the buyer should decrease every time a new bid enters the set of provisional winners. In both simulation studies presented in Chapter 8 the decrement was fixed, although the size of the decrement was different (5% of current total cost in the first simulation and 2% in the second). In section 8.2.1 I discussed the effects the size of the decrement can have on the convergence of the auctions. However, so far I have not discussed the effects the fact that the decrement is fixed has on the QSM and thereby on the convergence and outcome of the auction.

In its original form, the QSP forces the incoming bid to bear the whole burden of decreasing the total cost to the buyer by the fixed decrement δ (defined as a percentage of current lowest total cost). Especially, when the bid quantities are small, and therefore also the price attached to the bid is low, the burden is unreasonable. The price for that one bid has to go very low in order for the total cost – which possibly consists of several other bids – to decrease by a fixed decrement. This leads easily to unacceptable bid suggestions from the QSP and a definite bias towards large bid suggestions. Choosing a small δ would make it somewhat easier for small bids to be acceptable, but a small δ slows down the convergence of the auction, and it would not change the fact that large bids would still be favored.

[18] The ideas presented in this chapter are based on a brainstorming session with Professors Hannele Wallenius, Jyrki Wallenius and Murat Köksalan.

Thus, we developed the idea to make the price decrement "dynamic", i.e. dependent on the bid size. Following the idea presented by auction scholars in relation to Federal Communications Commission (FCC, 2000) auction # 31 (a combinatorial auction), we thought it might be a good idea to have the decrement depend on the items and item quantities in the bid. Only if the incoming bid is for the whole demand should the decrement be effective fully. For example, assume the full decrement is set at Δ = 5%, then, if the bid is for only "half" of the whole bundle, the required decrease would be only 2.5%.

The formulation of such dependence in terms of items and item quantities in a combinatorial auction is difficult. The items might be very different from each other, so that for example half of the total cost could be accrued from one single item, and the rest from all other items. Therefore the definition of "a half bundle" is not trivial, unless the item quantities were 50% of the total demand for each item. The auction owner (buyer) could assign relative weights to the items, so that the size of each bundle in monetary terms could be evaluated. Also, instead of buyer defined weights, the dual prices could be used as relative weights to estimate the value (monetary size) of any given bundle. The total cost of the bid could be estimated using the dual prices, and the cost of the bid would then be compared to the total demand, to see how large of a portion the bid represents. A potential problem can arise if some dual price is zero, because that can underestimate the costs drastically, and allow for too small a delta.

Even a simpler idea is to approximate the "size" of the bid by comparing the bid price in the new bid to the total lowest cost to the buyer. Let Δ denote the decrement by which a bid for the whole demand is required to lower the total cost, and let δ denote the decrement by which an incoming bid should decrease the total cost. Then δ would be defined as

$$\delta = \frac{p_{new}}{\sum_{i=1}^{N}\sum_{j=1}^{n_i} p_{ij}x_{ij} + p_{new}}\Delta \tag{28}$$

where p_{new} is the price of the incoming bid, p_{ij} the price of bidder i's j^{th} bid in the bid stream, and $x_{ij} \in \{0,1\}$ indicates which bids are among the current winners. However, defining δ in this manner results in a nonlinear constraint in the QSP. Thus, we decided to approximate the new total cost in the denominator with the old total cost C* (the solution of the WDP without the incoming bid), which is easily available. Because C* is by definition an upper bound for the denominator, it results in a slightly smaller δ than would be appropriate. However, we do not believe the difference would be significant.

The constraint (7) in the QSP (the one requiring that the total cost to the buyer decreases) would then take the form of

$$
\begin{aligned}
&\sum_{i=1}^{N}\sum_{j=1}^{n_i} p_{ij}x_{ij} + p_{new} \le \left(1-\frac{p_{new}}{C^*}\Delta\right)C^* \\
&or \\
&\sum_{i=1}^{N}\sum_{j=1}^{n_i} p_{ij}x_{ij} + (1+\Delta)p_{new} \le C^*
\end{aligned}
\tag{29}
$$

When testing the dynamic delta as suggested above in (29), it was discovered that the use of a dynamic delta led to the QSM to suggest drastically smaller bids, and the efficiency of the final allocation deteriorated. There are two reasons for this. First, the original QSM using dual prices already has a bias towards small bids, as the linearized cost function underestimates the cost of small bids and conversely overestimates the cost of large bids. The dual prices are consistently above the true variable cost parameters and there is no fixed cost element. Thus, small bids appear cheaper than they should, and large bids more expensive. This bias is aggravated with the dynamic delta, which requires less reduction in the total cost from small bids.

Thus, the fact that a fixed delta introduces a bias towards larger bids conveniently counteracts the bias towards small bids inherent in the QSM. That is why the original QSM worked as well as it did. What was puzzling though, was that when the approximated (linear) cost function in the QSM was replaced by the bidder's true cost function, the dynamic delta did not perform better than the fixed delta – in fact, it did worse. Naturally, the relative sizes of the fixed and dynamic deltas, affect the results. We tested two different levels for the variable delta, and the results were the same. We did not pursue the research into the dynamic delta because it seemed clear that it would not significantly improve the efficiency of the final allocation.

9.2 Initial Bids in the Form of Ranges

One of the reasons the efficient allocation was not found in many of the simulated auctions was that the initial bid stream did not contain the bids from the efficient allocation, and therefore also the QSM cannot find the missing efficient bids. Thus, one way to try to improve the performance of the QSM would be to increase the initial bid stream so that it would encompass more combinations. One idea would be to express the initial bids in the form of ranges. The bidders would give lower and upper bounds on the item quantities. Because the bids are linear within the bounds, also the price needs to be expressed linearly. Providing a per-unit price for each item is somewhat against the idea of combinatorial auctions, but since the price is only valid within the quantity ranges, it is only a minor compromise. The subsequent bids and the bid suggestions of the QSM would still be fixed (without ranges) and with only one price for the whole package, just as before. However, with the initial bid quantities expressed as ranges and per-unit prices the WDP needs to be adjusted from the original formulation (Eq. (2)). Now the q_{ijk}'s are also variables to be determined in the solution, hence the objective function takes the form

$$\min \sum_{i=1}^{N}\sum_{k=1}^{K} x_{i1} p_{i1k} q_{i1k} + \sum_{i=1}^{N}\sum_{j=2}^{n_i} x_{ij} p_{ij} \qquad (30)$$

where q_{i1k} is the quantity for item k in bidder i's initial (j = 1) bid, p_{i1k} is the per-unit price of item k in bidder i's initial bid, and p_{ij} is the bundle price for bidder i's subsequent bids. The following constraints need to be added so that the ranges specified by the bidders in their initial bids are taken into consideration when the q_{i1k}'s are solved for:

$$L_{ik} x_{i1} \le q_{i1k} \le U_{ik} x_{i1} \qquad (31)$$

Where L_{ik} is the lower bound and U_{ik} the upper bound for the quantity of item k in bidder i's initial bid. Similar adjustments need to be made to the QSP as well.

In the initial tests we ran, this modification of having bidders place the initial bids in the form of ranges did not have much of an effect on the efficiency of the final outcome. What happened was that the non-zero bid quantities were chosen to be either at the upper or lower bound. The bid space covered by the initial bid stream did expand, but only limitedly. Towards the end of the auction, the initial bids are edged out from the winning allocation. The bid quantities in the bids suggested by the QSM did change as a result of the new initial bids compared to the original situation in which the initial bids were formulated as bundle bids. However, the WDP's and the QSP's repeatedly chose the same upper or lower bound values for the initial item quantities. Hence, the effect of the modification remained modest. Of course, our findings are affected by our experiment set-up, hence it is possible that in another setting the results could be better.

9.3 Allowing for Shortages and Excesses in the Supply

To make the QSM more flexible in finding good bid suggestions it might make sense to allow for a shortage in the supply of some of the items, and conversely allow for supply exceeding demand for some other items. This way the exact complements for the bid suggestions do not have to exist in the bid stream.

The formulations of the WDP and QSM already allow for excess supply, but it is hardly ever present because having extra units of some items increases bidders' costs and thereby also the total cost to the buyer making such solutions rarely optimal. However, in reality, such solutions might not be all that undesirable for the buyer. In real life, procurement situations are rarely one time events. Rather, companies purchase the same items over an over again. Thus, in order to truly allow for the WDP or QSM to find solutions in which there is excess supply, some compensation should be given from the excess units. This compensation reflects the fact that now the buyer is in fact getting more for her money, and it makes the comparison with smaller packages more equitable.

The original formulations of the WDP and QSM do not allow shortages in any items. However, relaxing this constraint would give more flexibility in the solution of the problem than simply allowing excess supply. The shortages should be penalized, though, to indicate that it is not desirable to have a shortage but that it could be tolerable if the overall deal is then better. Allowing for shortages makes sense in practice, because one auction is hardly the only opportunity to procure these items. The missing units of the items can most likely be procured in some other auctions or then directly from the market. Naturally, if these conditions do not apply for some of the items, excesses and shortages can be allow for only a subset of the items in the auction.

Denote the excess in the supply of item *k* with e_k and the shortage with s_k. The per-unit penalty (i.e. extra cost) from shortage of item *k* is denoted with S_k and the compensation from excess supply with E_k. The WDP now takes the form

$$\begin{aligned}
\min \quad & \sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij} p_{ij} - \sum_{k=1}^{K}(E_k e_k - S_k s_k) \\
s.t. \quad & \sum_{i=1}^{N}\sum_{j=1}^{n_i} x_{ij} q_{ijk} = d_k + e_k - s_k \qquad \forall\ k = 1,\dots,K \\
& x_{ij} \in \{0,1\}
\end{aligned} \tag{32}$$

and similar additions need to be made to the QSP as well.

When testing the new formulation we soon realized that it is sensitive to the choice of E_k and S_k. If they were not chosen correctly, there would never be excesses or shortages. In addition, we noticed that the QSM very quickly clears away all the shortages that may have appeared in the initial stages of the auction. It cannot be a profit maximizing solution for any bidder to leave some shortage in the supply whenever the penalty from shortage S_k is larger than the bidder's variable cost c_k for item *k* (unless restricted by a capacity constraint). And, if $S_k < c_k$, this effectively means that the buyer can acquire the item cheaper somewhere else. Hence there would be no point in buying anything from this bidder, and the auction would become useless.

Based on our findings we concluded that the excesses and shortages formulation would not help in improving the efficiency of the final allocations in the quantity support auctions discussed in the previous chapter. However, we think that the excesses and shortages formulation could prove useful in one-shot auctions. By one-shot auctions we mean all non-iterative auctions, in which the WDP is solved only once during the entire auction. This means that the bidders cannot receive any feedback on how to improve upon their bids, and it is possible that the bids placed in the auction are not good complements to each other. By allowing the possibility of excesses and shortages in the formulation of the WDP, the buyer could potentially find different and perhaps more efficient solutions.

V THE GROUP SUPPORT MECHANISM

After testing the improvement ideas in Chapter 9 we concluded that minor changes in the QSM most likely would not improve the efficiency of the final allocations. The further observations from the results of the second simulation study (section 8.2.4.6) gave indication that one problem with the QSM is that it optimizes only one incoming bid at a time. First of all, when you can adjust only one bid at a time, it is not possible to overcome the threshold problem. As was explained already earlier, the threshold problem arises when the bidders should jointly revise prices in their existing bids so that together (possibly with one new bid) they could beat the current winner(s). Thus, by giving support to only one bidder at a time, the threshold problem remains. Also, it would appear that the "puzzle problem" is broader than we anticipated. The QSM was designed to alleviate the puzzle problem, in other words, to solve for the shape of the "last missing piece to the puzzle". Undoubtedly, in this capacity the QSM was successful: it was able to find the "missing pieces" allowing the quantity support auctions to continue much further than the price support auctions. However, oftentimes, if striving for the efficient allocation, it is not enough to solve only for the last missing piece. When the other bids that form the efficient allocation are not in the bid stream, the QSM cannot find the last missing efficient allocation bid. In fact, in this case there are several pieces missing from the puzzle, and the shape of just one of them cannot be solved. In order to significantly improve the allocative efficiency, the QSM should be modified to address this broader puzzle problem as well as the threshold problem.

10 THE GROUP SUPPORT MECHANISM[19]

The contribution of this chapter is to introduce another bidder decision support tool we have designed for a semi-sealed-bid multi-unit combinatorial auction. I also present a detailed example auction to explicate the use of the GSM. The design follows the logic of the QSM, but has significant improvements. Thus, one contribution is that we

[19] Material presented in this chapter will be published in Köksalan, Leskelä, Wallenius and Wallenius (2009), forthcoming in Decision Support Systems.

have designed a way to use information in bidders' bids to improve the approximation of the bidders' cost functions.

We call the new tool the Group Support Mechanism (GSM). The purpose of the GSM is to circumvent the problems arising from optimizing only one incoming bid at a time, which make it difficult for the QSM to lead auctions to an efficient allocation. Following the logic of the QSM, the GSM looks for bids that would become provisional winners. The difference is that the GSM suggests a combination of bids that either together satisfy the entire demand, or team up with one or more of the existing bids in the bid stream to become active as a group. This should improve the efficiency of the final allocation, because the GSM also chooses how many new bids it suggests, and is therefore not as dependent on the existing bids as the QSM. The QSM is a special case of the GSM, because the GSM will also support only a single bidder when finding it the optimal course of action. However, the GSM is free to suggest any number of incoming bids at a time, and the added flexibility should result in more efficient outcomes than the use of QSM.

10.1 The Group Support Problem

At the heart of the GSM is the Group Support Problem (GSP). The formulation of the GSP is related to the QSP – which is only natural since the underlying logic behind the two mechanisms is similar. The QSP was designed to maximize the profit of the incoming bidder. Conversely, the GSM attempts to maximize the joint profit of the provisionally winning bidders, that is, the profits of the new bids and existing bids to be included in the provisionally winning combination. The inclusion of the profit of existing bids in the GSM may appear somewhat counterintuitive, because the purpose of the QSM was to serve the incoming bidders and find new, profitable bids for them. However, not including the profit of the existing bids would create a bias towards filling the set of provisional winners with new bids. The existing bids would not be teamed-up with, even if they were good matches. Thus, the profits of both old and new bids are included in the objective function of the GSP.

The notation in the formulation of the GSP is similar to that used in the QSP: there are K items, d_k ($k = 1, \dots, K$) units of each item requested by the buyer, N bidders, and

n_i bids from each bidder i. Again, x_{ij}'s indicate the statuses of the old bids, $\tilde{c}_i(.)$ is the approximated cost function of bidder i, and $q_{i,new,k}$'s indicate the item quantities (elements of $Q_{i,new}$) and $p_{i,new}$'s the prices in the bid suggestions to be solved in the GSP. C^* is the current lowest total cost to the buyer, δ the percentage by which the total cost is required to decrease, and a_{ik} represent the bidders' capacities. At any point in time the bidders can be divided into two sets: active bidders, who are among the provisional winners, and inactive bidders who are not. Let I denote the set of inactive bidders. The GSP solves for a combination of new bids for the inactive bidders which, as a group (and possibly together with some of the existing bids in the bid stream), become active:

(GSP)
$$\max \ -\sum_{i\in I} s_i + \varepsilon\left(\sum_{i=1}^{N}\sum_{j=1}^{n_i}\left(p_{ij} - \tilde{c}_i(Q_{ij})\right)x_{ij}\right) + \varepsilon\left(\sum_{i\in I} e_i\right) \tag{33}$$

$$s.t. \quad \sum_{i=1}^{N}\sum_{j=1}^{n_i} p_{ij}x_{ij} + \sum_{i\in I} p_{i,new} \le (1-\delta)C^* \tag{34}$$

$$\sum_{i=1}^{N}\sum_{j=1}^{n_i} q_{ijk}x_{ij} + \sum_{i\in I} q_{i,new,k} \ge d_k \quad \forall\ k = 1,\ldots,K \tag{35}$$

$$p_{i,new} - \tilde{c}_i\left(Q_{i,new}\right) - e_i + s_i = 0 \quad \forall\ i\in I \tag{36}$$

$$\sum_{k=1}^{K} q_{i,new,k} - x_{i,new} \ge 0 \quad \forall\ i\in I \tag{37}$$

$$Mx_{i,new} - p_{i,new} \ge 0 \quad \forall\ i\in I \tag{38}$$

$$\sum_{j=1}^{n_i} x_{ij} \le 1 \quad \forall\ i\notin I \tag{39}$$

$$\sum_{j=1}^{n_i} x_{ij} + x_{i,new} \le 1 \quad \forall\ i\in I \tag{40}$$

$$q_{i,new,k} \le a_{ik} \quad k = 1,\ldots,K, \ \forall\ i\in I \tag{41}$$

$$x_{ij}, x_{i,new} \in \{0,1\} \quad \forall\ i, j$$

$$s_i, e_i, p_{i,new}, q_{i,new,k} \ge 0$$

The objective (33) is to maximize the combined (approximated) profit of all the bidders by choosing the statuses of the old bids (x_{ij}), and the quantities ($q_{i,new,k}$) and prices ($p_{i,new}$) in the new bids (i.e. sum of new profits e_i defined in (36)) suggested for the inactive bidders. The constraints (34), (35), (39), and (40) are the same as in the QSP. Notice that since the objective function maximizes the total profit of the active bidders, it does not discriminate against solutions in which some bidders accrue a loss. Because the new bids in the combination suggested by the GSM will become active only if all the bids are accepted by the bidders, such solutions are not desirable. However, we acknowledge that the profits we use in the formulations are approximations, and if we require that all profits be at least zero, we might not get a feasible solution even if one existed. This could be the case if we have overestimated the bidders' costs. Thus, we allow losses, but add a penalty for losses (s_i) into the objective function and formulate corresponding constraints (36). The addition of ε (a very small positive constant) in the objective function implies that the primary objective is to avoid losses and profit maximization is secondary. Essentially, the GSP chooses the most profitable combination of bids from combinations that do not result in losses for any bidder, if possible. If losses cannot be avoided, they will still be minimized. Constraints (37) and (38) (where M is any large number) ensure that a price in a new bid can be positive only if at least one item assumes a positive value and that the status variable for the new bid, $x_{i,new} = 1$, when the bid is not empty.

10.2 Customizing Cost Functions Approximations

The formulation of the GSP – just like the formulation of the QSP – requires an approximation of the bidders' cost functions. Thus, the second building block of the GSM is the procedure through which the cost functions are approximated. In the QSM, the cost function approximation was a linear cost function using dual prices as per-unit cost parameters. The QSM used the same cost function estimate for all bidders, and thereby the bid suggestions of QSM were anonymous (i.e. it gave the same suggested bids regardless of the bidder). However, it does not make sense for the GSM to be anonymous. If each bidder did not have a unique approximation of the cost function, the optimization algorithm would assign the bid suggestions randomly to some bidders, (as there would be numerable alternative optima) and it would pool the quantities into one large bid or few large bids (as allowed by the capacity constraints). We hold on to the assumption that bidders do not want to disclose any cost information, and thus the only information we have on the bidders' costs is information from the bidders' bids in the bid stream. We use the bid information to customize the cost function approximations for each bidder.

To start off, the form of the cost function to be used in the approximation was chosen. We decided to try out the same functional form as was used in the two simulation studies in Chapter 8. This way we could compare the results of the GSM auctions to the QSM simulations. To approximate the bidders' cost functions we designed an inverse optimization problem (in the spirit of Beil and Wein, 2003; see also Zionts and Wallenius, 1976), which utilizes the information we get in the form of bids in the bid stream. We assume that bidders do not place bids in which costs exceed the price. Thus, the task of the inverse optimization problem is to find a set of cost function parameters, which are consistent with a bidder's bidding behavior (see (42)). Also, we made some assumptions that allowed us to pose some constraints on the cost parameters. First, we assumed that by including an additional item into the bundle should not decrease the total cost, i.e. $F_{12} \geq F_1$ and $F_{12} \geq F_2$, (see constraints (45)). Secondly, we assumed that there would be economies of scope between the items, i.e. $F_{12} \leq F_1 + F_2$, (see constraints (46)). Also, in order to constrain the feasible set a little more, we assumed that upper and lower limits for all the cost parameters can be derived for any particular industry (constraints (43) and (44)).

The constraints of the Cost Estimation Problem (CEP) are:

$$p_{ij} - \tilde{c}_i(Q_{ij}) \geq \varepsilon \quad \forall \; j = 1,\ldots,n_i \tag{42}$$

$$l_{iL} \leq F_{iL} \leq u_{iL} \quad \forall \; L \subseteq \Gamma \tag{43}$$

$$l_{ik} \leq c_{ik} \leq u_{ik} \quad \forall \; k = 1,\ldots,K \tag{44}$$

$$F_{iL'} \leq F_{iL} \quad \forall \; L' \subset L, \forall \, L \subseteq \Gamma \tag{45}$$

$$F_{iL} \leq \sum_{L_t \in Par[L]} F_{iL_t} \quad \forall \; L \subset \Gamma \tag{46}$$

where ε is a small positive scalar, Γ is the set of all possible item combinations, and L, L' are subsets of Γ, Par[L] refers to all possible partitions of L, and L_t is an element of Par[L]. A partition of set L is the group of disjoint sets, which together form L.

These constraints (42) – (46), however, still leave a vast range of feasible options to choose from. Therefore, the choice of the objective function to a large extent determines the values for the cost parameters. Thus, the question becomes how to choose the objective function.

The use of different objective functions will lead the CEP to choose different points in the feasible set. Without any further information on the bidders' cost functions besides the constraints that make up the feasible set, there is no way of knowing, which point would be better than some other point. In other words, it is impossible to say which objective function would provide the best – or even good – approximations of the bidders' cost functions. Thus, we chose simply to maximize the sum of the cost function parameters as the objective of the CEP, that is,

$$\max \sum_{L \subseteq \Gamma} F_{iL} + \sum_{k=1}^{K} c_{ik} \, p_{ij} \tag{47}$$

and designed the following iterative scheme to approximate the bidders' cost function parameters and to narrow down the feasible set as the auction progresses.

First, lower and upper bound estimates for the cost function parameters are set. At the beginning of the auction the bounds coincide with "industry estimates", or in our experiments, the ranges of the distributions. The objective function will drive all the parameters to their upper bounds in the absence of any bid information to provide contradicting evidence. This may naturally be an over estimation of the cost functions, but it will not prohibit the GSP from finding bid suggestions, since losses are allowed in the formulation. Thus, it can still suggest the bids to the bidders even though the cost function approximations suggest that the bids would result in losses, and it is possible that the bidders will actually find them profitable. If the bidders accept what appeared to be unprofitable bids, it has the added benefit that now we get contradicting information and can update the estimates for the cost function parameters. The updated estimates are then set as the new upper bounds for the parameters, and the auction continues. If we started from a lower bound estimate for the cost parameters, and the GSP suggests bids in which all the estimated profits are positive, the acceptance of the bids is expected and would not give us any new information on the cost function parameters (the old estimates are still consistent with the new evidence). Naturally, in this case, if the bidders declined the bid suggestions we would get new information. However, it is desirable that the bidders accept the bid suggestions because that way the auction progresses. In order to maximize the information obtained from new bids and to speed up the auction process it makes more sense to start from the upper bound estimates.

It is worth noting that the GSP will not offer bid suggestions to all inactive bidders unless there is room for everybody in the set of provisional winners, which is more unlikely the more there are participants in the auction. Thus, some bidders are not suggested a bid by the GSP, and without any bid information the cost function estimate would not be lowered, which decreases the probability that the bidder would be offered a bid suggestion the next time. Some bidders could get stuck in this loop, and never be suggested anything. We could consider adding a constraint requiring that the bidder requesting support would be guaranteed a bid suggestion (not to upset the bidder), but so far we have not added any such constraints. Instead, recognizing that the estimates are above the true parameters, we decided to decrease the estimates by 1% for each bidder who is not suggested anything to improve their chances to receive a suggestion in the next round. We chose the decrement to be 1% in order to make only small adjustments in the estimates. Increasing the decrement could reduce the number of iterations needed to get the GSM to suggest a bid for the bidder, but a smaller decrement allows us to get closer to the true estimates. If, in the next round, the bidder receives a bid suggestion and accepts it, the upper bounds are replaced by the 1% lower estimates. If it turns out that the new estimate was too optimistic (the GSM offers a bid it thinks is profitable, but the bidder declines), we solve the cost function estimation problem again with the rejected bid added to the constraints, and receive an updated estimate of the parameters.

10.3 Example of an Auction with GSM

In order to present, how the GSM actually works, we designed an example auction, which is described below in detail. The example aims at clarifying how the GSP and CEP are formulated and solved, and how the auction advance. The outcome of the auction was also analyzed in order to evaluate the performance of the GSM.

