

Applying Revenue Management in Auction design

A.M.A. Sman
BMI-paper

Applying Revenue Management in Auction design

A.M.A. Sman
BMI-paper

Vrije Universiteit Amsterdam
Faculty of Sciences
Business Mathematics and Informatics
De Boelelaan 1081a
1081 HV Amsterdam
www.few.vu.nl

May 2007

Preface

A part of the master's programme of Business Mathematics and Informatics at the Vrije Universiteit Amsterdam is writing a BMI-paper. In the BMI-paper, a problem from the field of BMI is assessed using existing literature.

In this paper, the challenges in auction design are considered through an analysis of existing literature in the field of revenue management and game theory. Also, the findings are tested with a series of auction simulations.

I would like to express my gratitude to Dr. Auke Pot from the Vrije Universiteit Amsterdam for his guidance with writing this paper.

Abstract

In this paper we consider the most important literature from the field of revenue management and game theory on the subject of auctions. We start with a short description of auction to familiarize with the terms involved, and for some general information on the subject. In Section 2 we elaborate on the various types of auction models, adjustments to these models and the differences between them. Section 3 will explain more about certain general problems in auction design that we should be wary of. In the second part of this paper, starting with Section 4, we will use simulations to get a general idea of the revenue that we can expect from certain auction types and the influence of adjusting some settings. Finally we discuss the results we found and compare them to the theory from the first part of the paper.

Contents

1	Introduction to auctions	3
1.1	What is an auction?	3
1.2	Where do we find auctions?	3
1.3	Why are auctions used?	6
1.4	What are the types of auctions?	7
1.5	Which auction type is best?	8
2	Modeling auctions	11
2.1	Independent Private Value Model	12
2.1.1	Sealed-bid second-price auction	12
2.1.2	Sealed-bid first-price auction	14
2.1.3	Comparing the open- and sealed-bid auctions	15
2.2	Common Value Model	16
2.3	Extensions of the models	18
2.3.1	Customer asymmetry	18
2.3.2	Collusion	19
2.3.3	Risk aversity	19
2.3.4	Multiple objects	20
3	Auction pitfalls	23
3.1	Collusion	23
3.2	Entry Deterrence	23
3.3	Loopholes	24
3.4	Creditability of the Rules	24
3.5	Solutions	24
3.5.1	Anglo-Dutch Auction	25
4	Research question	27
5	Approach	29
5.1	Simulating the IPV Model	29
5.2	Simulating the CV Model	30
6	Results	33
6.1	The IPV Simulations	33
6.1.1	Second-price	33
6.1.2	First-price	34
6.2	The CV Simulations	34
6.2.1	Second-price	35
6.2.2	First-price	35
6.3	Comparing the two models	35
6.4	Determining the best bid	36
7	Conclusion	37

1 Introduction to auctions

1.1 What is an auction?

The English dictionary describes an auction as: "A public sale in which property or items of merchandise are sold to the highest bidder.", and *Wikipedia* [17] describes auctions as: "An auction is the process of buying and selling things by offering them up for bid, taking bids, and then selling the item to the highest bidder.". Now, both of these definitions describe the most important aspect of auctions: a limited amount of items that are sold through a set of predetermined rules to the bidders that are willing to pay the highest amount for them. Consider [2] for a brief introduction to auctions as well.

1.2 Where do we find auctions?

We can find auctions in a variety of markets. These include for instance industrial as well as consumer and financial markets. The type of market the auctions are used for determine to great extent the rules that are used, which in their turn determine the revenue to the bidder(s) and seller(s). In [15], a distinction is made between five different types of markets. A short description is given below:

- **Traditional Auction Houses:** the two largest and most famous traditional auction houses are Christie's and Sotheby's. They typically provide auctions for selling fine art, antiques, books, jewelry, toys, dolls, and other collectible memorabilia. Sotheby's was founded in 1744 by book dealer Samuel Baker. Twenty-two years later, in 1766, Christie's is founded by James Christie, who resigns from a Navy commission to take up auctioning. In both Christie's and Sotheby's, a variation of the so-called, ascending, open-priced (English) is used, which will be explained later.
- **Financial-Market Auctions:** the government uses auctions to sell bonds and bills regularly, to finance national debts. Institutional and individual investors bid for the minimum interest rate they are willing to pay. Also, securities exchanges use *double auctions* for trading stocks, bonds and foreign exchange. In double auctions, bid offers are made by buyers and ask offers are made by sellers. Double auctions are left out of the scope of this paper, for more information on double auctions, consider [18].
- **Government Auctions:** many public assets, such as public lands and natural resource rights are sold by the government using auctions. For instance, the radio spectrum auctions for third-generation (3G) cellular phone services in many European countries and the United States. These auctions were in fact such a success to the sellers that many of the companies eventually winning these auctions came into financial trouble due to the staggering sale prices.

- **Industrial-Procurement Auctions:** in many industries auctions are used for the procurement of materials, services and general subcontracting of production. For instance when several companies want to land a construction contract, the bids are *selling prices* at which the company is willing to do this specific construction. In contrast to many other types of auctions, factors other than price, such as quality levels and delivery schedules are important in the final selection as well.
- **Online Consumer Auctions:** online consumer auctions are types of auctions like eBay [6], where consumers can conduct online auctions for selling and buying any kind of items. However, not only individual consumers use online auctions, also small business use them. A different type of auction site is priceline.com, where the auction mechanism is in some way the exact opposite of that of eBay. Customers declare what they are willing to pay for an item, and supplying sellers accept or reject these offers.

Under certain circumstances it can be very profitable to use auctions instead of list pricing for instance. There are many success examples of the application of auctions mechanisms, especially through the use of internet. A few recent examples are:

- **eBay.com:** As was mentioned before, eBay is an American Internet company that manages ebay.com, an online auction and shopping web-site where goods and services are sold worldwide. It was founded in 1995 by computer programmer Pierre Omidyar as AuctionWeb, and the first item sold was a laser pointer for \$14.83. In September of 1997 the website changed its name from AuctionWeb to eBay. After it went public in 1998, both Pierre Omidyar and the company's first president Jeff Skoll became instant billionaires.

On eBay, three auction types are used: Auction-style listings, Fixed Price format and Dutch Auctions. The first, auction-style listing allows the seller to offer one or more items for sale for a specified number of days. In the Fixed Price format the seller offers one or more items for sale at a Buy It Now price, and in the Dutch auction two or more identical items are sold in the same auction, and bidders can bid for any number of items. On *Wikipedia* [17] is explained how the first and last of these auction types work in practice:

"For Auction-style listings, the first bid must be at least the amount of the minimum bid set by the seller. Regardless of the amount the first bidder actually bids, until a second bid is made, eBay will then display the auction's minimum bid as the current highest bid. After the first bid is made, each subsequent bid must be equal to at least the current highest bid displayed plus one bidding increment. The bidding increment is established by eBay based on the size of the current highest displayed bid. For example, when the current highest bid is less than or equal

to \$0.99, the bidding increment is \$0.05; when the current highest bid is at least \$1.00 but less than or equal to \$4.99, the bidding increment is \$0.25. Regardless of the amount each subsequent bidder bids, eBay will display the lesser of the bidder's actual bid and the amount equal to the previous highest bidder's actual bid plus one bidding increment. For example, suppose the current second-highest bid is \$2.05 and the highest bid is \$2.40. eBay will display the highest bid as \$2.30, which equals the second-highest bid (\$2.05) plus the bidding increment (\$0.25). In this case, eBay will require the next bid to be at least \$2.55, which equals the highest displayed bid (\$2.30) plus one bidding increment (\$0.25). The next bid will display as the actual amount bid or \$2.65, whichever is less. The figure of \$2.65 in this case comes from the then-second-highest actual bid of \$2.40 plus the bidding increment of \$0.25. The winning bidder pays the bid that eBay displays, not the amount actually bid. Following this example, if the next bidder is the final bidder, and bids \$2.55, the winner pays \$2.55, even though it is less than the second-highest bid (\$2.40) plus one bidding increment (\$0.25). However, if the next bidder is the final bidder and bids an arbitrarily large amount, for example \$10.00 or even more, the winner pays \$2.65, which equals the second-highest bid plus one bidding increment."

For Dutch Auctions, which are auctions of two or more identical items sold in one auction, each bidder enters both a bid and the number of items desired. Until the total number of items desired by all bidders equals the total number of items offered, bidders can bid any amount greater than or equal to the minimum bid. Once the total numbers of items desired by all bidders is greater than or equal to the total number offered, each bidder is required to bid one full bidding increment above the currently-displayed winning bid. All winning bidders pay the same lowest winning bid.

