Auctions for Targeted Television Advertising

by

Mark Steven Zschocke

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Management Sciences

Waterloo, Ontario, Canada, 2008

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

ABSTRACT

Television advertising is a billion-dollar industry in the United States. Currently, advertisers place their messages in television programs that are estimated to have a high proportion of their target demographic viewers. The advertising spots are usually purchased months in advance at set list prices or at negotiated prices. Technologies that can place advertisements at the cable box level, instead of the program level, will provide advertisers with the ability to target any demographic group directly and in real-time. This thesis explores the new decision-making required by this new technology and how the television advertisement space can be sold more effectively. In particular, it compares a list price system to a number of new auction models.

The structure of the auctions for the new targeted television advertising system is unique and has not been previously studied in the literature. This thesis explores new auction models that can capture these unique features and lead to desirable results for the seller of the advertisement space. A simplified analytical model shows how these features impact advertisers' bidding behavior and how a list price system compares to the auction models in the ability to raise revenue for the seller of the advertising space. These issues are then explored under various market settings with differing numbers of advertisers and value distributions that these advertisers have for the advertising space.

Since sequential first price auctions have undesirable consequences such as strong price fluctuations, this work focuses on second price auctions. The Vickrey-Clarke-Groves (VCG) mechanism is customized for this problem by developing an optimization formulation that determines the best set of advertisers for a particular advertisement space. Because execution time may be an issue, other auction models are developed that lead to similar outcomes as the VCG mechanism but require less computational effort.

This thesis provides guidance on when a list price system will lead to higher expected revenue than an auction model and vice versa in a targeted television advertising system. It also demonstrates why some of the standard auction models cannot be applied to this problem and what type of new models are required to lead to desirable advertising outcomes.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude and appreciation for my supervisor, Professor Beth Jewkes, for her support, encouragement and kind guidance. Her open-mindedness combined with constructive critiques provided an environment under which it was a pleasure for me to conduct this research.

I would also like to thank my readers, Benny Mantin and Selcuk Onay, for their valuable feedback and comments. Many thanks to the Department of Management Sciences, University of Waterloo for their support throughout my degree. I am also very grateful to NSERC for their generous funding through the Canada Graduate Scholarship Master's award.

A very special thanks to Dan Wilson from Invidi Technologies Corp. for his tireless effort to help make this thesis happen. His explanations of the new technology and our stimulating discussions of its implications were immensely valuable to me.

Thank you to my dear wife Heather for her love and care and her belief in my abilities. She makes every day wonderful just by being with me.

Many thanks to my father Jörg, my mother Marjorie, and my sisters Christina and Sandra for their love and support. Although they are far away, they are always close to my heart.

TABLE OF CONTENTS

1. INTRODUCTION........1

1.1. TELEVISION ADVERTISING........1

1.1.1. The Current Advertising System........1

1.1.2. Targeted Advertising........3

1.2. NEW TECHNOLOGY FOR TELEVISION ADVERTISING........4

1.2.1. Implementation........4

1.2.2. Pricing Decision........5

1.3. CONTRIBUTION OF THIS THESIS........7

2. LITERATURE REVIEW........8

2.1. AUCTION THEORY........8

2.1.1. Overview........8

2.1.2. Single-item Auctions........8

2.1.3. Multi-item Auctions........10

2.1.4. Private versus Common Value Auctions........13

2.1.5. Revenue Comparison........14

2.2. AUCTION IMPLEMENTATIONS ISSUES........16

2.2.1. Introduction........16

2.2.2. Avoiding Collusion........17

2.2.3. Information Sharing........18

2.2.4. Attracting a Sufficient Number of Bidders........18

2.2.5. Perception of Fairness........19

2.2.6. Training........19

2.2.7. Challenges with Vickrey-Clarke-Groves (VCG) Mechanism........20

2.3. RELEVANT APPLICATIONS........21

2.3.1. Smart Markets........21

2.3.2. Internet Advertising Auctions........21

2.3.3. The Television Advertising Market........23

2.4. TARGETED TELEVISION ADVERTISING SYSTEM........24

3. SIMPLIFIED ANALYTICAL MODEL........25

3.1. PROBLEM FRAMEWORK........25

3.2. FIRST PRICE AUCTION........27

3.2.1. Optimal Bidding Strategy........27

3.2.2. Expected Revenue for Seller........30

3.3. SECOND PRICE AUCTION........31

3.3.1. Optimal Bidding Strategy 31
3.3.2. Expected Revenue for Seller 31
3.4. LIST PRICE – OPTIMIZATION 32
3.4.1. Formulation 32
3.4.2. Setting an Optimal List Price 33
3.5. SECOND PRICE AUCTION WITH RESERVATION PRICE 34
3.5.1. Optimal Bidding Strategy 34
3.5.2. Expected Revenue for Seller 34
3.6. NUMERICAL EXAMPLE 36
3.6.1. Introduction 36
3.6.2. First Price Auction 37
3.6.3. Second Price Auction 39
3.6.4. List Price 39
3.6.5. Second Price Auction with Reservation Price 40
3.7. COMPARISONS & INSIGHTS 44
3.7.1. Introduction 44
3.7.2. Bidding Functions 45
3.7.3. Seller Revenue 50
3.8. CONCLUSIONS AND CAVEATS 54
3.8.1. Auction Model versus List Price System 54
3.8.2. Number of Types 55
3.8.3. Value Distributions 56
3.8.4. Additional Complexities 57

4. SMART MARKET AUCTION DESIGN 58

4.1. INTRODUCTION 58
4.1.1. Series of Auctions 58
4.1.2. Link between Impressions and Flotilla Slot 58
4.1.3. Overlapping Types 59
4.1.4. Motivating Example 59
4.2. PROPOSED SMART MARKET AUCTIONS 67
4.2.1. Introduction and Notation 67
4.2.2. Second Largest Ad 68
4.2.3. Reimburse 69
4.2.4. Reimburse Revised 71
4.2.5. Vickrey-Clarke-Groves Mechanism 72
4.2.6. Reservation Prices 75
4.3. LIST PRICE ALTERNATIVE 75

4.4. COMPUTATIONAL ANALYSIS 77
4.4.1. Introduction and Parameters 77
4.4.2. Base Case Scenario 78
4.4.3. Number of Advertisers 80
4.4.4. Setting the Correct List Price 81
4.4.5. Setting a Reservation Price 82
4.4.6. Value Distribution Functions 83
4.5. LIMITATIONS 84
4.5.1. Large-scale Optimization 84
4.5.2. Main Assumptions 85
4.6. CONCLUSIONS 86
4.6.1. Algorithm Reimburse vs. VCG Mechanism 86
4.6.2. List Price System vs. VCG Mechanism 86

5. CONCLUSIONS AND FUTURE RESEARCH 88

5.1. CONCLUSIONS 88
5.1.1. List Price versus Auction Model 88
5.1.2. Type of Auction Model 88
5.2. FUTURE RESEARCH 89

APPENDIX 90

APPENDIX A – EXPECTED VALUES 90
APPENDIX B – INTEGRATION BY PARTS 90
APPENDIX C – AUCTION VS LIST PRICE BY VALUE DISTRIBUTION 91

REFERENCES 92

List of Tables

Table 4.1: Example 1 – Full Information Matrix. 60
Table 4.2: Example 1 – Condensed matrix with updated bids. 61
Table 4.3: Example 1 – Revised full information matrix. 63
Table 4.4: Example 2 – Full information matrix. 66
Table 4.5: Example 2 – Full information matrix with shill bid. 66
Table 4.6: Sample data for demographic composition of viewers. 77
Table 4.7: Sample data for advertiser composition in terms of target demographics....................78

List of Figures

Figure 1.1: Timeline of new targeted television advertising system. 4
Figure 1.2: Sample two-slot Flotilla. 5
Figure 3.1: Flotilla construction example assuming 3 bidders with 2 types......................................26
Figure 3.2: Expected revenue of auction vs. list price dependent on the set list price. 40
Figure 3.3: Optimal first price bids with and without types dependent on value (Λ=1). 47
Figure 3.4: Optimal first price bids given four different value distributions. 49
Figure 3.5: Expected revenue of auction vs. list price (Λ=1). 52
Figure 3.6: Expected revenue of second price auction for three value distributions. 53
Figure 3.7: Expected revenue of second price auction for three normal value distributions. 54
Figure 4.1: Expected revenue of auctions and list price in base case. 79
Figure 4.2: Expected revenue of auctions and list price by number of advertisers. 80
Figure 4.3: Expected revenue of list price dependent on list price set. 81
Figure 4.4: Expected revenue of auctions with reservation price and list price in base case. 82
Figure 4.5: Expected revenue of VCG auction with reservation price. 83
Figure 4.6: Expected revenue of auctions depending on value distributions. 84
Figure 6.1: Second price auction vs. list price assuming [0.5,1.5]-Uniform distribution. 91
Figure 6.2: Second price auction vs. list price assuming [1,0.1]-Normal distribution.........................91

1. INTRODUCTION

1.1. TELEVISION ADVERTISING

An estimated $153.7 billion was spent on advertising in the U.S. in 2007 and over 40% of this amount was spent on television advertising (TNS Media Intelligence, 2007). Television is a very attractive advertising medium because it can use both sight and sound for dynamic selling that is very close in nature to personal selling. Furthermore, it allows the advertiser to target both selective and mass markets by placing advertisements in shows that are viewed by the desired demographic group. In addition, television advertising can be cost-efficient by selecting the appropriate time of day to broadcast the advertisements (Sissors and Baron, 2002).

Advertisers who are planning a television advertising campaign want to know how many people will see the ad (defined as *reach*), how often they will see the ad (defined as *frequency*), and the total number of times that the ad will be seen (defined as *gross impressions*). The information required to make these estimations is provided by Nielsen Media Research (http://www.nielsenmedia.com) who track a sample audience in order to estimate broadcast ratings. Broadcast ratings "represent an estimate of the audience that has viewed a program or has tuned in during a specific time period" (Sissors and Baron, 2002). More specifically, one rating point is equal to one percent of the audience. To calculate *gross rating points* (GRP), the broadcast ratings of each time the advertisement was shown are added together. By multiplying the GRPs with the target audience base, gross impressions, the actual size of the audiences, can be calculated. Even though both GRP and gross impressions contain duplicate audiences, they are the most commonly used measures for purchase and pricing decisions for advertisers.

1.1.1. The Current Advertising System

Television advertising space is typically bought in three ways (Katz 1995). In an *up-front buy* (or *long-term buy*) advertising space is bought for the entire upcoming broadcast year. This offers the advertiser insurance against sellouts, in which all available advertising spots are sold, but provides poor cancellation options. In *scatter buys* (or *short-term buys*) advertising space is bought for the coming quarter, which results in greater flexibility but comes with the risks that the best inventory (advertising spots) may already have been sold. Finally, *opportunistic buys* are last-minute purchases of available advertising space which usually go for a low price but have a high risk of sellout. As a result, the advertiser may not be able to broadcast his messages in his desired television shows.

When making purchasing decisions, advertisers typically try to reach their advertising target while minimizing the *cost per thousand* (CPM). This is the cost to achieve 1,000 gross impressions (Sissors and Baron, 2002). However, other measures also influence the price that advertisers are willing to pay. For example, advertisements (ads) placed in *prime time* (between 8 and 11 p.m.) will result in a larger reach than ads placed at other times because each time the ad plays, it is seen by many distinct viewers at once. More reach is a desirable trait and is one of the reasons that prime time advertising is so expensive. Advertisers must take all this information into account when trying to reach the intended audience size of their target demographic, with the desired frequency, and for a reasonable cost. Furthermore, they must reach these goals by choosing specific *programs* (television shows) that will carry their advertising.

The cost of advertising in these programs depends on three factors (Sissors and Baron, 2002). The first factor, advertisers' demand for particular programs, has the biggest effect on the final advertising prices. Secondly, the price is influenced by the buyers' and sellers' estimates of audience sizes for particular programs. Lastly, expensive productions (such as the "Super Bowl") are priced at a premium as networks try to recover the cost of programming and business overhead. The final advertising prices are determined through an elaborate negotiation process in which buyers and sellers must agree on when and where the advertisements will be placed.

In spite of its many benefits, television advertising has some limitations such as low viewer attention during advertisements and clutter resulting from the large number of advertisements broadcasted within any given television hour. One of the largest downsides, however, is the lack of *targeted advertising* possibilities. An advertiser who is trying to reach a particular demographic cannot do this directly but must instead choose programs that will likely have a significant proportion of his target demographic. A brief example will more clearly demonstrate this inefficiency.

Imagine a particular program, such as "The Amazing Race", which in this hypothetical example is seen by a viewer base consisting of 70% female and 30% male. An advertiser who is targeting exclusively females may choose to purchase some advertising space in this particular show and reach the large female audience that he is seeking; however, if the ad is targeted at females only, it is "wasted" on the smaller – yet substantial – number of male viewers. A technology that allows advertisers to truly target demographic groups, rather than to target programs with merely a high proportion of the target demographic group, would be of great value to advertisers.

1.1.2. Targeted Advertising

By inserting advertisements into the digital cable box of individual households, rather than the general program feed, Invidi Technologies Corporation (http://www.invidi.com) has developed the technology to allow targeted television advertising. Their advanced technology is capable of determining the demographic viewer composition of each household with great accuracy using previous television viewing patterns. Their comprehensive privacy framework does not allow Invidi to access this information; instead, a computer program will automatically, and in real time, indicate appropriate ads for the respective households.