10.3.1 Initiation Phase

In this example auction there are three items, 600 units of each item demanded by the buyer, and 10 bidders. All bidders are assumed to be "glocal" meaning that they have production capacity for all three items but they are willing to settle for any item and unit combination as long as it is profitable. The bidders are assigned capacities. Because the second simulation study demonstrated that the case of unequal capacities was more challenging, the capacities in this example are unequal as well. The three possible levels for the bidders' capacities are the same as in the simulation study (150, 225 and 300). The bidders are also assigned cost functions. Assigning specific cost functions allows for the evaluation of the profitability of the bid suggestions, and for the solution of the efficient allocation, against which the auction outcome can be evaluated. The cost functions are of the same for as in the simulation studies presented in Chapter 8. The parameters are chosen randomly from uniform distributions, and for the sake of simplicity no difference is made between the items: variable costs c_{ik} vary from [53.3; 66.7], fixed costs F_{iL} for single items from [1,777.8; 2,222.2], for pairs of items from [2,666.7; 3,333.3], and for combinations of three items from [3,555.6; 4,444.4]. It is made sure that the cost functions exhibit economies of scope, and the ranges for the parameters ensure that the fixed cost cannot decline as a result of an additional item being included into the bundle.

Because the GSM cannot be used unless there is already a feasible solution to the WDP (and thereby a total cost to the buyer), we assume that at the beginning each bidder places one bid. The bidders place a bid on the combination for which their fixed cost is the lowest compared to the expected value of the cost. The item quantities are set at the upper bounds of the capacity constraints. The initial bid stream is reproduced in Table 22.

Table 22 The initial bids and provisional winners

Bid	Item 1	Item 2	Item 3	Price (€)	Status*
x_{11}	0	225	150	27,199	0
x_{21}	150	0	150	22,954	1
x_{31}	225	225	150	45,990	0
x_{41}	0	225	0	17,613	0
x_{51}	0	225	150	25,594	0
x_{61}	225	150	300	49,356	0
x_{71}	300	0	0	22,007	0
x_{81}	0	225	225	33,210	1
x_{91}	225	150	0	27,831	1
$x_{10,1}$	225	225	225	50,076	1
* 1 = active, 0 = inactive; Total cost to buyer = 134,071€					

The next step is to estimate the cost functions of the bidders. The CEP is unique for each bidder, but only with respect to the constraints derived from the existing bids (42). The constraints resulting from the economies of scope assumption are the same for each bidder. Also initially, the upper and lower bounds for the parameters (which are set to be the upper and lower bounds of the uniform distributions from which the bidders' true cost function parameters were drawn) are identical for all bidders. As the auction progresses, the upper and lower bounds will differ across bidders. The objective function, which was chosen to be the sum of the parameters, is the same across bidders throughout the auction. For example, for Bidder 1 the CEP assumes the following form. For the sake of simplifying the notation, the index for the bidder has been left out.

$$
\begin{aligned}
\max \quad & F_1 + F_2 + F_3 + F_{12} + F_{13} + F_{23} + F_{123} + c_1 + c_2 + c_3 \\
s.t. \quad & 27199 - F_{23} - 225c_2 - 150c_3 \geq \varepsilon \\
& 1777.8 \leq F_i \leq 2222.2 \qquad i = 1,2,3 \\
& 2666.7 \leq F_{1j} \leq 3333.3 \qquad j = 2,3 \\
& 2666.7 \leq F_{23} \leq 3333.3 \\
& 3555.6 \leq F_{123} \leq 4444.4 \\
& 53.3 \leq c_i \leq 66.7 \quad \forall\ i = 1,\ldots,3 \\
& F_{12} \geq F_1,\ F_{12} \geq F_2 \\
& F_{13} \geq F_1,\ F_{13} \geq F_3 \\
& F_{23} \geq F_2,\ F_{23} \geq F_3 \\
& F_{123} \geq F_{12},\ F_{123} \geq F_{13},\ F_{123} \geq F_{23} \\
& F_1 + F_2 \geq F_{12},\ F_1 + F_3 \geq F_{13},\ F_2 + F_3 \geq F_{23} \\
& F_{12} + F_3 \geq F_{123},\ F_{13} + F_2 \geq F_{123},\ F_{23} + F_1 \geq F_{123}
\end{aligned}
\tag{48}
$$

where ε is a small positive constant. Once the CEP has been solved for each bidder, we have the following initial estimates for the cost function parameters:

Table 23 Initial estimates for the cost function parameters

	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
F_i	2222.2	2222.2	2222.2	2222.2	2222.2	2222.2	2222.2	2222.2	2222.2	2222.2
F_{ij}	3333.3	3333.3	3333.3	3333.3	3333.3	3333.3	3333.3	3333.3	3333.3	3333.3
F_{123}	4444.4	4444.4	4444.4	4444.4	4444.4	4444.4	4444.4	4444.4	4444.4	4444.4
c_1	66.7	66.7	66.7	66.7	66.7	66.7	65.9	66.7	64.4	66.7
c_2	61.6	66.7	66.7	66.7	54.5	66.7	66.7	66.7	66.7	66.7
c_3	66.7	64.1	66.7	66.7	66.7	66.4	66.7	66.1	66.7	66.7

As can be seen from the table, the CEP chooses the upper bounds of the fixed cost parameters for all bidders, and for most variable cost parameters, except for those bidders, whose initial bid indicates that it is not possible that all the cost parameters would be at the upper bound.

The auction is a continuous auction, that is, the WDP is solved after each incoming bid. However, since in this example, bids only enter based on the suggestions of the

GSM, it is easier to present the example “round by round”, in which one round continues as long as there is a change in the solution of the WDP (i.e., the set of active bidders changes).

10.3.2 The Auction

Round 1

It is now time to solve the GSP for the first time. The decrement δ by which the total cost to the buyer is required to decrease from round to round is set at 2%. The cost function parameters in Table 23 are inserted in the cost function in the GSP (33). However, because the cost function is discontinuous (the fixed cost term depends on the combination of items in the bid), the formulation becomes somewhat more complex, and it bears resemblance to the formulation of the QSP with the true cost function (Appendix 2) and of the efficient allocation problem (Appendix 3). A set of auxiliary variables $y_{i,jkl}$ is defined to construct constraints that guarantee that the correct fixed cost is taken into consideration in the objective function ($y_{i,jkl} = 1$ if the items in combination j,k and l all have a non-zero value, $y_{i,jkl} = 0$ otherwise). Also, for the same reason we need to create new variables for the item quantities in the incoming bids: one variable per item *per combination* it is in. If there are K items in the auction, each item is in 2^{K-1} combinations, so in this case each item is in four combinations.

Thus, the cost function $\tilde{c}(Q)$ for bidder i to be inserted into (33) takes the form

$$\begin{aligned}\tilde{c}_i(Q_{i,new}) = {} & c_{i,1}\sum_{s=1}^{4} q_{i,new,1,s} + c_{i,2}\sum_{s=1}^{4} q_{i,new,2,s} + c_{i,3}\sum_{s=1}^{4} q_{i,new,3,s} + F_{i,1}y_{i,1} + F_{i,2}y_{i,2} \\ & + F_{i,3}y_{i,3} + F_{i,12}y_{i,12} + F_{i,13}y_{i,13} + F_{i,23}y_{i,23} + F_{i,123}y_{i,123}\end{aligned} \tag{49}$$

To ensure that the correct fixed cost is taken into consideration in the cost function and that at most one combination is chosen per bidder, the following constraints are added to the GSP for each inactive bidder $i \in I$,

$$
\begin{aligned}
&My_{i,k} - q_{i,new,k,1} \geq 0 \quad k = 1,2,3 \\
&My_{i,12} - q_{i,new,1,2} - q_{i,new,2,2} \geq 0 \\
&My_{i,13} - q_{i,new,1,3} - q_{i,new,3,2} \geq 0 \\
&My_{i,23} - q_{i,new,2,3} - q_{i,new,3,3} \geq 0 \\
&My_{i,123} - q_{i,new,1,4} - q_{i,new,2,4} - q_{i,new,3,4} \geq 0
\end{aligned} \tag{50}
$$

$$y_{i,1} + y_{i,2} + y_{i,3} + y_{i,12} + y_{i,13} + y_{i,23} + y_{i,123} \leq 1$$

$$y_{i,1},\ y_{i,2},\ y_{i,3},\ y_{i,12},\ y_{i,13},\ y_{i,23},\ y_{i,123} \in \{0,1\}$$

where M is a constant larger than any conceivable item quantity. In this example, M = 1000 was used. Note that these constraints leave open the possibility that $y_{i,jkl}$ assumes the value of one, even though one or more of the items in the combination are zero. For example, when items one and two assume a nonzero value for bidder i, either $y_{i,12}$ = 1 or $y_{i,123}$ = 1. However, since the fixed cost for a bundle including more items is always larger than for a bundle with less items, the objective function will ensure that the $y_{i,jkl}$, which coincides exactly with the combination of nonzero item quantities, assumes the value one.

The bundle of bids suggested by the GSM in the first round is presented in Table 24. There are always multiple solutions, because the profit or loss in the bids (e_i or s_i) can be divided in an infinite number of ways between the bidders who are offered a bid suggestion. We chose the solution where the estimated profit/loss is divided equally among the bidders. Of course the true profits/losses of the bidders are not equal; we are dividing the profit/loss calculated with the cost function estimates equally among the bidders. Another approach would have been to divide the profit proportionately to the size of the bid. The main effect that the choice of solution has in the auction, is that in some cases it can cause the bidder to accept or reject a bid suggestion. E.g. if our cost estimate is a bit too low, offering a bid suggestion in which the GSP thinks the bidder will just break even, will not be good enough for the bidder and she will reject. However, if some of the surplus in the auction is allocated to this bidder, it may increase the bid price high enough so that she will accept the bid suggestion. Because we assume we do not know the bidders' true costs, it is difficult to know which cost estimates are underestimated, and which way of dividing the surplus would be best. Thus, we decided to start with dividing the (estimated) profit equally.

Table 24 Bids suggested by the GSM

Bid	Item 1	Item 2	Item 3	Price (€)
$x_{1,new}$	0	225	0	16,658
$x_{5,new}$	75	225	0	21,426
$x_{7,new}$	300	0	300	43,950

Bidders 1, 5 and 7 are suggested a bid, and their bids would team up with Bidder 6's initial bid (225; 150; 300; 49,356 €). The bidders all accept the new suggestions, so the set of provisional winners is now Bidder 1, Bidder 5, Bidder 6 and Bidder 7, and the total cost to the buyer is reduced to 131,390 €.

Next the cost functions are updated. The cost estimates for Bidders 3 and 4, who were inactive but still not offered a new bid, are now lowered by 1%. The cost estimates of the active bidders remain the same. The accepted bids offer no new information on the bidders' cost functions, because the bid prices are high enough to cover current estimated costs. The reason for the decrease in Bidder 3's and Bidder 4's estimates is that it is possible that we have overestimated their cost and that is the reason why they were not offered a bid. Bidders 2, 8, 9 and 10 were active so they could not have received bid suggestions regardless of their cost function estimates, and thus nothing is done to the estimates of their cost functions.

Round 2

The GSP with the updated cost information is solved for the new set of inactive bidders. The solution is presented in Table 25.

Table 25 Bids suggested by the GSM in Round 2

Bid	Item 1	Item 2	Item 3	Price (€)
$x_{4,new}$	300	225	225	53,426
$x_{8,new}$	0	150	225	27,736

This time the GSP suggests bids to Bidder 4 and Bidder 8, which would team up with the initial bids of Bidder 5 ($x_{51} = 1$) and Bidder 7 ($x_{71} = 1$). The GSP thinks the bids will result in a loss of 474 € for both bidders. However, both bidders accept the suggested bids. The GSP has overestimated their costs, so the true cost of producing the proposed bundles is below the bid prices. Now we have new information on the costs of Bidders 4 and 8, and their cost estimates are updated. The cost estimates for inactive bidders who did not receive a bid suggestion (Bidders 2, 3, 9 and 10) are lowered by 1%, and the estimates for bidders who previously were active (Bidder 1, 5, 6 and 7), out of which Bidders 5 and 7 are still active, remain untouched. The new cost function parameters are depicted in Table 26.

Table 26 Cost function parameters after Round 2

	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5	Bidder 6	Bidder 7	Bidder 8	Bidder 9	Bidder 10
F_i	2222.2	2200.0	2178.0	2200.0	2222.2	2222.2	2222.2	2222.2	2200.0	2200.0
F_{ij}	3333.3	3300.0	3267.0	3300.0	3333.3	3333.3	3333.3	3333.3	3300.0	3300.0
F_{123}	4444.4	4400.0	4356.0	4400.0	4444.4	4444.4	4444.4	4444.4	4400.0	4400.0
c_1	66.7	66.0	65.3	64.4	66.7	66.7	65.9	66.7	63.8	66.0
c_2	61.6	66.0	65.3	66.0	54.5	66.7	66.7	66.7	66.0	66.0
c_3	66.7	63.5	65.3	66.0	66.7	66.4	66.7	59.1	66.0	66.0

Round 3

First the GSP suggests a bid only for Bidder 10, and that bid (0; 150; 225; 25,160 €) is not acceptable to her. Also according to the estimated cost the bid is not profitable, and therefore we do not receive more information on B10's cost function. The cost estimates of all the other inactive Bidders (1, 2, 3, 6 and 9) are lowered by 1% and the GSP is solved again. This time the GSP suggests bids for Bidder 1: (150; 225; 0; 26,139 €), Bidder 6: (225; 0; 300; 37,075 €) and Bidder 9: (225; 150; 150; 37,379 €), which would team up with Bidder 5's initial bid resulting in a total cost to the buyer of 126,187 €. All the three bidders find the suggestions profitable, even though the GSP again thinks the bids would result in losses. The cost estimates are updated for all bidders, except the ones who are active.

Round 4

The first solution of the GSP provides only a suggestion for Bidder 3 (225; 150; 150; 34,855 €), which is not accepted. Thereafter, the GSP is solved seven times, and each time at least one bidder who is suggested a bid declines the suggestion, vetoing the "group" bid. The accepted bids are added to the bid stream, but the total cost to the buyer does not decline by the required 2% because the complementary bid(s) was not accepted. After each GSP solution the cost functions are updated. Finally, the ninth iteration produces bid suggestions to Bidder 4: (300; 0; 150; 30,515 €) and Bidder 7: (300; 150; 300; 50,897 €), which are both profitable for the bidders.

After this round, the GSP does not find bid combinations that would be profitable for all bidders. The auction ends. Going back to the bidders' cost functions we can conclude that the inactive bidders have such high costs that they could not afford to decrease the total cost to the buyer by the required 2%. The active bidders could have afforded to, but they did not have an incentive to do so, as they were already among the provisional winners. The final bid stream and the winning bids are presented in Table 27.

Table 27 Bid stream and solution of the WDP after Round 4

Bid	Item 1	Item 2	Item 3	Price (€)	Status
x_{11}	0	225	150	27,199	1
x_{21}	150	0	150	22,954	0
x_{31}	225	225	150	45,990	0
x_{41}	0	225	0	17,613	0
x_{51}	0	225	150	25,594	1
x_{61}	225	150	300	49,356	0
x_{71}	300	0	0	22,007	0
x_{81}	0	225	225	33,210	0
x_{91}	225	150	0	27,831	0
$x_{10,1}$	225	225	225	50,076	0
x_{12}	0	225	0	16,658	0
x_{52}	75	225	0	21,426	0
x_{72}	300	0	300	43,950	0
x_{42}	300	225	225	53,426	0
x_{82}	0	150	225	27,736	0
x_{13}	150	225	0	26,139	0
x_{62}	225	0	300	37,075	0
x_{92}	225	150	150	37,379	0
x_{43}	300	150	225	46,781	0
x_{44}	300	225	225	50,969	0
x_{73}	300	150	0	31,938	0
x_{45}	300	0	0	19,666	1
x_{74}	300	150	300	52,385	0
x_{46}	300	150	0	30,561	0
$x_{10,2}$	225	225	225	46,199	0
x_{47}	300	0	0	19,778	0
x_{48}	300	0	0	19,824	0
x_{49}	300	0	150	30,515	0
x_{75}	300	150	300	50,897	1

The total cost to the buyer is 123,356 €. In fact, the total cost decreased by more than 2% from Round 3. The previously inactive bid x_{45} made by Bidder 4 was more advantageous to the buyer now that a good match entered (x_{75}) the auction. The estimated profit for Bidder 4 was larger from the newest bid (x_{49}) and that is why the GSP favored that one (it looks at the auction from the bidders' perspective), but it is the WDP that determines the set of winners. The mark-up in Bidder 4's bid x_{45} is higher than in bid x_{49}. Hence, the bidder may be happier with this outcome.

10.3.3 Comparison with the Efficient Allocation

In order to see how well the GSM performed, we wanted to compare the final allocation of the auction to the efficient auction. The efficient allocation can be solved, because we know the bidders' cost functions. It is presented in Table 28.

Table 28 The efficient allocation of the example auction

Bidder	Item 1	Item 2	Item 3
B1	0	75	150
B4	300	0	0
B5	0	225	150
B7	300	300	300

Comparing the efficient allocation with the actual auction outcome (in Table 27) it is evident that the auction came close to the efficient allocation. All the bidders are the same, as are the items they bid on. The only differences are in two item quantities. If we compare the production costs of the efficient allocation (113,666 €) and the winning allocation (114,057 €), the difference (391 €) is very small.

10.3.4 Comparison with the QSM Auction

The auction presented in the example above was also run through using the QSM to support the bidders instead of the GSM. All the auction parameters (demand, number of bidders, bidders' cost functions and initial bids) were kept the same, but instead of using the GSM, the inactive bidders used the QSM, and placed bids based on the bid suggestions made by the QSM. The final allocation of the QSM auction is presented in Table 29.

Table 29 The final allocation of the auction with QSM

Bidder	Item 1	Item 2	Item 3
B4	225	150	0
B5	0	150	150
B7	0	225	150
B8	150	0	150
B9	225	75	150

The combined production cost of the winning bidders is 121,385 €, which is 7,719 €, or 6.8% higher than the efficient production cost, and 6.4% higher than the production cost of the winning allocation of the GSM auction. There are some bidders among the winners, who are not efficient and should not be there (Bidders 8 and 9), and one efficient bidder (Bidder 1) is not among the winners. Also, the number of winning bids in the final allocation exceeds that of the efficient allocation, which causes the bidders to incur unnecessary fixed costs. The difference between the GSM and the QSM is very clear, even though the efficient allocation consisted of only four bids. When the number of bids increases, one can expect the difference in the efficiency of the QSM auction and the GSM auction to increase.

In a way, the GSM can be considered as a generalization of the QSM. As a special case, the GSM will support a single bidder when it finds the optimal course of action, but it also provides the possibility of supporting any combination of bidders in each iteration. Hence, it has a substantially higher flexibility to improve the allocations. The benefits that can be obtained from this flexibility are, naturally, expected to increase as the number of bids in the efficient allocation increases. Also, the GSM alleviates the threshold problem and the extended puzzle problem, since it can offer bid suggestions to a group of "local" or "glocal" bidders to help them outbid a "global" bidder. Thus, the GSM should improve the allocative efficiency of the final allocation compared to the QSM in cases in which the efficient allocation consists of three bids or more.

VI COMBIAUCTION – IMPLEMENTING THE QUANTITY SUPPORT MECHANISM IN PRACTICE

The ultimate goal of our research project is to implement the bidder support tools in practice. Thus, we have designed and developed an Internet-based auction system, the CombiAuction, which enables combinatorial bidding, and to which the Quantity Support Mechanism (presented in Chapter 7) and the Price Support tool of Teich et al. (2001 and 2006) (formulation in section 8.2.3) have been integrated. Although per se insufficient, the Price Support tool is a good complement for the QSM. Price Support is faster to use, since it usually gives fewer suggestions than the QSM. Also, it is a helpful tool in situations in which the bidder has placed bids for good item combinations but the prices are too high.

The following chapter presents the CombiAuction system and the user interface. The auctions system has been tested with human subjects. The purpose of the tests was to study the feasibility of our auction system and the usability of the user interface, but also to study how easily people grasp the somewhat complex idea of combinatorial auction, and what kind of bidding strategies they use. Chapter 12 presents the experiment design and the results of the experiment. The contribution of this part of the thesis includes the implementation of the QSM and the design of the laboratory experiment. The most important contributions, though, come from the observations related bidder behavior. I present a classification of bidders according to their bidding strategies, and discuss the cognitive challenges of combinatorial auctions. Especially, analyzing the strategies used by the bidders in choosing the bids (prices and item quantities) is interesting, since identification of such strategies has not been done before.

11 THE COMBIAUCTION SYSTEM

The CombiAuction System is an Internet application designed by our research group[20] and developed by Valtteri Ervasti. An earlier application, NegotiAuction, had been developed by Teich et al. (2001, 2006) and used to run single-item, multi-unit auctions. In the development of CombiAuction, the database and the underlying data model were taken from NegotiAuction and developed further to suit the needs of combinatorial auctions. Aside from this, improvements in technology since NegotiAuction, and the fundamental differences between combinatorial and single-item auctions made it necessary make major changes in the application.

Combiauction is designed as a web application and written in Java according to the MVC (Model-View-Controller) architectural pattern. At the center of the application, a calculation package is written to handle all the calculations of the WDP and the QSP. A commercial software package (LINDO) is used to solve the problems. The data is stored in a MySQL database, and for access to the database, Hibernate is used in the persistence layer. For the Control and View parts of the architectural pattern, CombiAuction implements the widely used Apache Struts framework. The user interface consists of pages written in JSP, with the Struts tag libraries heavily utilized. The finished application is deployed in Apache Tomcat, a web container. Finally, the web container and the database were installed on a virtual server. With this configuration, CombiAuction can be accessed through the Internet using any web browser.

11.1 Organizing an Auction in the CombiAuction System

All auctions run on the CombiAuction system are continuous auctions, that is, the WDP is solved after each incoming bid. Some of the auction design parameters and auction rules are predetermined, but some the auction owner can specify. CombiAuction supports the organization of multiple-unit combinatorial auctions in both forward and reverse formats. Thus far, the only attributes the bidders can bid on

[20] Main contributors in the design apart from myself were professors Hannele Wallenius and Jyrki Wallenius and Valtteri Ervasti.

are the item quantities and bid price. Multiple attributes (quality, delivery terms, etc.) could be incorporated into the auction through "pricing out" as in NegotiAuction (Teich et al., 2001 and 2006). However, "pricing out" has not been implemented yet, since the focus has been on the combinatorial aspects of the auctions. In the following I will present the kinds of auctions one can organize in the CombiAuction system.

A new auction can be created by clicking on the link "Create a new Auction" in the menu bar. A new auction is created in four steps. First the owner provides the auction a name and a brief description, chooses the direction of the auction (reverse or forward), and specifies the number of items to be sold/bought (see Figure 3). Assume that in this example, the owner wants to create a reverse auction.

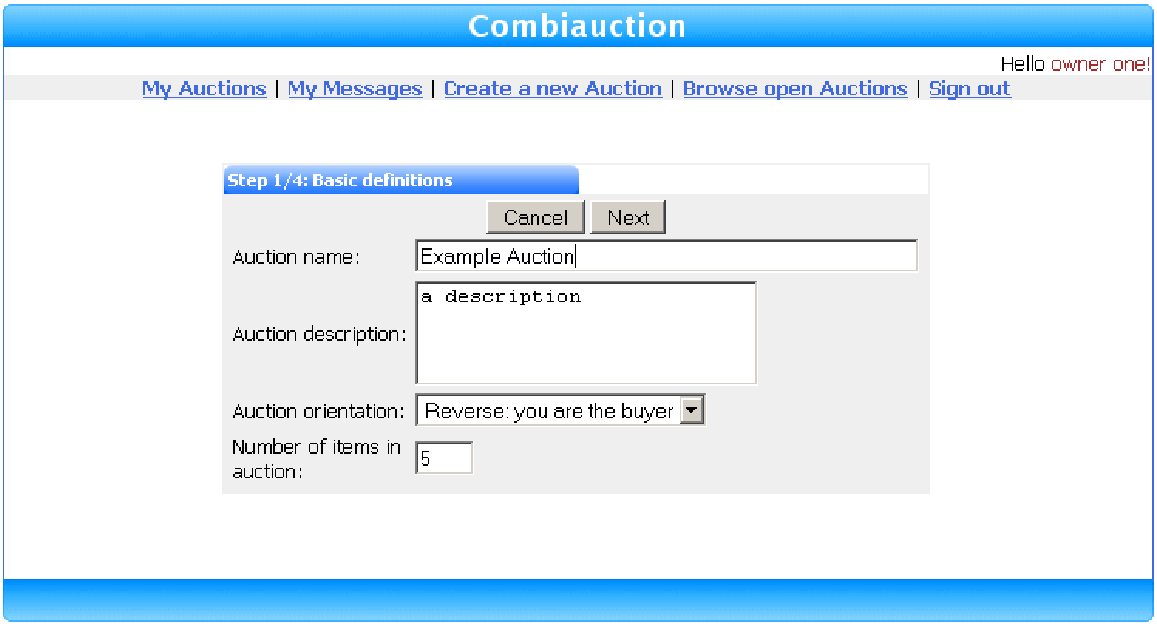

Figure 3 Screen view from Step 1 of the auction creation process

In the second step (Figure 4) the auction owner specifies the items, and the number and type (kilograms, liters, tones, units, etc.) of units of each item to be bought. The owner can give a description for each item, if she wants to, and specify a reservation price for the whole bundle. The reservation price is not revealed to the bidders. The owner can also specify a maximum number of units of each item that is allowed in each bid. If the owner restricts the maximum quantity offered by each bidder to be less than the total demand, the owner can make sure that at least two bidders are chosen as suppliers for that item. This way the buyer can decrease, for instance, her dependence on any particular supplier.