This formula has proven to be very successful in attracting many bidders to the auctions and thereby adding to the success of the site. For the first quarter of 2007, eBay has announced a \$1.77 billion revenue.

- **Google adWords:** AdWords ([1]) is Google's main source of revenue. AdWords offers pay-per-click advertising, and site-targeted advertising for both text and banner advertisements (ads), through local, national, and international distribution. Advertisers specify the words that should trigger their ads and the maximum amount they are willing to pay per click. When a user searches Google's search engine on www.google.com, ads for relevant words are shown as "sponsored link" on the right side of the screen, and sometimes above the main search results. This has the advantage that these ads are only shown to those users who are actually interested in the subject.

Based on the other advertisers' bids and the so called *quality score* of all ads the order of the paid listings is determined. The quality score is calculated by historical click-through rates and the relevance of an adver-

tiser's ad text, keyword, and landing page to the search, as determined by Google, this quality score is also used to determine the reserve price of the keywords, the minimum bid amount.

In the first quarter of 2007, Google had a 63% revenue increase over the same period last year to \$3.66 billion ([14]), the greatest part of which can be addressed to their advertisements income. According to [5], advertising accounts for 99% of Google's revenues.

- **Online tv commercial auctions:** In 2006, a group of 10 marketers including Wal-Mart, Hewlett-Packard Co., Philips and Microsoft championed an initiative for a test to buy and sell TV advertising through online auctions. According to them this can add to the transparency and efficiency of media buying. The online auction site eBay reacted to this by presenting the "Media Marketplace", an online trading platform that could be used. An advertiser can pick from a set of various different networks, ages, genders and time slots, the system would then return the available inventory and the advertiser can enter a bid. Within a matter of hours or even minutes the advertiser could receive the message that they are the winner of the auction. Tests with this type of online advertising auctions began March this year ([8]), and although by that time no networks had yet signed up to participate, it had already spawned a lot of interest from groups of marketers, who were so dedicated to making this system work that they would consider sitting out the second quarter market entirely.

1.3 Why are auctions used?

There are many reasons why auctions have certain benefits that other methods simply don't have, but perhaps the most important reason can be found on the web-site of *Addis Equine Auctions* ([3]): "The auction method of marketing anything is the only way to establish true market value.", which means that if an item is sold by competitive bidding it often draws prices higher than the seller's expectations. This can be explained with the concept of consumer surplus.

Consumer surplus is the difference between what consumers are willing to pay, and the actual selling price. Consider this example from *Wikipedia*: "For example, a person is willing to pay a tremendous amount for water since he needs it to survive, however since there are competing suppliers of water he is able to purchase it for less than he is willing to pay. The difference between the two prices is the consumer surplus." So in some way, the consumer surplus is the advantage consumers get from the market functioning the way it does.

There are other advantages as well to auctions to sellers and buyers as well. For instance, the seller knows exactly when his property will be sold, and the item can be sold faster, decreasing holding costs. One very important advantage to the seller is that he takes no part in the negotiating process, which allows the seller to sell the property under his/her terms and conditions.

To the buyers there are advantages as well, to some extent, they can set their own purchase price. It is an advantage to the buyer as well that the property is sold in a short amount of time. And also, very importantly, buyers feel better about their purchase, knowing that there is a contending bid just one increment below their purchase price, creating a feeling that they purchased the property at a fair market value. All in all, auctions provide a good alternative to fixed price sales.

1.4 What are the types of auctions?

There is a wide range of types of auctions. Auctions where one single item is sold are perhaps best known, but for this type of auction many different sets of rules can be used. For instance, auctions can be either open or sealed, indicating the ability of buyers to see each others' bids. And they can be ascending or descending. Consider below the most common types of auctions for a single item:

- **English auction:** The English auction is an open ascending auction, also known as the open-outcry auction, and is the mechanism that is used in the auctions at the traditional auction houses Sotheby's and Christie's. In this type of auction the auctioneer begins the auction at a certain price, the lowest acceptable price, and then gradually increases this price until there is only one bidder left willing to purchase the item for that price. The item is then sold to that bidder for that price. Because of the high level of competition in this type of auction buyers can sometimes get carried away and bid more than they otherwise would have. This can lead to a phenomenon called the *Winner's Curse*, which will be explained later on.
- **Dutch auction:** The Dutch auction is an open descending auction, in this type of auction the auctioneer begins with an extremely high asking price which is progressively lowered until a price is reached that someone is willing to pay. It is named after its most famous application of this auction type, the Dutch tulip auctions. The English system is in some way considered to be inferior to the Dutch mechanism in the fashion that in the English auction, the price that is eventually paid could be a great deal lower than the winner's actual valuation. If the second highest bidder doesn't drive the price up too much, the winner could end up paying a lot less than he was willing to. However, in the Dutch auction, if someone really wants an item, he cannot afford to wait too long to make his bid, causing him to bid near to his actual valuation of the item.
- **Sealed-bid first price auction:** In this type of auction, the bidding is sealed, which means that it is hidden from the other bidders. Every bidder makes a sealed bid, and the winner is bidder who bid the highest amount, and pays exactly that amount. Because bidders can't derive any information from the contending bidders, and because bidders usually get to submit only one bid, a lot of effort must go into the preparation of

the bid. If we consider this type of auction from a bidder's point of view, making a high bid raises the probability of winning but lowers the profit if the bidder is victorious. Because of this, bidders choose to shade their bids a bid downward (they do not bid their true valuation of the item), which also helps to avoid the winner's curse mentioned before.

- **Sealed-bid second price auction:** This auction type is also called the Vickrey auction. This auction type is named after William Vickrey, who in his article [16], first showed that this type of auction has certain desirable theoretical properties, namely that, in contrast to the sealed-bid first price auction, a bidder's best strategy is to bid their true valuation of the item (See Section 2.1.1). Like in the sealed-bid first price auction, the bids are sealed, and the bidder with the highest bid wins. However, the winner doesn't pay his own bid, but the second-highest bid.

Many variations can be applied to these standard auction types, such as reserve prices or minimum bid increments. Also, an auction can be *reversed*. In this type of auction, sellers compete to obtain business, so the role of buyer and seller are reversed, mainly to lower purchase prices. The four auction types mentioned were explained for the single-item variant, however, there are also auction types where multiple (identical or unique) units are sold to multiple bidders. In these so called *multi-unit auctions* C items are sold to N bidders, these items can be, but are not necessarily, homogeneous. For each of the auction types mentioned above, the multi-unit version can be easily derived. Consider the English auction: the seller has C homogeneous items for sale, and each of the N customers only wants one item. In this case, the price is increased until only C bidders remain willing to pay the price p , and each of these bidders is awarded an item at price p . In the open descending auction the price is lowered until someone is willing to pay the current price, after which the price is lowered again until another bidder claims the item, this is continued until all items are sold or the lowest selling price is reached. For the sealed-bid auctions, the C highest bids are accepted, and bidders pay either their own bid in case of the sealed-bid first price auction, or the $(C + 1)^{st}$ highest bid in case of the sealed-bid second price auction. Other types of multi-unit auction exist as well, for instance where customers are allowed to bid for multiple items, in which case bidders submit bids for the prices they are willing to pay for any number of items. We will elaborate on multi-unit auctions in section 2. There is also a subtype of multi-unit auctions called the *combinatorial auction*. Combinatorial auctions are auctions in which bidders can place bids on combinations of items, or packages, rather than just individual items. This is typically used when customers are willing to pay more for a certain combination of items, than the sum of what they would be willing to pay for the items alone.

1.5 Which auction type is best?

There are many different auction types, but which one is best? This greatly depends on the circumstances of the auction. It's not easy to say a certain

auction type is better than another, however, it is possible to say how well an auction type would fit a certain set of circumstances. Sometimes, it might be difficult to have all bidders physically present at the auctions, in which case an English or Dutch auction is not the best option. In some auction types, the risk of collusion is greater than in others. Some auction types might take longer for the item to be sold, making the auction type unfit for selling items like fresh fish or flowers. For reasons like this the question of which auction type is best is very difficult to answer. However, we can try optimally fitting the auction rules to the circumstances. For more information on fine-tuning the rules and mechanism of an auction to maximize revenues, consider [12].

2 Modeling auctions

This section will deal with the models that were used for our research, the independent private value model and the common value model. The description of these models and their extensions is based on their descriptions in the chapter on auction in [15] and on article [10]. The names of these models, the independent private value model and the common value model, refer to the most important underlying premises of the models, concerning the valuations of the bidders. The valuations of the bidders, or the price they feel or perceive the object in the auction is worth, can differ strongly between different types of auctions. These differences in the perceived value of the object influence the outcome of the auction, and thereby the profit of the seller.