The implementation of this new targeted television advertising technology raises many questions and requires new decision-making. The focus of this thesis is how to price advertisements when this technology is available to advertisers. The advertisement space may be sold by setting individual list prices for each demographic or by implementing an auction.

While this is the first implementation of such a system on television, targeted advertising is well established in other fields. One of the leaders in internet advertising, Google, uses search terms entered by the users to determine which ads to place next to the search results. For example, if a user searches for the word "camera" then the advertisements placed around the search results will be for camera stores, camera magazines and other relevant websites.

To sell their advertising space, Google has implemented an online auction in which advertisers submit bids for the available advertisement space. These bids are then used in combination with the search terms entered by the user to determine which advertisements are shown and how much they are charged (Edelman, Ostrovsky and Schwarz, 2007). Even though there are significant similarities between the established field of targeted internet advertising and the new targeted television advertising (as will be examined in more detail later), some critical differences exist. Consequently, a potential auction implementation for the new targeted television advertising field requires new auction models.

1.2. NEW TECHNOLOGY FOR TELEVISION ADVERTISING

1.2.1. Implementation

A technology that can target viewers at the cable box level, instead of the program level, removes many of the inefficiencies of the current television advertising market. Using the previous example, the female viewers of "The Amazing Race" could see a different ad than their male counterparts (at the exact same time) as long as they are watching the show on separate digital cable boxes. Furthermore, viewers of cable TV channels with very low ratings, who currently cannot be captured due to the large uncertainty regarding the audience size and demographic composition, can now be captured just as easily as viewers of the top rated programs.

There are many aspects of targeted television advertising that are new and unique; however, this thesis will focus specifically on the pricing decision: What is the most effective way of selling available television advertising space? The question will be analyzed under the framework of a current implementation plan of targeted television advertising by Invidi, hereafter referred to as *the seller*. Without specifying technical details, the implementation will operate on the timeline displayed in Figure 1.1.

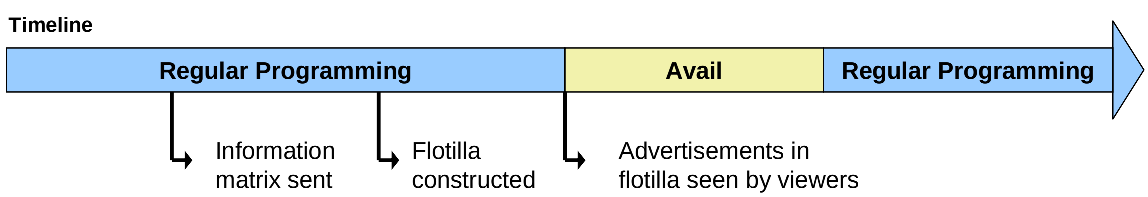

Figure 1.1: Timeline of new targeted television advertising system.

A few times each hour, advertising timeslots become available (defined as *avails*) for sale to advertisers. Immediately preceding these avails, the seller will receive an information matrix that lists, for a sample of the viewing households (or viewers), all the appropriate ads that are currently in the seller's system. Even though this is an automated procedure without human involvement, this process could be described as the households "voting" for which advertisements appeal to them.

Due to technological constraints, the seller can only make a small number of advertisements available to the households simultaneously. More specifically, the seller must construct a *flotilla,* which is a block of advertising space for the length of the avail with as many slots as the number of advertisements that can be shown simultaneously on different digital cable boxes. For example, Figure 1.2 shows a flotilla that lasts over the full length of an avail and consists of two slots. In this case, the top slot was filled with Advertisement X and the bottom slot with Advertisement Y. Once the avail occurs, viewers will see either one of these two advertisements.

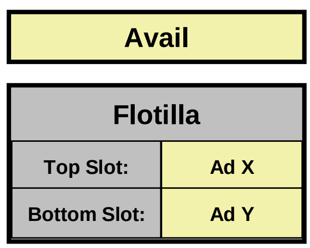

Figure 1.2: Sample two-slot Flotilla.

The challenge is to find a mechanism that determines which advertisements receive flotilla slots and how much the advertisers should be charged for this advertising space.

1.2.2. Pricing Decision

The seller could operate under a pricing system similar to the current one, in which its clients can buy impressions based on a predetermined or negotiated price. In this scenario, there are two main challenges: First, what are the optimal prices they should set in order to maximize revenue? Since the new technology is far more accurate than the regular markets in reaching the desired demographic group through its advertising decision system, advertisers should be willing to pay a premium price for this service; however, the amount of this premium remains to be determined.

The second challenge lies in the fact that fixed prices imply a promise of service. Once an advertiser has bought a certain number of impressions, the seller must provide these impressions. This obligation can become a problem if, for example, advertisers are trying to reach a large number of the same demographic group within a relatively small time window. If the number of actual viewers is not high enough to satisfy the advertisers' demand, the seller would have to find some way of compensating advertisers that did not receive their desired number of impressions. To avoid this scenario, the seller could stop selling advertising spots targeting this particular demographic group as soon as sales have reached the forecasted level of viewers for that demographic. However, even with this approach the impression obligation may not be met if viewer forecasts do not match actual viewer numbers.

As an alternative to a list price system, the seller could run an auction in which advertisers specify not only the desired quantity of impressions for a particular demographic group but also the price that they are willing to pay. Their success in achieving the desired number of impressions and the actual price they will pay per impression is then determined by the auction. Under this system, the seller is no longer responsible for ensuring that advertisers receive their desired number of impressions because the advertisers' success in reaching their impression goal is solely a function of their bids. In other words, unsatisfied bidders must simply raise their bids if they want to reach more viewers. Of course, this statement only remains true as long as the advertiser does not demand an audience size that is larger than the complete television advertising market that is under the seller's control.

Another advantage of an auction based system is the fact that it is more efficient in allocating impressions to the advertisers that value them the most (see literature review below). Auctions are not without drawbacks, however. Advertisers that are used to the traditional system may be uncomfortable with this new way of purchasing impressions. An auction burdens the advertisers with the additional challenge of determining an appropriate bid and it creates uncertainty with regards to the number of viewers that will be reached. Given that advertisers may have a fixed target number of impressions they need to reach, some may not be willing to participate in an auction.

1.3. Contribution of This Thesis

In this thesis, new auction models are presented which can handle the complexity of a new targeted television advertising system. The new technology contains some unique features which have not been previously studied in the literature. Through a simplified analytical model, insights are gained into the implications of these features on bidding behavior and revenue generation using the standard auction models. Furthermore, new auction models are presented and analyzed with respect to their effectiveness in raising revenue for the seller.

Within the context of the discussed implementation of the new technology, this thesis answers the following questions:

- How does the market structure (in terms of number of advertisers and their values for ads) influence the optimal list prices that the seller should set for their impressions?
- What should the rules of an auction model be if it were to be put in place instead of a list price system?
- How does the revenue generated by the proposed auction models compare to the list price system and how does the comparison depend on the market structure?

Following this introductory section, Chapter 2 presents a literature review of general auction theory, specific implementation issues, and relevant real world applications of auctions. The simplified analytical model of Chapter 3 shows the effect of the unique features of the new technology on the standard auction models and compares the resulting revenue to that of a list price system. Chapter 4 relaxes the assumptions and uses Monte Carlo simulation to derive results for the full model. Finally, extensions and future research are discussed in Chapter 5.

2. LITERATURE REVIEW

2.1. AUCTION THEORY

2.1.1. Overview

Auctions have been used to transfer goods from sellers to buyers since as early as 500 B.C. when they were used in Babylon (Krishna, 2000). Surprisingly, it took until the 1960s until auctions were first studied in the research literature. It was William Vickrey with his paper on "Counterspeculation, Auctions, and Competitive Sealed Tenders" in 1961 who started what is now a very active field of research. In fact, the topic has become so broad that it is investigated by researchers in many different disciplines, including economics, operations research / management science, sociology, and computer science (Rothkopf and Park, 2001). According to Rothkopf and Park, each discipline has had its own focus, with economics using simplifying assumptions to solve game theoretic models mathematically, operations research paying closer attention to details by loosening the constraint on strict optimality, sociology investigating the effect of repeated auctions on participant behaviour, and computer science designing automated auction agents and using auction concepts in general application designs. This thesis will review work from all of these fields with an emphasis on papers from the operations research discipline.

2.1.2. Single-item Auctions

There are four distinct auction forms that are commonly used to sell a single-item. Each of these auction types are identified by Vickrey (1961) and can be found in real-life applications.

English Auction – The English auction is frequently used in the sale of artwork and other collectibles, and is the type of auction commonly featured in Hollywood movies. During such an auction, bids are freely made or called out by the auctioneer (which also gives it the name "open ascending price auction") until no bidder is willing to increase the last announced bid. At this point, the item is sold to the last (and highest) bidder for the amount of his bid. The format of this auction promotes *truth-telling*; i.e., bidders will reveal their true valuation of the item (Vickrey, 1961).

A simple example will illustrate the concept of truth-telling: imagine a bidder A, who values an item at $100, and who has just observed bidder B announce a bid of $90; clearly, bidder A should continue bidding with an amount above $90 (but below $100) because he will else forgo the possibility of securing the item for a price that is below his value. However, if bidder B announces a bid of $100 or higher, bidder A should no longer participate in the auction because he would have to pay an amount that is beyond his value of the item. In effect, all bidders announce increasing bids up to their personal value and the bidder with the highest value receives the item for the last price given by the bidder with the second highest value (which will be equal to that second highest value).

Dutch Auction – Another, perhaps less well known, auction form is the Dutch auction. It is used extensively in the sale of flowers in the Netherlands (McAfee and McMillan, 1987). The auctioneer begins with a very high price and then continuously decreases this price until the first bidder accepts, receiving the item for the last announced price. From a bidding strategy perspective, this auction is far more challenging than the English auction. Truth-telling is *not* the optimal strategy because any bidder who accepts the auctioneer's price as soon as it reaches the bidder's value will always make a profit of zero (where profit is his value for the item minus the price paid). In comparison, bidders who allow the price to decrease to an amount below their value for the item create profit opportunities for themselves. The lower they let the price fall before they bid, the higher their potential profit; however, the further they allow the price to fall, the higher the probability that another bidder will capture the item before them. Balancing this trade-off, bidders will engage in *bid shading*, which means that they only accept prices that are below their respective values. The degree of bid shading depends on a number of factors and will be explored more closely in a later section.

Sealed Bid First Price Auction – In a sealed bid first price auction, bidders cannot observe each other's bids and can only submit one bid. The auctioneer observes all bids and sells the item to the highest bidder for the amount of his bid. This auction form is often used when large contracts are tendered, and most people will engage in this type of auction when purchasing a new home. Bidders face the same challenges as in the Dutch auction because the lower they bid, the higher their potential profit, but the higher also the chance of being overbid.

Sealed Bid Second Price Auction – This auction form is identical to the sealed bid first price auction except that the winner must only pay the second highest bid. This slight change induces truth-telling behaviour: the price that the winning bidder will pay for the item is *independent* of his own bid; instead, the price is the second highest bid. His own bid will only affect his probability of winning the item.

A small example demonstrates this clearly: bidder A values the item at $100 and only bids $90 while bidder B bids $95, leaving bidder A without the item and a profit of zero. Had bidder A bid his true value, he would have received the item for $95, making a profit of $5. Had bidder B only offered $80, bidder A would have made a profit of $20 regardless of whether he had bid $100 or $90. Clearly, bidding $100 – or his true value – is the superior choice.

While these are the four most commonly used auctions, this is certainly not an exclusive list of single-item auctions. Extensions are the inclusion of entry fees and reservation prices, which are set by the seller. An entry fee must be paid by all bidders regardless of whether they are successful in the auction and a reservation price is a minimum amount that the winning bidder has to pay. Further examples of alternative auctions include the *all-pay auction*, in which each bidder must pay his bid regardless of whether he wins the object or not. In the *third-price auction*, the winning bidder must only pay the third highest bid. Riley and Samuelson (1981) describe an even more eccentric two player auction, which they termed *sad loser auction*: both bidders pay an entry fee, the higher bidder receives the item for free, but the losing bidder must pay his bid. They show that even this auction model has an equilibrium bidding strategy. Many of these unusual auctions are only theoretical constructs and are rarely used in practice.

2.1.3. Multi-item Auctions

While single-item auctions have been studied extensively and are well understood, most of today's auction literature focuses on multi-item auctions. Multi-item auctions are significantly more complex than single-item auctions and fall into different categories. In the simplest setting, identical items are sold to bidders that each only desire one item. The auction becomes more complicated if bidders want to acquire more than one of these identical items, and matters are significantly complicated if these bidders experience increasing or decreasing utility for particular *packages* of items or if the items are no longer identical.

The auction models that are available to the seller are endless: he may choose to sell the items one by one in sequential auctions or simultaneously in one auction; furthermore, there are a vast number of possible auction rules that he may employ to sell his items. Following is a list of some of these auction models to give the reader a flavour of the variety.

Identical Items – In a setting with bidders desiring multiple items from a total of N identical items, Krishna (2002) describes the following sealed bid auction models: in a *discriminatory* auction, the top N bids win an item and must pay the amount of their bid; bidders that submit multiple bids of different amounts will end up paying these differing amounts for each of the items. In *a uniform price* auction, all N items are sold at the price at which the total amount demanded equals the total amount supplied. And finally in a *Vickrey* auction, bidders must pay the amount required to secure that particular item; i.e., for the first item the bidder must pay the N^{th} highest competing bid, for his 2nd item he must pay the N-1th highest bid, etc.