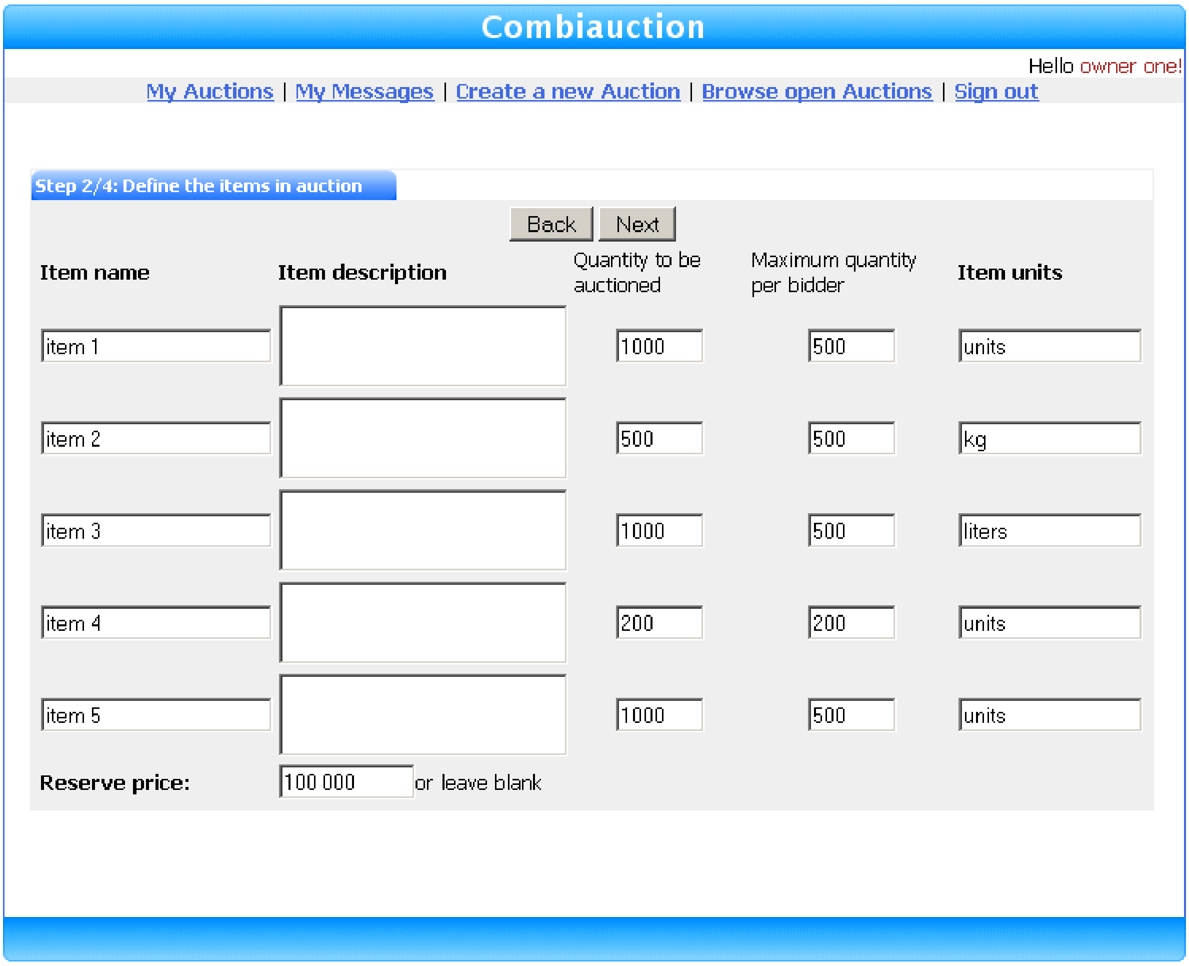

Figure 4 Screen view from Step 2 of the auction creation process

In the third step (Figure 5) the auction owner sets the opening and closing times for the auction. Although the owner specifies a closing time, the auction closing rule is flexible. Every time someone places a bid within the last 10 minutes of the auction, the closing time is pushed back by 10 minutes. The closing time is pushed back until bidding stops. This "soft" closing rule is often used to prevent bidders from sniping.

Besides the auction opening and closing times, the auction owner specifies the minimum bid decrement (increment in a forward auction) by which the total cost to the buyer must decrease when the provisionally winning allocation changes. As all the auctions are continuous, a round is defined as a change in the provisionally winning allocation. The auction owner can choose whether the bid support tools are available or not. More specifically, choosing "full quantity support" means that both quantity and price support are available. The two other alternatives are to enable only price support, or to offer no support at all. The owner can either organize an open-entry auction to which everyone registered in the CombiAuction system can participate in, or restrict participation by requiring that the owner has to confirm the bidders' request to participate ("Entry by confirmation"). The last design element to choose is to specify what information is revealed to the bidders (bid visibility). "Closed bidding" here refers to what I have called a semi-sealed-bid auction in this thesis. That means that bidders see their own bids and their statuses (active or inactive), but not the number or content of other bids. "Open bidding" means that all bids placed in the auction and their statuses are visible to all participants. If choosing an open auction, the owner can select whether the bidders' identities are kept hidden (second alternative), or revealed to everybody (third alternative). Regardless of the type of bid visibility chosen, the auction owner will always see all the bids and know which bidder placed them.

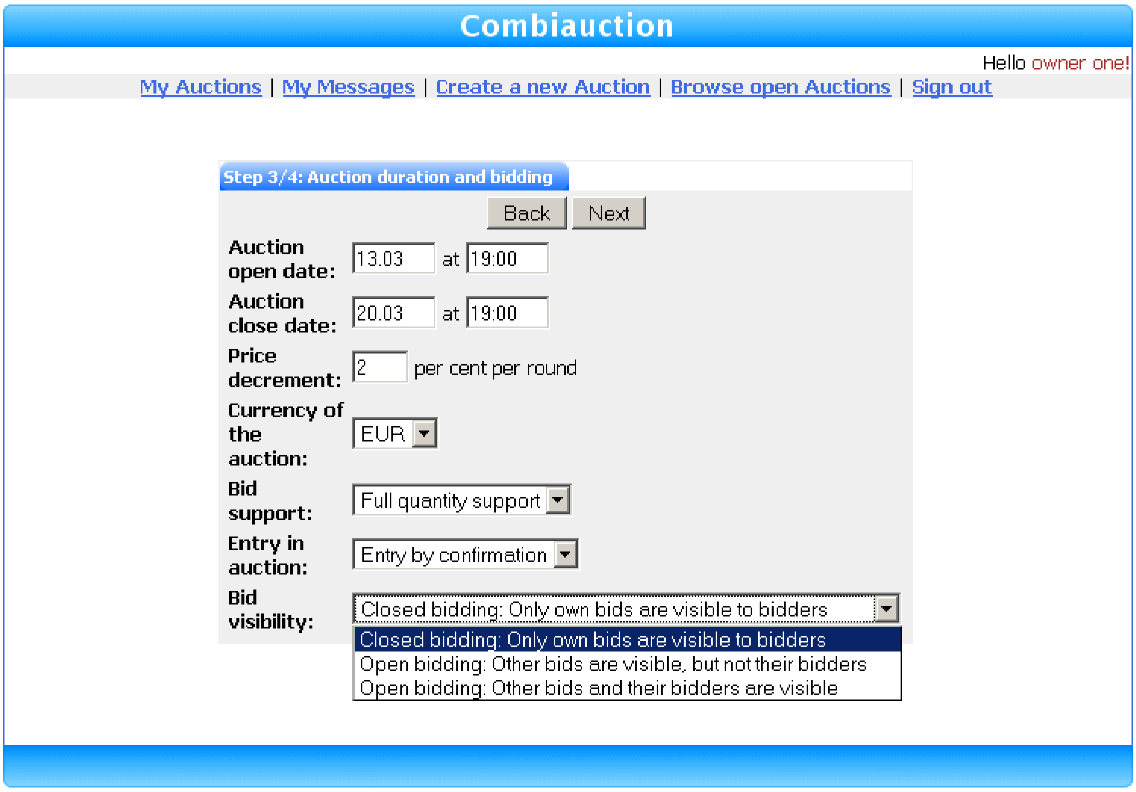

Figure 5 Screen view from Step 3 of the auction creation process

The last step simply pools together all the design information specified by the owner for a final confirmation before the auction is created (Figure 6).

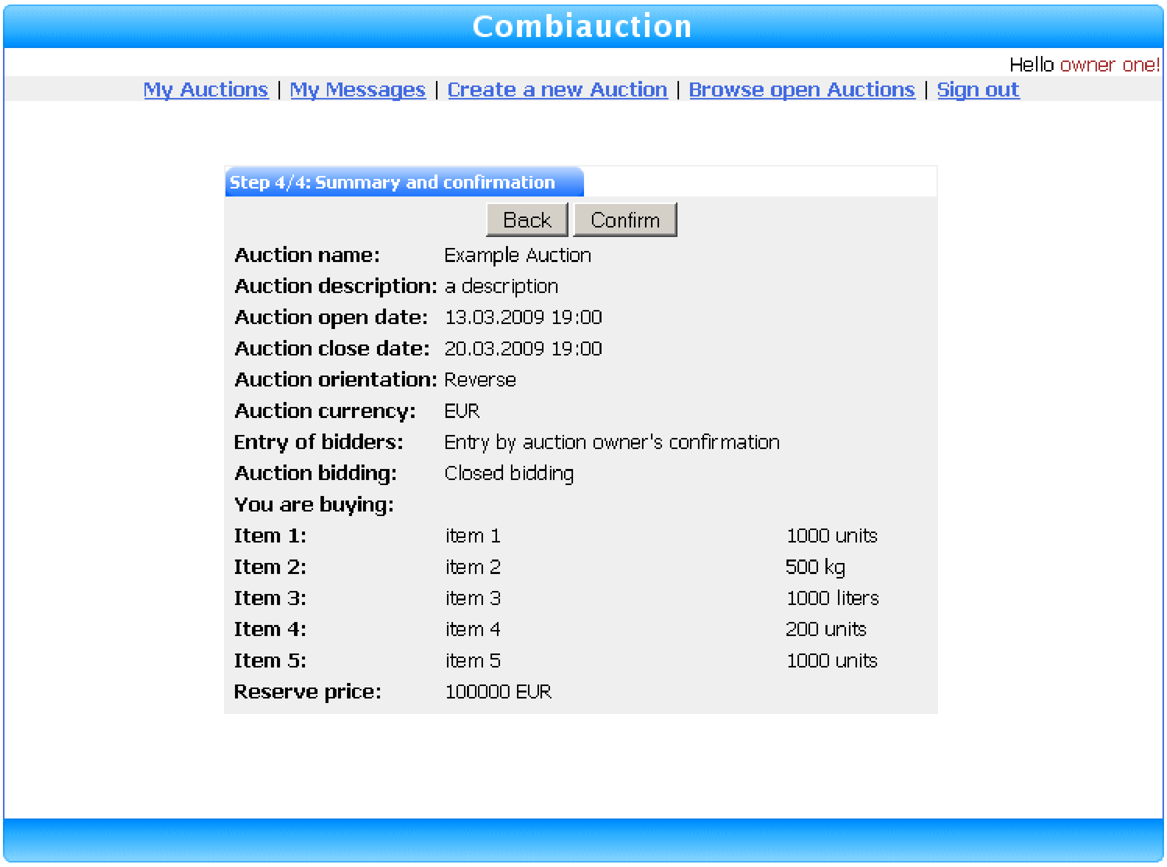

Figure 6 Screen view from Step 4 of the auction creation process

After the auction is created, the auction owner can monitor the progress of the auction from the owner's auction home page (Figure 7). The owner's auction home page contains four boxes, each containing information related to the auction.

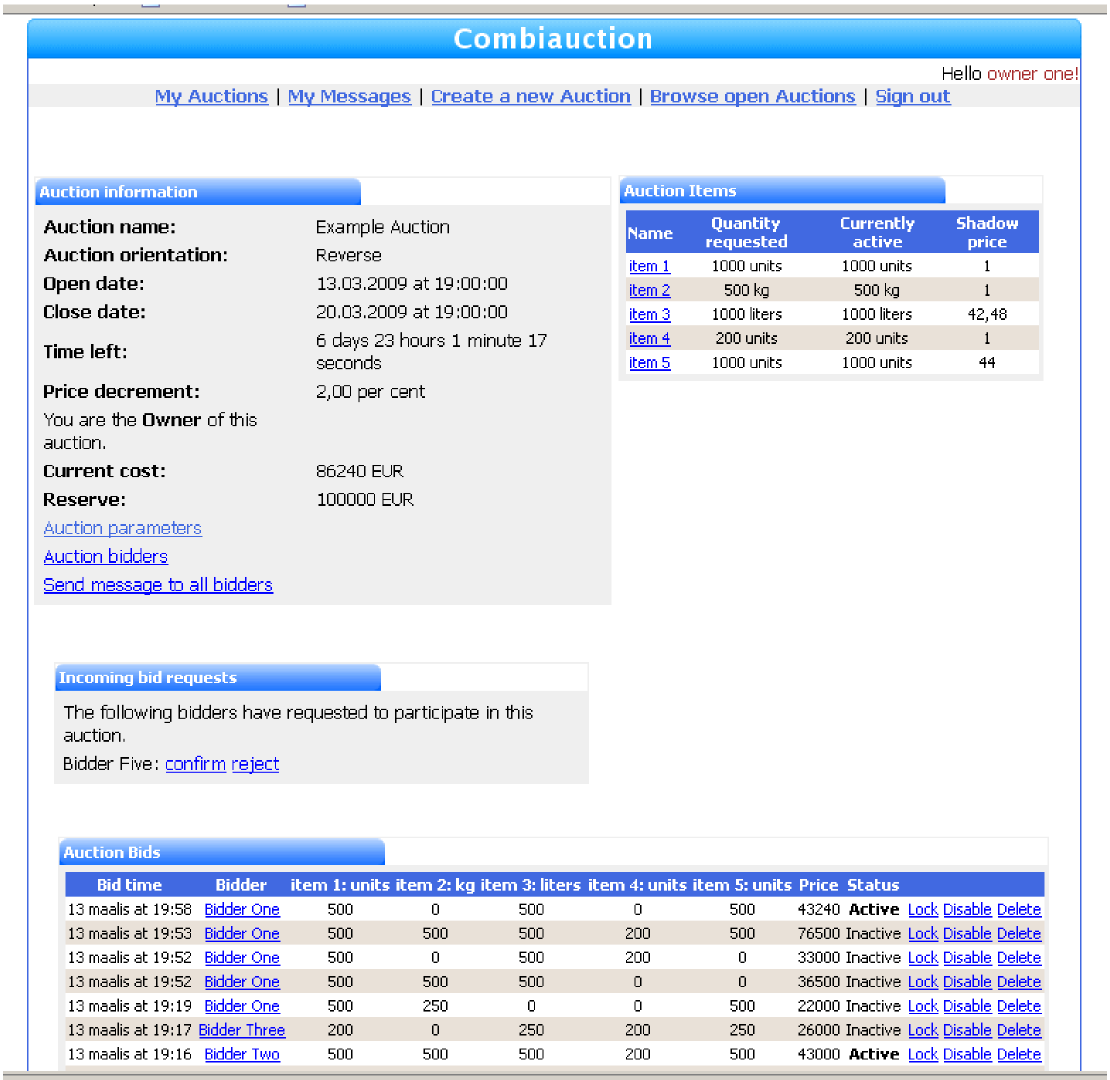

Figure 7 Screen view of the auction owner's auction home page

The “Auction Information” box reminds the auction owner of the opening and closing times, and the new closing time will also be updated there in case the closing time is pushed back due to late bidding. The owner can also check the current total cost from the “Auction Information” box. The link “Auction parameters” lets the owner review some of the auction parameters specified in Steps 2 and 3 of the auction creation process. Through the link the owner can make changes in the bid decrement, the reservation price, the auction closing time, and the minimum and maximum item quantities allowed in each bid. The link “Auction bidders” returns a list of bidders whose request to participate the owner has accepted. The owner can also conveniently send a message to all participants through the messaging system built in the CombiAuction system by clicking on “Send message to all bidders.”

The “Auction Items” box lists the items in the auction and item quantities demanded by the owner (buyer). By clicking on the item names the owner can check the item description given originally in Step 2 of the auction creation process. The box also informs the owner of the sum of item quantities in the currently active bids. As can be seen from screen shot in Figure 7, currently the demand is met, and there is no excess supply. The auction system does not allow shortages of item supply, but it is acceptable to have excess supply of the items in the active bids as long as the total cost is the lowest possible and the reserve price is not exceeded. The “Auction Items” box also reports the current shadow prices the QSM uses to solve for the bid suggestions. In case the shadow price for some item is zero, the shadow price is set to one because otherwise the item would have a zero cost and disappear from the objective function of the QSP.

The “Incoming bid requests” box lists all the bidders wishing to participate in the auction. The owner can either confirm or reject the bidder’s request. The “Incoming bid requests” box disappears when there are no requests.

The “Auction Bids” box lists all the bids placed in the auction with the newest bid on top. The owner can see the time the bid was placed, the bidders’ names, the content of the bids, and the bid statuses. The statuses (active or inactive) will change to “won” or “lost” once the auction closes. The owner can “lock” some of the bids, which means that those bids are guaranteed winners regardless of what other bids may enter the auction. This feature was incorporated already in *NegotiAuction*, and has been included in the *CombiAuction* as well. The ability to lock some bids prior to the end of the auction brings the auction closer to a negotiation. For instance, a supplier could be pressed for time and promise to make a good offer under the condition that the buyer makes an immediate decision whether to accept it or not. The owner can also “disable” bids, which prevents the bids from becoming active, or “delete” them entirely from the bid stream. By clicking on the names of the bidders the auction owner can send messages to the bidders.

11.2 Bidding in the CombiAuction System

A bidder registered in the CombiAuction system can browse through the list of open auctions and auctions scheduled to open in the future (Figure 8).

Figure 8 Screen view of the list of open and scheduled auctions.

By clicking on the name of the auction, the bidder can look at the details of the auction (Figure 9).

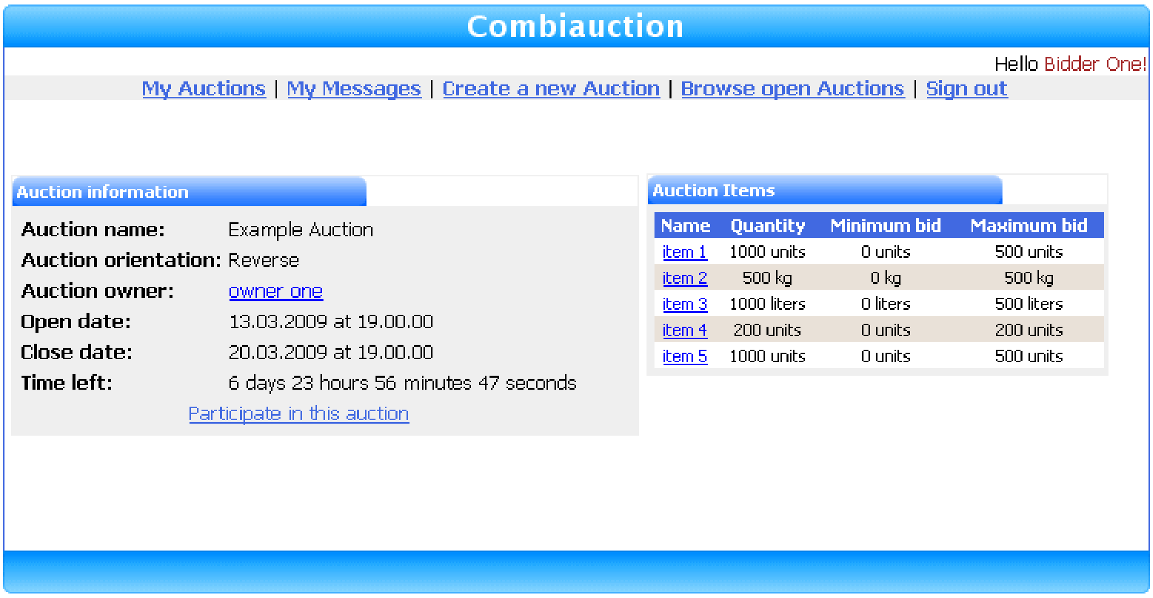

Figure 9 Screen view of the auction information page

The bidder can see the items and item quantities demanded by the owner (buyer), and by clicking on the names of the items, the bidder can see the item descriptions given by the owner. If the bidder wishes to participate in the auction, she will click on the "Participate in the auction". If it is an open-entry auction, the auction will be directly added to the bidder's "My Auctions" page (Figure 10). If participation requires the auction owner's participation, clicking on the link will send a request to the owner. After the owner confirms the request, the auction is added to the bidder's "My Auctions" page. Bidders can enter open auctions at any point in time; they are not required to be present from the beginning to the end. Also, because of this free entry, no activity rules are enforced in the CombiAuction.

Figure 10 Screen view of the bidder's "My Auctions" site

Once a participant in an auction, the bidder is directed to the bidder's auction home page (Figure 11), where she can review the details of the auction (which are the same presented in on the auction information page for prospective participants in Figure 9), monitor her status in the auction, and submit new bids.

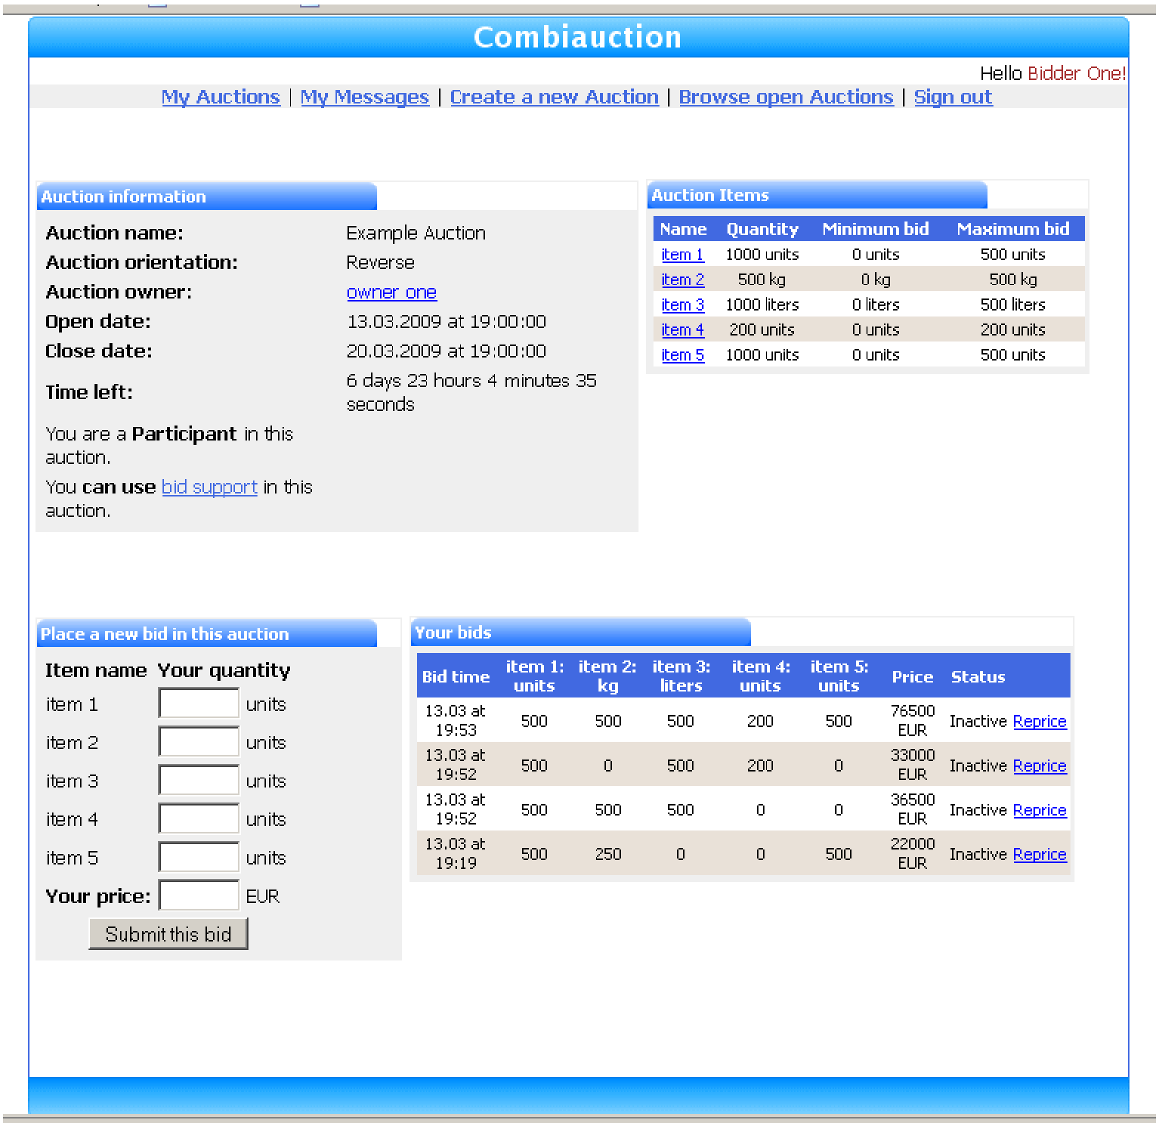

Figure 11 **Screen view of the bidder's auction home page**

The bidder can submit bids in three ways. First, the bidder can place a new, “self-made” bid by filling out the item quantities and the bid price in the dialog box in the lower left hand corner. In CombiAuction, all bids are allowed to enter the bid stream. In other words, the bidder is not required to place a bid that becomes active upon submission. Also inactive bids are accepted. This is to help overcome the threshold problem. If the bidder only wants to change the price on one of her existing bids, she can use the “Reprice” option next to the bid. Thirdly, the bidders can use the two support tools, the Quantity Support Mechanism and the Price Support tool to help them submit bids – if the auction owner has included the tools in the auction. Notice, however, that the bidder support tools are never available at the beginning of the auction. There has to be some bids in the bid stream before the support tools can find bids that can become active.