Let us consider the meaning of the words private and independent with respect to the valuations of the bidders. The independent private value model says that the bidders have an independent and private valuation of the object in the auction. Here, independent means that whatever bidder i thinks the object is worth tells us nothing about another bidder j 's valuation of the object. That the valuations are private means that the bidders don't know the other bidders' valuations. The combination of these assumptions is realistic when the object for sale is something with a value dependent of personal taste, such as a piece of art or consumption, for personal use (not with the intent of selling). One bidder might be willing to pay over \$100000 for a painting when another wouldn't think of paying half of that. The bidders have valuations that are independent of each other, and they don't know the others' valuations.

The second model that is described in [15] is called the Common Value Model, and as the name says, the valuations of the bidders are common. Common in the sense that the object has one common value to all bidders, and the difference between bidders lies in the fact that they all have a different perception of the value. This is the case when for example a construction contract is auctioned. In this type of auction, the bidders are companies that want to land the construction contract and the bids they make is the amount of money they ask for their services. Logically, the bidders try to make a bid as low as possible. The bidders make a bid based on the value they think the contract is worth, based on an approximation of the costs involved. When this guess is lower than the actual costs, the bidder will make a bid that is too low, not being able to cover the costs involved in the contract. When his guess is higher, the bidder will make a bid that is so high they will probably not win the auction. Either way, information on other bidders' valuation is very interesting to the bidders, because it tells them something about their own perception.

So when deciding on a model for an auction we first look at the circumstances and then decide which one fits best. Only after that we start looking at which auction type we should best apply for that situation, a question that we will address in the following subsections.

2.1 Independent Private Value Model

The Independent Private Value Model is discussed in great detail in [15], this section is a summary of that description focussed on the performance of the open and sealed bid auction types of the first price and second price auction. We will try to answer the question of which of the four auction types performs best when the bidders have private and independent valuations. To model the valuations v_i of the customers we will draw them from a continuous density f , with cumulative distribution function F , defined between 0, the lowest possible valuation, and a v_{max} , the highest possible valuation.

So for N bidders we randomly draw $v_1 \dots v_N$ from F , resulting in N independent valuations. In the IPV model, each bidder knows their own valuation, and they know the distribution F their valuation is drawn from. This means that the bidders know if their valuation is high or low compared to the others. They can use this information in determining the bid b_i they make.

Another assumption the writers make in [15] is that the bidders will behave rationally. This means that considering their valuation v_i and the distribution F , along with the auction rules, the bidders will determine a bid according to a *bidding strategy* that will maximize their *expected* revenue. The revenue of a bidder that wins is the price they pay minus their valuation. This optimal bidding strategy is determined in consideration of the others' bidding strategies.

When every bidder has a bidding strategy that optimizes their expected return, none of them have an incentive to change their strategy, because it would decrease their expected return and the bidding strategy would become suboptimal. This balance of bidding strategies among bidders is called an *equilibrium set of strategies*, or the *Nash equilibrium* in game theory. Now that we have set the rules and further defined the model we will start looking at the auction types individually and watch how they perform compared to each other.

2.1.1 Sealed-bid second-price auction

In [15], the IPV model is first used in a sealed-bid second-price type auction. In this auction type, the winner is the highest bidder and he will pay an amount equal to the second highest bid. In the auction, there are N bidders, each with their own independent private valuation v_i randomly drawn from F , and they will use this valuation and their knowledge of the underlying distribution F to determine their bid b_i .

As mentioned before, in determining the winner of the auction we are interested in the bidder that placed the highest bid, and in determining the payment we want to know the second highest bid. Therefore, $b_{[i]}$ is introduced as the i^{th} highest bid: $b_{[1]} \geq b_{[2]} \geq \dots \geq b_{[N]}$. Using this notation, the winner is the bidder that made bid $b_{[1]}$ and will pay an amount of $b_{[2]}$. We will denote the bidding strategy that the bidders apply as $b_i(v_i)$: the bid as a function of the valuation, and the optimal bidding strategy is denoted as $b_i^*(v_i)$.

When we look at the optimal strategy for this type of auction, the optimal bidding strategy is to bid your own valuation, or: $b_i^*(v_i) = v_i$. In order to prove

this point, first consider that in second-price type auctions the eventual winner cannot influence the price they end up paying. This is because the winner made bid b_1 and the payment is b_2 , a bid made by someone else on whom the bidder has no influence. So by making a bid higher or lower they can only adjust their winning chance, not the payment.

To prove the optimal strategy is indeed $b_i^*(v_i) = v_i$ consider the following analysis from [15] of the possible outcomes from bidding the true valuation: the bidder either wins the auction, or loses it. Next is determined what would have happened if they had bid more or less, to see if that would have improved the expected return.

- **Bidder wins:** The bidder ends up winning the auction with a bid equal to his valuation ($b_i = v_i = b_{[1]}$). Apparently, he has made the highest bid, and will pay an amount equal to the second highest bid $b_{[2]}$, which is, naturally, lower than his valuation so he has a positive return. If we suppose he would have made a bid higher than his valuation he still would have won, and still would have paid the same amount, so bidding higher than the valuation shows no improvement over the original strategy. If the bidder had bid less than his valuation he could have made a bid that was below the second highest bid $b_{[2]}$, and end up losing the auction, generating a revenue of 0, when in fact he could have won. So when $b_i = v_i$ is a winning strategy, it is the optimal strategy.
- **Bidder loses:** The bidder loses the auction with a bid equal to his valuation ($b_i = v_i < b_{[1]}$). Someone has made a higher bid, and we know that that bid is also higher than our bidder's valuation. If the bidder had made a higher bid than his valuation he could have ended up winning the auction. However, he would have had to pay a price equal to the current best bid $b_{[1]}$, which we know is higher than the bidder's valuation, so his return would be negative, he would pay more for the item than he actually wants to pay for it. Finally, lowering the bid wouldn't have increased the return either because the bidder didn't win for bid v_i , much less for any bid lower than that. This shows that even when $b_i = v_i$ is a losing strategy, it is still the optimal strategy.

The analysis above shows that regardless of the eventual outcome of the auction, the optimal strategy a bidder in a sealed-bid second-price auction can apply is to bid their true valuation.

In the analysis, no assumption were made on the strategies of the other bidders, other than that the valuations are independent and private and that the bidders behave rationally. When this is the case, the strategies are called *dominant*, because they are optimal regardless of the strategies of the others. Analogously, the equilibrium is called a *dominant-strategy equilibrium*. To calculate the expected revenue to the seller in this auction is now rather simple, to illustrate this consider the following example from [15].

Example 2.1 There are N customers with valuations uniformly distributed on $[0, 1]$, so $F(v) = v$ on this interval. Under the second-price auction, the optimal strategy is $b^*(v_i) = v_i$.

The expected revenue earned by the seller is simply $E[v_{[2]}]$: the expected value of the second highest bid. It is not hard to show that for $U(0, 1)$ distributed valuations:

$$E[v_{[2]}] = \frac{N-1}{N+1}$$

Now we know the expected revenue to the seller for the situation described above, and we also know that the expected revenue increases with the number of bidders N . The more bidders, the better it is for the seller, more competition means a higher selling price.

2.1.2 Sealed-bid first-price auction

After the sealed-bid second-price auction, the IPV model is next applied to the sealed-bid first-price auction type in [15]. In a sealed-bid first-price auction, bidders submit a sealed bid, and the winner pays an amount equal to his own bid. Unlike for second-price auctions, a bidder does in fact determine the price he pays, should he win. This changes the bidding strategies, because if you bid your true valuation in a first-price auction, and you win you pay exactly what you think the item is worth, creating a revenue of 0.

The writers make two additional assumptions made to the model for this auction type. The first is that all bidders have the same bidding strategy (assumption 1), and the second is that in the bidding strategy, a higher valuation leads to a higher bid, so $b_i(v_i)$ is strictly increasing (assumption 2).

To determine the expected return for the seller in a first-price auction, we will look at the winner of an auction, and determine the best strategy for a bidder. If bidder i competes in the auction and wins, we know that his bid was highest, higher than all other $N - 1$ bids. And according to the second assumption, if his bid was highest, then so was his valuation, because a higher valuation leads to a higher bid. So when someone wins an auction, his valuation was highest, higher than the other $N - 1$.

To calculate the probability of this happening we can use the distribution F of the valuations. The probability of someone having a lower valuation than v_i is $F(v_i)$, so the probability of all $N - 1$ other bidders having a lower valuation is: $F^{N-1}(v_i)$. In [15], this is denoted as $P(v_i)$, the probability of winning with valuation v_i .