Krishna also lists the corresponding open auction models: in the *multiunit Dutch* auction, bidders buy items one by one – as the auctioneer keeps lowering the price – until all items are sold (the outcome is equivalent to the discriminatory auction). In a *multiunit English* auction, bidders indicate the number of items they are willing to purchase at the current price. The auctioneer keeps increasing the price until the number of demanded items equals the supply of items (the outcome is equivalent to the uniform price auction). Lastly in the *Ausubel* auction, the auctioneer keeps increasing the price and the bidders purchase items at the price at which the residual supply (items that no other bidders are willing to buy at the current price) becomes available (outcome equivalent to the Vickrey auction).

Non-identical items – In an auction with non-identical items, bidders must be able to specify the price that they are willing to pay for certain groups, or packages, of items. As bidders make offers for overlapping packages, a second problem besides determining the price that each bidder should pay emerges. Which bidders have won which items? This is known as the *winner determination problem*. For example, if bidder A bids $6 for item X and bidder B makes individual $5 bids for item X and Y, as well as a combined bid of $12 for both items, item X may go to bidder B even if bidder A values the item more individually. However, if another bidder places a bid on item Y that exceeds \$6, then bidder A will receive item X. This is the realm of *combinatorial auctions*.

The paper by Rassenti et al (1982) was one of the first papers on the topic of combinatorial auctions. It was set in the context of developing a mechanism to allocate arrival and departure airport time slots. In a more recent paper, Pekec and Rothkopf (2003) list the current strategies that are used to cope with the mathematically hard winner determination problem: the *algorithmic approach* uses either integer programming techniques, which work particularly well for smaller problems and guarantee optimality, or approximation algorithms, which can handle larger problems and are usually close to optimality. Another solution is to *relegate computational complexity* by forcing bidders to indicate other complementary bids that make their offer part of the optimal set. This does not remove the complexity of the winner determination problem but merely transfers it to the bidder. Therefore, bidders with superior computational technology available to them are favoured by this method. A clever solution to *maintaining fairness in the face of computational limits* is to provide bidders with an opportunity to provide a different set of winners if their suggested set results in higher revenue for the seller. In the context of auctions, fairness is achieved when all bidders are treated equally by the seller and when sub-optimal solutions, which favour some bidders, are rejected. Alternatively, complexity reduction can be achieved by *limiting biddable combinations* which makes the algorithmic approach feasible for larger problem sets.

In addition to the solution to the winner determination problem, rules regarding the payment structure must be put in place. Such rules could specify a first price system, in which each bidder must pay his bid if he wins his desired item(s). As an alternative solution, Clarke (1971) and Groves (1973) generalized the Vickrey auction to a more general solution named the Vickrey-Clarke-Groves (VCG) mechanism. The mechanism is famous for its dominant truth-telling strategy and perfect efficiency – even in combinatorial auctions (Rothkopf, 2007). Efficiency is achieved when the total value of the winners is maximized (i.e. when the bidders who value the items the most receive the items). The VCG mechanism achieves these properties by optimizing the revenue to the seller based on the existing bids but reimbursing each successful bidder by the amount of additional revenue that the particular bidder has added to the system (i.e., the decrease in revenue that would result if that bidder had not submitted his bid). The precise way in which the VCG mechanism determines payments of bidders will be demonstrated in Chapter 4.

A further discussion on combinatorial auction mechanisms is outside of the scope of this thesis but Pekec and Rothkopf's paper provides a good starting point for the interested reader. For a full review, Crampton, Shoham and Steinberg (2006) have published an extensive book on combinatorial auctions.

Given that the auctions for targeted television advertisements will offer multiple flotilla slots simultaneously, this system clearly falls into the realm of multi-item auctions for non-identical objects (different demographic groups are worth different amounts to advertisers). Advertisers may bid for particular groups and combination of groups which – superficially – gives it the appearance of a typical combinatorial auction. However, as will be discussed later, the unique way in which impressions (the actual desired item) are linked to flotilla slots (the means to receive the desired item) gives this system a structure very distinct from combinatorial auctions.

2.1.4. Private versus Common Value Auctions

Aside from the number of items for sale, another important characteristic of an auction is whether the items have private or common values. Vickrey (1961) worked under the private value model, where each bidder has his own independent value for the item. This assumption is usually valid for items such as artwork and collectibles in which individual preferences determine the personal value.

In contrast, some items have a common value to all bidders but there is uncertainty regarding the amount of this value. For example, in an auction for oil extraction rights, the value of the oil will be the same for all bidders (world market price); however, the exact size of the oilfield is unknown at the time of the auction. Milgrom and Weber (1982) describe how this can lead to the *winner's curse*: if the value for the item is not only common – but also deterministic – then all bidders would place exactly the same bid for the precise value of the item. But since the value is stochastic, the bids will differ according to the individual estimates of the value. The bidder with the highest estimate will place the highest bid and win the item; however, the bidder with the highest estimate out of a group of bidders will likely have overestimated the value of the item. As a result, he will overpay. This is the winner's curse. Anticipating the winner's curse, savvy bidders will bid below their estimated value (shading) by an increasing amount the more uncertain the value of the item. Consequently, the seller should provide as much information as possible about the item to avoid excessive bid shading.

In this thesis, the commodity of the seller is advertising space. Since organizations value advertisements differently, as demonstrated by large variations in advertising efforts, and cannot resell the impressions for some type of market price, this system will be analyzed under a private value assumption. Therefore, the winner's curse will not be an issue in the choice of auction model for this new targeted television advertising system.

2.1.5. Revenue Comparison

The most important feature of any auction model to the seller is their expected revenue. However, the most profitable auction model depends on the characteristics of the bidders and the auction structure. McAfee and McMillan (1987) defined the *benchmark model* by the following assumptions:

1. All bidders are risk neutral.
2. All bidders have their own private and independent value for the item.
3. The bidders are symmetric, drawing their values from the same common distribution.
4. Payment is dictated by the bids alone (no other fees).

Based on these four assumptions, the four standard auctions described in Section 2.1.2 result in the same revenue. This is known as the *revenue equivalence theorem*. This may surprise the reader since the only difference between the sealed bid first price and the sealed bid second price auction is whether the winner pays the highest or second highest bid. The reason that both models will still result in the same revenue is that bidders will bid lower in the sealed bid first price auction. While the benchmark model provides a striking result, its assumptions are often violated in real world applications. Therefore, the effect on revenue will be examined as some of the benchmark assumptions are relaxed.

Krishna (2002) illustrates that if bidders are risk averse, first price auctions will lead to higher revenue than second price auctions. The reasoning is that, while risk aversion does not affect bidding behavior in the second price auction, it does influence the first price bidding behavior. As described earlier, bidders in a first price auction will engage in bid shading by balancing their probability of winning (by bidding high) with their expected profit (by bidding low). As their risk aversion grows, their willingness to gamble decreases and they would rather make a small profit with a high probability than a large profit with a low probability. Therefore, their bid shading decreases, resulting in higher bids and making first price auctions more profitable than second price auctions.

If the bidders' values for the item are common, instead of private, the highest revenue will be generated by the English auction, followed by the sealed bid second price auction, and the lowest revenue will come from the Dutch and the sealed bid first price auction (Milgrom and Weber, 1982). As mentioned before, uncertainty about the value of the item increases the fear of the winner's curse in a common value auction. Information gained from the seller before the auction and information gained from other bidders during the auction alleviates this fear. In the English auction, bidders can observe each other's bids and gain information regarding how other bidders value the item. Even though bidders cannot obtain the same information in a sealed bid second price auction, they can at least rest assured that, even if they grossly overestimated the value, they will only have to pay the second highest bid. In contrast, bidders have neither advantage in the Dutch or the sealed bid first price auction and will therefore decrease their bids to a greater extent.

Relaxing the constraint on bidder symmetry does not lead to a clear preference of one auction model; instead, which auction model will generate the highest revenue depends on the particular value distributions of the bidders (Krishna, 2002). With regards to bidder payments, Menezes and Monteiro (2005) show that while an entry fee has no positive impact on revenue, a reservation price can further increase the expected revenue for the seller. Intuitively, setting a reservation price is comparable to a monopolist setting the price above the marginal cost.

Not only is the comparison *between* auctions interesting but a seller may contemplate his expected revenue of an auction compared to a *list price system*. Wang (1993) states that with certain cost parameters auctions will outperform list price systems. Going a step further, Caldentey and Vulcano (2007) explore the challenges of managing a dual channel system. In their model, buyers can either purchase the item immediately for the list price or opt to participate in a terminal auction over time *T*. Their results show that the auction is superior only if the seller has a sufficient quantity of items and the discount factor (of the value of items over time) is relatively small.

The immense size of the auction theory literature does not allow for a complete discussion of these topics. The interested reader can find a more complete discussion in the excellent literature guide to auctions by Klemperer (1999).

With regards to the new targeted television advertising system, an interesting question is the circumstances under which the revenue equivalence theorem holds or whether one auction model theoretically dominates the others in terms of revenue. Furthermore, as stated in Chapter 1, the auction models will be compared to a list price system for the purpose of helping the seller decide whether they should use an auction. Since this thesis involves a real world application, we now turn to some practical auction implementation issues.

2.2. AUCTION IMPLEMENTATIONS ISSUES

2.2.1. Introduction

There have been a number of high profile cases in which auctions were used to sell single or multiple items. One such case made it into the Wall Street Journal: in 1994, the Federal Communications Commission (FCC) used a series of auctions to sell radio spectra for personal communications services. The auction was rated as a great success generating \$617 million in revenue for ten licenses (Cramton, 1995). At the heart of the auction mechanism was an ascending bid auction, but a team of economists had worked out detailed rules to ensure the success of this high stake venture. As described by Cramton, strict payment rules were put into effect (such as upfront payments), minimum bid increments were specified, bidders had to maintain an active bidding history to remain in the auction, and a detailed set of rules was put into place to specify when the auction would end. All these rules were aimed at avoiding some of the pitfalls of earlier auction implementations and at creating an environment in which a successful auction could be completed.

Pekec and Rothkopf (2003) summarized the desirable properties of auction mechanisms: *allocative efficiency* is achieved by maximizing the total value to the winners. From the seller's perspective, *revenue maximization* is a very important property of any auction mechanism; and, both seller and buyer can benefit from *low transaction costs* by avoiding entry fees and limiting the extent of time and effort necessary to calculate and submit an appropriate bid. In addition, *fairness* is required to ensure bidder participation and the auction should be *failure free*, avoiding possible disastrous outcomes of extremely low revenue which may result from lack of participation both in terms of the quantity and the size of bids. The following sections will discuss some of the main implementation issues and how they play a role in achieving the desirable auction properties suggested by Pekec and Rothkopf.

2.2.2. Avoiding Collusion

One of the most costly mistakes that a seller can make is to implement an auction mechanism that makes it easy for bidders to collude. Klemperer (2002) lists collusion as the "first major set of concerns for practical auction design" (p. 170). Some auction mechanisms are more vulnerable to collusion than others.

To illustrate this point, imagine a situation with two bidders competing for one item and agreeing that bidder A should receive the item for a very low price, x, and compensate bidder B with a side payment y. In a sealed bid second price auction, they can quite easily collude. Bidder A can submit a high bid (close to his true value) while bidder B submits a bid equal to x, which would win the item for bidder A at price x. Should bidder B try to cheat bidder A by bidding an amount slightly above x, he will not succeed since bidder A submitted a relatively large bid; therefore, the collusion will likely be successful. In contrast, in a sealed bid first price auction, the collusion could only be successful if bidder A submits the low bid x and bidder B submits an even lower bid; however, bidder B now has a strong incentive to bid just above x and win the item for himself at a very low price.

Klemperer lists some high profile cases in which participants colluded successfully in an ascending auction and notes that "a frequently repeated auction market […] is particularly vulnerable to collusion" (p. 172). Since the targeted television advertising system will feature frequently repeated auctions, collusion is an issue that needs to be addressed. Even though collusion is a larger problem in an auction of very expensive items (such as radio spectra, oil fields, etc.) than in less expensive items (such as viewer impressions), the proposed auction should be constructed in a way that does not permit systematic collusion between the main advertisers.

2.2.3. Information Sharing

The extent of information sharing that will be present in the auction is another aspect that will have an effect on collusive behaviour. As noted by Gallien and Wein (2005), in a multi-round auction, full disclosure of previous bids can lead to collusion and strategic behaviour. For example, using peculiar numbers as bids can be a method of signaling competitors (Klemperer, 2002) and collusive agreements can more easily be enforced by the participants if they have access to full information. However, in a multiunit auction, a complete lack of information can lead to inefficiencies in the price formation process as bidders are unsure which bids to increase or decrease (Gallien and Wein, 2005).

Another aspect that needs to be considered with respect to information sharing is how potential bidders may feel about the revelation of their bids because this may lead to other bidders gaining competitive information. In an experiment conducted by Millet et al. (2004) through an online procurement auction, they found that the revelation of bidding information leads to lower acceptance rates (percent of invited suppliers that accept the invitation and log into the auction event). In an even more troubling finding, they note that auctions with high information sharing may attract suppliers looking for market intelligence and may prevent bidders who are concerned about confidentiality from participating in the auction. This can lead to fewer serious bids and can increase the extent of bid shading – both can depress revenues substantially.

In the new targeted television advertising system, advertisers will expect some type of information sharing about their current advertising campaign so that they can increase or decrease bids as necessary to achieve their goals. The proposed system will need to balance this information need with the privacy concerns of advertisers.