When the support tools are available, the link appears in the "Auction Information" box. The link leads to the support dialog page (Figure 12).

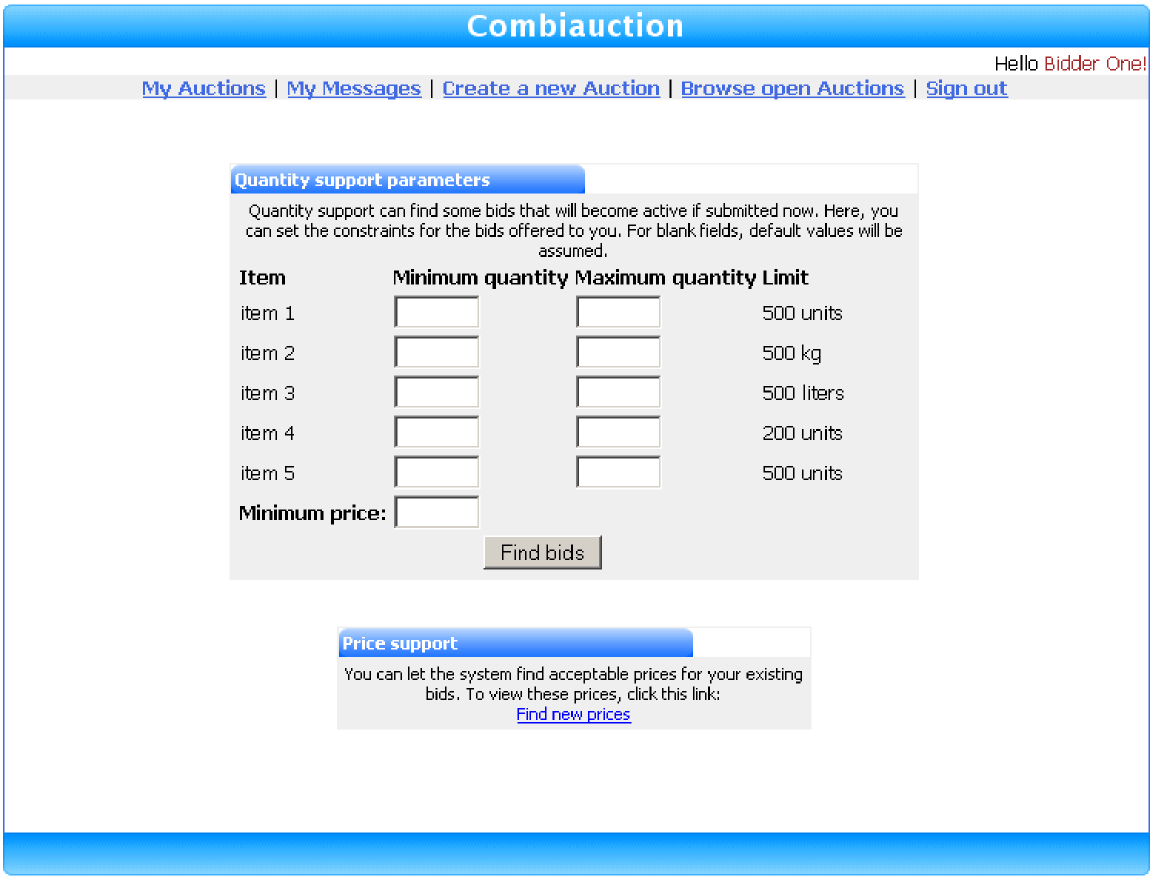

Figure 12 Screen view from the support tool dialog page

The link to the price support tool is at the bottom of the page. Simply by clicking on the link "Find new prices", the auction system will offer the bidder a list of those of her inactive bids that can become active with a lower price (Figure 13).

Figure 13 Screen view of the Price Support tool

In this example, only one of the bidder's existing bids can become active. The bidder has the option of submitting this bid suggestion, or returning to the auction home page without placing a bid. If the bidder submits the bid, it is added to the bid stream and it will appear under "Your Bids" on the bidder's auction home page.

The Quantity Support tool (upper box in Figure 12) allows the bidder to specify some constraints on the suggestions the QSM will suggest. The bidder can restrict the maximum amount of each item, which is convenient if the bidder is operating under a capacity constraint. The bidder can also specify a lower limit for item quantities and the bid price in case she is interested only in bids of a certain size, or wants to bid on some items in particular. If the bidder does not insert any specifications, the QSM will use zero as the default value for the lower limits and the bid price, and the maximum allowed item quantity defined by the auction owner as the default value for the upper limit. The shortlist in the QSM is the "full" short list described in section 8.2.2.

The bid suggestions offered by the QSM are offered to the bidder (Figure 14), and the bidder can evaluate the suggestions and submit the one that is the most profitable to her (if any). If the bidder submits a bid, it is added to the bid stream and it will appear under "Your Bids" on the bidder's auction home page.

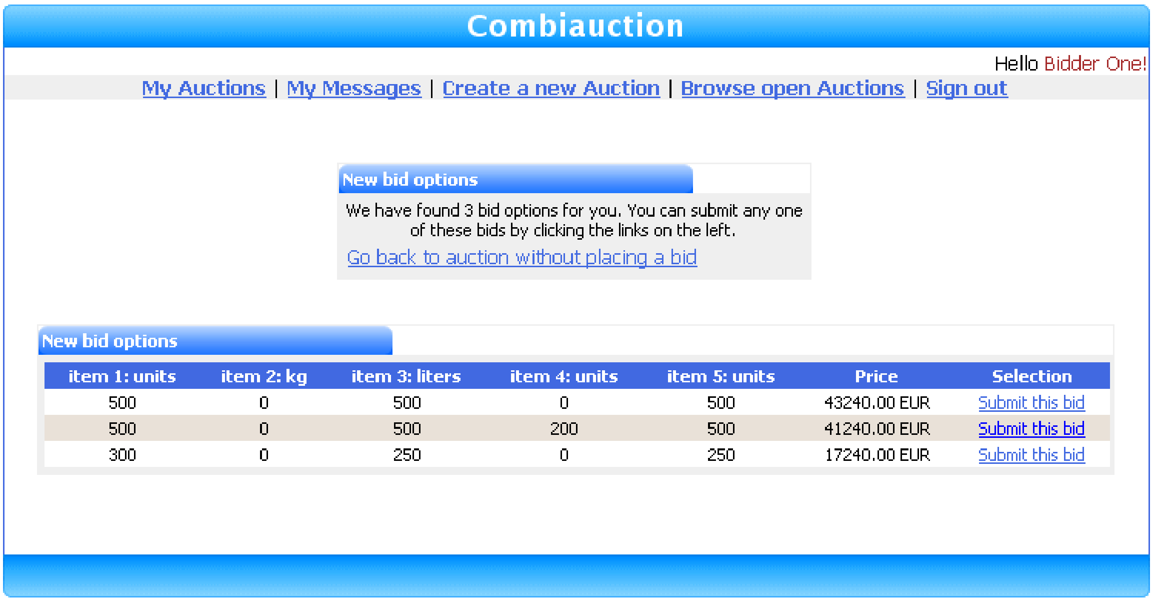

Figure 14 Screen view of the bid suggestion made by the QSM

The auction continues until the closing time, or until no new bids have been submitted for 10 minutes.

12 TESTING THE COMBIAUCTION SYSTEM[21]

The CombiAuction system was tested in an experiment with human subjects. The objectives of the experiment were to 1) test the feasibility and usability of the CombiAuction system, 2) study how the auction outcomes (efficiency and total cost to buyer) compare with those obtained in the second simulation study, 3) study what kind of strategies bidders use, and 4) how human users understand the concept of a combinatorial auction and the support tools.

12.1 Experiment Set-Up

In the experiment, two different auction settings were used. The designs for both settings were chosen from the second simulation study (presented in section 8.2) to allow for direct comparison to the simulated auctions. Also all the parameter values were set to be the same as in the simulations studies (5 items, 600 units of each item, 2% decrement). In the simulation study we tested 48 different designs, which could not have been reproduced with a limited number of human subjects. Thus, I chose two of the designs, one design with equal capacities, and one with unequal capacities, since based on the simulation study it was the bidders' capacities which had the biggest impact on the efficiency of the auction outcome.

For the first experiment setting – auction A – I chose an equal capacities auction with 15 bidders, normal economies of scope, and full shortlist. The initial bids were created based on advantage in variable costs (Bid1). In the simulation study, the efficient allocation changed from replication to another (there were 50 replications of each design). For the laboratory experiment I chose only one replication, that is, one set of cost functions. Using the same cost functions for each group of participants allowed for better comparisons between the groups. I deliberately chose a replication which had ended in the efficient allocation in the simulation in order to see if the auctions would end in the same, efficient allocation with real users as well.

[21] I am solely responsible for the design and conduct of the laboratory experiment as well as for the analysis of the results.

For the second setting (auction B) I chose the corresponding unequal capacities auction (15 bidders, normal economies of scope, full shortlist and initial bids created based on Bid1). However, this time I chose a replication in which the winning allocation was inefficient (winners' combined production costs were 9.2% above the costs of the efficient allocation), in order to study whether human users could direct the auction to a more efficient allocation or not.

Experiment Participants

In total 74 students – both undergraduate and graduate students – participated in the experiment. The students were all participants of a course on managerial economics at Helsinki University of Technology. Thus, the students were already familiar with the concepts of production and cost functions, and economies of scale and scope. However, only 14 of them had participated in online auctions before, and the combinatorial auction was an unfamiliar concept to all of them. Each student was required to participate in two auctions: first in one A auction and then in one B auction. All 74 participants bid in the first (A) auction, but only 69 of them bid in the second (B) auction. The students were rewarded by giving them points that counted towards their final grade from the course. In the beginning each student received 5 points, and they were rewarded extra points for playing well (winning with a profit larger than their counterparts), and for answering a post-experiment questionnaire. In case the students did not perform as expected (for example, placed unprofitable bids or failed to win when they had the chance to), points were deducted from them. The maximum score attainable was 11 points, which is 11% of the final grade from the course. Giving credit for participation is an easy and cheap way of motivating students to participate in laboratory experiments.[22]

Prior to the experiment, all participants were required to participate in a briefing session. The briefing session contained some general theory on combinatorial auctions. Also, the students were briefed on how to use the CombiAuction system, how the support tools work, and how to use the cost function parameters given to them on an

---

[22] Also Bichler (2000) gave students credit for participating in his experiments.

Excel sheet to calculate their costs for each bundle. It was explained to the students how the winners are determined in a reverse combinatorial auction, and what principles the QSM is based on (and why it is not available at the beginning of the auction). However, no mathematical notation or formulations of the WDP or the QSP were presented.

The Organization of the Experiment

The experiment was organized in four sessions during four consecutive days so that there were two sessions (days) for A auctions: A1 and A2, and two sessions (days) for B auctions: B1 and B2. The reason for having two sessions for both auctions was to offer the students the possibility to choose the days that suited their schedule the best. Each auction lasted for 23,5h (or longer if the closing time was extended). The participants were physically in different locations during the auction, but participated over the Internet. The duration of 23,5h is longer than the 1-4h duration usually used in laboratory experiments (in which the participants usually are in the same place at the same time). However, I chose such a long duration to better simulate an actual online auction in which the endogenous arrival of bidders is a key characteristic. Real online auctions usually last for several days or weeks, partly in order to give time for potential buyers to find the auction. However, since in this experiment the bidder pool was predefined and the prospective bidders were all knowledgeable about the auction, there was no need to extend the auction longer than was needed to be able to observe different bidding strategies. The auctions began at 5pm and the scheduled closing time was 4.30 pm the following day. This way the bidders had essentially two days time to place bids, which I anticipated to be enough to separate early bidders from late bidders.

Even though the simulated auctions, from which I borrowed the design parameters, cost functions and capacities, all had 15 bidders, I chose to divide the students into smaller groups. There were two reasons for this. Firstly, that way I could get more replications of the auctions. Then, if for some reason some of the auction outcomes were distorted (e.g. due to mistakes made by the bidders) there would be enough data left to analyze. Secondly, I thought it would be more rewarding for the students if a bigger portion of them could win at least once during the experiment. For the A auctions, 46 students enrolled in session A1, and 28 students in A2. Thus, I divided the A1 participants into 9 groups of 5 students, and one group of 6 students, and A2 participants into 4 groups of 7 students. Five students were enough to guarantee sufficient competition in equal capacities auctions in which the efficient allocation consisted of two bidders. Even if bidders placed bids for smaller bundles than maximum capacity, it would be very unlikely that all five bidders would be provisional winners simultaneously. For the unequal capacities auctions five bidders might not have been enough, since the bidders' capacities were smaller than in the equal capacities auctions, and more bidders would always be needed in the winning allocations. Thus, the 40 participants in the B1 session were divided into 4 groups of 7 students and 2 groups of 6 students, and the 24 participants in the B2 session were divided into 4 groups of 7 students and one group of 6 students.

The students were assigned cost functions to identify them. From the 15 bidders in the simulations I chose 5-7 bidders to be assigned as identities to the students. The same bidder identities (cost functions) were used in all the groups to allow for direct comparisons between groups. This way I could compare the auction outcomes (efficiency and total cost to the buyer) to see if they differed from one group to another – even though the starting points were equal. Also, when the same bidder identities were used in each group, the bidders' performance could be compared with corresponding bidders in other groups. This comparison determined the participants reward points. For the equal capacities (A) auctions with five bidders I chose the cost functions of the two bidders, who formed the efficient allocation (Bidders 12 and 15)[23]. In addition I chose the cost function of one bidder (Bidder 10) whose cost for the efficient allocation bid (300, 300, 300, 300, 300) was close enough to the efficient bidders so that a "pseudo efficient" allocation could be reached. A pseudo efficient allocation was defined as an allocation which was not efficient, but in which the total cost to the buyer was within the 2% decrement of the efficient production costs. In the pseudo efficient allocation the efficient bidder(s) cannot afford to submit new bids because in order to reduce the buyer's cost by 2% they would have to incur a loss. The

---

[23] The bidders' numbers refer to their number in the simulated auction, which makes is easy to keep track of the bidders' costs.

two remaining bidders whom I chose (Bidders 5 and 8) had higher costs, and they should not be among the winners. When there were 6 or 7 participants in the auction, the extra bidders were assigned costs that could also result in pseudo efficient outcomes (Bidders 1 and 11), but who had higher costs than Bidders 12, 15 and 10. The bidders' cost function parameters are presented in Appendix 4.

In the unequal capacities (B) auctions I included the bidders from the efficient allocation (Bidders 3, 8 and 13), the bidders who won in the simulated auction (Bidders 4 and 15 in addition to Bidder 3), and the remaining one or two bidders were chosen at random to be Bidders 1 and 9. In the unequal capacities case it would have been impossible to try to deduce which bidders could create a pseudo efficient allocation, since the item combinations in the winning bids also would have to change due to the different capacities (in the equal capacities case the winning bids would almost always be for half the total demand). The bidders' cost function parameters and capacities are presented in Appendix 4.

The organization of the experiment is summarized in Figure 15. The participants were divided into A1, A2, B1 and B1 sessions according to their preferences. I then further divided them into smaller auction groups within each session. In each auction there were the same set of cost functions (bidder identities) given to the participants. The cost functions were assigned randomly in the A auctions, but in the B auctions I tried to give better cost function to those who had received the worst ones in the A auction. This way I tried to give everyone equal chances to obtain extra points. Everybody participated first in one A auction and then in one B auction. The equal capacities (A) auctions were slightly simpler bidding environments, hence they also served as a practice session for the B auctions. The participants were grouped differently in the A and B auctions. The idea was to have bidders bid against new competitors – who maybe used a different strategy – in the second auction. Also, due to the participants' diverse preferences and different group sizes, it would have been impossible to maintain the same groups in both A and B auctions.

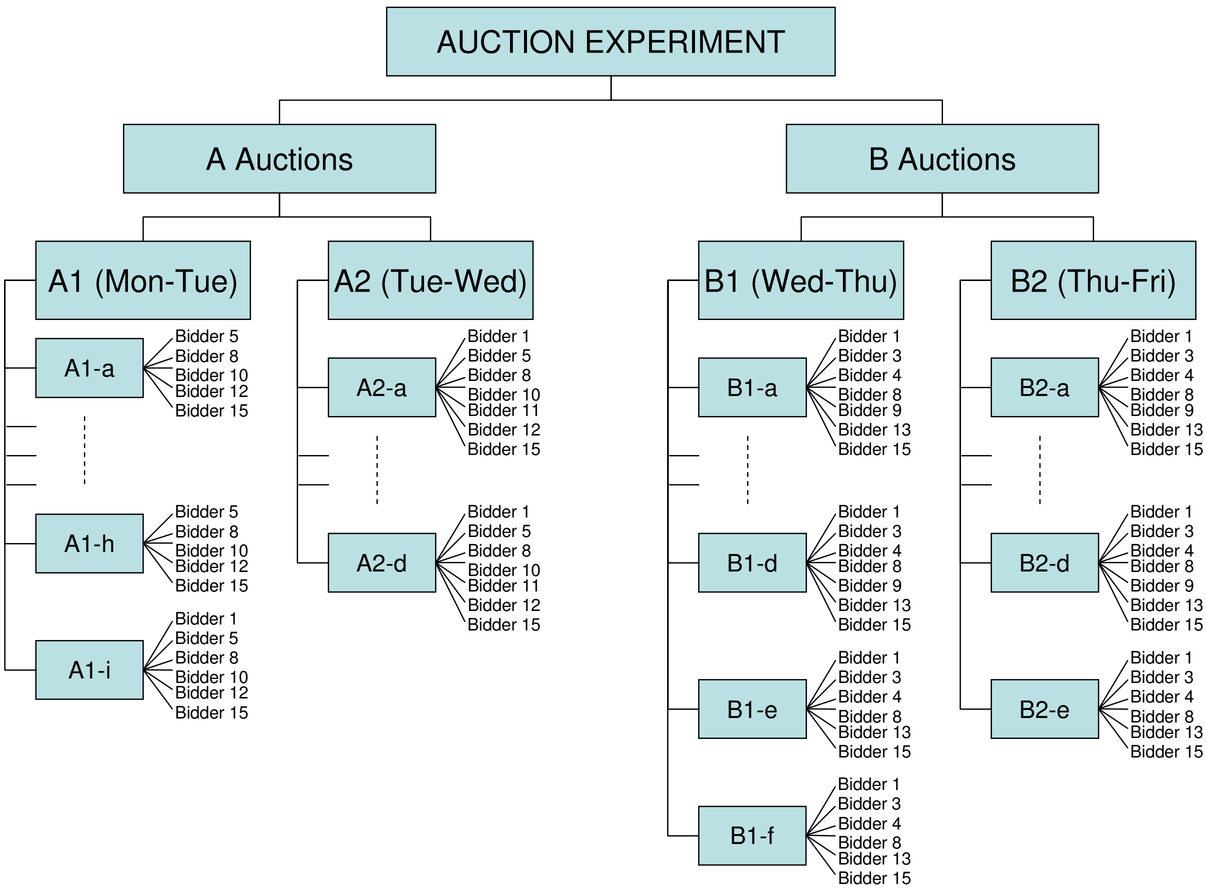

Figure 15 The organization of the laboratory experiment

12.2 Results of the Auction Experiment and Observations

In this chapter I will first compare the outcomes of the laboratory auctions to the simulated auctions. Thereafter I will discuss bidder behavior (strategies), bidders' understanding of the combinatorial auction and the support tools, and the usability of the auction system. The discussion is based on my observations from the auctions and the answers the participants gave to a post-experiment questionnaire after completing the auctions. The post-experiment questionnaire can be found in Appendix 5. Out of the 74 participants, 66 returned the post-experiment questionnaire.

12.2.1 Efficiency of the Final Allocation and Total Cost to Buyer

The efficiency of the final allocation and the total cost to the buyer were the key performance measures studied in the second simulation study. Thus, they served as a convenient starting point for the analysis of the laboratory experiment as well.

Equal Capacities (A) Auctions

The simulated version of the equal capacities (A) auction ended in an efficient allocation. Table 30 presents the winning allocations of the simulation, and of the laboratory experiments. Auctions A1-d, A2-a, and A2-c were left out from the table, because in those auctions one of the winning bidders had mistakenly placed a bid with negative profit, which distorted the final outcome. Thus, studying the efficiency or the total cost to the buyer in these auctions does not make sense. The efficiency could be really bad and the total cost to the buyer still low because of the unprofitable bids. These auctions are taken into consideration in the next sections, which discuss bidder behavior.

Table 30 Winning bidders and bids in the equal capacities (A) auctions

	A1-a		A1-b		A1-c		A1-e		A1-f		A1-g	
Bidder	B 12	B 15	B 8	B 15	B 10	B 15	B 12	B 15	B 12	B 15	B 5	B 12
Item 1	300	300	300	300	300	300	300	300	300	300	300	300
Item 2	300	300	300	300	300	300	300	300	300	300	300	300
Item 3	300	300	300	300	300	300	300	300	300	300	300	300
Item 4	300	300	300	300	300	300	300	300	300	300	300	300
Item 5	300	300	300	300	300	300	300	300	300	300	300	300
Price	142,214	138,000	145,100	137,349	141,979	140,138	158,232	140,000	150,000	139,000	147,000	140,000
Prod. cost	139,417	137,348	145,092	137,348	141,979	137,348	139,417	137,348	139,417	137,348	146,881	139,417
Profit	2,797	652	8	1	0	2,790	18,815	2,652	10,583	1,652	119	583

	A1-h		A1-i		A2-b			A2-d		Simulated	
Bidder	B 5	B 15	B 12	B 15	B 11	B 12	B 15	B 10	B 15	B 12	B 15
Item 1	300	300	300	300	300	300	0	300	300	300	300
Item 2	300	300	300	300	0	300	300	300	300	300	300
Item 3	300	300	300	300	300	0	300	300	300	300	300
Item 4	300	300	300	300	300	300	0	300	300	300	300
Item 5	300	300	300	300	300	300	0	300	300	300	300
Price	147,924	144,042	143,000	137,493	118,141	110,643	59,000	141,979	138,344	142,982	142,152
Prod. cost	146,881	137,348	139,417	137,348	117,003	109,521	58,071	141,979	137,348	139,417	137,348
Profit	1,043	6,693	3,583	145	1,138	1,122	929	0	996	3,565	4,804

Out of the 10 auctions presented in Table 30, four auctions ended in the efficient allocation (A1-a, A1-e, A1-f and A1-i). In addition, two auctions ended in a "pseudo efficient" allocation (A1-c and A2-d). In these two auctions the inactive efficient bidder (Bidder 12) should have placed an unprofitable bid in order to decrease the total cost to the buyer by the required 2%. Had the decrement been smaller, the true efficient allocation might have been reached.

Out of the remaining four auctions, which ended up in an inefficient allocation, three auctions (A1-b, A1-g and A1-h) were such that at least the efficient bidder would have been able to become a winner with a profitable bid. It is possible that the reason for them not placing the bid is that the QSM did not suggest it to them. Another possible reason is that the “shortlist” provided by the QSM was too long, and the bidders did not have time to calculate the profits for all the bid suggestions – or did not want to go through the trouble. Some participants complained that towards the end of the auction, when most bidders were logged into the system, the QSM became slow, and that it took several minutes for it to provide a list of suggested bids. This most likely has affected the auction outcomes. This could also explain why the efficient bidders failed to place their winning bids. The auction A2-b on the other hand, seems to have suffered from the threshold problem. Bidder 15 placed her bid early in the auction, and directly with a relatively low profit. There bids belonging to the efficient allocation were also in the bid stream, but because the price was very high, the incoming bidders could not afford to team up with them, but rather teamed up with Bidder 15’s cheap bid.