We know that all bidders have the same strategy (assumption 1), so this must hold for all bidders, so we can lose the subscript i . Now that we know the probability of winning, we can calculate the expected revenue to the bidder: it is his probability of winning times the revenue if he does:

$$S(b(v), v) = P(v)(v - b(v))$$

The optimal bidding strategy $b^*(v)$ can be determined by differentiation over v

(for proof consider [15]):

$$b^*(v) = v - \frac{\int_0^v P(s)ds}{P(v)}$$

with $P(v) = F^{N-1}(v)$. This optimal bidding strategy tells us that the optimal bid is always v minus a fraction, dependent of the distribution function F , and this fraction decreases as the number of bidders increases. Apparently, the bidders shade their bids but as the number of bidders increases they shade their bid by less and less, trying to increase their probability of winning by lowering the revenue if they win. Now that we have found the optimal bidding strategy for the sealed-bid first-price auction we can make a comparison with the second-price type by applying the same valuation density function. Consider Example 2.2 from [15]:

Example 2.2 Again N bidders in an auction with valuations uniformly distributed on $[0, 1]$, so $F(v) = v$. In this case, $P(v) = v^{N-1}$ and the equilibrium bidding strategy is

$$b^*(v) = v - \frac{\int_0^v P(s)ds}{P(v)} = v - \frac{\int_0^v s^{N-1}ds}{v^{N-1}} = v(1 - \frac{1}{N})$$

This means that each customer bids a fraction $1 - \frac{1}{N}$ of his own valuation, and the more bidders N are competing in the auction, the closer the bids are to the true valuation of the customer.

The bidder with the highest valuation wins the auction. The expected valuation of the winner $E(v_{[1]})$ is not hard to determine since all valuations are drawn from the uniform distribution, it is simply the highest realization among N draws from a uniform distribution: $E(v_{[1]}) = \frac{N}{N+1}$.

With this expected valuation of the winner we can determine the expected bid, and thus the expected return of the seller. The optimal bid is $v(1 - \frac{1}{N})$ and with $v = \frac{N}{N+1}$ the expected highest bid is:

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

The expected revenue found here is equal to the expected return found in Example 2.1. This leads us to the conclusion that the expected return in a sealed-bid first-price and a sealed-bid second-price auction are equal.

2.1.3 Comparing the open- and sealed-bid auctions

Next, the writers of [15] make a comparison between the sealed-bid first-price and second-price auctions and the open auction types, the Dutch and English auction. The results from the sealed-bid versions are considered and translated to the open-bid versions. As it turns out, there is no difference in expected revenue in a sealed-bid or in an open-bid auction. In other words, the sealed-bid first-price auctions have the same expected revenue as the Dutch auction

and the the same holds for the sealed-bid second-price auction and the English auction type.

To demonstrate this, take a closer look at the optimal strategies that can be applied in the open-bid type auctions. We will find that the optimal strategies for the sealed-bid and open-bid types are exactly the same, generating the same expected revenue.

First consider the open (ascending) English auction. A bidder i is in the bidding process, suppose our bidder has a valuation v . For any price p below v the bidder will continue to participate in the bidding. Only when the price p increases above the valuation of the bidder the bidder will drop out of the bidding. In a way, the highest bid the bidder was willing to make was a bid equal to his valuation v and nothing more. This is equal to saying that his bidding strategy was $b^*(v) = v$.

To show that the expected payment is the same, consider the following: if the price had stayed below the valuation of the bidder until the end, the bidder had continued upping the bid, causing all other bidders to drop out at their respective highest bid. The price our bidder would then have to pay would be that of the second highest valuation, the highest possible bid his last competitor was willing to make. So the payment would be $v_{[2]}$. This tells us that the English auction and the sealed-bid second-price auction have the same expected revenue. This is due to the fact that the optimal bidding strategy is the same. When two auction types have the same optimal bidding strategy they are called *strategically equivalent*.

When we consider the (descending) Dutch auction, every bidder in the auction makes a calculation beforehand at which price to claim the object. The bidder does this by comparing his chances of winning with the revenue he will receive. This provides him with an optimal strategy for the auction; the bid b that maximizes his expected revenue. Now he will wait until the selling price drops below the maximum price he is willing to pay for the item and claim it. This strategy is equivalent to that of the sealed-bid first-price auction, resulting in the same expected revenue.

This section shows that for the expected revenue an open-bid type auction makes no difference with a sealed-bid type auction. The sealed-bid second-price auction and the English auctions perform the same, as well as the sealed-bid first-price auction and the Dutch auction and thus are strategically equivalent and have the same expected revenue.

2.2 Common Value Model

Now that we know how the auction types perform when the valuations of the bidders are independent and private, we want to make a step towards another type of situation where the valuations of the bidders have one common component, using the Common Value Model from [15]. This is the case when the object for sale is for example a construction contract, and the bidder with the lowest bid (the one that agrees to do the contract for the least amount of money) is awarded the contract. If the bid you make doesn't cover all costs involved,

you might land the contract, but it will only cost you money. The trick is to make an estimate of the costs involved and make a bid that covers the costs but is still low enough to make a real chance in winning the auction. It is for this reason that bidder's are now suddenly very interested in other bidders' valuations (or perception of the contract's value), because it tells them something about how well their own estimate of the costs is.

The bidders in the auction all try to make a good estimate of the value of the contract, but some of them will end up on the low end of the real value, and others on the higher end. The writers describe the Common Value Model as following: the common value is denoted a , which is what the contract is actually worth (the minimal costs involved in the project), and add to this an uncertainty factor ε with mean = 0, which tells us how much the estimate is off. So for bidder i the perceived value of the auctioned item is:

$$p_i = a + \varepsilon_i$$

with ε_i the random noise for that bidder with mean 0. Furthermore, the bidders in the auction are aware of the distribution of ε , so they know the degree of uncertainty of their own valuation.

Now suppose we have a sealed-bid first-price auction under the aforementioned assumptions. All bidders have their own perception of the value of the contract p_i . The bidder that wins the auction is the one that places the *lowest* bid, the one that agrees to do the construction job for the least amount of money. This is in all cases also the bidder that made the lowest cost estimation. The bidder that made the lowest cost estimation is the bidder with the lowest p_i . For N bidders, the lowest perceived value is $\min\{p_1, \dots, p_N\}$. Because $p_i = a + \varepsilon_i$, and ε_i has a mean of 0, the lowest perceived value p_i has a negative uncertainty component, and so $p_i < a$. This means that the winner of the auction is the one that has made the lowest estimation of the costs, which is less than the actual costs. So you might have won the auction, but it will cost you more than you gain. This is called the *winner's curse* (see also [11]).

The winner's curse is a well known phenomenon, also among the bidders, and they react to this by adjusting their bids upward, making sure that even if they make a too low estimate of the costs, their bid will still be high enough to cover the costs if they win. This of course alters the bidding strategy, and thereby the expected return to the seller.

The effect of this change in bidding strategy and uncertainty on their own estimate can have different impacts depending on the type of auction. In sealed-bid type auctions, the bidders can learn nothing from each other, so they are not influenced by each other. However, in the Dutch and English type auctions seeing how many bidders persist on bidding (English) or refrain from claiming the item (Dutch) may cause them to start doubting their estimate and make last minute changes to their bid.

For this reason it can be shown that for example the English type auction is better for the seller than the sealed-bid second-price auction under common value model, when in fact, under the independent private value model they

perform exactly the same. How much this difference is between the types and under the models is a question that will be addressed later in this document.

2.3 Extensions of the models

After the basic models are described in [15], the writers take some of the underlying assumption in the IPV model and the CV models, and change them to see what happens to the revenue and the bidding strategies if the circumstances change.

For example, in the previous section we implicitly assumed that all bidders are the same, when it is very well possible that there are customers that are wealthier than others or have more expertise. We assumed that bidders don't work together, when in fact they sometimes do. Some customer might be more afraid of risk than others, and in some auction types multiple objects are sold at the same time.

In this section, we will discuss these extensions to the models, and describe the new situation. We will then use the knowledge of the previous sections in maximizing the expected revenue to the seller.

2.3.1 Customer asymmetry

So far we have considered all bidders to be individual and autonomous, but essentially the same. They have the same bidding strategy and, more importantly, the same valuation distribution. It could occur however, that at an auction there are different types of bidders present, some might be more experienced than others, some might have more expertise on the object for sale, and some might simply have more money to spend. In [15], a way to model this is proposed, by making groups of bidders, all with different valuation distributions.