2.2.4. Attracting a Sufficient Number of Bidders

Information sharing is not the only aspect of an auction mechanism that can reduce the number of bidders. Ascending auctions can also lead to fewer bidders because everyone expects the bidder with the highest value for the item to win (Klemperer, 2002). In particular with auctions for expensive items – which make the process of developing an appropriate bid, in itself, expensive – smaller entities may be discouraged from participating in the auction if they feel that the large entities will always have the chance to overbid them in the end. In comparison, a first price auction allows small entities to bid aggressively in hopes of outbidding a large entity that shaded their bid too heavily.

An interesting finding that came from the experiment by Millet et al. (2004) was that acceptance rates decline as more suppliers were invited. Auction participants are well aware that their chance of success declines as the pool of participants increases. Therefore, the seller must choose a fine line between encouraging extensive participation while making each potential bidder feel that they have an opportunity to make a profit.

The seller can accommodate by segmenting the advertising market. While there may be significant competition for nationwide, prime time flotilla slots, which may deter smaller advertisers from participating in the auction, other local and off-peak hours flotilla slots will attract less fierce competition and will thus be more attractive to smaller advertisers. In such a way, a wide participation of advertisers in the new targeted television advertising system can be achieved.

2.2.5. Perception of Fairness

Not only must the seller assure the bidders that their participation is worthwhile, but the seller must also ensure that the process and the outcome of the auction are perceived as fair. In their discussion on the procurement auctions of Mars Inc., Hohner et al. (2003) point out that "it is very important that the bidders perceive Mars' auctions to be fair" (p. 30) to ensure that their suppliers do not refuse to participate in the auction. Fairness can be achieved by having explicit – and credible – rules, as well as a transparent process. For example, for a sealed bid second price auction to work, bidders must be assured that the seller will not make a fake bid to increase the selling price. However, as mentioned earlier, information sharing has some negative effects and a trade-off between protecting bidder privacy and a transparent process must be reached.

As a new technology in the market place, the seller will have to build a reputation as a credible auction broker with clear rules that are enforced equally on all players.

2.2.6. Training

Another important aspect of auction implementation is the extent of training offered by the seller (Hohner et al., 2003 and Millet et al., 2004). Information sessions, or even mock auctions, may make potential bidders feel more confident about participating in the auction and can therefore lead to a larger pool of bidders.

Given the extensive number of advertisers that will potentially participate in the seller's auction, information sessions are not a feasible option. However, clear instructions on how the auction system works and how advertisers can best achieve their goals should be easily accessible to all potential clients.

2.2.7. Challenges with Vickrey-Clarke-Groves (VCG) Mechanism

As previously discussed, the VCG mechanism has some very attractive theoretical attributes, such as allocative efficiency and enticement of truth-telling, but it has some significant implementation issues that need to be discussed. Rothkopf dedicated a whole article, "Thirteen Reasons Why the Vickrey-Clarke-Groves Process Is Not Practical" (2007), to discuss the serious drawbacks of this method. He points out that truth-telling bidding behaviour results in *weak equilibria* and that in a case where bidders draw their values from different value distributions, bidders who are sure to lose may not participate. In other words, bidders who are convinced that others value the item(s) higher will not participate thus leading to lower revenues. Other concerns are the *bid preparation and bid communication costs* which will occur as bidders try to estimate their value for the item. In the general VCG process for n items, bidders will have to submit $2^n - 1$ bids (for each combination of items) which implies that they need to estimate their value for each of those bids. This large number of bids also makes the *winner determination effort* more demanding. Another practical concern, which is not addressed sufficiently by theoretical work, is the effect of *budget constraints* of the bidders which can destroy truthful bidding. Even if truthful bidding persists, this in itself causes the problem of *information revelation*. While some auction mechanisms (such as the sealed bid first price auction) do not force bidders to reveal their true value, the VCG mechanism certainly does. This may be unappealing to bidders, as discussed above. Furthermore, the VCG mechanism is susceptible to *cheating* because it is easier for colluding partners to enforce their collusion. In a *sequence of auctions* the VCG mechanism may also no longer be strategy-proof because the bidders' values in any particular auction will depend on their assessment of the level of competition in the future auctions. Finally, the process is also *revenue deficient* because it can result in zero revenue for the seller.

While only some of these issues may be applicable to any particular auction implementation, these issues must be carefully considered before proceeding with a potential VCG implementation. To complete this literature review, the following section will provide an overview of some real world auction implementations in fields that have similar attributes to the targeted television advertising field.

2.3. Relevant Applications

2.3.1. Smart Markets

The auction format that the seller chooses to sell targeted television advertisements may belong to the family of market mechanisms known as *smart markets*. Gallien and Wein (2005) define smart markets as "exchange institutions supported by a computer executing an optimization algorithm to solve the allocation problem associated with each given set of bids" (p. 77). In their paper, a linear program provides competitive information feedback to the bidders, who then have the chance to update their bids. Their updated values create new inputs to the linear program which is subsequently re-solved. Using a game-theoretic framework, they develop an upper bound for the winning bids and show some structural and convergence properties of the bidders' myopic best-response dynamics.

The targeted television advertising system can also be described as a smart market in which individual auctions are executed electronically through an algorithmic mechanism. However, this new system is quite different from the Gallien and Wein model of a producer-supplier setting in which suppliers work under production capacities and experience linear production costs.

2.3.2. Internet Advertising Auctions

Users of Google and Yahoo! search engines know that there is an area at the top and right of the screen reserved for advertisements relevant to the search terms they have used. Some users, however, are not aware that these advertisement slots were auctioned off to potential advertisers. In fact, Google's current auction system has some strong resemblance to a potential future auction for the targeted television advertising spots. Both are targeting viewer impressions and will be comprised of many mini-auctions in which there is no direct bidding but where instead a smart market determines the winners and their required payments. The timing and wording of search terms, as well as the timing and viewer composition of avails, are so unpredictable and frequent that holding individual (non-automated) auctions is not possible. A further similarity is the fact that items (advertisement space and flotilla slots) are related to each other in the sense that advertisements located higher in the search results will experience more impressions than lower ranked advertisements, just as advertisements in the top flotilla slots have an advantage over the ads in the bottom flotilla slots.

The two differences between the systems stem from the way consumers are targeted and from technical configurations. For Google and Yahoo!, consumers are grouped and characterized by their search terms. Yahoo! employs a single measure when selecting appropriate advertisements and ranks advertisers by their respective bid amounts. Google goes one step further by factoring in a *quality score* which indicates how appropriate each advertisement is for the given search term (Edelman, Ostrovsky and Schwarz, 2007). In contrast, the targeted television advertising system targets demographic groups and ranks potential advertisers by their bids as well as the current size of the respective advertiser demographic target group. Also, under Google and Yahoo!, higher ranked advertisements will experience a higher click-rate but will not capture all appropriate viewers exclusively. In comparison, the highest ranked advertisement that targets any particular demographic group in the new television system will capture that demographic group exclusively. For example, should the top flotilla slot go to an advertiser targeting all males, then all advertisers in lower flotilla slots, who are also targeting males or male subgroups, will not be able to capture *any* impressions. Despite these differences, the similarities between the two systems warrant a closer look at the auction format chosen by Google and Yahoo!.

One of the earliest auction models used to sell Internet advertising was developed by Overture (who were acquired by Yahoo! at the end of 2003). Edelman and Ostrovsky (2007) examined Overture's experiences with this first price auction model and conclude that it was not an optimal auction choice. In particular, they demonstrated that a first price auction, in the context of repeated auctions, creates a very unstable system in which bidders act strategically and try to game the system. After any given auction, some bidders may realize that they overpaid (the next highest bid was substantially lower than theirs) and will therefore decrease their bid in the next round. Simultaneously, other bidders may realize that they could have won the item by bidding higher and will update their bids accordingly in the next round. The change in bids for the next auction will present further opportunities for improvement and an infinite cycle of bidding adjustments occurs. As a result, bidders will invest significant effort into gaming the system despite the fact that no value is created by these actions. The costs of this effort is then passed on to the seller. Edelman and Ostrovsky show empirically that this auction format did indeed lower advertisement revenue for Overture.

Other search engine companies learned from this experience and chose a general second price auction to promote truth-telling and discourage strategic bidding. Google, for example, realized that any bidder in position j is only willing to pay the small amount necessary to beat bidder in position $j + 1$. As a result, the highest bidder will pay the second highest bid, the second highest bidder will pay the third highest bid, etc. However, as pointed out by Edelman, Ostrovsky and Schwarz (2007), this is not a VCG mechanism and truth-telling is not the dominant strategy in all situations. Under some scenarios, bidders can achieve a higher profit by sacrificing their rank for lower costs by bidding smaller amounts. Edelman et al. demonstrate the resulting difficulty in characterizing the bidding equilibrium.

Google's auction model is increasingly creating academic interest as demonstrated by another recent paper by Varian (2007) who compares Google's system to the assignment game and calculates the equilibria of the ad auction. He then uses his method to empirically determine the relationship between bids and values. While the state of internet advertising auctions may be described as more progressed than that of television advertising, some work has also been done in the latter field.

2.3.3. The Television Advertising Market

Due the unique technological capabilities of the new targeted television advertising system, no comparable system exists or has been studied by the research community. However, there has been some work regarding the way that television advertisement airtime is currently sold in the up-front market. Jones and Koehler (2002) show how computer-supported online auctions could simplify the sale of television airtime. They propose a combinatorial auction that accepts rule-based bids which does not require the bidder to specify bids for all combinations of possible airtime slots. Their main contribution is a winner determination heuristic that is fast enough to be practical in the proposed setting. Some of their work is based on Jones's dissertation (2000), who also provides an excellent overview of network television advertising practices.

Another paper on the up-front television airtime market was written by Bollapragada et al. (2002) and describes how NBC used an optimization-based sales system to be able to react quickly to customer requests for advertising placements in order to maximize NBC's revenue. However, they discuss a list price system – not an auction model.

2.4. TARGETED TELEVISION ADVERTISING SYSTEM

The new targeted television advertising system contains a unique environment in which to implement an auction. In the literature, auctions have been well studied but, to the best of my knowledge, they have never been studied under the special structure in which bidders are bidding for objects (in this case flotilla slots) that will enable them to capture their actual desired items (impressions) and in which the objects have different success rates of leading to items (ads in higher flotilla slots can receive more impressions). The academic papers on Google's system perhaps come closest to analyzing this situation but in their system viewers will consider any advertisement (with higher ranked ads being more *likely* to be seen) while in the new system only the highest ranked ad that is appropriate to a particular viewer will be seen.

The system is very complex to analyze mathematically so to gain some initial insight into the structure, the following chapter presents a simplified analytical model.

3. SIMPLIFIED ANALYTICAL MODEL

3.1. PROBLEM FRAMEWORK

In this chapter, analytical models are used to compare multiple auction formats to a list price system in a framework that is very close to the new targeted television advertising system described in Chapter 1. In particular, there are two items for sale (two flotilla slots) for which bids are submitted by *n* bidders, or advertisers. The new system recognizes which advertisements are appropriate for the individual households. Since more than one advertisement may be appropriate for some households – but only one of them can be viewed– attaining a slot in the flotilla may not be sufficient to guarantee impressions from all appropriate households; instead, the position within the flotilla will determine which ad receives the impression by assigning the impression to the top slot should a household be appropriate for both the ad in the top and bottom slot. This situation will arise commonly as advertisers target similar demographics and their ads are therefore either both suitable or both unsuitable for a particular household.

To capture some of this complexity, while maintaining a tractable mathematical model, each advertisement is assigned one of two "types", type A or type B. One could think of these types as, for example, indicating whether an advertisement is targeting females (type A) or males (type B). For this model, it was assumed that advertisers are of type A or B with equal probability. In reality, there are many more types since genders are further broken down by age cohorts and possibly income but the restriction to two types is sufficient to show the impact of types on the bidding behaviour and revenue generation of auctions.

Since the advertisement that attains the top slot in the flotilla will always be seen by all appropriate households, all advertisers desire this slot regardless of their type. However, advertisers only desire the bottom slot if they are not of the same type as the winner of the top slot. Using the example of types that indicate gender, Figure 3.1 illustrates this point. Assume that there are three bidders, Ad X, Y and Z, competing for the flotilla slots. Ad X and Y are both of type A (targeting females) and Ad Z is of type B (targeting males). Furthermore, Ad X has placed the largest bid, followed by Ad Y and Ad Z submitted the smallest bid. Since Ad X placed the largest bid, it wins the top flotilla slot as shown in both Scenario 1 and 2 of Figure 3.1. However, Scenario 1 shows why the second largest bidder, Ad Y, is not interested in the 2nd flotilla slot. Since Ad Y is of the same type as Ad X, it will not be able to capture any viewers: no male viewers will be interested in either advertisement and all female viewers will find both suitable which by the default rule described above will allocate the impression to the top slot. As a result, the 2nd flotilla slot will instead go to the lower bidder, Ad Z, as shown in Scenario 2. In this type of system, it is also obvious that each bidder only desires one flotilla slot since no benefit can be realized by filling both flotilla slots with the same advertisement.