Besides efficiency, also the total cost to the buyer is of interest in the auction outcome – especially for the buyer. Table 31 summarizes the total cost to the buyer (as the ratio of total cost to buyer and efficient production cost), as well as the efficiency indicator (see definition in section 8.2.4.2) and the efficiency status of the auctions.

Table 31 Total cost to buyer and efficiency indicators from the A auctions

	A1-a	A1-b	A1-c	A1-e	A1-f	A1-g	A1-h	A1-i	A2-b	A2-d	Simul.
Total cost to buyer (ratio to efficient cost)	1.012	1.021	1.019	1.078	1.044	1.037	1.055	1.013	1.040	1.013	1.030
Efficiency indicator	1	1.021	1.009	1	1	1.034	1.027	1	1.028	1.009	1
Efficiency status	eff.	ineff.	pseudo eff.	eff.	eff.	ineff.	ineff.	eff.	ineff.	pseudo eff.	eff.

According to Table 31 two of the auctions (A1-e and A1-f), which ended in the efficient allocation, resulted in a cost to the buyer higher than in the simulated auction. The explanation to this is that actually in these two auctions (contrary to the simulated auction), there were still some bidders who could have afforded to submit a provisionally winning bid, but for some reason they did not. Their answers to the questionnaire indicate that one problem was the slow performance of the QSM at the end of the auction. One of the bidders had not even tried using the QSM, but resorted to placing self-made bids, which did not become active. Interestingly, the two remaining efficient auctions (in which other bidders could not have afforded to place a winning bid), both pseudo efficient auctions and even one inefficient auction ended in a total cost to the buyer which was lower than in the simulated auction. This is because in the experiment, some bidders bid unnecessarily low. They placed bids for very low profit margins already quite early in the auction, when it would not have been necessary to become a provisional winner. Either the bidders were very risk averse and tried to maximize their probability of winning, or they did not want to constantly monitor the auction and keep bidding, or then they did not want to use the QSM (or it was not available yet). The bidders' behavior is discussed further in the following sections (12.2.2 and 12.2.3).

All in all, the QSM seemed to have performed quite well in the hands of human users in the equal capacities auctions. The efficient and pseudo efficient allocation was reached 6 times out of 10, and the efficiency indicators of the inefficient final allocations are quite small (Table 31). The heavy bidding activity right before the auction closing caused some problems for some of the bidders. This, however, is a computational or a server-related issue, and not due to the structure of the QSM. This observation does indicate though that the shortlist should be kept relatively short. Even with five bidders there can be dozens of bids in the bid stream, and hence the shortlist can become quite long. The threshold problem in auction A2-b shows that the content of bids in bid stream and the order in which bids are submitted affects the auction outcome.

One source of discrepancy in the results of the experiment auctions and the simulations is the fact that the real auctions differed from the simulations on a few design issues. First of all, in the simulated auctions the bidders did not place self-made bids after the QSM became available. This affects the content of the bid stream a lot – especially since the bidders in the laboratory experiment were more eager to submit self-made bids throughout the auction than we had anticipated. Secondly, in the simulations the bidders always submitted the most profitable bid on the shortlist. This is not necessarily always the case in the laboratory experiment, especially when the shortlist is long. Unfortunately, the shortlists are not recorded in the auction system, and hence there is no way to verify afterwards whether the bidders chose the profit maximizing bid or not. The system will be changed for the next experiment. Thirdly, all the bidders were present at the same time in the simulated auctions, hence all the inactive bidders were as likely to be the one submitting the next bid. In the laboratory experiment this was not the case, as some of the bidders waited as long as until the last hour to start bidding. However, one should also keep in mind that there is a random element in the simulated auction as well (the inefficient bidder using the QSM is chosen randomly), but the auctions were not repeated with the same cost functions. Had the same auction been simulated several times, the outcomes might not have been the same every time.

Unequal Capacities (B) Auctions

For the unequal capacities (B) auctions I had chosen an auction which in the simulation study had ended up in an inefficient allocation. Table 32 presents the winning allocation of the simulation and of the laboratory experiments, and also the efficient allocation. Auctions B1-e, B1-f, B2-a, B2-d. and B2-e were left out from the table, because in those auctions one of the winning bidders had mistakenly placed a bid with negative profit, which distorted the final outcome of these auctions. Also auction B1-a was left out because one participant, who had been given the identity of Bidder 3 (one of the efficient bidders) did not bid in the auction, even though she had requested for participation, and had been accepted as a participant. These auctions are considered in the sections on bidder behavior (sections 12.2.2 and 12.2.3).

Table 32 Winning bidders and bids in the unequal capacities (B) auctions

	B1-a			B1-b			B1-c			B1-d		
Bidder	B 4	B 8	B 13	B 1	B 4	B 8	B 3	B 8	B 15	B 1	B 4	B 15
Item 1	300	300	0	0	300	300	150	300	150	0	300	300
Item 2	150	150	300	225	150	225	0	300	300	200	150	250
Item 3	300	0	300	225	300	75	300	150	150	200	300	100
Item 4	300	0	300	0	300	300	0	300	300	0	300	300
Item 5	300	0	300	300	300	0	150	300	150	300	300	0
Price	135,000	56,454	120,000	80,000	134,000	100,047	73,000	163,467	123,892	73,000	133,300	109,554
Prod. cost	133,270	53,429	116,897	73,899	133,270	97,340	71,619	137,291	115,615	71,179	133,270	101,048
Profit	1730	3,025	3,103	6,101	730	2,707	1,381	26,176	8,277	1,821	30	8,506

	B2-b			B2-c			Simulation			Efficient allocation		
Bidder	B 3	B 8	B 13	B 3	B 8	B 13	B 3	B 4	B 15	B 3	B 8	B 13
Item 1	300	300	0	300	300	0	300	0	300	300	300	0
Item 2	300	0	300	300	0	300	300	0	300	300	0	300
Item 3	300	0	300	150	150	300	300	150	150	300	0	300
Item 4	0	300	300	0	300	300	150	150	300	0	300	300
Item 5	0	300	300	0	300	300	150	300	150	0	300	300
Price	91,437	87,266	130,138	79,067	106,000	130,000	133,401	69,491	130,067			
Prod. cost	85,432	86,943	116,897	76,980	100,546	116,897	124,101	67,690	124,077	85,432	86,943	116,897
Profit	6,005	322	13,241	2,087	5,454	13,103	9,300	1,801	5,990			

Interestingly, in two auctions (B2-b and B2-c) the winning bidders were the same as in the efficient allocation; and in B2-b even the bids were identical to the efficient allocation bids. According to the bidders' answers in the questionnaire, it was the support tools that helped guide the auction to the efficient allocation. One of the winning bidders in the B2-b auction had used price support, and the other two quantity support to find the winning bids. Looking at the other auctions it is easy to see that the winning bids vary a lot from one auction to another. This, and the fact that in the simulated auction the QSM did not find the efficient allocation, demonstrates the significance of the bids the bidders place without the help of the QSM in shaping the progress and ultimately the outcome of the auction. Also the fact that two of the three winning bids in auction B2-c were placed before the QSM was available attests to the significance of the initial bids in the bid stream.

A better comparison between the experiment auctions and the simulation can be done by studying the total cost to the buyer and the efficiency of the final allocations (see Table 33).

Table 33 Total cost to buyer and efficiency indicators from the B auctions

	B1-a	B1-b	B1-c	B1-d	B2-b	B2-c	Simul.
Total cost to buyer (ratio to efficient cost)	1.077	1.086	1.246	1.092	1.067	1.089	1.151
Efficiency indicator	1.050	1.053	1.122	1.056	1	1.018	1.092

Both the total cost to the buyer and the efficiency of the winning allocation are better in the laboratory experiments than in the simulated auction with the exception of auction B1-c. The inefficient winning allocation and high cost to buyer in auction B1-c can be explained through lack of competition. There were seven bidders registered in the auction, but only five of them placed bids. In all the other auctions in Table 33 all bidders registered for the auctions also placed bids. In addition, out of those five bidders who submitted bids in the B1-c auction, one bidder (Bidder 13) bid only at the beginning of the auction. She placed several bids before anyone else had entered the auction, but did not bid at all after other bidders started bidding. Thus, as the closing approached, there were three active bidders and only one inactive bidder attempting to submit bids.

An explanation for the better outcomes of the laboratory auctions compared to the simulated auction can be found by observing the bidders' bidding behavior in the auctions. In auctions B1-a, B1-b, and B1-d two out of the three winning bids were placed without the help of the support tools. Also, these bids were placed relatively early in the auction (sometimes already before the demand was fulfilled), and the profits in these bids were relatively low. In other words, the bidders were selling themselves short not knowing that much higher profits could be gained. In B1-d, the highest profit was made by the bidder, who was the last to bid – and who was the only one of the winning bidders to have used the QSM. In B2-c the situation was the same in the sense that again two of the three winning bids were placed among the first bids in the auction. However, in this auction the profit margins in these two bids were high. Luckily for these two bidders, they were both efficient bidders and their bids were close to the efficient allocation bids – and the third efficient bidder who entered the auction later was offered a good complementary bid by the QSM. Thus, even with high profits in the winning bids, the total cost to the buyer remained at a reasonable level, and much lower than in the simulated auction.

Based on the above-mentioned observations, it would seem that one explanation for the better outcomes of the B auctions compared to the simulated auction is the fact that bidders placed a lot more self-made bids than in the simulated auction. This widened the possibilities for the QSM to find profitable bid suggestions. It would appear, though, that a more powerful explanation is bidders' unnecessarily low profits in their initial bids. It is not easy to try to figure out a good strategy in combinatorial auctions – especially in auctions where the efficient allocation is not necessarily reached, since there is an element of luck involved. Figuring out a strategy is even more difficult for inexperienced bidders, who are not familiar with the characteristics of combinatorial auctions. In this experiment, the experience from the A auctions, in which profits were lower in the winning bids, can have guided the bidders' behavior in the B auctions.

12.2.2 Bidding Strategies

Bidding strategies in online auctions have been studied by Bapna et al. (2000, 2003 and 2004), Shah et al. (2003) and Puro (2009). They identify attributes that can be used to categorize bidders' behavior into distinct strategies. I will first review the strategies they have identified. These strategies are not directly applicable in the combinatorial auctions held in CombiAuction, and hence I will discuss the peculiarities of the CombiAuction auctions before analyzing the bidders' strategies in the laboratory experiments.

12.2.2.1 Strategies Identified in Literature

Bapna et al. (2004) identify four bidder types each following their own strategy: *evaluators*, *opportunists*, *participants*, and *sip-and-dippers*. The timing of the bids is used as the defining attribute in the categorization of the bidder types. The *evaluators* submit only one bid in the auction. The bid is placed either in the beginning or towards the middle of the auction. The evaluator knows her valuation for the item, and minimizes her effort cost from participating in the auction. The downside is that she may end up paying too much (in a forward auction) or selling too cheap (in a reverse auction). The *opportunists* are bargain-hunters. They enter the auction late and hope that by bidding late they will leave little time for competitors to act. When the ending rule is flexible, as in CombiAuction, the opportunist's strategy is not as effective as in an auction with a fixed closing rule. However, considering a wide variety of auctions (and not a single isolated auction), the opportunist's strategy of waiting to see how fierce the competition in each auction is before participating, can be effective even with flexible closing rule. *Participants* spend a lot of time bidding in the auction. They start bidding early on, and continue bidding until the closing time approaches. The *sip-and-dippers* participate in the beginning and at the end, but not in between. A typical sip-and-dipper places two bids: one bid at the beginning to establish her presence and to assess competition, and one bid that reveals their valuation when the closing time approaches.

Shah et al. (2003) also identify four bidder strategies: evaluator strategy, skeptic strategy, sniping strategy and unmasking strategy. Instead of using the timing of bids as the basis of categorization, Shah et al. examine whether bidders bid above the required increment (in a forward auction). *Evaluators* usually bid significantly higher than what the increment suggests – in addition to bidding early in the auction. *Skeptics* bid from the beginning to the end, as do the *participants* of Bapna et al. (2004), and they typically increase their bid exactly by the required increment. The name “skeptic” is derived from the fact that a proxy agent was available in the auctions Shah et al. (2003) studied, but these bidders did not use it. They rather bid manually every time they became outbid. The *sniping strategy* was equivalent to opportunist’s strategy of Bapna et al. (2004). The *unmasking strategy* consisted of a series of bids submitted close to each other. Shah et al. (2003) suspect that the purpose of such a strategy is to try to reveal other bidders’ proxy bids – and hence the name. This strategy is specific to auctions that enable proxy bidding, and thereby was not considered by Bapna et al. (2004).

Based on previous literature, Puro (2009) develops a categorization of bidders’ strategies that combines the timing of bidding, the number of bids placed, and the size of the bid decrement below the minimum required decrement. Puro studies online people-to-people (P2P) auctions in which are conducted as reverse auctions. Puro draws a distinction between strategies in which bidder places only one bid in the entire auction and strategies with multiple bids. He identifies five single-bid strategies: *sniping*, *late bidding*, *opportunist*, *evaluator* and *portfolio bidding*, and four multi-bid strategies: *all late*, *all skeptic*, *last bid late* and *stepped bidding*. In the single-bid strategies *late bidding* refers to bidding within the last 12 hours of the auction. *Sniping* and *opportunistic* bidding are special cases of late bidding. *Snipers* place their single bid within the last 30 minutes of the auction. *Opportunists* place their single bid before the snipers, and always with a price which is close to the required decrement. Thus, Puro’s definition of an opportunist is a little different from that of Bapna et al. (2004). In other words, the opportunists try their luck with one bid which becomes the leading bid when placed, but which maximizes the probability of being outbid. *Evaluators* place their bid before the last day of the auction, and always with a decrement much larger than what would have been required. *Portfolio bidding* is specific to the Prosper.com auction site Puro uses in his study. It refers to the possibility for bidders to specify their preferred products and valuations, and a bidding agent places automatically a bid in an auction that meets the bidder's criteria. In the multi-bid strategies *all skeptic* refers to a strategy similar to Shah et al.'s (2003) skeptic strategy. *All late* strategy is the multi-bid version of late bidding, that is, all bids are placed within the last 12 hours of the auction. *Last bid late* refers to a strategy in which bidder places at least one bid within the last 12 hours, but has already bid before that as well. *Stepped bidding* is the multi-bid equivalent of evaluator strategy. In stepped bidding the bidder places several bids early in the auction with decrements much lower than required.

All the studies reviewed above use data from real life auctions. One can then assume that a significant portion of the participants have prior experience from online auction. Also, all of the previous studies focus on single-item (single- or multiple-unit) auctions, which are conceptually much simpler bidding environments than combinatorial auctions. Thus, the impact of inexperience of the participants cannot be forgotten in analyzing the auctions in this experiment.

When combinatorial bidding is involved, strategies need to be evaluated somewhat differently. The timing of bids is still a relevant attribute, but when inactive bids are allowed to enter the bid stream (as is the case with CombiAuction) the bid decrement/increment is not directly applicable any more. Of course, when bids are placed with the help of support tools, the decrement is clearly defined. However, whenever bidder places a self-made bid, she does not know what price is required to make the bid active. Thus, in those cases, one cannot really talk about deliberately bidding below the required decrement. In combinatorial auction bidders rarely place only one bid, so the first five strategies identified by Puro (2009) are not applicable. Also, in combinatorial auctions the combination of items and item quantities become a defining characteristic in bidding strategies.

12.2.2.2 Bidders' Strategies in the Laboratory Experiment

I analyzed the auctions in order to identify distinct strategies using the categorizations in existing literature as guidelines. The relevant strategies are the evaluator/stepped bidding strategy of placing very competitive bids early on in the auction, the opportunist/sniping/all late strategy of entering the auction very late, the participant/skeptic strategy of bidding throughout the auction, and the sip-and-dip strategy of bidding at the beginning and at the end. The identification of these strategies is based mainly on the timing of bidding, and the competitiveness of the bid (profit margin below required decrement). However, an interesting aspect in combinatorial auctions is the content and number of the bids (what item combinations, which quantities, and how many different combinations one bidder bids for). My goal is to identify strategies for forming bids, as this has not been done in literature.

Thus, in order to identify bidding strategies, I studied the following characteristics in the bidders' behavior: 1) what time bidders place their first bid, 2) how often they bid and whether they were actively monitoring the auction around the closing time, 3) what item combinations and quantities they bid for in self-made bids, and 4) what prices they attach to the self-made bids. In addition, I analyzed how all of these attributes of the bidders' behavior changed from the A auction to the B auction. I used my observations from the auctions and the bidders' comments from the post-experiment questionnaire to answer these questions.

Time of Placing First Bid

The time a bidder chooses to enter the auction is one common attribute used to categorize bidders. Bidding early is part of the evaluator/stepped bidding strategy, participant/skeptic strategy as well as the sip-and dip strategy. Entering the auction late, on the other hand, is the essence of the opportunist/sniping/all late strategy. I defined early bidders as those, who placed their bid within the first hour of the auction, because these bidders are clearly eager to start bidding. In the A auctions, 16 bidders (21.6%) qualified as early bidders. In B auctions, only 6 bidders (8.7%) can be categorized as early bidders, and 5 of those also were early bidders in the A auctions. Even if the definition of early bidders were extended to bidding within the first two hours, the numbers would rise only to 17 (A auctions) and 11 (B auctions). I defined the late bidders as those who submitted their first bid during the last two hours of the auction. It is clear that they have deliberately bid late, because having had 21.5 hours time to bid they cannot plead a busy schedule as a reason for not bidding. In A auctions, 9 bidders (12.2%) were late bidders, and in B auctions 18 bidders (26.1%) bid late. One explanation for this is that because bidders realize they need to monitor the auction when the closing time approaches anyway, they can minimize their effort by starting to bid first then. Also, the opportunist strategy is the best strategy in simpler auctions, like the ones in eBay. Perhaps the bidders learned it from there. A cruder categorization is to look how many bidders began bidding during first day, and how many waited until the second day[24]. In A auctions, 42 bidders (56.8%) bid during the first day, but only 33 bidders (47.8%) bid during the first day in the B auctions.

A big portion of the bidders changed the timing of their first bids from A auction to B auction. Out of the 69 bidders 43 (62.3%) bid either at least two hours later or two hours earlier in the B auction than in the A auction, and 21 bidders (30.4%) even changed the day they started bidding. Of course the participants' schedules affect when they have time to log into the auction system and focus on bidding. However, it is also very likely that due to the complexity of the auction, the bidders experiment with different strategies in hopes of finding a good one. The fact that the outcome of the A auction (whether a bidder was a winner or a loser) does not correlate with the decision to bid at a different time in the A auction supports the latter explanation. Out of the losers of the A auctions 64.3% and out of the winners 59.3% changed the timing of their first bid by more than 2 hours. It is also possible, that some bidders' strategy was to wait until the support tools became active, and started bidding only after that. This hypothesis cannot be verified though, because there is no way of knowing if a bidder had visited the auction site, unless she also placed a bid.

Frequency and Timing of Bidding

By observing the time of entry into the auction alone, only the opportunists/snipers can be identified. In order to identify the other strategies, the entire bidding pattern (frequency and timing) needs to be studied.

[24] The auctions started at 5pm during the first day and the scheduled closing was at 4:30pm the following day.

The frequency of bidding varied a lot. Naturally, if your initial bid is active until the very last minute, you have no incentive of placing more bids. Also, I cannot know how many times the bidders used the support tools in order to place a bid but found no feasible suggestion and therefore did not place a bid. Therefore, seeming inactivity of a bidder (no bids in the bid stream) does not mean she would not have participated in the auction actively and visited the auction site often. This is where I used the answers bidders gave in the questionnaire to get a better understanding of the bidders' strategies.

The 9 late bidders/opportunists in A auctions were identified above. Late bidding was defined as entering the auction 2h before scheduled closing or later. If the definition is extended to include bidders who entered within 3h of closing, the number increases to 11 (14.9%). I defined participant/skeptic strategy as bidding soon after auction opened on the first day and then in the morning and afternoon of the second day all the way until closing. There were 9 bidders, who followed a pure participant strategy. In addition, there were 8 bidders who exhibited participant behavior, but did not enter the auction until after 9 pm on the first day (over 4h after auction began), and 2 bidders who participated from early on, but failed to monitor the closing of the auction. Thus, depending on how strictly one wants to define the participant/skeptic strategy, either 12.2% or 25.7% of bidders followed it. There was also a group of 7 bidders, who I call "partial participants", who started bidding on the morning of the second day and bid until the end.

There were only 5 sip-and-dippers (6.8%), if defined as bidding within the first 3 hours and then the 2 last hours of the auction. If this definition is extended to include bidding later on the first day, and returning in the afternoon the following day (when 4h were left), the number of sip-and-dippers increases to 16 (21.6%) of bidders.

The behavior of only 5 bidders (6.8%) could be interpreted as that of an evaluator. Characteristic of these bidders was that they placed only few bids either at the beginning or middle of the auction, and at a price close to production costs. Also a common characteristic is that they did not use the support tools to place the bids. The small number of evaluators is not surprising, because it is difficult to know what combinations to bid on if one does not monitor the auction and try out different combinations. By bidding on only a few combinations – albeit with a low price – and trusting that other bidders will then place complementing bids, the bidder takes a big risk. Interestingly, though, four of these evaluators ended up winning (although one actually made a loss with her bid). One explanation is that by placing a highly competitive bid relatively early, their bid became the bid the QSM would use as a complementing bid in most suggestions. However, the profits made by the winning evaluators were much lower than the average profit among the winners.

The remaining 16 bidders did not follow any identifiable strategy. Some of them did not monitor the auction until the end, some of them entered the auction rather late but not late enough to be an opportunist, and some did not provide enough information on their strategy in the questionnaire.

There were winning bidders among all the bidder categories. This is because strategy is not the only determinant of winning; cost matters as does the content of the bids, and the kind of bids the competitors have made.

The 18 late bidders/opportunists in the B auctions were identified already earlier. Just as in A auctions, if the definition of late bidding is extended to 3 hours before closing, the category is extended by two bidders to 20 (30.0%). In this second set of auctions, only 5 (7.2%) of the bidders could be categorized as participants. If the category is extended to those entering the auction later on the first day, 4 bidders can be added. If also the bidder who started early but did not participate at the end is added, the total number of bidders following the participant strategy is 10 (14.5%). This is clearly less than in the A auction. Perhaps based on their experience in the A auctions the bidders thought that the participant strategy is not the best one, and switched to other strategies. Many bidders seemed to have learned the importance of the last hour of bidding, which led them to reduce their effort prior to that. One indication of this is that 6 bidders categorized as participants in the A auction were now using the late bidding strategy. Another possibility is that the excitement of the game had worn off already during the first auction, and the bidders did not want to spend so much time in the second auction. One indication of this is that 5 bidders failed to bid in the second auction altogether. The number of partial participants was 8 (11.2%), which is almost the same as in the A auction.

There were 6 strictly sip-and-dipper bidders (8.7%) in the B auctions, and extending the definition increased the number to 11 (15.9%). As was the case with participants, also from the sip-and-dippers in the A auctions some bidders switched to late bidding. However, also some bidders switched to the participant strategy. The number of evaluators was again 5 (7.2%), and three of them were winners – although one of them had miscalculated their cost and actually won with an unprofitable bid.

The rest of the bidders (20 = 29.0%) did not fit into these categories, or they had not provided enough information on their bidding strategy in the questionnaire. Like in the A auctions, it is difficult to say, which strategy would have been the best. Bidders from all categories were among the winners.

Table 34 summarizes the bidder strategies observed in the all the experiment auctions.

Table 34 Frequencies of bidder strategies in the experiment auctions

	Late Bidders	Participants	Partial Participants	Sip & Dippers	Evaluators	Other
A Auctions*	11 (14.9%)	19 (25.7%)	7 (9.5%)	16 (21.6%)	5 (6.8%)	16 (21.6%)
B Auctions*	20 (30.0%)	10 (14.5%)	8 (11.2%)	11 (15.9%)	5 (7.2%)	20 (29.0%)

*There were 74 bidders in the A auctions and 69 in the B auctions.

Of course, one must keep in mind that strategy is not the only thing determining the frequency and timing of bidders' bids. The bidders (students) have busy schedules, and perhaps cannot participate in the auctions as much as they would like to. Also, some students are more motivated than others to make time for participation.

Item Combinations and Quantities

One interesting aspect of bidding behavior in combinatorial auctions is the kind of bids the bidders place, namely which items and what quantities the bidders bid for. In CombiAuction the availability of the QSM muddles the data somewhat, because the QSM gives bidders suggestions on item combinations and quantities. However, at least in the beginning of the auctions, when the support tools were not available, it was possible to observe the contents of the bidders' self-made bids.

One division can be made between *proactive* and *reactive* bidders. Proactive bidders place a wide array of different item combinations - "hooks in the water", as one bidder put it – so that other bidders' bids could team up with them. A few bidders submitted bids for all 31 item combinations (at maximum capacity). Reactive bidders resorted to the support tools, and placed bids based on the suggestions from the tools. As many as 13 bidders in both auctions report having placed bids only with the help of support tools, once they became available. Some bidders report using mainly the support tools to submit new bids, but placed self-made bids when the support tools suggested only unprofitable bids. The majority of the bidders were somewhere in between the two extremes; they placed self-made bids but also used the support tools.