Let's suppose we have $N = 10$ bidders, 5 of which are of type 1, with distribution function F_1 , and the other 5 are type 2, with distribution function F_2 . Now let's assume that type 1 customers generally have lower valuations, and type 2 customers have higher valuations. The normal rules would have the highest bidder win the item, so the winner will oftentimes be a bidder from group 2, and only every now and then a type 1 customer. With one group being vastly superior to the other, competition is not very high in the bidding. By increasing competition among the bidders the seller can drive up the bids and have a higher revenue.

One way to do this is by favoring the 'lower valuation' group over the other group. By giving group 1 a higher probability of winning, group 2 will react by bidding higher to increase their winning chances. That way, by occasionally even accepting a lower bid from group 1, the expected revenue will increase. Sometimes it can be very useful to discriminate among bidders.

2.3.2 Collusion

One other implicit assumption in the IPV and CV model is that customers do not cooperate, they only try to maximize their own profit. In reality, sometimes bidders do work together, and form a so-called *bidding ring*. A bidding ring takes away the most crucial factor in a successful auction: competition.

One example of collusion is when the bidders work together in a bidding ring by agreeing to submit bids that will minimize the eventual payment that has to be made by the winner (See section 3.1 for more collusion types). We will use a simple example to demonstrate the effect of collusion on the revenue.

As we saw before, the number of bidders in an auction are important to the eventual bidding price, because more bidders means more competition and the more competition in an auction the higher the bids. There are $N = 10$ bidders in an auction, when 5 of them work together, the five of them are in fact just one bidder. So we have not 10 bidders, but only six. This decreases the competition and thereby the expected revenue. It gets even worse when all bidders in an auction are colluding, there is only 1 bidder, so any bid above the reserve price wins.

So if the seller suspects collusion among the bidder it is sensible to adjust the reserve price accordingly. When there are more bidders in a bidding ring, the more the ring is willing to pay for the item, since the highest amount the ring is willing to pay is $\max\{v_1, \dots, v_N\}$, or the highest valuation of anyone in the ring. This shows that the more colluders there are in the auction, the higher the reserve price should be.

2.3.3 Risk aversity

Another implicit assumption that is adjusted in [15] is that bidders are risk-neutral. It is possible, for example, that some bidders in the auction are risk-averse, in the sense that they would rather accept a lower, yet more certain payout.

Consider the sealed-bid type auctions, with risk-averse bidders, and a risk-neutral seller. In a second-price auction the bidder has no influence on their return if they win, because it is determined by the second highest bid. This causes the risk-averse bidders to favor the sealed-bid first-price type auctions, because they have full control over the payment they have to make. Their risk aversity also causes them to shade their bids by less than the risk-neutral customers. They increase their probability of winning, by decreasing their return. From a seller's perspective, it is better to hold a first-price type auction in an environment of risk-averse bidders.

There is also the possibility of risk-neutral bidders and a risk-averse seller. Comparing the first to the second-price type auctions, the first-price type auctions has a more certain revenue, so the risk-averse seller will favor this type over the second-price auction. So in a risk-neutral environment the first- and second-price type auctions perform the same under the IPV model, but in an environment of risk aversity (either bidders or seller) the first-price type auc-

tion is preferred. This could be the reason for the more prominent use of the first-price type auctions in practice.

2.3.4 Multiple objects

The last extension of the model we will discuss is the auction for multiple objects. The simultaneous selling of multiple objects requires a change of auction rules compared to the ones we have discussed before. This section will describe some of them.

The first auction type we will discuss is the *multi-unit Vickrey auction*, based on the second-price type (or single-unit Vickrey auction). In this type of auction, the bidders don't submit one single bid, but a demand curve, indicating how much they are willing to pay for 1, 2, ..., M items. The best demand curves are then selected to maximize the profit to the seller and the payment is as follows: the bidder that is awarded k items he will pay the sum of the 1st + ... + k^{th} *rejected* bids. This mechanism has the same advantage as the single-unit Vickrey auction, which is that it motivates the bidders to bid their true valuation. However, the rules are rather complex, and as a result, it is rarely ever applied in real-world auctions.

Another type of multi-unit auction is the *multi-unit uniform price auction* ([7]). This auction type is more commonly used and is a relatively simple concept. The bidders submit a maximum price p_i they are willing to pay for one item, and the number of item k_i they are willing to buy. Then, the highest bidder is awarded the number of requested item, and so is second highest bidder and the third and so on until the supply is exhausted. The price each has to pay then is that of the highest *winning* bid. One downside to this auction type however, is that it is highly vulnerable to collusion.

In an article by Lawrence M. Ausubel [4], an auction type is proposed that combines the incentive of bidding your true valuation from the multi-unit Vickrey auction with the relative simplicity of the multi-unit uniform price auction. In his article Ausubel explains that all successful auction types have two characteristics in common: the price paid by a winning bidder should be independent from his own bid, and the auction should be designed to maximize the available information to all bidders in the auction. The first causes bidders to bid their true valuation, the second creates more competition and causes the bidders to bid more aggressively, which is positive for the expected return.

The auction starts at a reserve price p_r , and all bidders submit how many items they are willing to buy at this price. As the price goes up they continue to do so, and for every price p , the total demand of all bidders except the highest (the one that claims most items for that price) is compared to the total supply. If that total demand does not exceed the supply the remaining items are *clinched*, which means that they are sold at the highest bidder at the current price p . After that the price is raised again, bidders submit the number of items they are willing to buy and so on until all items are sold. For an example of this type of auction consider the following example from [4]:

Example 2.3 Suppose that two identical objects are available and that three bidders (A , B and C) initially bid for quantities of 2, 1 and 1 respectively. The total of the highest bidder's (A) competitor's demand is not less than the total supply, so the price goes up.

The bidders continue to bid these quantities until price p , when bidder C reduces his bid from 1 unit to 0 units, dropping out of the auction. While there continues to be excess demand, bidder A 's opponents now collectively demand only one unit, while the total supply (M) is still 2. Therefore, bidder A clinches 1 unit at price p , and the auction (for the remaining object) continues.

This example shows the workings of the third and last multi-unit auction type that will be discussed in this section. This concludes the section about the extensions to the IPV and CV model, the next session will discuss possible pitfalls in designing an auction.

3 Auction pitfalls

In the article [9] by Paul Klemperer, many pitfalls in auction design are discussed and emphasized with examples of auctions going dramatically wrong, in this section we will consider the most important ones. As we have shown before, there is no clear answer to the question of what the best auction type is. What's important is that the circumstances are taken into consideration when designing the auction mechanism that is to be used, and that it fits these circumstances to guarantee an optimal result. Many examples can be given of auctions that went wrong due to any kind of reason. In this section we discuss some of these pitfalls and mention possible countermeasures to avoid them.

3.1 Collusion

Collusion is a common problem in auctions that we discussed briefly in the previous section. Bidders colluding in an auction to avoid the price from going up too high is one of the greatest risks to the seller. For instance, in an English auction bidders can use the early stages of the bidding by signalling each other about who is supposed to win the item, if the others pick up this signal they can stop bidding to avoid the price from going up too high. This signalling can occur by means of using the last and least significant digit of the bid to send a message. This can also lead to punishment from other bidders if one of these doesn't comply to their rules and compete in an auction they 'weren't supposed' to bid in. The ring will then retaliate by driving up the price in an auction the bidder was supposed to win. This shows that bidding rings even have their ways of assuring the compliance of bidders in their strategies by punishing and/or awarding bidders.

3.2 Entry Deterrence

Another important practical risk in auctions is when there are too few bidders competing in the auction. We have shown before that the seller profits from a larger number of bidders due to the higher degree of competition. In the ascending auction, the probability of the bidder with the highest valuation eventually winning is very high, because at any moment in the bidding they can probably top the current price. This means that when bidders know they will probably be outbid by the competition they will probably not enter the bidding in the first place, especially when they have to pay bidding costs.

Other auction types can suffer from lack of entry as well, in one specific case in the 1991 U.K. sale of television franchises by a sealed-bid auction, regions were being sold for as much as 10 to 16 pounds per head of the population. But for one region, the company that was bidding learned that there were no competing companies bidding and they then placed the only and highest bid of one-twentieth of one penny.

Another issue discussed before as well is the *winner's curse*. When bidders know about the winner's curse they will adjust their bids downward leading to

less profit for the seller. This problem also leads to fewer bidders entering in the auction.