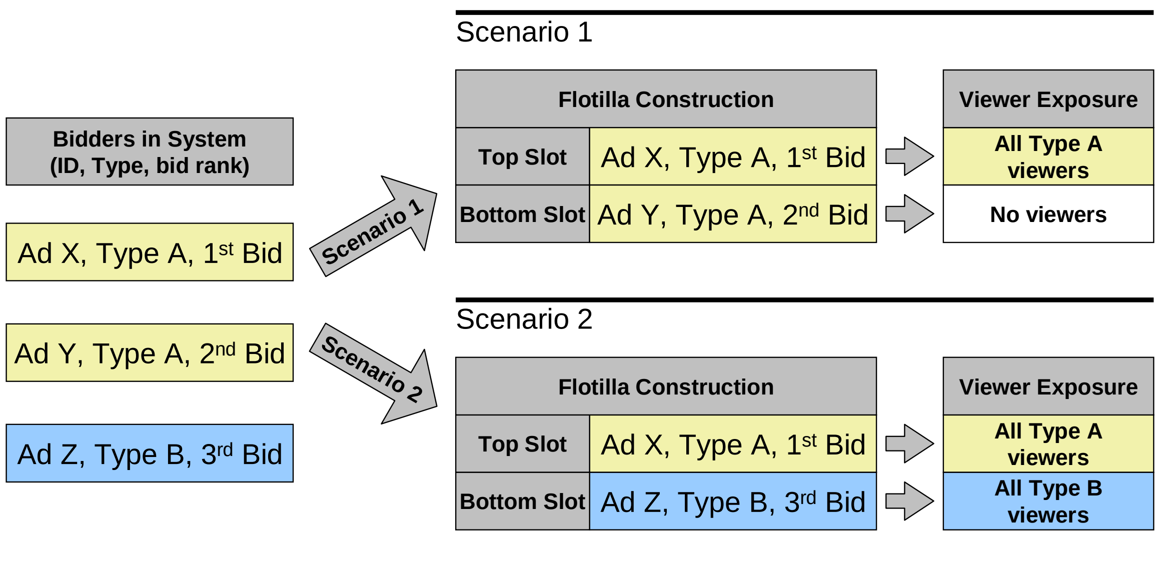

Figure 3.1: Flotilla construction example assuming 3 bidders with 2 types.

We assume that the values that bidders place on a flotilla slot follow a particular value distribution known to all bidders, $F(\cdot)$. Each bidder also knows his own valuation but he does not know the valuations of the other bidders, nor would his valuation of the item change if he knew. This is a reasonable assumption because impressions do not have a common value to all advertisers (such as oil or gold) but instead corporations have valuations that depend on the worth of the impressions to that particular advertiser. The validity of this assumption can be seen in the varying degrees of advertisement engagements by different organizations. Hence, this is a perfect private value auction. Furthermore, we assume that advertisers are risk-neutral and are trying to maximize their expected profit.

The following sections compare two auction models (a first price and a second price auction), a list price system, and a second price auction with reservation price and shows the effect of "types" on bidding behaviour and revenue generation, as well as the impact of the value distributions.

3.2. First Price Auction

3.2.1. Optimal Bidding Strategy

As described in Section 2.1.2, the first price auction is defined by the highest bidder winning the item and paying the amount that he bid. A bidder in this type of auction model faces the tradeoff between increasing his chance of winning the item by bidding a high value and decreasing the amount he needs to pay by bidding a low value. Using a similar approach to Menezes and Monteiro (2005), we analyze this situation from the point of view of one of the players, say Advertiser 1. Let us assume that this advertiser has a valuation $v = v_1$ and believes that all other advertisers follow a bidding strategy $b(\cdot)$. The profit π that Advertiser 1 can expect to derive from a particular auction is his valuation of the item v minus his bid b_1 if he secures a flotilla slot, and his profit is zero if he does not secure a slot. Defining $P_W(b_1)$ as the probability of winning a flotilla slot with a bid b_1, we have:

$$\pi(b_1) = P_W(b_1)(v - b_1) \tag{1a}$$

Assuming that the bid function $b(\cdot)$ is strictly increasing and differentiable (to be verified later) and that the range of $b(\cdot)$ is $[0,v]$ (since the bidder will never bid above his value and pay more than he values the item, nor bid less than zero), there exists $x \in [0, v]$ such that $b_1 = b(x)$. Therefore, the problem of Advertiser 1 is equivalent to choosing $x \in [0, v]$ to maximize the expected utility:

$$\pi(x) = \pi(b(x)) = P_W(b(x))(v - b(x)) \tag{1b}$$

In the absence of "types" and with only one object for sale, P_W is simply the probability of having the highest valuation of all bidders, $F(v)^{n-1}$. In the framework of this section, winning an item can be achieved by having the highest valuation of all bidders, or by having the second highest valuation and not being of the same type of the winner of the top slot (since bidders are of either type with equal probability, the probability of being of a different type is equal to one half), or by having the third highest valuation and being of a different type than the winner of the top slot while the second highest bidder is of the same type as the winner (occurring with a probability of one quarter, this scenario leads to the second highest bidder losing interest in the auction and the third highest bidder having the chance to capture the second slot), and so forth. Under this scenario, even the lowest bidder has the chance of winning if all other bidders are of the other type. This occurs with probability $2^{-(n-1)}$.

Using order statistics (Ross, 2003), the cumulative distribution function of the i^{th} smallest random variable (RV), X_i, of n i.i.d. RVs is:

$$P\left(X_{(i)} \leq x\right) = \sum_{k=i}^{n} \binom{n}{k} \left(F(x)\right)^{k} \left(1 - F(x)\right)^{n-k}$$

Therefore, we get:

$$\begin{aligned}
P_W(b(x)) = P_W(x) &= 2^0 P\left(x \geq X_{(n-1)}\right) + 2^{-1}\left[P\left(x \geq X_{(n-2)}\right) - P\left(x \geq X_{(n-1)}\right)\right] + \ldots \\
&\quad + 2^{-(n-2)}\left[P\left(x \geq X_{(1)}\right) - P\left(x \geq X_{(2)}\right)\right] + 2^{-(n-1)}\left[1 - P\left(x \geq X_{(1)}\right)\right] \\
&= \binom{n-1}{n-1}\left(F(x)\right)^{n-1}\left(1-F(x)\right)^{n-1-n+1} + 2^{-1}\left[\binom{n-1}{n-2}\left(F(x)\right)^{n-2}\left(1-F(x)\right)^{n-1-n+2}\right] + \ldots \\
&\quad + 2^{-(n-2)}\left[\binom{n-1}{1}\left(F(x)\right)^{1}\left(1-F(x)\right)^{n-1-1}\right] + 2^{-(n-1)}\left[\binom{n-1}{0}\left(F(x)\right)^{0}\left(1-F(x)\right)^{n-1}\right] \\
&= \sum_{k=0}^{n-1} 2^{-(n-1-k)}\binom{n-1}{k}\left(F(x)\right)^{k}\left(1-F(x)\right)^{n-1-k} \\
&= \left(\frac{F(x)+1}{2}\right)^{n-1} \qquad (2)
\end{aligned}$$

To find the maximum utility that Advertiser 1 can achieve, we take the derivative of equation (1b) with respect to x and set it equal to zero:

$$\pi'(x) = \left(v - b(x)\right)P_W'(x) - b'(x)P_W(x) \qquad (3a)$$

As stated by Menezes and Monteiro (2005), the expected profit is maximized at $x = v$ in a symmetric equilibrium. Setting $\pi'(x) = 0$ yields:

$$b'(v)P_W(v) = \left(v - b(v)\right)P_W'(v) \qquad (3b)$$

This can be solved by noticing:

$$\begin{aligned}\left(b(v)P_W(v)\right)' &= b'(v)P_W(v) + b(v)P_W'(v)\\ &= vP_W'(v)\end{aligned}$$

Using the Fundamental Theorem of Calculus (Stewart, 1999):

$$b(v)P_W(v) = \int_0^v xP_W'(x)dx + k \tag{3c}$$

As v goes to zero, the left-hand side of equation (3c) goes to zero because $b(.)$ is bounded and therefore $k = 0$. By substituting equation (2) into equation (3c) and solving for the optimal bid, b^*, we get for $v > 0$:

$$\begin{aligned}b^*(v) &= \frac{\int_0^v xP_W'(x)dx}{P_W(v)}\\ &= \frac{\int_0^v \left(x\frac{d}{dx}\left[\left(\frac{F(x)+1}{2}\right)^{n-1}\right]\right)dx}{\left(\frac{F(v)+1}{2}\right)^{n-1}}\\ &= \frac{\int_0^v \left(\frac{n\text{-}1}{2}\left(\frac{F(x)+1}{2}\right)^{n-2} xf(x)\right)dx}{\left(\frac{F(v)+1}{2}\right)^{n-1}}\\ &= \frac{\int_0^v \left((n-1)(F(x)+1)^{n-2} xf(x)\right)dx}{(F(v)+1)^{n-1}}\end{aligned} \tag{4}$$

And $b^*(v) = 0$ for $v = 0$.

To check whether $b(v)$ is continuous, we notice that when $v > 0$,

$$b^*(v) = \frac{\int_0^v \left((n-1)(F(x)+1)^{n-2} xf(x)\right)dx}{(F(v)+1)^{n-1}}$$

$$< \frac{\int_0^v \left((n-1)(F(x)+1)^{n-2} vf(x)\right)dx}{(F(v)+1)^{n-1}} < v$$

And thus $b(v)$ is continuous everywhere. In addition, from equation (3a) and (3b) we get:

$$\pi'(x) = (v-x)P_W'(x)$$

If $x < v$, $\pi'(x) > 0$ and if $x > v$, $\pi'(x) < 0$. Therefore $x = v$ maximizes the expected payoff and b^* is an equilibrium.

The optimal bidding function of equation (4) maps any possible bidder value into an optimal bid that will maximize that bidder's expected payoff. Assuming that all bidders follow this optimal bid strategy, we can now calculate the expected revenue of the seller.

3.2.2. Expected Revenue for Seller

The seller will receive from each bidder the revenue R equal to the bid multiplied with the probability of winning a slot in the flotilla of that particular bidder.

$$R = b^*(x) * P_W(x)$$

$$= \frac{\int_0^v xP_W' dx}{P_W} * P_W$$

$$= \int_0^v xP_W' dx$$

$$= \int_0^v \left(\frac{n\text{-}1}{2} \left(\frac{F(x)+1}{2} \right)^{n-2} xf(x) \right) dx \tag{5}$$

To find the expected revenue of the seller from this auction, we multiply the number of bidders n with the expected revenue $E[R]$ (expectation of equation (5)) from each bidder:

$$nE[R] = n\int_0^{\infty}\left[\int_0^{v}\left(\frac{n\text{-}1}{2}\left(\frac{F(x)+1}{2}\right)^{n-2} xf(x)\right)dx\right]f(v)dv \tag{6}$$

A numerical example of a first price auction will be provided later.

3.3. Second Price Auction

3.3.1. Optimal Bidding Strategy

As described in Section 2.1.2, the second price auction is characterized by the highest bidder receiving the item but only paying the second highest bid. In the framework of this section, the highest bidder of each type receives one slot and must pay the next highest bid of the same respective type. Similar to the standard second price auction, this mechanism promotes truth-telling, i.e. bidding the precise personal value of the object.

3.3.2. Expected Revenue for Seller

To find the expected revenue of the seller, we start by calculating the probability mass function (pmf) of the revenue that the seller will receive from the highest bidder, M_1. The highest bidder will have to pay the second highest bid, $X_{(n-1)}$, if the second highest bidder happens to be of the same type (probability of one half); in contrast, he will only have to pay the third highest bid, $X_{(n-2)}$, if the second highest bidder is of the other type but the third highest bidder is of the same type as him (probability of one quarter), and so forth. We define $X_{(0)} = 0$.

$$P\left(M_1 = X_{(i)}\right) = \begin{cases} 2^{i-n} & 0 < i < n \\ 2^{1-n} & i = 0 \end{cases}$$

Similarly, the second highest bidder will have to pay the third highest bid if the highest bidder is of the other type (else the second bidder has no more interest in the auction) and the third highest bidder is of the same type (probability of one quarter). With similar reasoning as before, we get:

$$P\left(M_2 = X_{(i)}\right) = \begin{cases} 2^{i-n} & 0 < i < n-1 \\ 2^{1-n} & i = 0 \end{cases}$$

And in general, the k^{th} bidder has the pmf of payment:

$$P\left(M_k = X_{(i)}\right) = \begin{cases} 2^{i-n} & 0 < i < n-k+1 \\ 2^{1-n} & i = 0 \end{cases}$$

The total expected revenue of the seller is the sum of the expected payments of all bidders, namely:

$$E[R] = \sum_{k=1}^{n-1} 2^{k-n} E\left[X_{(k)}\right] + \sum_{k=1}^{n-2} 2^{k-n} E\left[X_{(k)}\right] + \ldots + \sum_{k=1}^{1} 2^{k-n} E\left[X_{(k)}\right]$$

which can be rewritten as:

$$E[R] = \sum_{k=1}^{n-1} (n-k) 2^{k-n} E\left[X_{(k)}\right] \tag{7}$$

3.4. List Price – Optimization

3.4.1. Formulation

As an alternative to an auction model, the seller could specify a list price at which flotilla slots are sold to any buyer willing to pay the posted price. If demand surpasses supply, items are typically awarded on a first-come-first-serve basis; however, in the automated system that the seller will be running, this rule will not be practical and it is conceivable that, as an alternative, an optimization mechanism could pick the two advertisements that will maximize the seller's revenue. Following is a formulation of such an optimization model, assuming that the seller knows the actual valuations of the advertisers V_a.