Another division can be made between *intelligent* and *random* bidders. Intelligent bidders placed bids on particular combinations for justified reasons. In total 16 bidders said they tried to bid for combinations for which they had the lowest cost. It is not trivial, how to determine the combinations with the largest cost advantage in a combinatorial auction, when you do not have much information on your competitors' cost. Because in the set-up of this experiment, all items were treated equally (costs drawn from the same distributions), many bidders cleverly calculated per-unit average costs for bundles. If the items in the bundle had very different variable and fixed costs, such per-unit calculations across items could not have been done. In addition, 5 bidders said they tried to bid for large quantities to minimize the average cost. Where the intelligent bidders placed only a few well thought bids, the random bidders did not seem to have any logic in which items and quantities they placed. At the extreme, a random bidder would place bids for every item combination, and several bids with the same items, but different quantities.

The most popular bid was to place a bid containing all five items at full capacity, which was to be expected. Bidding for full capacity seemed to be the most popular choice in item quantities overall, which makes sense because of the economies of scale. In addition to bidding for all five items, bidders bid for various subcombinations of the full capacity bid (including bids for single items). We anticipated such behavior already when designing the "less support" benchmark cases in the second simulation study. Also the bid for all items, but with ½ of maximum capacity as quantities (which we also used in the second simulation study) was submitted in some of the equal capacities (A) auctions, but not very often. In the unequal capacities auction no such bids were placed. The bidders turned out to be much more creative than we had anticipated. A vast array of different bids (other than full capacity) was placed. A common denominator in the choice of bid quantities was to use fractions of the total demand (e.g. 100, 100, 100, 100, 100)[25] or (200, 200, 200, 200, 200). The bidders had a tendency to bid for equal quantities of all items also in the unequal capacities auction. Another popular choice was to bid for quantities that might complement other bids (e.g. 75, 75, 75, 75, 0) in the unequal capacities auctions where some bidders had capacity limits of 225. The bidders did not restrict themselves to even quantities across items; bids like (300, 100, 200, 300, 200) and (0, 300, 150, 150, 0) were placed also in the equal capacities (A) auctions, but were more common in the unequal capacities (B) auctions. The bidders seemed to have a tendency to place "desperate" bids (= bids with virtually zero profit) on small quantities of just one or two items when they could not find any active bids in the auction. Evidently they did not understand that due to economies of scope and scale the smaller bids had an even lower probability of becoming active.

In 8 of the 13 A auctions and 2 of the 11 B auctions some bidders placed bids which did not follow any of the bid creating logics described above. Bidders could place bids like (134, 200, 85, 93, 240), (220, 220, 220, 220, 220), and even (0, 0, 0, 1, 0). The fact that there were fewer odd bids in the B auctions leads me to believe that inexperience and frustration caused some of the bidders to bid completely randomly. Also, these odd bids were placed by bidders, who had difficulties in grasping the auction concepts (several of them placed bids that exceeded their capacity or resulted in a loss for them), or who placed dozens of bids during the auction.

Bid Prices

Bidders followed very different pricing strategies. Five distinct strategies could be defined: opportunists, evaluators, satisficers, support users, and desperate bidders. *Opportunists* tried their luck by submitting bids with very high margins (over 100%,

[25] I have left the bid prices out from the bids because they are irrelevant in this discussion.

sometimes even 1000%) at the beginning of the auctions. As the auction proceeded, they were forced to lower their prices drastically in order to stay competitive. The highest initial margins were observed only in the A auctions; bidders soon learned that such bids would not become active. However, some bidders did manage maintain profit margins as high as 20% in some of he B auctions. *Evaluators* were more concerned about winning than making a large profit. Moreover, they did not want to spend much time monitoring the auction and repeatedly submitting bids. Thus, they placed a few bids with very low margins before the heated competition around the closing of the auction. Usually the evaluators also study their cost function to identify the combinations where they should have a cost advantage. *Satisficers* have a profit target in mind (either profit margin or amount of euros) when they set out to bid. They place bids – both self-made and obtained with the help of a support tool – in which the target is achieved. Sometimes they even place bids with a lower price than suggested by the support tools, if the suggestion contains a higher profit. The *support users* put the least effort into the bidding process in the sense that they did not make any preparations prior to the auction. They simply used the support tools and evaluated the suggestions. This way their time cost from participation was smaller than for other bidders. Of course, sometimes they needed to put a lot of effort in the evaluation process – e.g. when the shortlist offered by the QSM was long – but it was rather mechanical a task.

The fifth strategy, *desperate bidding*, is a strategy any of the bidders can switch to when the support tools are not helpful, and the bidder does not have any active bids. Usually this happens near the closing time, but it can happen at any point in the auction when a bidder desperately wants to become active. Desperate bidding entails placing bids on seemingly random combinations of items and with very low profit margins (even lower than in the evaluator's bids). Usually the bidder places several such bids within a very short time period.

12.2.3 Other Observations on Bidders' Behavior and the CombiAuction System

The last objective of the laboratory experiment was to study the usability of the CombiAuction website. This means studying the bidders' perceptions of the site and the support tools. Also, studying how well the bidders learned to use the auction system and how well they understood the auction concepts is relevant, as it can validate or undermine the results presented in the two previous sections. There is little point in studying the outcomes of the auctions or the bidders' strategies, if the bidders had no idea of what they were supposed to do. Also, studying the bidders' perceptions and understanding of the auction system has implications on the organization of future auctions.

Usability of the User Interface and Support Tools

The participants were asked to rate how easy the user interface of the CombiAuction was to use, and how helpful the price and quantity support tools were (see Appendix 5). The average rating for the easiness of use (scale 1-5, 5 being the easiest) was 4.23 (standard deviation 0.80). The main complaint from the participants was that towards the end, when all bidders were logged in at the same time, the auction system became slow.

The average score for the helpfulness of price support was 3.81 (st. dev. 1.20) and for the helpfulness of quantity support 4.05 (st. dev. 1.10) on a scale 1-5 with 5 being most helpful. Most participants (45 out of 66) rated the price support and quantity support tools equally (un)helpful, which I had not anticipated. I imagined the bidders would find the quantity support tool more helpful. Moreover, a few participants (6 out of 66) even rated the price support tool more helpful than the quantity support. It could be that the combination of price and quantity support is much better than either of them alone. The QSM provides the bidders with good item combinations and they can then use the price support to keep the bids active as the auction proceeds.

Another explanation is that bidders do not care about the efficiency of the auctions. Knowing that they were not supposed to win does not make them feel any better about losing the auction – and without knowledge of other bidders' cost functions they cannot even know that they were not supposed to win. When they fail to find profitable bids they are equally unsatisfied with both support tools. This would also explain why 8 bidders answered the question on the helpfulness of the quantity support either "I somewhat disagree" (= 2) or "I totally disagree" (= 1). If at the end of the auction a bidder cannot find profitable bids from the shortlist, she is frustrated. It does not matter to her what the reason for that is. If the auction has ended in an efficient allocation with the total cost so low that others cannot afford to bring it down by the required decrement, the auction owner is pleased. The losing bidders, however, do not know what the reason for unprofitable suggestions is; they just see that the auction is closing and they are not among the winners.

Observations on Bidder Behavior and Understanding

The participants claimed that they understood the rules of the auctions and what their goal was. The average ratings for these questions in the questionnaire were 4.45 and 4.49 respectively (st. dev. 0.93 and 0.68) on a scale 1-5 with 5 indicating they understood well. However, the bidders made a lot of mistakes in the auctions. At least 14 bidders placed bids that exceeded their capacities, and at least 19 bidders submitted bids with a negative profit. Some of these mistakes were simply the result of carelessness; bidders forgot one zero from the price, they accidentally clicked on an undesirable bid in the support tools, had a mistake in their excel calculations, or made a typo when submitting a self-made bid (even though a confirmation window opens every time a bid is submitted). Several mistakes took place when the auction was about to end. Bidders were under pressure to act quickly in order to submit their bid before the auction closed, which led to carelessness. Also, it is possible that some bidders had so strong a desire to win, they deliberately made unprofitable bids. However, I did not find any indication of this in the questionnaires or in the emails I exchanged with some students.

Many times the bidder noticed the mistakes themselves and contacted me asking if I could delete the bids. However, among the winners there were 9 bidders, who won the auction with an unprofitable bid. There were also other signs of bidders not totally understanding how a combinatorial auction and the support tools worked. Bidders sometimes interpreted the inactive status to be the result of too high a price, when in fact it was due to lack of other bids in the auction. This led the bidders to unnecessarily lower the bid prices already at the beginning of the auction, when there were no other bidders present. A few bidders were also bewildered when the support tools were not available, and contacted me asking why they were not available. Also, in the participant questionnaire, three students answered that they did not understand how the QSM worked, and one even wrote that he "didn't understand the tools so good [well] that [he] could rely on them". Clearly, the briefing session was not enough for some of the participants to understand the concept of a combinatorial auction and of the support tools. However, these were isolated instances, and they have not affected my analysis of the auction outcomes or the analysis of the bidders' strategies.

The challenge of combinatorial auctions is that participating in them requires an understanding of the combinatorial aspects of the auction. In order to understand the winner determination and the quantity support, the bidders should have knowledge of the basic concepts of optimization. If not, the auction appears to them as a black box, and they cannot discern the link between their actions and the auction outcomes. Also, not understanding how the support tools work, bidders may not use them – or expect too much from them and become frustrated. What this experiment clearly shows is that briefing the participants in advance is a crucial phase in holding combinatorial auctions.

It is relatively easy to understand that cost advantages matter – and several participants in the experiments had understood this and tried to act accordingly. However, it is not trivial to understand that cost advantage is ultimately a relative concept: it is enough to have relatively low costs on items for which there are good complements. Thus, other bidders' costs affect what is the best bid for you. Unfortunately, even if a bidder understands this, there is not much she can do, because she does no know the other bidders' cost functions. All the bidders can try to do is to place some self-made bids on combinations for which her costs are low, and to use the support tools in order to find combinations to team up with other bids. This is not sufficient to guarantee a winning bid in the auction – but not using any support makes winning more difficult, or at least decreases the profits in the winning bids.

VII CONCLUSIONS

Combinatorial auctions have become an interesting subject of research. The literature focused on different aspects of combinatorial auctions has increased significantly during the past decade. There have also been some applications of combinatorial auctions into practice. However, there are still many ways combinatorial auctions could be improved to make them easier to use and applicable to a wider array situations.

The underlying problem with combinatorial auctions is that they are complex in many ways. Not only are combinatorial auctions computationally difficult, but they are also challenging environments for bidders. In this thesis I have introduced the *puzzle problem*. The puzzle problem refers to the situation in which bidders should coordinate the items and item quantities they bid for (in addition to price), in order to overcome the current provisional winners. Usually all communication among bidders is forbidden, as bid takers try to prevent collusion among bidders. With collusion I mean bidders' attempts to benefit at the expense of the bid taker. Also, reverse auctions are often sealed-bid (or semi-sealed-bid) auctions meaning that the bidders do not know the contents of their competitors bids. The problem facing a bidder in such an auction can be compared to the task of trying to complete a puzzle without knowing the size and shape of the missing piece. Hence the name "puzzle problem."

In our research project we developed a bidder support tool called the Quantity Support Mechanism (QSM) to help bidders overcome the puzzle problem in continuous, semi-sealed-bid combinatorial auctions. The QSM can be used equally in forward and reverse auctions, but in this thesis I have presented it in the reverse auction setting. At the heart of the QSM there is an IP problem (QSP), which maximizes the approximated profit (difference between the price in the new bid and the approximated cost of producing the suggested item quantities) of a currently non-winning bidder subject to the total cost to the buyer decreasing by a required decrement and the buyer's demand being fulfilled. The bidders' profit is approximated with a linear cost function. The estimates for the per-unit costs for each item are obtained from the dual prices of the linear relaxation on the Winner Determination Problem. Because the approximation is anonymous (= each bidder is assumed to have the same costs), and because the linear approximation may be far off from the true cost function (which presumably exhibits economies of scope since it is a combinatorial auction), a shortlist of alternative bid suggestions is created. The shortlist is created by forcing non-basic variables into the basis (i.e. forcing inactive bids among the provisional winners) and by forcing different item combinations to zero, each at a time. The shortlist is offered to the bidder, who then decides which bid to submit, if any.

We ran simulations to test the QSM. The simulations showed that the QSM helped solve the puzzle problem. The first simulation study showed that the QSM found good suggestions for the bidders, and that it produced better suggestions than a QSM which used random cost parameters instead of the dual prices. The second simulation study showed that the QSM improved the efficiency of the auction outcomes compared to the situation in which only price support was available.

The QSM was implemented in an online auction system called the CombiAuction. The auction system is designed for continuous combinatorial auctions, but it is up to the auction owner to decide whether the auction is a reverse or a forward auction and whether it is a semi-sealed-bid or an open-cry auction. The auction owner can also decide, whether the bidders have access to the QSM and the price support. The CombiAuction system was tested in an experiment with human subjects. The objective of the experiment was threefold. Firstly, the objective was to study the usability of the user interface of the CombiAuction system as well as the QSM. Secondly, I wanted to study the bidders' behavior and strategies in the auction. The third objective was to compare the outcomes of the experiment auctions to those of the simulations.

From the equal capacities auctions about half ended in an efficient (or pseudo-efficient) allocation, as had the simulated auction. The rest fell prey to bidders who made unprofitable bids or bidders who could not monitor the auction at the end. In addition, one auction fell prey to the threshold problem. Also in several of the unequal capacities auctions many bidders won with an unprofitable bid, which distorted the final allocation. Out of those auctions in which bidders did not win with unprofitable bids, one third ended in the efficient allocation. This is a good result seeing that the corresponding simulated auction ended in a very inefficient allocation (the production cost of the winning allocation was 9.2% above efficient cost). The good outcomes of the laboratory experiment auctions compared to the simulated auctions also helped validate the simulation results in general. Whenever simulating human behavior, there is the risk that by making simplifying assumptions one creates a more favorable setting, which leads to too good results. However, since the auctions with real bidders ended up in similar or better outcomes than the simulated auctions, the assumptions we made in the simulations seem to have been reasonable.

In the literature, there are several classifications of bidder strategies in online auctions. Using these categorizations as guidelines I could identify four bidding strategies: late bidders (opportunists, snipers), participants (skeptics), sip-and-dippers, and evaluators. However, a significant portion of the bidders did not fit into these categories, nor did they form new categories. This was partially due to not having enough information on their strategy, and partly due to not being able to detect any patterns from their bidding behavior. The above-mentioned strategies refer to the timing of bidding, and also the size of the mark-up in the bids (with evaluators). The strategies bidders use to form their combinatorial bids have not been studied in the literature. Based on the bid data I formed two categorizations: *proactive* vs. *reactive* bidders, and *intelligent* vs. *random* bidders. *Proactive* bidders place a lot of bids on different combinations for other bidders to team up with, whereas *reactive* bidders mainly use the support tools and try to find bids that complement existing bids. When creating self-made bids, *intelligent* bidders choose the item combinations based on their cost function. They try to place bids where they think they may have a cost advantage. Intelligent bidders also bid at full capacity in order to benefit from the economies of scale. *Random* bidders bid on many different combinations with no seeming logic, and they typically bid on more combinations than the intelligent bidders. Random bidders also often choose item quantities below their maximum capacity. Also intelligent bidders sometimes choose quantities below their maximum capacity when attempting to place bids that they expect to team up well with other bids or sum up easily to total demand. For example, if the bidder's capacity is 225 units for a particular item, and the demand is 600 units, the intelligent bidder may choose to bid for 200 units. All in all, the bidders proved to be more creative when creating self-made bids than I had anticipated.

Based on my observations and feedback given by the participants the interface seemed to be good and easy to use. The only complaints were related to the technical difficulties encountered during the last session (B2-auctions). However, the bidders made some mistakes: they placed unprofitable bids and bids that exceeded their capacity. We should redesign the interface to minimize the possibility of making such mistakes. For instance, whenever selecting a suggested bid from the shortlist, a confirmation window should pop up. Also, making it easier to export data from the auction system to Excel and to import it back might reduce mistakes. Based on my observations, some bidders did not quite grasp the concept of a combinatorial auction; perhaps a more thorough briefing session combined with a quiz would be called for.

The second simulation study revealed that the QSM did not always guide the auctions to the efficient allocation. Moreover, whenever the efficient allocation consisted of more than two bids, the efficient allocation was hardly ever found. The QSM suffers from the *extended puzzle problem*, which is in way an extension of the threshold problem well recognized in the literature. The QSM is good at finding the last missing piece to the puzzle, but unless the other pieces are the ones from the efficient allocation, it will not find the efficient bids. This is because it relies on the existing bids to find good complements for the new bid. Just as in the threshold problem it is useless for one bidder to decrease bid price alone, it is useless for the bidder to place an efficient bid unless its complements are present. It would not become active (a provisional winner), and thereby the QSM will not find it, because the QSM is designed to find bids that will become active immediately upon submission.

In order to help bidders overcome the extended puzzle problem we designed a new support tool, the Group Support Mechanism (GSM). The GSM is based on the same principles as the QSM: it aims at maximizing the profit of the incoming bidders while decreasing the total cost to the buyer and fulfilling the total demand. The biggest difference is that it suggests bids for several bidders at a time, and together the group of bidders would become provisional winners. Another major difference is that in the

GSM the cost function approximations are customized for each bidder based on the information obtainable from their bids. Based on preliminary tests the GSM seems to improve the efficiency of auction beyond what the QSM could do. This result is logical, since the GSM is essentially an extension of the QSM. As a special case the GSM could also suggest a bid for only one bidder – like the QSM would – if it was the optimal course of action. The weakness of the GSM is that whenever one of the bidders does not submit her bid suggestion, none of the bidders become active. This slows down the convergence of the auction and can cause frustration in the bidders.

Implications on Applying Combinatorial Auctions to Practice

The simulation studies already indicated that support is beneficial for the auction owner in semi-sealed-bid auctions. The laboratory experiment strengthened this impression. Moreover, the QSM seemed to produce better (= more efficient) outcomes for the unequal capacities auctions in the laboratory experiments than was expected based on the simulations. However, bidders made a lot of mistakes in the auctions. On the one hand this calls for training and experience, but also challenges the design of the user interface. When submitting a self-made bid in CombiAuction, a confirmation window always opens. However, no such window opened if placing a bid through the support tools. In both price and quantity support the bid was immediately submitted, if the bidder clicked on the "Submit this bid" link. Consequently, bidders made more mistakes with bids from the support tools than self-made bids. It is also very important that all numbers – especially large ones – are easy and quick to read. The bidder needs to be able to differentiate between a million and hundred thousand at a glance without having to count the zeros. The experiment showed that bidders made a lot of mistakes towards the end of the auctions when they were in a hurry. Therefore, the closing rule should not be fixed, but flexible (as it was in the experiment), and the extension time should be long enough that bidders can choose their course of action and estimate the profitability of all the bid suggestions. The 10 minutes used in the experiment may not have been long enough for that. The bidders' behavior in the experiment indicated that with more experience bidders shift their bidding closer to the end of the auction. Hence the auction owner should anticipate heavy bidding activity during the last hour of the auction. Naturally, this sets some requirements for the robustness of the system as well; it needs to be able to handle heavy traffic.

The laboratory experiment showed that combinatorial auctions are indeed difficult environments for bidders to bid in – especially for inexperienced bidders. One of the reasons is that a combinatorial auction does not advance linearly: your more recent bids can become inactive and older ones active. When the bidders did not understand the logic of winner determination and auction progression, it resulted in confusion and frustration. The participants felt that the support tools were easy to use, but some of the bidders decided not to use them because they did not understand how they produced the bid suggestions. Clearly, the bidders need to be briefed thoroughly on the intricacies of combinatorial auctions, and the support tools. What is sufficient information and how it is best conveyed is still an open issue. The need for training effectively rules out the possibility of organizing online combinatorial auctions open for everyone with access to the Internet. Fortunately, combinatorial auctions have many applications in the B2B markets, in which the bidding is done by professional sellers/buyers, who can be trained and who can accumulate the required experience.

Future Research

In this thesis I have brought up new insights into the challenges of combinatorial auctions. I have also described the tools we have developed to tackle these challenges. However, there are still many issues that need further studying.

First of all, the GSM should be developed further. The estimation of the bidders' cost functions, which is a crucial element in the GSM, should be studied further. Different alternatives for the estimation and updating procedure will be developed and compared with each other. Also the possibility of using the "pseudo dual prices" used in many existing combinatorial auction mechanisms could be explored. The current form of the cost function used in the Cost Estimation Problem has too many parameters to be estimated from the limited bid information. Hence, the first step will be to explore alternative, simpler forms for the cost function, which would still exhibit economies of scope. The GSM should also be tested, first through simulations and then in laboratory experiments with human subjects. Through simulations it is possible to compare different configurations of the GSM. The laboratory experiments on the other hand provide knowledge of how real bidders perceive and understand the GSM, and how convenient it is to use.

Secondly, the effect of bidders' learning and experience in combinatorial auctions on strategies and auction outcomes should be studied. In the laboratory experiment described in this thesis the bidders did not have prior experience of combinatorial auctions. Everybody participated in two auctions, but one auction hardly gives enough experience in such a complex bidding environment so that one could call the bidders experienced in the second auction. However, because combinatorial auctions are complex bidding environments, I expect experience to have a significant effect on bidders' strategies and thereby potentially on the auction outcomes. Already there were some indications of learning taking place. For instance, some bidders learned wait until close to the end to bid.

The QSM should also be fine-tuned before the next experiments. At least the length of the shortlist should be restricted to ease the bidders' burden of evaluating bid suggestions. It is not trivial, what is the optimal way of shortening the shortlist, and hence some alternative methods should be tested. Also, the system should be redesigned to record more information on bidders' strategies. For example, it would be good to know afterwards, when the bidders were logged in the system, and what the contents of the shortlists offered by the QSM were. It would also be interesting to compare the QSM and GSM in a laboratory setting to get a better understanding of what kind of allocations they lead to with human users, and which tool the bidders prefer.

References

Adomavicius, G., and Gupta, A. (2005), Toward Comprehensive Real-Time Bidder Support in Iterative Combinatorial Auctions, *Information Systems Research* 16: 169–185.

An, N., Elmaghraby, W., and Keskinocak, P. (2005), Bidding Strategies and Their Impact on Revenues in Combinatorial Auctions, *Journal of Revenue and Pricing Management* 3(4): 337–357.

Andersson, A., Tenhunen, M., and Ygge, F. (2000), Integer Programming for Combinatorial Auction Winner Determination, in: Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS), Boston, MA.

Ausubel, L.M. (2004), An Efficient Ascending-Bid Auction for Multiple Objects, *The American Economic Review* 94: 1452–1475.

Ausubel, L.M., Cramton, P., and Milgrom, P. (2006), The Clock-Proxy Auction: A Practical Combinatorial Auction Design, Chapter 5 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press, 115–138.

Ausubel, L.M, and Milgrom, P. (2002), Ascending Auctions with Package Bidding, *Frontiers of Theoretical Economics* 1: 1–42.

Ausubel, L.M., and Milgrom, P. (2006), The Lovely but Lonely Vickrey Auction, Chapter 1 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press, 17–40.

Bajari, P., and Hortaçsu, A. (2003), The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions, *RAND Journal of Economics* 34(2): 329–355.

Banks, J.S., Ledyard, J.O., and Porter, D.P. (1989), Allocating Uncertain and Unresponsive Resources: An Experimental Approach, *RAND Journal of Economics* 20: 1–25.

Banks, J., Olson, M., Porter, D., Rassenti, S., and Smith, V. (2003), Theory, Experiment and the Federal Communications Commission Spectrum Auctions", *Journal of Economic Behavior and Organization* 51: 303–350.

Bapna, R., Goes, P., and Gupta, A. (2000), A Theoretical and Empirical Investigation of Multi-Item On-line Auctions, *Information Technology and Management* 1(1–2): 1–23.

Bapna, R., Goes, P., and Gupta, A. (2003), Analysis of Design of Business-to-Consumer Online Auctions, *Management Science* 49(1): 85–101.

Bapna, R., Goes, P., Gupta, A., and Jin, Y. (2004), User Heterogeneity and its Impact on Electronic Auction Market Design: An Empirical Exploration, *MIS Quarterly* 28(1): 21–43.

Bazerman, M.H. and Samuelson, W.F. (1983), I Won the Auction but Don't Want the Prize, *Journal of Conflict Resolution* 27: 618–634.

Beil, D.R., and Wein, L.M. (2003), An Inverse-Optimization-Based Auction for Multiattribute RFQs, *Mangement Science* 49(11): 1529–1545.

Bichler, M. (2000), An Experimental Analysis of Multi-Attribute Auctions", *Decision Support Systems* 29(3): 249–268.

Bichler, M. (2001), The Future of e-markets: Multidimensional Market Mechanisms, Cambridge University Press, Cambridge, UK.

Binmore, K. and Klemperer, P. (2002), The Biggest Auction Ever: The Sale of the British 3G Telecom Licenses, *The Economic Journal* 112: C74–C96.