3.3 Loopholes

When determining the auction mechanism, one must be aware of the risks of bidders detecting loopholes in the rules. To illustrate this, consider the following example from "What really matters in auction design":

In 2000, Turkey announced two telecom licenses sequentially, with an additional twist that set the reserve price for the second license equal to the selling price of the first. One firm then bid far more for the first license than it could possibly be worth if the firm had to compete in the telecom market with a rival holding the second license. But the firm had rightly figured out that no rival would be willing to bid that high for the second license, which therefore remained unsold, leaving the firm without a rival holding the second license!

3.4 Creditability of the Rules

It is very important that the rules of the auction are abided by. However, in some instances, the auctioneer can find himself in a situation where following the rules might mean the end of the auction. For example when there are N bidders competing for $N - 1$ items, all it takes is for one bidder to drop out for the auction to fail. In this case, the auctioneer might not be too prone to undertake action against one of the bidders in case of misconduct.

In case of the Turkish telecom license loophole, the government may consider holding a new auction to sell the second license. However, the impact of this on the creditability of the rules can be disastrous for future auctions. If an item is sold even though the reserve price hasn't been met, the reserve price has no real meaning, and the bidders will react to this by ignoring it in future auctions.

3.5 Solutions

Some of the aforementioned problems can cause an auction to fail dramatically, which usually comes down to producing very little revenue. By simply adjusting the rules a little, these problems can be prevented. For instance, setting a proper reserve price is very important. If the auctioneer sets a low reserve price, the inclination of bidders to form a bidding ring is much higher, since so much more can be gained from it, in contrast to the case that a reasonable reserve price has been set. The problem with setting a high reserve price however is that it increases the probability of the item not selling, in which case the auction is seen as a failure, a very embarrassing outcome to the seller. This political problem makes setting a proper reserve price all the more difficult.

Making it more difficult for bidders to signal each other can also be a useful method in preventing collusion. The problem of using the last digit to send a message can be solved by forcing bidders to bid round numbers, and making the

bids anonymous. Also, pre-specifying the minimum bid increment can thwart the bidders' efforts to send messages.

Another way is to choose a different auction type all together, in a sealed-bid auction all bidders make simultaneous bids, making it impossible for bidders to react to each other, so bidders can't retaliate against other bidders who fail to cooperate with them. In his paper [13], Marc Robinson studied auction for oil leases to put this theory to the test. Another advantage of this auction type over the ascending auction for instance, is that the outcome is less certain. From this follows that 'weaker' bidders are more likely to enter the auction, this improves the competitiveness of the auction, which increases the expected revenue to the seller. This effect can also be reached by not ensuing any bidding costs, as this will stop the weaker bidders from entering the auction.

3.5.1 Anglo-Dutch Auction

One solution that Klemperer proposes to combine both the advantages from the open and sealed bid auctions is the Anglo-Dutch Auction. This auction consists of two phases: in the first phase, an ascending auction is being held, the price is increased until only two bidders remain. Then, in phase two, both bidders will make one final sealed bid which has to be higher than the price at the end of the first phase, and the winner pays his bid.

According to Klemperer, the greatest advantage from running the Anglo-Dutch type auction arises when one bidder is considered to be much stronger than the others. In a classical English auction this bidder would be a sure winner, making potential rivals unwilling to enter. The sealed bid phase however induces some uncertainty about which of the final two bidders will win and that attracts much more potential bidders. This can cause the bidding price at the end of phase one to be higher than when simply an English auction has been used.

4 Research question

In Section 2 of this paper we took a close look at some of the basic auction models and certain adjustments to model more specific auction/customer types. In this part of the paper we will use this knowledge to simulate auctions and study the results that different auction mechanisms and settings produce.

The research question that we try to address is to what extent do the results from auction simulations match up with the results the theory would predict. More specifically, are second- and first-price auctions revenue equivalent, are the English (open) auction and the sealed-bid second-price auction basically the same in terms of revenue produced, and does this also hold for the Dutch (open) auction and the sealed-bid first-price auction mechanisms. With use of this simulation program we will study the effect on the auction mechanisms of changing the parameter settings. Running the simulations with different amounts of bidders, other valuation distributions and such we can see how this affects the auction mechanisms performance. While this gives as an idea of the robustness of the mechanisms to changes in the parameters settings, it also provides us with a more thorough understanding of the auction theory from Section 2.

After that we take the opportunity to study a method that was used in the simulations to determine an optimal bid given the valuation of a certain customer and its distribution. This shows the optimal bid for various valuation distributions and parameter settings.

5 Approach

We can get an idea of the revenue produced by an auction type by simulating this auction and every aspect of it, if we then change some its settings and run it again we can see the effect this has on for instance the revenue it produces. Section 2 of this paper is on modeling auctions, and the main focus of that section is on two different types of models. The Independent Private Value Model (section 2.1) and the Common Value Model (section 2.2). For the simulation we used two different programs, one for each model. We will explain the workings of each of these briefly.

5.1 Simulating the IPV Model

In the Independent Private Value Model, no difference is made between the open- and the sealed types of first-price and second-price auctions. Therefore, neither does the program in simulating these auctions, there are two settings: first-price (Dutch / sealed-bid first-price) and second-price auctions (English / sealed-bid second-price). For both types, there are 4 different valuation distributions, these are the distributions that determine the valuation of the customers. The four types are the uniform distribution (a, b) , the exponential distribution (λ) , the normal distribution (μ, σ^2) and the triangular distribution (a, b, c) . Furthermore, the number of bidders N in the auction can be chosen. Now that we have determined the settings, we can run M auctions, and determine the mean and variance of the revenue these auctions produce.

The workings of the second-price type of auctions are rather simple: N valuations are drawn from the distribution. Because bidders in the second type auctions bid their true valuation this is also immediately their bid in the auction. The N bids are considered, but we are interested in the second-highest bid only, because this is the price paid by the winner and the revenue produced by the auction. This revenue r is stored in an array, and another auction is simulated and its revenue is also stored, until M auctions are completed. With this array of M revenues we can determine the mean $\bar{r} = \frac{1}{M} \sum_{i=0}^M r_i$ and the variance $\sigma_r^2 = \frac{1}{M} \sum_{i=0}^M (\bar{r} - r_i)^2$ for these settings of the auction. We can now alter the number of bidders N or the distribution type or its parameters to see what effect this has on the expected revenue the auction produces.

The first-price auction is simulated differently. Again, N valuations, one for each bidder, are drawn from the valuation distribution. We saw in Section 2 that in first-price auction bidders tend to shade their bids in order to make a positive surplus. How much though, should they shade their bid such that they maximize their expected surplus. Given their valuation v and a certain bid $b \leq v$ it is possible to determine the expected surplus. If they win, their surplus is $v - b$, if they don't the surplus is 0. The expected surplus is then $P(win)(v - b)$, this has a maximum for a certain b^* which can be calculated given the valuation distribution and number of bidders. This is done by a method in the simulation program, it calculates the expected surplus for each possible bid and returns the bid for which it is highest. Now, the (optimal) bids are determined for each

customer, the N bids are compared and the highest bid is the price paid by the winner, and also, the revenue to the seller. Like in the second-price auction the revenue is stored, and after M simulations \bar{r} and σ_r^2 are calculated.

5.2 Simulating the CV Model

The main difference between the Independent Private Value Model and the Common Value Model is the fact that in the later model the bidders can draw information and conclusions from the bids (or absence of bids) of their competitors. For obvious reasons, this effect is canceled out in the case of sealed bids. Therefore, in the CV model auctions simulations we consider only the open bid auctions, the English and Dutch auctions.

For the simulation of the English auctions, again N numbers are drawn from a probability distribution that represent the valuations of the customers. The valuation of customer i is the sum of the actual value of the item a and a random noise term ε_i (for more information, see section 2.2). Now that we know the valuations of the customers the bidding can begin at a predetermined reserve price p_r . Every customer that has a valuation higher than p_r is willing to bid for the item, so a random customer is picked, and this customer makes the bid for price p . Now that the price has gone up to p , any customer that still has a valuation higher than p plus the minimum bid increment is willing to make a bet, so again a random customer is picked that makes a bet. Because in the CV model the bidders use the information from the others' bids to re-examine their own valuations, the valuations are updated after every bid. Suppose customer i has valuation v_i , as soon as bids are being made that are higher than $v_i - \sigma$, the valuation v_i of customer i is increased. By how much it is increased depends on two things: the difference between the valuation and the highest bid, and the *information value*, this information value represents how much the bidders rely on what they see their competitors bidding. If this information value is set to zero, they pay no attention to it, in which case one would get the same results as in the IPV model. After all valuations are updated to the latest bid, another bid can be made. This process is repeated until only one bidder remains, this is the winner and pays the current highest bid (p) for the item. This auction is repeated M times, and just like before, \bar{r} and σ_r^2 can rather easily be determined.