Linear Integer Programming Optimization Model – List Price System

Model parameters:

V_a = value per impression of advertiser a

$T_{t,a}$ = 1 if ad a is of type $t \in \{A, B\}$, else 0

L = list price charged per impression

Decision variables:

$$x_{t,a} = \begin{cases} 1 & \text{ad } a \text{ of type } t \text{ has attained the item (flotilla slot)} \\ 0 & \text{else} \end{cases}$$

Objective function:

$\max_{x} \; P = L\sum_{t,a} x_{t,a}$ (total revenue for seller)

Constraints:

$\sum_{t} x_{t,a} \le 1 \quad \forall a$ (each ad can only receive one flotilla slot or none)

$\sum_{t} x_{t,a} \le 1 \quad \forall t$ (only one ad of each type can attain a flotilla slot)

$L\sum_{t} x_{t,a} \le V_a \quad \forall a$ (each ad can only receive a slot if its V is higher than L)

$x_{t,a} \le T_{t,a} \quad \forall t, a$ (each ad can only claim a slot in congruence with its type)

$x_{t,a} \in \{0,1\} \quad \forall t, a$

A close examination of the problem reveals that its complexity does not require an optimization model and that a few lines of code, or even visual inspection, will find the solution more efficiently. To solve this problem it is merely required to find any two advertisements of different types with a personal valuation of the item above the listed price. A more interesting question, therefore, is how to set the optimal list price.

3.4.2. Setting an Optimal List Price

To find the optimal list price, we need to find an expression of the seller's profit as a function of the list price charged, L. The seller will receive the list price at least once if the highest bidder has a personal valuation of the item above the list price, which has a probability of $1 - F(L)^n$. He will receive the list price a second time if the second bidder is above the list price and the top two bidders are not of the same type, with a probability of $\left(1-2^{-1}\right)\binom{n}{n-2}\left(F(L)\right)^{n-2}\left(1-F(L)\right)^2$, or if the third bidder is above the list price and the top three bidders are not of the same type, with probability of $\left(1-2^{-3}\right)\binom{n}{n-3}\left(F(L)\right)^{n-3}\left(1-F(L)\right)^3$, and so forth. The profit of the seller is therefore:

$$\pi(L) = L\left[1-\left(F(L)\right)^n + \sum_{k=0}^{n-2}\left(1-2^{k+1-n}\right)\binom{n}{k}\left(F(L)\right)^k\left(1-F(L)\right)^{n-k}\right] \tag{8}$$

By setting the first order derivative of the profit with respect to L to zero and solving for L, one could find the optimal list price. A simple way to find the optimal L is to use Solver in Microsoft Excel by maximizing equation (8) subject to L.

3.5. Second Price Auction with Reservation Price

3.5.1. Optimal Bidding Strategy

One way to combine an auction model with a list price system is to implement an auction with a reservation price. In this type of auction, the seller is decreasing the probability of receiving a very low payment but is increasing the number of potential bidders who do not participate in the auction. The higher he sets the reservation price, the higher his guaranteed payment (if the item sells) but the lower his probability of selling the item.

The optimal bidding strategy in a second price auction with a reservation price is the same as the optimal bidding strategy in a second price auction without a reservation price except that bidders with values below the reservation price do not participate in the auction.

3.5.2. Expected Revenue for Seller

As before, we find the expected revenue of the seller by calculating the probability mass function (pmf) of the revenue (payment) that the seller will receive from the highest bidder, M_1. The highest bidder will have to pay the second highest bid, $X_{(n-1)}$, if the second highest bidder happens to be of the same type (probability of one half) and has a value above the reservation price; in contrast, he will only have to pay the third highest bid, $X_{(n\text{-}2)}$, if the second highest bidder is of the other type but the third highest bidder is of the same type as he is (probability of one quarter) and his value is above the reservation price, and so forth. He will have to pay the reservation price r if all bidders with values above the reservation price are of the other type than him (without the reservation price, he would have had to pay zero). This leads to:

$$P\left(M_1 = X_{(i)} \middle| X_{(i)} \geq r\right) = 2^{i-n} P\left(X_{(i)} \geq r\right) \qquad 0 > i > n$$

$$P\left(M_1 = r\right) = \sum_{j=1}^{n} 2^{j-n} \left[P\left(X_{(j-1)} < r\right) - P\left(X_{(j)} < r\right)\right]$$

which results in:

$$P\left(M_1 = X_{(i)} \middle| X_{(i)} \geq r\right) = 2^{i-n} \left[1 - \sum_{k=i}^{n} \binom{n}{k} \left(F(r)\right)^k \left(1 - F(r)\right)^{n-k}\right] \qquad 0 > i > n$$

$$P\left(M_1 = r\right) = \sum_{j=1}^{n} 2^{j-n} \binom{n}{j-1} \left(F(r)\right)^{j-1} \left(1 - F(r)\right)^{n-j+1}$$

And in general, the k^{th} bidder has the pmf of payment:

$$P\left(M_k = X_{(i)} \middle| X_{(i)} \geq r\right) = 2^{i-n} \left[1 - \sum_{k=i}^{n} \binom{n}{k} \left(F(r)\right)^k \left(1 - F(r)\right)^{n-k}\right] \qquad 0 > i > n-k+1$$

$$P\left(M_k = r\right) = \sum_{j=1}^{n-k+1} 2^{j-n} \binom{n}{j-1} \left(F(r)\right)^{j-1} \left(1 - F(r)\right)^{n-j+1}$$

The total expected revenue of the seller is the sum of the expected payments of all bidders, namely:

$$
\begin{aligned}
E[R] = &\sum_{j=1}^{n-1} 2^{j-n}\left[1-\sum_{k=j}^{n}\binom{n}{k}(F(r))^k(1-F(r))^{n-k}\right]E\left[X_{(t)}\middle|X_{(t)} \geq r\right] \\
&+ r\sum_{j=1}^{n} 2^{j-n}\binom{n}{j-1}(F(r))^{j-1}(1-F(r))^{n-j+1} \\
&+ \sum_{j=1}^{n-2} 2^{j-n}\left[1-\sum_{k=j}^{n}\binom{n}{k}(F(r))^k(1-F(r))^{n-k}\right]E\left[X_{(t)}\middle|X_{(t)} \geq r\right] \\
&+ r\sum_{j=1}^{n-1} 2^{j-n}\binom{n}{j-1}(F(r))^{j-1}(1-F(r))^{n-j+1} \\
&+ \ldots \\
&+ \sum_{j=1}^{1} 2^{j-n}\left[1-\sum_{k=j}^{n}\binom{n}{k}(F(r))^k(1-F(r))^{n-k}\right]E\left[X_{(t)}\middle|X_{(t)} \geq r\right] \\
&+ r\sum_{j=1}^{2} 2^{j-n}\binom{n}{j-1}(F(r))^{j-1}(1-F(r))^{n-j+1} \\
&+ r\sum_{j=1}^{1} 2^{j-n}\binom{n}{j-1}(F(r))^{j-1}(1-F(r))^{n-j+1}
\end{aligned}
$$

which can be rewritten as:

$$
E[R] = \sum_{j=1}^{n}\left[\begin{array}{l}(n-j)2^{j-n}\left[1-\sum_{k=j}^{n}\binom{n}{k}(F(r))^k(1-F(r))^{n-k}\right]E\left[X_{(j)}\middle|X_{(j)} \geq r\right] \\ +r(n-j+1)2^{j-n}\binom{n}{j-1}(F(r))^{j-1}\left(1-F(r)^{n-j+1}\right)\end{array}\right] \tag{9}
$$

To give the reader a more concrete example, the next section provides a numerical example and compares the introduced auction models and a list price system.

3.6. Numerical Example

3.6.1. Introduction

In this numerical example, there are three bidders who each draw their value for a flotilla slot from a [λ]-exponential distribution so $n = 3$, $F(x) = 1 - \exp(-\lambda*x)$, and $f(x) = \lambda*\exp(-\lambda*x)$. The exponential distribution is one of many distributions that could be used to model impression valuations of advertisers. One could argue that there are many advertisers who have a relatively low value per impression; perhaps they are targeting a broad range of customers and are willing to pay only a small amount per impression, or their market research indicates that advertising only provides a relatively small value to them. Conversely, some advertisers are willing to pay higher prices if they believe in a higher conversion rate from advertisement to purchase. Some advertisers may even value impressions extremely high if they can reach their highly specified target group, such as high income females. The exponential distribution models such a situation well because it encompasses a wide left-hand tail close to zero (representing the large amount of advertisers with relatively low values) and an open right-hand tail (representing the fact that some advertisers may value their impressions extremely high). To provide a complete analysis, however, a later section will provide results given some different value distributions. The following calculations show the expected revenue from a first price auction, a second price auction (with and without a reservation price) and a list price system given three bidders with [λ]-exponential distributions.

3.6.2. First Price Auction

Using equation (4), we derive the following bidding function:

$$b^*(v) = \frac{\int_0^v 2\left(1 - e^{-\lambda x} + 1\right)^1 x\lambda e^{-\lambda x} dx}{\left(2 - e^{-\lambda v}\right)^2}$$

$$= \frac{4\int_0^v x\lambda e^{-\lambda x} dx - 2\int_0^v x\lambda e^{-2\lambda x} dx}{4 - 4e^{-\lambda v} + e^{-2\lambda v}}$$

$$= \frac{4\left[-xe^{-\lambda x}\right]_0^v + 4\int_0^v e^{-\lambda x} dx + \left[xe^{-2\lambda x}\right]_0^v - \int_0^v e^{-2\lambda x} dx}{4 - 4e^{-\lambda v} + e^{-2\lambda v}}$$

$$= \frac{-4ve^{-\lambda v} + 4\left[-\frac{1}{\lambda}e^{-\lambda x}\right]_0^v + ve^{-2\lambda v} - \left[-\frac{1}{2\lambda}e^{-2\lambda x}\right]_0^v}{4 - 4e^{-\lambda v} + e^{-2\lambda v}}$$

$$= \frac{-4ve^{-\lambda v} - 4\frac{1}{\lambda}e^{-\lambda v} + 4\frac{1}{\lambda} + ve^{-2\lambda v} + \frac{1}{2\lambda}e^{-2\lambda v} - \frac{1}{2\lambda}}{4 - 4e^{-\lambda v} + e^{-2\lambda v}}$$

$$= \frac{\frac{7}{2\lambda} - \left(\frac{4}{\lambda} + 4v\right)e^{-\lambda v} + \left(\frac{1}{2\lambda} + v\right)e^{-2\lambda v}}{4 - 4e^{-\lambda v} + e^{-2\lambda v}}$$

$$= \frac{\frac{7}{2\lambda} - \left(\frac{4}{\lambda} + 4v\right)e^{-\lambda v} + \left(\frac{1}{2\lambda} + v\right)e^{-2\lambda v}}{4 - 4e^{-\lambda v} + e^{-2\lambda v}} \tag{10}$$

Expected revenue can be calculated using equation (6):

$$E[R] = 3\int_0^\infty \left[\int_0^v \left(\left(\frac{2 - e^{-\lambda x}}{2}\right)x\lambda e^{-\lambda x}\right)dx\right]\lambda e^{-\lambda v}dv$$

$$= 3\int_0^\infty \left[\int_0^v x\lambda e^{-\lambda x} - \frac{1}{2}x\lambda e^{-2\lambda x}dx\right]\lambda e^{-\lambda v}dv$$

$$= 3\int_0^\infty \left[\left[-xe^{-\lambda x}\right]_0^v + \int_0^v e^{-\lambda x}dx + \left[\frac{1}{4}xe^{-2\lambda x}\right]_0^v - \int_0^v \frac{1}{4}e^{-2\lambda x}dx\right]\lambda e^{-\lambda v}dv$$

$$= 3\int_0^\infty -\lambda ve^{-2\lambda v} - e^{-2\lambda v} + e^{-\lambda v} + \frac{1}{4}\lambda ve^{-3\lambda v} + \frac{1}{8}e^{-3\lambda v} - \frac{1}{8}e^{-\lambda v}dv$$

$$= 3\left(\begin{array}{l}\left[\frac{1}{2}ve^{-2\lambda v}\right]_0^\infty + \left[\frac{1}{4\lambda}e^{-2\lambda v}\right]_0^\infty + \left[\frac{1}{2\lambda}e^{-2\lambda v}\right]_0^\infty + \left[-\frac{7}{8\lambda}e^{-\lambda v}\right]_0^\infty + \\ \left[-\frac{1}{12}ve^{-3\lambda v}\right]_0^\infty + \left[-\frac{1}{36\lambda}e^{-3\lambda v}\right]_0^\infty + \left[-\frac{1}{24\lambda}e^{-3\lambda v}\right]_0^\infty\end{array}\right)$$

$$= 3\left(-\frac{1}{4\lambda} - \frac{1}{2\lambda} + \frac{7}{8\lambda} + \frac{1}{36\lambda} + \frac{1}{24\lambda}\right) = \frac{7}{12\lambda}$$

3.6.3. Second Price Auction

Since the optimal bidding behaviour is truth-telling, we can easily calculate the expected revenue using equation (7) and Appendix A:

$$E[R] = \sum_{k=1}^{2} (3-k) 2^{k-3} E\left[X_{(k)}\right]$$

$$= 2^{-1} E\left[X_{(1)}\right] + 2^{-1} E\left[X_{(2)}\right]$$

$$= 2^{-1} \frac{1}{3\lambda} + 2^{-1} \frac{5}{6\lambda} = \frac{7}{12\lambda}$$

The expected revenue is equal to the first price auction scenario! While this is not a conclusive proof, it does provide a strong indication that the revenue equivalence theorem does indeed hold in this framework and that any auction model will result in the same expected revenue. Revenue equivalence will be examined in more detail in a later section.