Borcherding, K., Eppel, T., and von Winterfeldt, D. (1991), Comparison of Weighting Judgments in Multiattribute Utility Measurement, *Management Science* (37): 1603–1619.

Branco, F. (1997), The Design of Multidimensional Auctions, *Rand Journal of Economics* (28): 63–81.

Cantillon, E., and Pesendorfer, M. (2006), Auctioning bus routes: the London experience, Chapter 23 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press, 573–592.

Capen, E., Clapp, R., and Campbell, W. (1971), Competitive Bidding in High Risk Situations, *Journal of Petroleum Technology* 23: 641–653.

Caplice, C. (2007), Electoronic Markets for Truckload Transportation, *Production and Operations Management* 16(4): 423–436.

Caplice, C., and Sheffi, Y. (2006), Combinatorial auctions for transportation services, Chapter 21 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press, 539–572.

Che, Y.-K. (1993), Design Competition through Multidimensional Auctions, *RAND Journal of Economics* 24: 668–679.

Chen-Ritzo, C.-H., Harrison, T.P., Kwasnica, A.M., and Thomas, D.J. (2005), Better, Faster, Cheaper: An Experimental Analysis of a Multiattribute Reverse Auction Mechanism with restricted Information Feedback, *Management Science* 51: 1753–1762.

Cho, Y. (2003), Economic Efficiency of Multi-Product Structure: The Evidence from Korean Housebuilding Firms, *Journal of Housing Economics* 12: 337–355.

Clarke, E.H. (1971), Multipart Pricing of Public Goods, *Public Choice* 11: 17–33.

Conen, W. and Sandholm, T. (2001), Minimal Preference Elicitation in Combinatorial Auctions, Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Workshop on Economic Agents, Models, and Mechanisms, Seattle, WA, August 6th (2001).

Coppinger, V.M., Smith, V.L., and Titus, J.A. (1980), Incentives and Behavior in English, Dutch and Sealed-Bid Auctions, *Economic Inquiry* 18(1): 1–22.

Cox, J.C., Smith, V.L., and Walker, J.M. (1988), Theory and Individual Behavior of First-Price Auctions, *Journal of Risk and Uncertainty* 1: 61–99.

Crescenzi, P., and Kann, V. (2006), A Compendium of NP Optimization Problems, (http://www.nada.kth.se/~viggo/wwwcompendium/node142.html, accessed Nov. 16th, 2006.

Cripps, M., and Ireland, N. (1994), The Design of Auctions and Tenders with Quality Thresholds: The Symmetric Case, *The Economic Journal* 104(423): 316–326.

Day, R.W., and Milgrom, P. (2008), Core-selecting package auctions, *International Journal of Game Theory* 36: 393–407.

Day, R.W., and Raghavan, S. (2007), Fair Payments for Efficient Allocation in Public Sector Combinatorial Auctions, *Management Science* 53(9): 1389–1406.

Day, R.W., and Raghavan, S. (2008), A combinatorial procurement auction featuring bundle price revelation without free-riding, *Decision Support Systems* 44: 621–640.

De Vries, S., Schummer, J., and Vohra, R. (2007), On ascending Vickrey auctions for heterogeneous objects, *Journal of Economic Theory* 132: 95–118.

De Vries, S., and Vohra, R. (2003), Combinatorial Auctions: A Survey, INFORMS *Journal on Computing* 15: 284–309.

Dunford, M., Hoffman, K., Menon, D., Sultana, R., and Wilson, T. (2004), Testing Linear Pricing Algorithms for use in Ascending Combinatorial Auctions, working paper, Department of Systems Engineering and Operations Research, George Mason University, Fairfax, VA.

Engelbrecht-Wiggans, R. (1988), Revenue Equivalence in Multiple Object Auctions, *Economics Letters* 26(1): 15–19.

Engelbrecht-Wiggans, R. (2001), The Effect of Entry and Information Cost on Oral versus Sealed-Bid Auctions, *Economics Letters* 70: 195–202.

Epstein, R., Henriques, L., Catalan, J., Weintraub, G.Y., and Martinez, C. (2002), A Combinatorial Auction Improves School Meals in Chile, *Interfaces* 32(6): 1–14.

Ervasti, V., and Leskelä, R.-L. (2009), Allocative Efficiency in Simulated Multiple-Unit Combinatorial Auctions with Quantity Support, forthcoming in *European Journal of Operational Research*.

FCC, The Federal Communications Commission (2000), Comment Sought on Modifying the Simultaneous Multiple Round Auction Design to Allow Combinatorial (Package) Bidding, Public Notice DA 00-1075, http://wireless.fcc.gov/auctions/31/releases/da001075.pdf, accessed Nov 30th, 2008.

FCC, The Federal Communications Commission (2007), Notice and Filing Requirements, Minimum Opening Bids, Reserve Prices, Upfront Payments, and Other Procedures for Auctions 73 and 76, Public Notice DA-4171, http://fjallfoss.fcc.gov/edocs_public/attachmatch/DA-07-4171A1.pdf, accessed Oct. 17, 2007.

FCC, The Federal Communications Commission (2008), Public Notice DA 08-595, Attachment A: 62-63. http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-08-595A2.pdf, accessed Oct. 1, 2008.

Friedman, L. (1956), A Competitive Bidding Strategy, *Operations Research* 4(1): 104–112.

Fujishima, Y., Leyton-Brown, K., and Shoham, Y. (1999), Taming the Computational Complexitiy of Combinatorial Auctions: Optimal and Approximate Approaches, Proceedings of IJCAI -99, Stockholm, Sweden: 548–553.

Gallien, J, and Wein, L.M. (2005), A Smart Market for Industrial Procurement with Capacity Constraints", *Management Science* 51: 76–91.

Garvin, S. and Kagel, J. (1994), Learning in Common Value Auctions, *Journal of Economic Behavior and Organization* 25: 351–372.

Goeree, J.K., and Offerman, T. (2002), Efficiency in Auctions with Private and Common Values: An Experimental Study, *The American Economic Review* 92(3): 625–643.

Gonen, R., and Lehmann, D. (2000), Optimal Solutions for Multi-Unit Combinatorial Auctions: Branch and Bound Heuristics, in Second ACM Conference on Electronic Commerce (EC-00), Minneapolis, Minnesota, October 2000.

Groves, T. (1973), "Incentives in Teams", *Econometrica* 41: 617–631.

Guo, Y., Lim, A., Rodrigues, B., and Tang, J. (2006), Using a Lagrangian Heuristic for a Combinatorial Auction Problem, *International Journal on Artificial Intelligence Tools* 15(3): 481–489.

Guttman, R.H., and Maes, P. (1998), Cooperative vs. Competitive Multi-agent Negotiations in Retail Electronic Commerce, MIT Media Lab Paper, in Proceedings of the Second International Workshop on Cooperative Information Agents, Paris, July 1998.

Günlük, O., Lászlo, L., and de Vries, S. (2005), A Branch-and-Price Algorithm and New Test Problems for Spectrum Auctions, *Management Science* 51(3): 391–406.

Hoffman, K., Menon, D., and van den Heever, S. (2004), A Bidder Aid Tool for Dynamic Package Creation in the FCC Spectrum Auction, working paper, Dept. of Systems Engineering and Operations Research, George Mason University.

Hoffman, K., Menon, D., and van den Heever, S., and Wilson, T. (2006), Observations and Near-Direct Implementation of the Ascending Proxy Auction, Chapter 17 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press, 415–450.

Hohner, G., Rich, J., Ng, E., Reid, G., Davenport, A.J., Kalaganam, J.R., Lee, H.S., and An, C. (2003), Combinatorial and Quantity-Discount Procurement Auctions Benefit Mars, Incorporated and Its Suppliers, *Interfaces* 33: 23–35.

Isaac, R.M., and James, D. (2000), Robustness of the Incentive Compatible Combinatorial Auction, *Experimental Economics*, 3: 31–53.

Jap, S.D. (2003), An Exploratory Study of the Introduction of Online Reverse Auctions, *Journal of Marketing* 67: 96–107.

Jones, J.L., and Koehler, G.J. (2002), Combinatorial Auctions Using Rule-Based Bids, *Decision Support Systems* 34: 59–74.

Jones, J.L., and Koehler, G.J. (2005), A Heuristic for Winner Determination in Rule-Based Combinatorial Auctions, *INFORMS Journal on Computing* 17(4): 475–489.

Kagel, J.H. (1995), Auctions: A Survey of Experimental Research, in: Kagel, J.H., and Roth, A.E. (eds.), Handbook of Experimental Economics, Princeton, NJ, Princeton University Press.

Kagel, J.H., Harstad, R.M., and Levin, D. (1987), Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study, *Econometrica* 55(6): 1275–1304.

Kagel, J.H., and Levin, D. (1986), The Winner's Curse and Public Information in Common Value Auctions, *American Economic Review* 76: 894–920.

Keeney, R., and Raiffa, H. (1976), Decisions with Multiple Objectives: Preferences and Value Tradeoffs, Wiley, New York.

Kelly, F., and Steinberg, R. (2000), A Combinatorial Auction with Multiple Winners for Universal Service, *Management Science* 46(4): 586–596.

Klemperer, P. (1998), Auctions with almost common values: the 'Wallet Game' and its applications, *European Economic Review* 42: 757–769.

Klemperer, P. (2002), What Really Matters in Auction Design, *The Journal of Economic Perspectives* 16(1): 169–189.

Koppius, O., and van Heck, E. (2002), Information Architecture and Electronic Market Performance in Multidimensional Auctions, working paper, Rotterdam School of Management, Erasmus University Rotterdam.

Korhonen, P., and Wallenius, J. (1996), Letter to the Editor: Behavioral Issues in MCDM: Neglected Research Questions, *Journal of Multi-Criteria Decision Analysis* 5: 178–182.

Krishna, V. (2002), Auction Theory, Academic Press, San Diego.

Kumar, V., Kumar, S., Tiwari, M.K., and Chan, F.T.S. (2006), Auction-Based Approach to Resolve the Scheduling Problem in the Steel Making Process, *International Journal of Production Research* 44: 1503–1522.

Kwasnica, A.M., Ledyard, J.O., Porter, D., and C. DeMartini (2005), A New and Improved Design for Multiobject Iterative Auctions, *Management Science* 51: 419–434.

Kwon, R.H., Anandalingam, G., and Ungar, L.H. (2005), Iterative Combinatorial Auctions with Bidder-Determined Combinations, *Management Science* 51(3): 407–418.

Köksalan, M, Leskelä, R.-L., Wallenius, H., and Wallenius, J. (2009), Improving Efficiency in Multiple-Unit Combinatorial Auctions: Bundling Bids from Multiple Bidders, forthcoming in *Decision Support Systems*.

Larichev, O. (1984), Psychological Validation of Decision Methods, *Journal of Applied Systems Analysis* 11: 37–46.

Ledyard, J.O., Olson, M., Porter, D., Swanson, J.A., and Torma, D.P. (2002), The First Use of a Combined-Value Auction for Transportation Services, *Interfaces* 32(5): 4–12.

Ledyard, J.O., Porter, C., and Rangel, A. (1997), Experiments Testing Multiobject Allocation Mechanisms, *Journal of Economics & Management Strategy* 6(3): 639–675.

Leskelä, R.-L., Teich, J., Wallenius, H., and Wallenius, J. (2007), Decision Support for Multi-Unit Combinatorial Auctions, *Decision Support Systems* 43: 420–434.

Levin, D., Kagel, J.H., and Richard, J.-F. (1996), Revenue Effects and Information Processing in English Common Value Auctions, *American Economic Review* 86(3): 442–460.

Lucking-Reiley, D., (1999), Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet, *American Economic Review* 89(5): 1063–1080.

Maasland, E., and Onderstal, S. (2006), Going, Going, Gone! A Swift Tour of Auction Theory and Its Applications, *De Economist* 154: 197–249.

Maskin, E., and Riley, J. (1989), Optimal Multi-Unit Auctions, in: Hahn, F. (ed.), The Economics of Missing Markets and Information, Oxford University Press, Oxford, UK.

McAfee, R., and McMillan, J. (1987), Auctions and Bidding, *Journal of Economic Literature* 25(2): 699–738.

McAfee, R., and McMillan, J. (1996), Analyzing the Airwaves Auction, *Journal of Economic Perspectives* 10(1): 159–175.

McMillan, J. (1994), Selling Spectrum Rights, *Journal of Economic Perspectives* 8(3): 145–162.

Metty, T., Harlan, R., Samelson, Q., Moore, T., Morris, T., Sorensen, R., Schneur, A., Raskina, O., Schneur, R., Kanner, J., Potts, K., and Robbins, J. (2005), Reinventing the Supplier Negotiation Process at Motorola, *Interfaces* 35: 7–23.

Milgrom, P.R. (1989), Auctions and Bidding: A Primer, *Journal of Economic Perspectives* 3(3): 3–22.

Milgrom, P. (1998), Game Theory and the Spectrum Auctions, *European Economic Review* 42: 771–778.

Milgrom, P. (2000), Putting Auction Tehory to Work: The Simultaneous Ascending Auction, *The Journal of Political Economy* 108(2): 245–272.

Milgrom, P.R., and Weber, R.J. (1982), A Theory of Auctions and Competitive Bidding, *Econometrica* 50(5): 1089–1122.

Mishra, D., and Parkes, D.C. (2007), Ascending price Vickrey auctions for general valuations, *Journal of Economic Theory* 132: 335–366.

Mishra, D., and Veeramani, D. (2007), Vickrey-Dutch procurement auction for multiple items, *European Journal of Operational Research* 180: 617–629.

Mito, M. and Fujita, S. (2004), On Heuristics for Solving Winner Determination Problem in Combinatorial Auctions, *Journal of Heuristics* 10: 507–523.

Murray, J.D., and White, R.W. (1983), Economies of Scale and Economies of Scope in Multiproduct Financial Institutions: A Study of British Columbia Credit Unions, *The Journal of Finance* 38(3): 887–902.

Myerson, R.B. (1981), Optimal Auction Design, *Mathematics of the Operations Research* 6(1): 58–73.

Nisan, N. (2000), Bidding and Allocation in Combinatorial Auctions, Proceedings of ACM Conference on Electronic Commerce (EC 00), ACM Press, New York: 1–12.

Nisan, N. (2006), Bidding Languages for Combinatorial Auctions, Chapter 9 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press: 215–232.

Ono, C., Nishiyama, S., and Horiuchi, H. (2003), Reducing complexity in winner determination for combinatorial ascending auctions, *Electronic Commerce Research and Applications* 2: 176–186.

Papadimitriou, C.H., and Steiglitz, K. (1982), Combinatorial Optimization: Algorithms and Complexity, Prentice Hall, New Jersey.

Park, S., and Rothkopf, M. (2005), Auctions with Bidder-Determined Allowable Combinations, *European Journal of Operational Research* 161(2): 399–415.

Parkes, D.C. (1999), *i*Bundle: An Efficient Ascending Price Bundle Auction, Proceedings of First ACM Conference on Electronic Commerce (EC-99), Denver, Colorado, 1999: 148–157.

Parkes, D.C. (2001), Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency, Doctoral dissertation, Computer and Information Science, University of Pennsylvania, Philadelphia.

Parkes, D.C. (2006), Iterative Combinatorial Auctions, Chapter 2 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press: 41–77.

Parkes, D., and Ungar, L.H. (2000), Iterative Combinatorial Auctions: Theory and Practice, Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-00): 74–81.

Pekeč, A., and Rothkopf, M. (2000), Making the FCC's First Combinatorial Auction Work Well, Comment on FCC Public Notice, DA 00-1075, http://wireless.fcc.gov/auctions/31/releases/workwell.pdf, accessed Nov 30th, 2008.

Pekeč, A., and Rothkopf, M. (2003), Combinatorial Auction Design, *Management Science* 49: 1485–1503.

Pinker, E.J., Seidmann, A., and Vakrat, Y. (2003), Managing Online Auctions: Current Business and Research Issues, *Management Science* 49(11): 1457–1484.

Porter, D., Rassenti, S., Roopnarine, A., and Smith, V. (2003), Combinatorial Auction Design, Proceedings of the National Academy of Sciences of the United States of America (PNAS), 100(19): 11153–11157.

Puro, L. (2009), Strategic Decision Making in Online Auctions: Prosper.com People-to-People Lending Site, Master's Thesis, Helsinki University of Technology, Dept. of Industrial Engineering and Management.

Rassenti, S.J., Smith, V.L., and Bulfin, R.L. (1982), A Combinatorial Auction Mechanism for Airport Time Slot Allocation, *The Bell Journal of Economics* 13(2): 402–417.

Riley, J.G., and Samuelson, W.F. (1981), Optimal Auctions, *The American Economic Review* 71(3):381–392.

Robinson, M.S. (1985), Collusion and the Choice of Auction, *RAND Journal of Economics* 16: 141–145.

Roth, A.E. and Ockenfels, A. (2002), Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet, *The American Economic Review* 92(4):1093–1103.

Rothkopf, M., and Harstad, R. (1994), Modeling Competitive Bidding: A Critical Essay, *Management Science* 40(3): 364–384.

Rothkopf, M., Pekeč, A., and Harstad, R. (1998), Computationally Manageable Combinatorial Auctions, *Management Science* 44(8): 1131–1147.

Sandholm, T. (2000), Approaches to Winner Determination in Combinatorial Auctions, *Decision Support Systems* 28(1-2): 165–176.

Sandholm, T. (2002), Algorithm for Optimal Winner Determination in Combinatorial Auctions, *Artificial Intelligence* 135: 1–54.

Sandholm, T. (2007), Expressive Commerce and Its Application to Sourcing: How We Conducted $35 Billion of Generalized Combinatorial Auctions, *AI Magazine* 28(3): 45–58.

Sandholm, T. and Boutlier, C. (2006), Preference Elicitation in Combinatorial Auctions, Chapter 10 in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), Combinatorial Auctions, MIT Press: 233–263.

Sandholm, T., Levine, D., Concordia, M., Martyn, P., Hughes, R., Jacobs, J., and Begg, D. (2006), Changing the Game in Strategic Sourcing at Procter & Gamble: Expressive Competition Enabled by Optimization, *Interfaces* 36(1): 55–68.

Sandholm, T., and Suri, S. (2003), BOB: Improved Winner Determination in Combinatorial Auctions and Generalizations, *Artificial Intelligence* 145: 33–58.

Sandholm, T., and Suri, S. (2006), Side Constraints and Non-Price Attributes in Markets, *Games and Economic Behavior* 55: 321–330.

Sandholm, T., Suri, S., Gilpin, A., and Levine, D. (2005), CABOB: A Fast Optimal Algorithm for Winner Determination in Combinatorial Auctions, *Management Science* 51(3): 374–390.

Schnizler, B., Neumann, D., Veit, D., and Weinhardt, C. (2008), Trading Grid services – a multi-attribute combinatorial approach, *European Journal of Operational Research* 187: 943–961.

Shah, H.S., Joshi, N.R., Sureka, A., and Wurman, P.R. (2003), Mining eBay: Bidding Strategies and Shill Detection, in: Zaïane, O.R., Srivastava, J., Spiliopoulou, M., and Masand, B. (eds.), WEBKDD 2002 – Mining Web Data for Discovering Usage Patterns and Profiles, Lecture Notes in Artificial Intelligence 2703: 17–34.

Sheffi, Y. (2004), Combinatorial Auctions in the Procurement of Transportation Services, *Interfaces* 34: 245–252.

Simon, H. (1955), A Behavioral Model of Rational Choice, *Quarterly Journal of Economics* 69: 7–19.

Srivinas, W., Tiwari, M.K., and Allada, V. (2004), Solving the Machine-Loading Problem in a Flexible Manufacturing System Using a Combinatorial Auction-Based Approach, *International Journal of Production Research* 42(9): 1879–1893.

Teich, J.E., Wallenius, H., and Wallenius, J. (1999), Multiple Issue Auction and Market Algorithms for the World Wide Web, *Decision Support Systems* 26(1): 49–66.

Teich, J.E., Wallenius, H., Wallenius, J., and Koppius, O.R. (2004), Emerging Multiple Issue E-auctions, *European Journal of Operational Research* 159: 1–16.

Teich, J.E., Wallenius, H., Wallenius, J., and Zaitsev, A. (2001), Designing Electronic Auctions: An Internet-Based Hybrid Procedure Combining Aspects of Negotiations and Auctions, *Electronic Commerce Research* 1: 301–314.

Teich, J.E., Wallenius, H., Wallenius, J., and Zaitsev, A. (2006), A Multi-Attribute e-Auction Mechanism for Procurement: Theoretical Foundations, *European Journal of Operational Research* 175: 90–100.

Tenorio, R. (1999), Multiple Unit Auctions with Strategic Price-Quantity Decisions, *Economic Theory* 13: 247–260.

Tukiainen, J. (2008), Testing for common cost in the City of Helsinki bus transit auctions, *International Journal of Industrial Organization* 26: 1308–1322.

Vakrat, Y. and Seidmann, A., 2000, Implications of the bidders' arrival process on the design of online auctions, Proceedings of the 33rd Hawaii International Conference on System Sciences, 2000: 1–10.

Vickrey, W. (1961), Counterspeculation, Auctions, and Competitive Sealed Tenders, *The Journal of Finance* 16(1): 8–37.

Vohra, R., and Weber, R.J. (2000), A 'Bids and Offers' Approach to Package Bidding, Comment on FCC Public Notice, DA 00-1075, http://wireless.fcc.gov/auctions/31/releases/bid_offr.pdf, accessed Nov 30th, 2008.

Wilson, R.B. (1969), Competitive Bidding with Disparate Information, *Management Science* 15(7): 446–448.

Wilson, R.B. (1977), A Bidding Model of Perfect Competition, *The Review of Economic Studies* 44(3): 511–518.

Wilson, R.B. (1979), Auctions of Shares, *Quarterly Journal of Economics* 93(4): 557–585.

Wilson, R.B. (1992), Strategic Analysis of Auctions, in: Aumann, R.J., and Hart, S. (eds.), Handbook of Game Theory Vol. I, North-Holland, Amsterdam: 227–271.

Wurman, P., and Wellman, M. (2000), A*k*BA: A Progressive, Anonymous-Price Combinatorial Auction, Proceedings of Second ACM Conference on Electronic Commerce (EC-00), Minneapolis, Minnesota, 2000: 21–29.

Yang, S., Segre, A.M., and Codenotti, B. (2009), An optimal multiprocessor combinatorial auction solver, *Computers and Operations Research* 36(1): 149–166.

Yokoo, M., Sakurai, Y. and Matsubara, S. (2004), The effect of false-name bids in combinatorial auctions: new fraud in internet auctions, *Games and Economic Behavior* 46: 174–188.

Zionts, S., and Wallenius, J. (1976), An Interactive Programming Method for Solving the Multiple Criteria Problem, *Management Science* 22(6): 652–663.

Özer, A.H., and Özturan, C. (2009), A model and heuristic algorithms for multi-unit nondiscriminatory combinatorial auction, *Computers & Operations Research* 36: 196–208.

APPENDIX 1: FIXED COST PARAMETERS IN THE FIRST SIMULATION STUDY

Table 35 Upper and lower limits for the uniform distributions from which the fixed cost parameters were drawn in the first simulation study

	"Low" Fixed Cost		"High" Fixed Cost	
	Lower	Upper	Lower	Upper
F_1	1000	1200	8000	10000
F_2	1000	1200	7000	11000
F_3	700	1000	5000	7000
F_4	1500	2000	9000	12000
F_5	1600	2500	10000	13000
F_{12}	1700	2300	13000	18000
F_{13}	1700	2200	12000	16000
F_{14}	2100	3000	17000	20000
F_{15}	2500	3500	18000	22000
F_{23}	1600	2100	11000	15000
F_{24}	2500	3000	16000	23000
F_{25}	2500	4000	17000	24000
F_{34}	2200	3000	13000	16000
F_{35}	2000	3000	15000	20000
F_{45}	2900	4000	17000	22000
F_{123}	2300	3200	18000	24000
F_{124}	3100	4200	21000	27000
F_{125}	3200	5000	22000	30000
F_{134}	3200	4200	20000	26000
F_{135}	3300	4600	22000	29000
F_{145}	3600	5200	23000	29000
F_{234}	3000	4000	23000	27000
F_{235}	3100	4500	24000	30000
F_{245}	3800	5400	25000	32000
F_{345}	3700	5500	22000	26000
F_{1234}	4100	5500	30000	38000
F_{1235}	3900	5400	31000	40000
F_{1245}	4600	6300	34000	41000
F_{1345}	4500	6200	28000	38000
F_{2345}	4600	6200	34000	40000
F_{12345}	5500	7500	42000	50000

F_{ijk} = the fixed cost of producing items *i*, j and *k* jointly

APPENDIX 2: FORMULATION OF THE QSP WITH TRUE COST FUNCTIONS

Here I will present the formulation of the QSP for the case of ten bidders and five items, which corresponds to the one of the settings in the first simulation study. The formulation in the case of three items is analogous but simpler due to fewer cost function parameters. The formulation is presented in the context of the simulation study instead of in a general form in order to make it easier to read. The QSP is formulated for bidder *m*. To simplify the notation, the index indicating bidder *m* is omitted from the cost function variables and parameters.