The Dutch auction is a descending auction type, so we start at a very high price p , one that no-one is willing to pay. Using the method mentioned before that determines the optimal bid given a certain valuation. Now the N customers all have an amount b that the maximum amount they are willing to pay for the item. As the price p declines, it comes closer to this amount. However, as long as nobody has claimed the item, this will tell the bidders more and more about the others' valuations. What is the value of the information they can draw from this, if anything, it will lower their valuations of the item. This is done in a similar way as in the English auctions, as soon as the price drops within a range of σ above the valuation v_i of customer i , this customer starts to adjust his valuation downward. By how much, again, is determined by the difference between the current price, and the information value. If this information is 0, the

mechanism is simply a first-price auction, with valuations $N(a + \mu, \sigma^2)$. These auctions will be repeated M times, after which the \bar{r} and σ_r^2 are calculated.

6 Results

Now we will take a look at the results of the simulations. First we will consider the Independent Private Value model and its results, and after that the results of the simulations of the Common Value model. After that we will take a moment to compare these two models. Finally we will also discuss some results that came up in the analysis of the method that determines the best bids for a certain valuation that was mentioned a few times before.

6.1 The IPV Simulations

Using the program that simulates auctions according to the Independent Private Value model we were able to plot the expected revenue an auction type produces as a function of the number of bidders N for various distribution (types) of their valuations. First, we will consider the results from the second price type auctions.

6.1.1 Second-price

To plot the expected revenue against the number of bidders N in the auction, the auctions were simulated for six different valuation distributions, all of which having a mean of 100. This means that in the simulations of the second price auctions, all bidders have an expected valuation of 100:

- A uniform distribution from 50 to 150.
- An exponential distribution with $\lambda = 0.01$.
- A normal distribution with $\mu = 100$ and $\sigma^2 = 35$.
- A triangular distribution with settings $a = 0, b = 300, c = 0$.
- A triangular distribution with settings $a = 50, b = 180, c = 70$.
- A triangular distribution with settings $a = 20, b = 150, c = 130$.

The simulations were run 10000 times for each distribution type for $N = 2, 3, \dots, 15, 20, 25, \dots, 50, 60, \dots, 100$. The results are shown in Figure 1 in the Appendix.

The most obvious conclusion one can draw from Figure 1 is that the expected revenue an auction type produces increases significantly with the number of bidders. Especially for small numbers of N , adding one extra bidder to the auction can result in a revenue increase of up to 10% depending on the distribution. It seems that the more outliers the valuation distribution has, the more revenue the auction produces, especially for high numbers of bidders (N), to illustrate this, consider Figure 2.

In this figure the valuation distributions for $N(100, 10)$ and $N(100, 35)$ are plotted against the auction revenue. The probability of high outliers amongst the valuation of the bidders plays a great role in the profitability of an auction.

The downside of this, is that when many outliers occur, the amount of revenue is less certain. For instance, in the latter example, the normal distribution with the lower σ^2 has a revenue variance of 5.42 for $N = 10$, whereas the one with many outliers and more revenue has a variance of 261.02. The revenue is higher, however it is much less certain.

6.1.2 First-price

Again, the simulation are run for the same numbers of bidders N , and for the same valuation distributions as before, but this time for the first price auction types. The results are plotted in Figure 3. At first glance the curves look exactly the same as for the second-price type auctions, however, if we take a closer look it seems that this doesn't hold when the valuations have an exponential distribution (see Figure 4). The revenue is on average 10 points higher in a first-price auction than in a second-price auction when the valuations are exponentially distributed. This also happens to a smaller degree when the valuations are triangularly(0,300,0) distributed, the difference is 1 or 2 points on average. The reason for this behavior is not explained in the literature, however, an explanation could lie in the fact that both these distributions have a high probability of very high outliers, especially the exponential distribution. Apparently, if in an auction some members can have a vastly higher valuation than others, like in the exponential distribution, the seller is better of using a first-price auction mechanism.

In all other cases, such as for the normal and uniform distributions, the first- and second price mechanisms produce the same revenue to the seller. One important difference however, between the two auction types, is indicated in the literature as well. The variance in first-price auctions is much lower. In the second price auctions with $N(100, 5)$ and $N(100, 35)$ valuations the variance is 5.42 and 261.02 respectively, for the first-price auction the variances are 3.30 and 153.06; the first-price and second-price have the same expected revenue, but the first-price types have a more certain return, a solid reason to prefer the first-price type auction over the second-price auctions.

6.2 The CV Simulations

If we simulate auctions according to the Common Value model, we expect a higher revenue for the second-price auction (English) and a lower revenue for the first-price auction (Dutch). If we exclude the information-gain factor however, the results should be exactly the same. Suppose we do a hundred-thousand runs of the English auction, with $a = 100$, $\mu = 0$, $\sigma^2 = 35$ and $N = 10$, without the bidders using the information they gain from the bids . The expected revenue is then 135.5 with variance 263.4. If we do the same in the IPV model for $\mu = 100$, $\sigma^2 = 35$ and $N = 10$, we get an expected revenue of 135.1 with variance 262.9, this is practically the same. It is clear from this example that the sealed bid auction, sealed or open, do in fact return the same revenue under this model,

and that the difference in the expected results from the CV model compared to the IPV model will only occur because of the behavior of the bidders.

6.2.1 Second-price

As expected, we do see an increase in revenue in the second-price type auctions when the bidders can learn from each others' bids. In Figure 5 the same valuation distributions are plotted for the two different models. For $\sigma^2 = 35$ the difference is clearly visible, with a revenue of approximately 10 points higher with the same valuations. For $\sigma^2 = 5$ this difference is hardly visible in the plot, but also for small differences between bidders the revenue increases due to bidders learning from each other.

6.2.2 First-price

In the first price type of auctions we expect to see a decrease in revenue, as long as the bidders haven't made a bid yet, they become more and more uncertain of their own valuations and adjust them downward, to illustrate this, again, the simulations are run for 10 bidders and for valuations $a = 100$, $\mu = 0$, and $\sigma^2 = 5$ and 35, for the Dutch auction. And indeed, the simulations show a decrease in revenue, especially when bidders have a high valuation variation.

6.3 Comparing the two models

In the table below the results, expected revenue ($E(r)$) and revenue variance ($Var(r)$), are summarized for the IPV model and the CV model for $\sigma^2 = 5$ and 35. The differences between the models and the auction types are most significant for a high variance, but also for the low variance there is a clear difference.

	IPV		CV	
	$E(r)$	$Var(r)$	$E(r)$	$Var(r)$
First	136.1	159.2	125.1	100.8
Second	135.6	263.4	144.7	189.9

Table 1: Simulation results for high valuation variance.

	IPV		CV	
	$E(r)$	$Var(r)$	$E(r)$	$Var(r)$
First	105.1	3.2	103.4	3.1
Second	105.0	5.4	106.4	4.7

Table 2: Simulation results for low valuation variance.

In Table 1 and 2 we can see that for the IPV model, the first- and second-price auction produce the same amount of revenue but, as mentioned before, the revenue variance is less for the first-price auctions. In the CV model, where we consider the open-bid auctions only, we see a decrease in revenue for the first-price types, and an increase for the second-price, this is due to the type of information the bidders gain from observing their competitors. In both cases, the revenue variance becomes less if the bidders' observe each-others' bids and try to use this information in re-examining their own valuations.

6.4 Determining the best bid

Finally we will use the methods that were used to determine the optimal bid to see how much margin the customers take on their valuation to determine their optimal bid depending on their valuation distribution and the number of bidders. In Figure 6 the optimal bids are plotted against the valuation for an auction with 5 bidders, when we compare this to Figure 7, where the number of bidders is 20, it is clear that bidders realize that when they have more competitors for the item, they have to adjust their bids upward. Also, the valuation distribution (type) has a lot of influence on the optimal bid. The shape of the valuation distribution and the number of customers together determine by how much customers can shade their bids.

7 Conclusion

This paper demonstrates how even though auctions seem very complicated at first glance, especially due to factors like customer bidding behavior, they can be broken down into smaller pieces and analyzed to get a grasp on auction performance. Using existing literature on auctions it is demonstrated that the performance of an auction type greatly depends on the circumstances they are held in. The question if customers observe and learn from the other bids, is a very important one. But also, the type of product and the risk of collusion in the auctions, determine to great extent the revenue to the seller.