3.6.4. List Price

By equation (8), the profit as a function of the list price is:

$$\pi(L) = L\left[1 - \left(1 - e^{-\lambda L}\right)^3 + \sum_{k=0}^{1} \left(1 - 2^{k-2}\right) \binom{3}{k} \left(1 - e^{-\lambda L}\right)^k \left(e^{-\lambda L}\right)^{3-k}\right]$$

$$= L\left[1 - \left(1 - e^{-\lambda L}\right)^3 + \tfrac{3}{4}\left(e^{-\lambda L}\right)^3 + \tfrac{3}{2}\left(1 - e^{-\lambda L}\right)\left(e^{-\lambda L}\right)^2\right]$$

Using Excel Solver, the maximum profit is achieved with the list price L set to 1.19 / λ which results in a profit of 0.93 / λ. Evidently, the seller can achieve a higher profit if he sets a list price instead of selling the flotilla slots through an auction. The seller's expected profit does, however, depend on him setting the correct list price. Assuming a [1]-exponential value distribution of the three bidders, Figure 3.2 shows the expected revenue as a function of L. The graph demonstrates that the list price could be set anywhere between 0.5 and 2.5 and still result in an expected revenue that is above the revenue that can be expected from an auction. In other words, even if the seller sets the list price 66% below or 116% above the optimal list price, he can still expect a higher profit than had he sold the items through an auction.

Intuitively, the auctions suffer from the low ratio of number of bidders to number of items. As shown in the previous section, the seller has a fifty-fifty chance of receiving payment equal to the smallest or the second smallest bidder valuation. As Appendix A showed, these are likely to be small values. Under the list price system, however, the seller has the opportunity to capture higher amounts from high-value bidders. By setting the price above the expected value, he risks having no revenue in some situations but when he does make a sale, payments will be large. A later section will explore under which circumstances this statement remains valid.

Figure 3.2: Expected revenue of auction vs. list price dependent on the set list price.

3.6.5. Second Price Auction with Reservation Price

In an auction with a reservation price we expect the revenue to be at least equal to the greater of a pure auction system or a pure list price system since these systems can be emulated by simply setting the reservation price to zero or the optimal list price, respectively. This section demonstrates the additional value that can be derived from an optimally set reservation price. Using equation (9), we find:

$$E[R]=\sum_{j=1}^{3}\left[\begin{array}{l}(3-j)2^{j-3}\left[1-\sum_{k=j}^{3}\binom{3}{k}(F(r))^{k}(1-F(r))^{3-k}\right]E\left[X_{(k)}\middle|X_{(k)}\geq r\right]\\ +r(3-j+1)2^{j-3}\binom{3}{j-1}(F(r))^{j-1}(1-F(r))^{3-j+1}\end{array}\right]$$

$$\begin{aligned}&=2^{-1}\left[1-\sum_{k=1}^{3}\binom{3}{k}(F(r))^{k}(1-F(r))^{3-k}\right]E\left[X_{(1)}\middle|X_{(1)}\geq r\right]+r(3)2^{-2}\binom{3}{0}(1-F(r))^{3}\\&+2^{-1}\left[1-\sum_{k=2}^{3}\binom{3}{k}(F(r))^{k}(1-F(r))^{3-k}\right]E\left[X_{(2)}\middle|X_{(2)}\geq r\right]+r(2)2^{-1}\binom{3}{1}(F(r))^{1}(1-F(r))^{2}\\&+r\binom{3}{2}(F(r))^{2}(1-F(r))\end{aligned}$$

$$\begin{aligned}&=2^{-1}\left[\binom{3}{0}(1-F(r))^{3}\right]E\left[X_{(1)}\middle|X_{(1)}\geq r\right]+r(3)2^{-2}\binom{3}{0}(1-F(r))^{3}\\&+2^{-1}\left[\binom{3}{0}(1-F(r))^{3}+\binom{3}{1}(F(r))^{1}(1-F(r))^{2}\right]E\left[X_{(2)}\middle|X_{(2)}\geq r\right]\\&+r(2)2^{-1}\binom{3}{1}(F(r))^{1}(1-F(r))^{2}+\ r\binom{3}{2}(F(r))^{2}(1-F(r))^{1}\end{aligned}$$

$$\begin{aligned}&=(1-F(r))^{3}\left[\frac{2E\left[X_{(1)}\middle|X_{(1)}\geq r\right]+3r}{4}\right]\\&+2^{-1}(1-F(r))^{3}E\left[X_{(2)}\middle|X_{(2)}\geq r\right]+2^{-1}3(F(r))(1-F(r))^{2}E\left[X_{(2)}\middle|X_{(2)}\geq r\right]\\&+r3(F(r))(1-F(r))^{2}+\ r3(F(r))^{2}(1-F(r))\end{aligned}$$

$$\begin{aligned}&=(1-F(r))^{3}\left[\frac{2E\left[X_{(1)}\middle|X_{(1)}\geq r\right]+2E\left[X_{(2)}\middle|X_{(2)}\geq r\right]+3r}{4}\right]\\&+3(F(r))(1-F(r))^{2}\left[\frac{1E\left[X_{(2)}\middle|X_{(2)}\geq r\right]+2r}{2}\right]+\ 3(F(r))^{2}(1-F(r))[r]\end{aligned}\qquad(11)$$

To maximize expression (11), we need to find $E\left[X_1\middle|X_1\geq r\right]$ and $E\left[X_2\middle|X_2\geq r\right]$. We start by finding $E\left[X_3\middle|X_3\geq r\right]$:

$$E\left[X_3 \middle| X_3 \geq r\right] = \int_r^\infty x f_{X_3|X_3 \geq r} dx = \int_r^\infty x\left(\frac{d}{dx}\frac{P(X_3 \geq r, X_3 \leq x)}{P(X_3 \geq r)}\right)dx$$

So we first need:

$$P(X_3 \geq r, X_3 \leq x) = P(r \leq X_3 \leq x) = \begin{cases} F(x)^3 & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$= \begin{cases} 1 - 3e^{-\lambda x} + 3e^{-2\lambda x} - e^{-3\lambda x} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$P\left(X_3 \leq x \middle| X_3 \geq r\right) = \begin{cases} \dfrac{P(X_3 \geq r, X_3 \leq x)}{P(X_3 \geq r)} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$= \begin{cases} \dfrac{1 - 3e^{-\lambda x} + 3e^{-2\lambda x} - e^{-3\lambda x}}{3e^{-\lambda r} - 3e^{-2\lambda r} + e^{-3\lambda r}} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$f_{X_3|X_3 \geq r} = \begin{cases} \dfrac{d}{dx}\dfrac{P(X_3 \geq r, X_3 \leq x)}{P(X_3 \geq r)} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$= \begin{cases} \dfrac{3\lambda e^{-\lambda x} - 6\lambda e^{-2\lambda x} + 3\lambda e^{-3\lambda x}}{3e^{-\lambda r} - 3e^{-2\lambda r} + e^{-3\lambda r}} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

which using Appendix B gives us:

$$E\left[X_3 \middle| X_3 \geq r\right] = \int_r^\infty x f_{X_3|X_3 \geq r} dx$$

$$= \int_r^\infty x \frac{3\lambda e^{-\lambda x} - 6\lambda e^{-2\lambda x} + 3\lambda e^{-3\lambda x}}{3e^{-\lambda r} - 3e^{-2\lambda r} + e^{-3\lambda r}} dx$$

$$= \frac{3\left(r + \dfrac{1}{\lambda}\right)e^{-\lambda r} - 6\left(\dfrac{1}{2}r + \dfrac{1}{4\lambda}\right)e^{-2\lambda r} + 3\left(\dfrac{1}{3}r + \dfrac{1}{9\lambda}\right)e^{-3\lambda r}}{3e^{-\lambda r} - 3e^{-2\lambda r} + e^{-3\lambda r}}$$

Similarly, for $E[X_1|X_1 \geq r]$ and $E[X|X \geq r]$, we get:

$$P(X_1 \geq r, X_1 \leq x) = \begin{cases} R(x)^3 & x \geq r \\ 0 & \text{otherwise} \end{cases}$$
$$= \begin{cases} e^{-3\lambda x} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$P(X \geq r, X \leq x) = \begin{cases} F(x) & x \geq r \\ 0 & \text{otherwise} \end{cases}$$
$$= \begin{cases} e^{-\lambda x} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$P(X_1 \leq x | X_1 \geq r) = \begin{cases} \dfrac{e^{-3\lambda x}}{e^{-3\lambda r}} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$
$$= \begin{cases} e^{-3\lambda(x-r)} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$P(X \leq x | X \geq r) = \begin{cases} \dfrac{e^{-\lambda x}}{e^{-\lambda r}} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$
$$= \begin{cases} e^{-\lambda(x-r)} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$f_{X_1|X_1 \geq r} = \begin{cases} 3\lambda e^{-3\lambda(x-r)} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$f_{X_1|X_1 \geq r} = \begin{cases} \lambda e^{-\lambda(x-r)} & x \geq r \\ 0 & \text{otherwise} \end{cases}$$

$$E[X_1|X_1 \geq r] = 3\left(\frac{1}{3}r + \frac{1}{9\lambda}\right)$$
$$= r + \frac{1}{3\lambda}$$

$$E[X|X \geq r] = r + \frac{1}{\lambda}$$

And since $E[X_2|X_2 \geq r] = 3E[X|X \geq r] - E[X_1|X_1 \geq r] - E[X_3|X_3 \geq r]$, we get:

$$E[X_2|X_2 \geq r] = 3r + \frac{3}{\lambda} - r - \frac{1}{3\lambda} - \frac{3\left(r + \frac{1}{\lambda}\right)e^{-\lambda r} - 6\left(\frac{1}{2}r + \frac{1}{4\lambda}\right)e^{-2\lambda r} + 3\left(\frac{1}{3}r + \frac{1}{9\lambda}\right)e^{-3\lambda r}}{3e^{-\lambda r} - 3e^{-2\lambda r} + e^{-3\lambda r}}$$
$$= 2r + \frac{8}{3\lambda} - \frac{3\left(r + \frac{1}{\lambda}\right)e^{-\lambda r} - 6\left(\frac{1}{2}r + \frac{1}{4\lambda}\right)e^{-2\lambda r} + 3\left(\frac{1}{3}r + \frac{1}{9\lambda}\right)e^{-3\lambda r}}{3e^{-\lambda r} - 3e^{-2\lambda r} + e^{-3\lambda r}}$$

Finally, substituting all these value into equation (11), we get the expected revenue of an auction with reservation price:

$$E[R]=(1-F(r))^3\left[\frac{2r+\frac{2}{3\lambda}+4r+\frac{16}{3\lambda}-\frac{6\left(r+\frac{1}{\lambda}\right)e^{-\lambda r}-12\left(\frac{1}{2}r+\frac{1}{4\lambda}\right)e^{-2\lambda r}+6\left(\frac{1}{3}r+\frac{1}{9\lambda}\right)e^{-3\lambda r}}{3e^{-\lambda r}-3e^{-2\lambda r}+e^{-3\lambda r}}+3r}{4}\right]$$

$$+3(F(r))(1-F(r))^2\left[\frac{2r+\frac{8}{3\lambda}-\frac{3\left(r+\frac{1}{\lambda}\right)e^{-\lambda r}-6\left(\frac{1}{2}r+\frac{1}{4\lambda}\right)e^{-2\lambda r}+3\left(\frac{1}{3}r+\frac{1}{9\lambda}\right)e^{-3\lambda r}}{3e^{-\lambda r}-3e^{-2\lambda r}+e^{-3\lambda r}}+2r}{2}\right]$$

$$+\ 3(F(r))^2(1-F(r))[r]$$

$$=(1-F(r))^3\left[\frac{9}{4}r+\frac{3}{2\lambda}-\frac{3\left(r+\frac{1}{\lambda}\right)e^{-\lambda r}-6\left(\frac{1}{2}r+\frac{1}{4\lambda}\right)e^{-2\lambda r}+3\left(\frac{1}{3}r+\frac{1}{9\lambda}\right)e^{-3\lambda r}}{2\left(3e^{-\lambda r}-3e^{-2\lambda r}+e^{-3\lambda r}\right)}\right]$$

$$+3(F(r))(1-F(r))^2\left[2r+\frac{4}{3\lambda}-\frac{3\left(r+\frac{1}{\lambda}\right)e^{-\lambda r}-6\left(\frac{1}{2}r+\frac{1}{4\lambda}\right)e^{-2\lambda r}+3\left(\frac{1}{3}r+\frac{1}{9\lambda}\right)e^{-3\lambda r}}{2\left(3e^{-\lambda r}-3e^{-2\lambda r}+e^{-3\lambda r}\right)}\right]$$

$$+\ 3(F(r))^2(1-F(r))[r]$$

As in the list price case, this expression can be maximized with any Solver. In the context of this numerical example, the optimal reservation price with λ=1 is 0.79 and the resulting expected revenue is 1.16. As expected, this auction with a reservation price creates the highest revenue for the seller. Chapter 4 will further explore the use of reservation prices in the proposed auction system for the targeted television advertising system but we now turn to some further analysis of this analytical model.