The objective is to find a bid price p_{new}, bid quantities $q_{new,kj}$ and the combination of complementing bids from other bidders that maximizes the profit for the bidder. The status of bidder *i*'s *j*[th] bid is indicated by x_{ij} ($x_{ij}=1$ indicates active status, $x_{ij}=0$ indicates inactive status). Let F_{ijk} denote bidder *m*'s fixed cost of producing items *i*, *j*, and *k* jointly, and c_k the variable cost of item *k*. Because of the discontinuous cost functions, the formulation becomes quite cumbersome and lengthy. Auxiliary variables y_i, y_{ij}, y_{ijk}, y_{ijkl} and $y_{ijklm} \in \{0,1\}$ ensure that the correct fixed cost parameter is applied. E.g. if in the optimal bid, items 1, 3, 4 and 5 assume a positive value, the system should set $y_{1345} = 1$, and others to zero. For the construction of the constraints that ensures that this in fact happens, we need several (mutually exclusive) variables for the quantities of each item. These variables are denoted $q_{new,kj}$, where *k* indicates the item ($k = 1, \ldots, 5$), and *j* the different variables for the same item ($j = 1, \ldots, 16$). In other words, there is one quantity variable for each combination that item *k* is a part of. Thus, for any item *k* only one $q_{new,kj}$ can assume a positive value.

The problem is to maximize:

$$\begin{aligned}
&p_{new} - \sum_{j=1}^{16}\sum_{k=1}^{5} c_k q_{new,kj} - \sum_{k=1}^{5} F_k y_k - \sum_{k=2}^{5} F_{1k} y_{1k} - \sum_{k=3}^{5} F_{2k} y_{2k} - \sum_{k=4}^{5} F_{3k} y_{3k} - F_{45} y_{45} \\
&- \sum_{k=3}^{5} F_{12k} y_{12k} - \sum_{k=4}^{5} F_{13k} y_{13k} - \sum_{k=4}^{5} F_{23k} y_{23k} - \sum_{k=1}^{3} F_{k45} y_{k45} - \sum_{k=4}^{5} F_{123k} y_{123k} - \sum_{k=2}^{3} F_{1k45} y_{1k45} \qquad (51) \\
&- F_{2345}\; y_{2345} - F_{12345} y_{12345}
\end{aligned}$$

subject to the constraints:

(i) Total cost to the buyer must decline 5% from the current best C^*:

$$\sum_{i=1}^{10} p_{ij} x_{ij} + p_{new} \leq 0.95C^* \tag{52}$$

Notice that in the first simulation study we are considering the point in time in the auction when there is only one bid from each bidder i in the existing bid stream (j = 1), so there is no need to sum over the j's.

(ii) The buyer's demand for each item (= 1000 units) must be fulfilled, q_{ijk} is the quantity of good k in bidder i's initial bid (again, no need to sum over j):

$$\sum_{i=1}^{10} q_{ijk} x_{ij} + \sum_{j=1}^{16} q_{new,kj} \geq 1000 \quad \forall\, k = 1,\ldots,5 \tag{53}$$

(iii) Exactly one fixed cost is chosen:

$$\begin{aligned} &\sum_{k=1}^{5} y_k + \sum_{k=2}^{5} y_{1k} + \sum_{k=3}^{5} y_{2k} + \sum_{k=4}^{5} y_{3k} + y_{45} + \sum_{k=3}^{5} y_{12k} + \sum_{k=4}^{5} y_{13k} + \sum_{k=4}^{5} y_{23k} + \sum_{k=1}^{3} y_{k45} \\ &+ \sum_{k=4}^{5} y_{123k} + \sum_{k=2}^{3} y_{1k45} + y_{2345} + y_{12345} = 1 \end{aligned} \tag{54}$$

(iv) The correct fixed cost should be chosen

The constraint (iii) assures that only one of constraints (iv) can be nonbinding. The objective function (maximization) assures that given the bid quantities, the minimum fixed cost is chosen. E.g. if the bid quantities are nonzero for only items 1 and 2, the algorithm will set y_{12} = 1 allowing $q_{new,1,2}$ and $q_{new,2,2}$ assume a positive value. The constraints would allow the algorithm to set for instance y_{12345} = 1 ($q_{new,1,16}$ and $q_{new,2,16}$ would assume positive values), but because $F_{12} < F_{12345}$ it would not be optimal.

$$q_{new,k,1} - My_k \le 0 \quad \forall\ k = 1,\ldots,5$$

$$
\begin{aligned}
& q_{new,1,2} + q_{new,2,2} - My_{12} \le 0 \\
& q_{new,1,3} + q_{new,3,2} - My_{13} \le 0 \\
& q_{new,1,4} + q_{new,4,2} - My_{14} \le 0 \\
& q_{new,1,5} + q_{new,5,2} - My_{15} \le 0 \\
& q_{new,2,3} + q_{new,3,3} - My_{23} \le 0 \\
& q_{new,2,4} + q_{new,4,3} - My_{24} \le 0 \\
& q_{new,2,5} + q_{new,5,3} - My_{25} \le 0 \\
& q_{new,3,4} + q_{new,4,4} - My_{34} \le 0 \\
& q_{new,3,5} + q_{new,5,4} - My_{35} \le 0 \\
& q_{new,4,5} + q_{new,5,5} - My_{45} \le 0 \\
& \\
& q_{new,1,6} + q_{new,2,6} + q_{new,3,6} - My_{123} \le 0 \\
& q_{new,1,7} + q_{new,2,7} + q_{new,4,6} - My_{124} \le 0 \\
& q_{new,1,8} + q_{new,2,8} + q_{new,5,6} - My_{125} \le 0 \\
& q_{new,1,9} + q_{new,3,7} + q_{new,4,7} - My_{134} \le 0 \\
& q_{new,1,10} + q_{new,3,8} + q_{new,5,7} - My_{135} \le 0 \\
& q_{new,1,11} + q_{new,4,8} + q_{new,5,8} - My_{145} \le 0 \\
& q_{new,2,9} + q_{new,3,9} + q_{new,4,9} - My_{234} \le 0 \\
& q_{new,2,10} + q_{new,3,10} + q_{new,5,9} - My_{235} \le 0 \\
& q_{new,2,11} + q_{new,4,10} + q_{new,5,10} - My_{245} \le 0 \\
& q_{new,3,11} + q_{new,4,11} + q_{new,5,11} - My_{345} \le 0 \\
& \\
& q_{new,1,12} + q_{new,2,12} + q_{new,3,12} + q_{new,4,12} - My_{1234} \le 0 \\
& q_{new,1,13} + q_{new,2,13} + q_{new,3,13} + q_{new,5,12} - My_{1235} \le 0 \\
& q_{new,1,14} + q_{new,2,14} + q_{new,4,13} + q_{new,5,13} - My_{1245} \le 0 \\
& q_{new,1,15} + q_{new,3,14} + q_{new,4,14} + q_{new,5,14} - My_{1345} \le 0 \\
& q_{new,2,15} + q_{new,3,15} + q_{new,5,15} + q_{new,5,15} - My_{2345} \le 0 \\
& \\
& q_{new,1,16} + q_{new,2,16} + q_{new,3,16} + q_{new,4,16} + q_{new,5,15} - My_{12345} \le 0
\end{aligned}
\qquad (55)
$$

(v) Capacity constraints as defined in the first simulation study:

$$q_{new,kj} \le 500 \quad \forall\ k = 1,\ldots,5,\ \ j = 1,\ldots 16 \qquad (56)$$

(vi) Only one bid can be active per bidder:

$$x_{m1} = 0 \tag{57}$$

Bidder *m*'s new bid will become active, so her initial bid cannot be active. There is only one bid from all the other bidders, at this point in time, hence there is no need for additional constraints yet.

(vii) The bid status variables of the initial bids are binary variables:

$$x_{ij} \in \{0,1\} \quad \forall\ i = 1,\ldots,10 \tag{58}$$

(viii) Every y_{ijklm} is binary:

$$\begin{aligned}
&y_i \in \{0,1\} \quad \forall\ i = 1,\ldots,5 \\
&y_{ij} \in \{0,1\} \quad \forall\ i = 1,\ldots,5,\ \ j = 2,\ldots,5 \\
&y_{ijk} \in \{0,1\} \quad \forall\ i = 1,\ldots,5,\ \ j = 2,\ldots,5,\ \ k = 3,\ldots,5 \\
&y_{ijkl} \in \{0,1\} \quad \forall\ i = 1,\ldots,5,\ \ j = 2,\ldots,5,\ \ k = 3,\ldots,5,\ \ l = 4,5 \\
&y_{12345} \in \{0,1\}
\end{aligned} \tag{59}$$

APPENDIX 3: FORMULATION OF THE EFFICIENT ALLOCATION PROBLEM

The efficient allocation problem in the second simulation study is in some respects similar to the formulation of the quantity support problem with true cost functions presented in Appendix 2. The logic by which the correct fixed cost is chosen is the same. Also the constraints assuring that buyer's demand is fulfilled, and that only one bid per bidder can be accepted are the same. The objective function is different, but it also contains similar elements. The efficient allocation problem is formulated for the case of five items and 15 bidders, which corresponds to one of the simulation designs. A general formulation of the efficient allocation problem would be much more difficult for readers to follow.

Let F_{jklmn} denote the fixed cost of producing items j, k, l, m and n jointly, and c_{ij} bidder i's variable cost of producing item j. The capacity constraints for bidder i's item j are denoted a_{ij}. The objective is to select the combination of bids from different bidders that minimize the total production cost. Because of the discontinuous cost functions, the formulation becomes quite cumbersome and lengthy. Auxiliary variables $y_{i,j}$, $y_{i,jk}$, $y_{i,jkl}$, $y_{i,jklm}$ and $y_{i,jklmn} \in \{0,1\}$ ensure that the correct fixed cost parameter is applied. E.g. if in the optimal allocation, items 1, 3, 4 and 5 assume a positive value for bidder 1, the system should set $y_{1,1345} = 1$, and other y_i's to zero. For the construction of the constraints that ensures that this in fact happens, we need several (mutually exclusive) variables for the quantities of each item. These variables are denoted q_{ijk}, where i indicates the bidder ($i = 1, \ldots, 15$), j the item ($j = 1, \ldots, 5$) and k enumerates the different variables for the same item ($k = 1, \ldots, 16$). For any bidder i and item j only one $q_{i,j,k}$ can assume a positive value.

The problem is to minimize:

$$\sum_{i=1}^{15} (\sum_{j=1}^{5}\sum_{k=1}^{16} c_{ij}q_{ijk} + \sum_{m=1}^{5} F_{i,m}y_{i,m} + \sum_{m=2}^{5} F_{i,1m}y_{i,1m} + \sum_{m=3}^{5} F_{i,2m}y_{i,2m} + \sum_{m=4}^{5} F_{i,3m}y_{i,3m} + F_{i,45}y_{i,45}$$
$$+ \sum_{m=3}^{5} F_{i,12m}y_{i,12m} + \sum_{m=4}^{5} F_{i,13m}y_{i,13m} + \sum_{m=4}^{5} F_{i,23m}y_{i,23m} + \sum_{m=1}^{3} F_{i,m45}y_{i,m45} + \sum_{m=4}^{5} F_{i,123m}y_{i,123m} \tag{60}$$
$$+ \sum_{m=2}^{3} F_{i,1m45}y_{i,1m45} + F_{i,2345}y_{i,2345} + F_{i,12345}y_{i,12345})$$

subject to the constraints:

(i) The buyer's demand for each item (= 600 units) must be fulfilled

$$\sum_{i=1}^{15}\sum_{k=1}^{16} q_{ijk} \geq 600 \qquad \forall\ j = 1,\ldots,5 \tag{61}$$

(ii) At most one fixed cost is chosen for each bidder

$$\sum_{m=1}^{5} y_{i,m} + \sum_{m=2}^{5} y_{i,1m} + \sum_{m=3}^{5} y_{i,2m} + \sum_{m=4}^{5} y_{i,3m} + y_{i,45} + \sum_{m=3}^{5} y_{i,12m} + \sum_{m=4}^{5} y_{i,13m} + \sum_{m=4}^{5} y_{i,23m}$$
$$+ \sum_{m=1}^{3} y_{i,m45} + \sum_{m=4}^{5} y_{i,123m} + \sum_{m=2}^{3} y_{i,1m45} + y_{i,2345} + y_{i,12345} < 1 \qquad \forall\ i = 1,\ldots,15 \tag{62}$$

(iii) The correct fixed cost should be chosen for each bidder $i = 1, \ldots, 15$. The following constraints allow a $y_{i,jklmn}$ to assume the value of one only if the right combination of item quantities assumes a positive value.

$$q_{i,j,1} - My_{i,j} < 0 \qquad \forall\ j = 1,\ldots,5$$

$$\begin{aligned}
&q_{i,1,2} + q_{i,2,2} - My_{i,12} \leq 0\\
&q_{i,1,3} + q_{i,3,2} - My_{i,13} \leq 0\\
&q_{i,1,4} + q_{i,4,2} - My_{i,14} \leq 0\\
&q_{i,1,5} + q_{i,5,2} - My_{i,15} \leq 0\\
&q_{i,2,3} + q_{i,3,3} - My_{i,23} \leq 0\\
&q_{i,2,4} + q_{i,4,3} - My_{i,24} \leq 0\\
&q_{i,2,5} + q_{i,5,3} - My_{i,25} \leq 0\\
&q_{i,3,4} + q_{i,4,4} - My_{i,34} \leq 0\\
&q_{i,3,5} + q_{i,5,4} - My_{i,35} \leq 0\\
&q_{i,4,5} + q_{i,5,5} - My_{i,45} \leq 0
\end{aligned} \tag{63}$$

$$q_{i,1,6}+q_{i,2,6}+q_{i,3,6}-My_{i,123}\le 0$$
$$q_{i,1,7}+q_{i,2,7}+q_{i,4,6}-My_{i,124}\le 0$$
$$q_{i,1,8}+q_{i,2,8}+q_{i,5,6}-My_{i,125}\le 0$$
$$q_{i,1,9}+q_{i,3,7}+q_{i,4,7}-My_{i,134}\le 0$$
$$q_{i,1,10}+q_{i,3,8}+q_{i,5,7}-My_{i,135}\le 0$$
$$q_{i,1,11}+q_{i,4,8}+q_{i,5,8}-My_{i,145}\le 0$$
$$q_{i,2,9}+q_{i,3,9}+q_{i,4,9}-My_{i,234}\le 0$$
$$q_{i,2,10}+q_{i,3,10}+q_{i,5,9}-My_{i,235}\le 0$$
$$q_{i,2,11}+q_{i,4,10}+q_{i,5,10}-My_{i,245}\le 0$$
$$q_{i,3,11}+q_{i,4,11}+q_{i,5,11}-My_{i,345}\le 0$$

$$q_{i,1,12}+q_{i,2,12}+q_{i,3,12}+q_{i,4,12}-My_{i,1234}\le 0$$
$$q_{i,1,13}+q_{i,2,13}+q_{i,3,13}+q_{i,5,12}-My_{i,1235}\le 0$$
$$q_{i,1,14}+q_{i,2,14}+q_{i,4,13}+q_{5,13}-My_{i,1245}\le 0$$
$$q_{i,1,15}+q_{i,3,14}+q_{i,4,14}+q_{i,5,14}-My_{i,1345}\le 0$$
$$q_{i,2,15}+q_{i,3,15}+q_{i,4,15}+q_{i,5,15}-My_{i,2345}\le 0$$
$$q_{i,1,16}+q_{i,2,16}+q_{i,3,16}+q_{i,4,16}+q_{i,5,16}-My_{i,12345}\le 0$$

$$\forall\ i=1,\ldots,15$$

The constraints above do not rule out the possibility that e.g. $y_{i,12345}=1$, even though only four items or less actually assume positive quantities. However, since $F_{i,12345}$ is set to be larger than any other fixed cost, the objective to minimize total cost will choose the lowest fixed cost allowed by the constraints. The same argument applies to any other F_i as well.

(iv) Capacity constraints

$$q_{ijk}\le a_{ij}\qquad \forall\ i=1,\ldots,15,\ \ j=1,\ldots,5,\ \ k=1,\ldots,16 \tag{64}$$

(v) Every auxiliary variable is binary

$$\begin{aligned}
&y_{i,j}\in\{0,1\} && \forall\ i=1,\ldots,15,\ \ j=1,\ldots,5\\
&y_{i,jk}\in\{0,1\} && \forall\ i=1,\ldots,5,\ \ j=1,\ldots,5,\ \ k=2,\ldots,5\\
&y_{i,jkl}\in\{0,1\} && \forall\ i=1,\ldots,15,\ \ j=1,\ldots,5,\ \ k=2,\ldots,5,\ \ l=3,\ldots,5\\
&y_{i,jklm}\in\{0,1\} && \forall\ i=1,\ldots,15,\ \ j=1,\ldots,5,\ \ k=2,\ldots,5,\ \ l=3,\ldots,5,\ \ m=4,5\\
&y_{i,12345}\in\{0,1\} && \forall\ i=1,\ldots,15
\end{aligned} \tag{65}$$

APPENDIX 4: BIDDERS' COST FUNCTIONS IN THE LABORATORY EXPERIMENTS

Table 36 presents the cost function parameters for bidders in equal capacities (A) auctions. In groups of five students, the cost functions used were those of Bidders 5, 8, 10, 12 and 15. Bidder 1 was added to the six person group, and Bidder 11 to the seven person groups. The capacities are equal (300, 300, 300, 300, 300) for all bidders.

Table 37 presents the cost function parameters for bidders in unequal capacities (B) auctions. In the groups with only six bidders, Bidder 9 was removed from the auction. The bidders' maximum capacities are presented in Table 38.

Table 36 Bidders' cost function parameters in the equal capacities (A) auctions

		Bidder 1	Bidder 5	Bidder 8	Bidder 10	Bidder 11	Bidder 12	Bidder 15
Variable Costs	c_1	59,38	63,47	61,53	56,31	65,53	53,51	55,14
	c_2	56,64	57,93	56,72	66,59	55,12	63,85	54,03
	c_3	66,51	59,30	58,00	65,03	66,23	65,32	56,86
	c_4	60,85	60,21	61,50	61,67	61,59	53,36	63,43
	c_5	56,27	55,43	57,83	56,15	61,87	60,21	59,14
Fixed Costs	F_1	16644	19329	19673	19523	17690	17823	19984
	F_2	18364	18454	19960	17703	16585	18183	16335
	F_3	17773	19657	19426	17371	17675	16475	17316
	F_4	16205	16037	17720	16041	16948	17606	17746
	F_5	18027	17700	16883	17122	18696	18393	19308
	F_{12}	26753	24299	27610	28980	24416	25395	24408
	F_{13}	24237	24180	28605	29348	26447	26261	29530
	F_{14}	26535	24491	27360	29895	26915	24002	25775
	F_{15}	29568	24169	27738	26707	26455	29454	28900
	F_{23}	26655	26067	29387	25488	24119	29981	24804
	F_{24}	27489	29978	27523	29698	28452	29236	26080
	F_{25}	27329	25960	26733	26804	27111	25219	27939
	F_{34}	28784	27738	26447	24977	24157	29108	28124
	F_{35}	26035	24982	26579	27734	29911	27822	24020
	F_{45}	25853	25066	26335	27644	27911	28575	25683
	F_{123}	32027	33187	33649	39980	38591	38147	35436
	F_{124}	38542	33124	35341	37485	38648	34275	34703
	F_{125}	35553	38147	32306	39640	34664	34620	34437
	F_{134}	36471	33870	37599	34281	33366	33626	35147
	F_{135}	35695	32205	32983	37035	33573	35710	32926
	F_{145}	38652	32273	37331	34045	33142	34218	36823
	F_{234}	34789	37650	32590	36189	37196	33867	38993
	F_{235}	33640	37321	39511	38119	33195	36955	34543
	F_{245}	34385	33144	38243	33370	34933	39823	33217
	F_{345}	38968	37475	35569	34928	37218	38972	39675
	F_{1234}	47512	42397	42200	41194	48537	41679	47694
	F_{1235}	43276	49508	43240	42877	47978	42148	49835
	F_{1245}	49030	43952	42911	44056	42098	40240	40949
	F_{1345}	42334	48547	49552	48188	40433	42921	49354
	F_{2345}	48607	49466	48132	43639	46775	41971	40184
	F_{12345}	54550	57982	56416	50253	50363	50539	50770

Table 37 Bidders' cost function parameters in the unequal capacities (B) auctions

		Bidder 1	Bidder 3	Bidder 4	Bidder 8	Bidder 9	Bidder 13	Bidder 15
Variable Costs	c_1	58,06	64,20	62,50	57,76	53,85	64,03	56,42
	c_2	54,92	53,61	57,95	66,27	55,01	65,53	65,93
	c_3	53,88	56,35	66,62	62,05	53,61	62,37	62,51
	c_4	65,52	59,30	56,61	56,37	65,19	54,31	63,24
	c_5	55,51	59,19	56,11	54,46	55,25	64,72	65,84
Fixed Costs	F_1	19635	17355	18239	16258	17196	16262	18466
	F_2	16576	18505	17497	19397	19719	19580	19094
	F_3	19135	17729	19955	17188	18652	16345	18284
	F_4	18093	18624	16583	16041	19937	18950	18681
	F_5	19704	16071	19752	19516	18491	16676	17829
	F_{12}	24964	27114	28477	26161	28601	28430	26739
	F_{13}	26764	26657	26962	24764	24926	25063	27071
	F_{14}	25847	25108	29133	27139	26637	25105	25626
	F_{15}	25042	28210	26492	24684	24773	27615	25549
	F_{23}	27501	29891	28301	29100	25777	26823	25422
	F_{24}	28442	24164	29082	24868	24889	24857	27556
	F_{25}	26582	26570	28532	24972	25423	25184	28519
	F_{34}	28385	29860	29385	27243	26194	26385	25532
	F_{35}	25112	24656	26220	26547	24586	27173	27953
	F_{45}	26151	28130	27062	28896	26993	28669	27715
	F_{123}	38070	33185	39646	36258	35136	36352	38739
	F_{124}	35861	36397	32763	39856	35548	34976	36750
	F_{125}	38723	37323	33967	34399	34256	36813	39910
	F_{134}	33403	36506	37369	38713	36166	37043	38794
	F_{135}	35990	36207	34658	33096	33748	32624	35264
	F_{145}	33484	39025	38708	36367	34471	37745	36521
	F_{234}	33749	36226	33008	38549	36967	32360	35033
	F_{235}	32767	34727	36100	38136	39249	33489	39010
	F_{245}	36642	38103	35639	33369	35321	32031	35594
	F_{345}	35093	32670	32373	39741	33992	39524	38494
	F_{1234}	49256	48888	47454	43536	47339	47114	42418
	F_{1235}	41993	44341	49214	49981	47720	49611	41734
	F_{1245}	45501	42292	48059	47062	49374	46941	47954
	F_{1345}	42764	41781	49930	40662	46926	41533	43732
	F_{2345}	48579	46659	45918	41613	49433	42818	41537
	F_{12345}	55647	54080	52026	57525	59721	56902	49149

Table 38 Bidders' capacities in the unequal capacities (B) auctions

	Bidder 1	Bidder 3	Bidder 4	Bidder 8	Bidder 9	Bidder 13	Bidder 15
Item 1	225	300	300	300	300	150	300
Item 2	225	300	150	300	300	300	300
Item 3	225	300	300	150	225	300	150
Item 4	225	150	300	300	150	300	300
Item 5	300	150	300	300	225	300	150

APPENDIX 5: AUCTION EXPERIMENT PARTICIPANT QUESTIONNAIRE

NAME: _______________
STUDENT ID: _________

1. Have you participated in any kind of online auctions before (yes/no)?

2. Did you place bids only with the help of the support tools after it became available (yes/no)?

Answer the following questions by choosing the appropriate number from 0-5.
0 = I don’t know
1 = I totally disagree
2 = I somewhat disagree
3 = I don’t agree or disagree
4 = I somewhat agree
5 = I totally agree

3. I understood the rules of the auction game	0	1	2	3	4	5
4. I understood what my goal was in the game	0	1	2	3	4	5
5. The CombiAuction site was easy to use	0	1	2	3	4	5
6. I understood how the price support tool works	0	1	2	3	4	5
7. I understood how the quantity support tool works	0	1	2	3	4	5
8. The price support tool was easy to use	0	1	2	3	4	5
9. The quantity support tool was easy to use	0	1	2	3	4	5
10. The price support tool was helpful	0	1	2	3	4	5
11. The quantity support tool was helpful	0	1	2	3	4	5

12. Please describe your bidding strategy during the two auctions (When did you bid? How often did you bid? How did you choose the items, quantities and price in your bids when you didn’t use the support tools? Etc.).

13. How did your experience in the A auction affect your bidding in the B auction?

14. Did you monitor the auctions within the 10 minutes before it was scheduled to close? If you did not have active bids, did you use the support tools? Were you successful in finding profitable bids?

15. How would you improve the usability of the CombiAuction system?

16. Is there anything else you want to comment on about the game?