We saw that, when customers don't observe each others' bids, as is usually the case with items for personal consumption, the revenue is equal for many different auction types, open or closed and ascending or descending. However, a customer that wants risk may prefer the auction types with less variance in return, such as the first-price types, when another customer may want to risk it and try for an even higher return in the second-price auction. Under the Common Value Model the bidders do in fact watch the other bids and draw conclusions from them. In this case, in the second price auctions, the bidders are inclined to increase their bids when they many competitors approaching their own valuation. They start questioning their own assessment of the item's value and start bidding more aggressively, resulting in a higher return for the seller than otherwise. The opposite is the case in the CV model for the first-price open, or Dutch, auction. As the price drops and no-one has claimed the item yet, bidders become less certain of their own assessment and do the opposite as in the English auction, they adjust their bids downward resulting in a lower return for the seller. Again, this is only the case when there is a common value factor, when bidders can draw information from the others' bids.

This is underlined by the results from the auction simulations, where we see that there are differences between auction types, depending on the circumstances. We also see that the number of bidders has influence on the expected revenue, as was also mentioned in the literature. The reason for this is simply that an increasing number of bidders increases the competition and the bidders realize that they have to compensate for this by bidding higher.

There has much been written about auction theory, many different auction types and many different models to describe them. In this paper the most common ones were discussed and represented in a simulation that did in fact match what we expected from the studied literature.

Appendix




Figure 1: Second-price auction simulation results

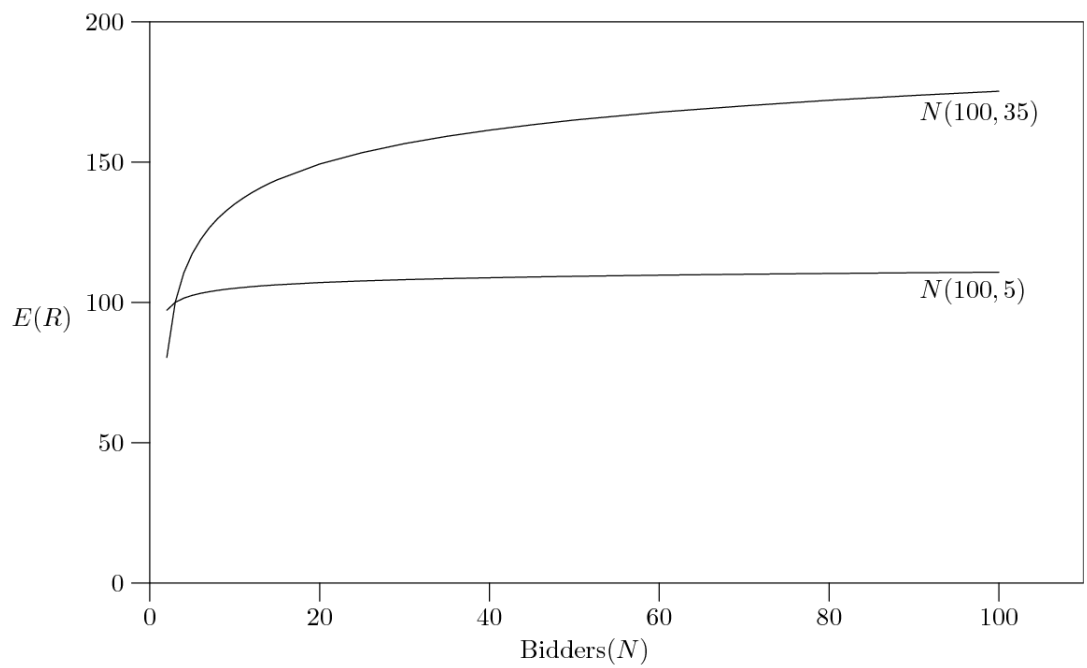


Figure 2: Second-price auction simulation results for normal distributions

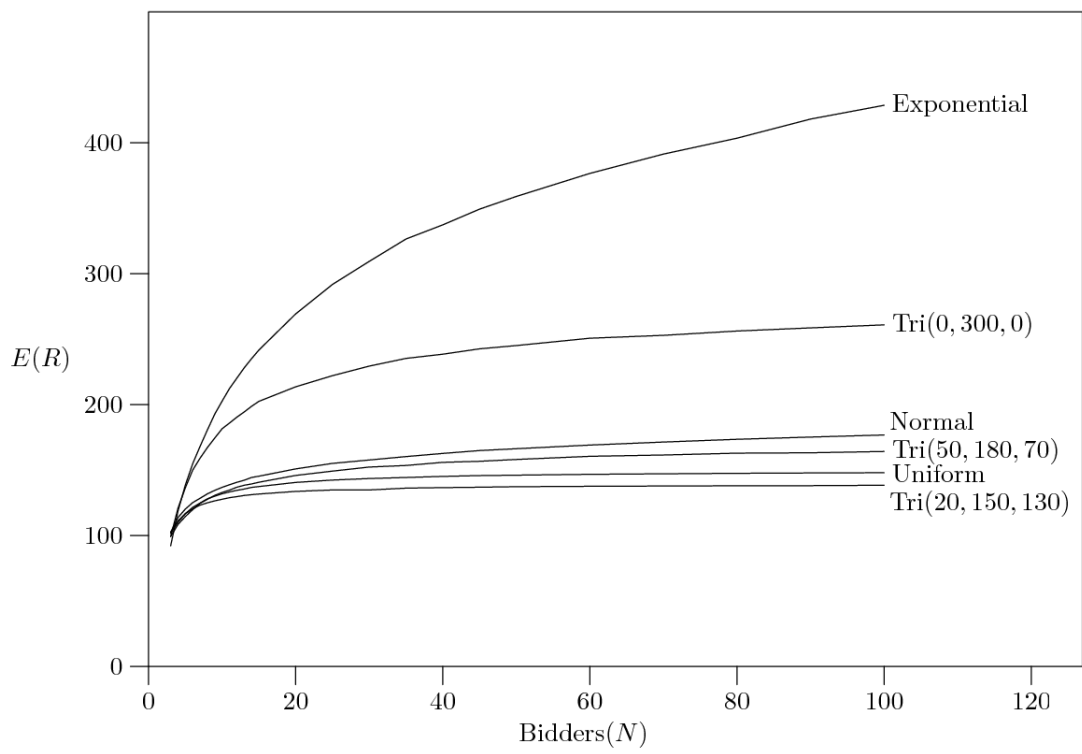


Figure 3: First-price auction simulation results




Figure 4: First-price auction (dashed) and second-price auction results




Figure 5: Comparison between IPV and CV model for normal distribution




Figure 6: Optimal bids for 4 distribution types for 5 bidders




Figure 7: Optimal bids for 4 distribution types for 20 bidders

References

- [1] Google Adwords. <http://adwords.google.nl>.
- [2] Inc. Agorics. Library: Auctions. www.agorics.com/Library/auctions.html, 1996.
- [3] Addis Equine Auctions. <http://www.addisequineauctions.com>.
- [4] L.M. Ausubel. An efficient ascending-bid auction for multiple objects. *The American Economic Review*, 94(5):1452–1475, Dec 1998.
- [5] P. Coy. The secret to google’s success: Its innovative auction system has ad revenues soaring. http://www.businessweek.com/magazine/content/06_10/b3974071.htm, 2006.
- [6] eBay. <http://www.ebay.com>.
- [7] R. Engelbrechts-Wiggans and C.M. Kahn. Multi-unit auctions with uniform prices. *Economic Theory*, 12(2), Aug 1998.
- [8] D. Goetzl. Advertisers set march launch for ebay ‘marketplace’, beta reveals details for automated tv buys. <http://publications.mediapost.com/index.cfm?fuseaction=Articles.san&s=55716&Nid=27594&p=117792>.
- [9] P. Klemperer. What really matters in auction design. *The Journal of Economic Perspectives*, 16(1):169–189, 2002.
- [10] E. Maskin. The unity of auction theory: Milgrom’s masterclass. *Journal of Economic Literature*, 42(4):1102–1115, Dec 2004.
- [11] P. Milgrom. Auctions and bidding: A primer. *The Journal of Economic Perspectives*, 3(3):3–22, 1989.
- [12] R.B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1978.
- [13] Marc S. Robinson. Oil lease auctions: Reconciling economic theory with practice. *Departments of Economics Working Paper, University of California - Los Angeles*, (292), May 1983.
- [14] K. Swenrud. Salesforce for google adwords blog. <http://blogs.salesforce.com/adwords>, 2007.
- [15] K.T. Talluri and G.J. van Ryzin. *The Theory and Practice of Revenue Management*. Springer-Verlag, New York, 2006.
- [16] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, Mar 1961.

- [17] Wikipedia. <http://en.wikipedia.org>.
- [18] R. Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1115, Sep 1985.