3.7. COMPARISONS & INSIGHTS

3.7.1. Introduction

The previous section demonstrated that, given three bidders with [λ]-exponential value distributions, setting an appropriate list price will result in higher revenue than an auction. The question remains whether this is true under all situations and how a change in the value distribution function or the number of bidders would impact this result. Furthermore, it is not yet clear how the introduction of "types" shaped the results of this section and how the results compare to a standard auction of two items without bidders falling into separate categories.

3.7.2. Bidding Functions

We begin the analysis of the impact of types by comparing the bidding functions of a first price auction for two items with types versus without types.

Without types, the probability of winning one of the two items is now simply the probability P_W:

$$P_W(x) = P\left(x \geq X_{(n-1)}\right) + \left[P\left(x \geq X_{(n-2)}\right) - P\left(x \geq X_{(n-1)}\right)\right]$$

$$= \binom{n-1}{n-1}\left(F(x)\right)^{n-1}\left(1-F(x)\right)^{n-1-n+1} + \binom{n-1}{n-2}\left(F(x)\right)^{n-2}\left(1-F(x)\right)^{n-1-n+2}$$

$$= \left(F(x)\right)^{n-1} + (n-1)\left(F(x)\right)^{n-2}\left(1-F(x)\right)$$

In the current example of three bidders with [λ]-exponential value distribution:

$$P_W(x) = \left(1-e^{-\lambda x}\right)^2 + 2\left(1-e^{-\lambda x}\right)\left(e^{-\lambda x}\right)$$

$$= 1 - 2e^{-\lambda x} + e^{-2\lambda x} + 2e^{-\lambda x} - 2e^{-2\lambda x}$$

$$= 1 - e^{-2\lambda x}$$

Using equation (3), the optimal bid b^* is:

$$b^*(v) = \frac{\int_0^v x P_W'(x)dx}{P_W(x)}$$

$$= \frac{\int_0^v 2\lambda x e^{-2\lambda x} dx}{1-e^{-2\lambda x}}$$

$$= \frac{1}{\left(1-e^{-2\lambda v}\right)}\left[\frac{1}{2\lambda} - \left(\frac{1}{\lambda}v + \frac{1}{2\lambda}\right)e^{-2v\lambda}\right]$$

$$= \frac{1-(2\lambda v+1)e^{-2\lambda v}}{2\lambda\left(1-e^{-2\lambda v}\right)} \tag{12}$$

We can now answer the question if the auction participants tend to bid higher or lower in a scenario *with* types versus a scenario *without* types. Figure 3.3 shows how the optimal bid depends on the particular value of the bidder (assuming all bidders derive their value from a [1]-exponential distribution) and demonstrates that the answer to the question above depends on how high their value is. The result can be explained intuitively as follows: In an auction without types, a bidder that has a low valuation for the item must bid a relatively high percentage of his value in order to stand a chance of winning an item that will likely go to a bidder with a higher valuation. With the introduction of types, however, the low-value bidder may have the good fortune that all bidders with higher values are of the same type, but different than his own type, and that he receives the second flotilla slot even though there are bidders with higher valuations for the slot. The low-value bidder therefore has an incentive to shade his bid more and, as a result, the dashed line is above the full line for low valuations.

In contrast, if a bidder has a high valuation and there are no types, the bidder will not bid a high percentage of his value because he is content with receiving the second slot. With multiple types, however, the high-value bidder must consider the possibility of entering the second highest bid and discovering that the highest bidder is of the same type and that he has lost the chance for an impression. Therefore, the high-value bidder will try to secure the top flotilla slot by bidding a higher percentage of his value, as indicated by the full line being above the dashed line for high values in Figure 3.3.

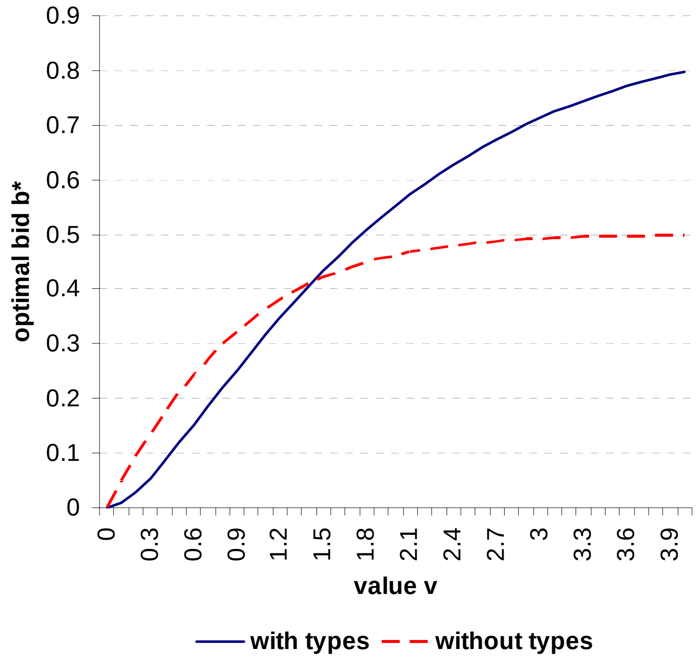

Figure 3.3: Optimal first price bids with and without types dependent on value (λ=1).

Another observation that can be made from Figure 3.3 is that the bidding functions flatten out for high values. In fact, by letting v go to infinity in equation (10) and (12), we can find the highest possible bid:

$$\text{As } v \to \infty,\ \frac{\frac{7}{2\lambda} - \left(\frac{4}{\lambda} + 4v\right)e^{-\lambda v} + \left(\frac{1}{2\lambda} + v\right)e^{-2\lambda v}}{4 - 4e^{-\lambda v} + e^{-2\lambda v}} \to \frac{\frac{7}{2\lambda}}{4} = \frac{7}{8\lambda}\text{; and similarly,}$$

$$\text{As } v \to \infty,\ \frac{1-(2\lambda v+1)e^{-2\lambda v}}{2\lambda\left(1-e^{-2\lambda v}\right)} \to \frac{1}{2\lambda}$$

Therefore, the highest possible bid with types is $7/(8\lambda)$ and without types is $1/(2\lambda)$. Since the expected value of a [λ]-exponential random variable is $1/\lambda$, we can conclude that bidders in auctions with and without types will never bid above 87.5% and 50% of the expected item value, respectively – even if their own value is far above the average.

It is worthwhile to note that the introduction of types does not change the bidding behaviour in a second price auction where truth-telling remains the optimal strategy. However, it does affect the expected revenue of the seller which will be demonstrated in a later section.

Bidding behaviour is also heavily impacted by the value distribution. To illustrate this point, the following calculations show the optimal bidding function if the bidders' values are drawn from a uniform distribution between *a* and *b*. According to equation (4):

$$b^*(v) = \frac{\int_a^v 2\left(\frac{b-2a+x}{b-a}\right)x\frac{1}{b\text{-}a}dx}{\left(\frac{v-a}{b-a}+1\right)^2}$$

$$= \frac{\int_a^v 2\left(\frac{(b-2a)x+x^2}{(b-a)^2}\right)dx}{\left(\frac{b-2a+v}{b-a}\right)^2}$$

$$= \frac{2\int_a^v \left((b-2a)x+x^2\right)dx}{(b-2a+v)^2}$$

$$= \frac{2\left[\left(\frac{b}{2}-a\right)x^2+\frac{1}{3}x^3\right]_a^v}{(b-2a+x)^2}$$

$$= \frac{2\left[\left(\frac{b}{2}-a\right)v^2+\frac{1}{3}v^3-\left(\frac{b}{2}-a\right)a^2-\frac{1}{3}a^3\right]}{(b-2a+v)^2}$$

$$= \frac{\frac{4}{3}a^3-ba^2+(b-2a)v^2+\frac{2}{3}v^3}{(b-2a+v)^2}$$

Figure 3.4 compares the bidding functions resulting from the previous [1]-exponential distribution with the bidding functions resulting from a [0.5]-exponential distribution, a [0.5,1.5]-uniform distribution and a [0.5,3.5]-uniform distribution. The thin lines and the thick lines in Figure 3.4 are bidding functions that stem from value distributions with the same expected value, respectively.

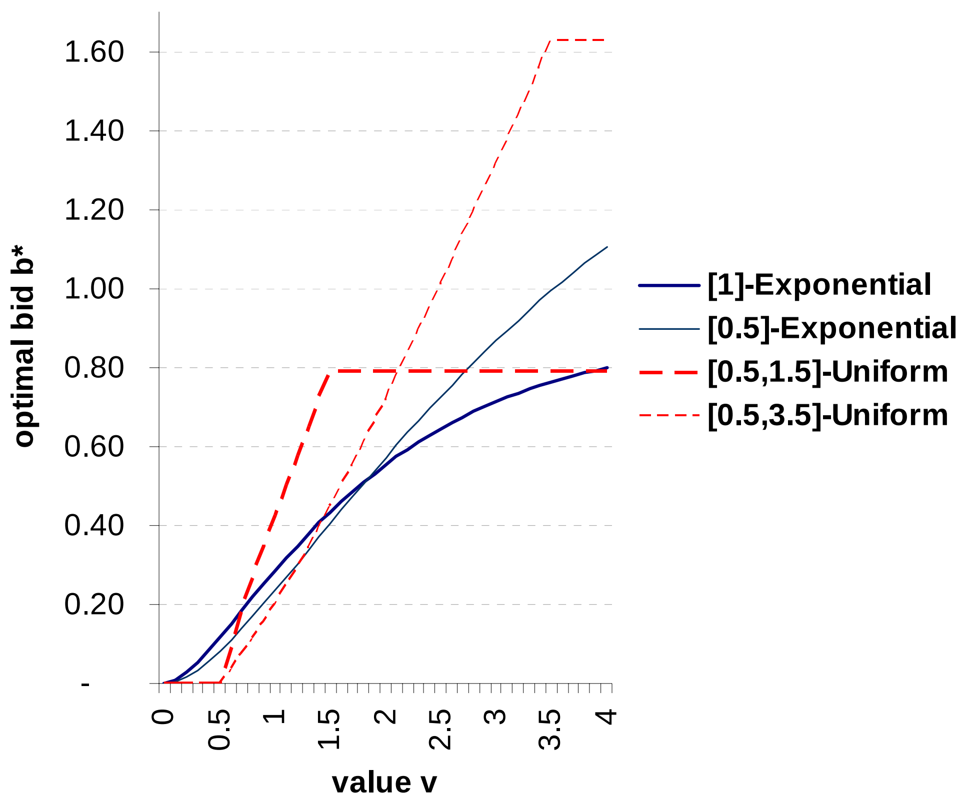

Figure 3.4: Optimal first price bids given four different value distributions.

Figure 3.4 provides two main insights: first, bidders whose valuation of the item is close to the mean bid more aggressively if the distribution has a lower standard deviation. This is intuitive because bid shading significantly reduces the probability of winning if most other bidders are likely to have values close to their own value. Second, and perhaps less intuitive, if the mean of the bidders' value distribution for the item increases, the bid function does not always move upwards. In particular, for lower values, the optimal bid function is now below the bid function given the lower mean value distribution function; only for higher values is the optimal bid function above the previous bid function. This is best explained by the fact that in the high-mean scenario, bidders that have low valuations are unlikely to attain the item unless the "types" work in their favour and they attain the second slot because all other bidders are of the same type but different from them. For bidders with higher valuations, the increase in the mean fuels the attempt to secure the top flotilla slot.

3.7.3. Seller Revenue

While the discussion on bidding behaviour has provided some interesting insights, the more important question is what impact the introduction of types has on the expected revenue and how the seller's expected revenue depends on the value distributions of the bidders as well as the number of bidders. Moreover, the ultimate question is whether the seller should set a list price or run an auction to sell the flotilla slots.

Using a similar approach as in the derivation of equation (5) and (6), the expected revenue, if no types exist, is:

$$nE[R] = n\int_0^{\infty} \left[b^*(x)P_W(x)\right] f(v)dv$$

Further, with 3 bidders for two items without types, the probability of winning is simply 1 - the probability of being the third highest bidder:

$$P_W(x) = 1 - R(x)^2 = 1 - e^{-2\lambda x}$$

Now, using equation (12):

$$nE[R] = 3\int_0^{\infty} \left[\frac{1-(2\lambda v+1)e^{-2\lambda v}}{2\lambda}\right] \lambda e^{-\lambda v} dv$$

$$= 3\int_0^{\infty} \left[\frac{e^{-\lambda v}}{2} - \lambda v e^{-3\lambda v} - \frac{e^{-3\lambda v}}{2}\right] dv$$

$$= 3\left(\left[-\frac{e^{-\lambda v}}{2\lambda}\right]_0^{\infty} + \left[\frac{1}{3} v e^{-3\lambda v}\right]_0^{\infty} + \left[\frac{1}{9\lambda} e^{-3\lambda v}\right]_0^{\infty} + \left[\frac{1}{6\lambda} e^{-3\lambda v}\right]_0^{\infty}\right)$$

$$= 3\left(\frac{1}{2\lambda} - \frac{1}{9\lambda} - \frac{1}{6\lambda}\right) = \frac{2}{3\lambda}$$

By comparing this amount to the expected revenue with types of $7/(12\lambda)$, it is evident that the introduction of types has decreased the expected revenue by $1/(12\lambda)$, or 12.5%, from the standard auction without types. This relatively small decrease in expected revenue may be surprising since types have introduced the possibility of only one of two items being sold.

However, this finding can be explained by the fact that types have also increased the competition for the top flotilla slot (since order matters).

With regards to the type of auction, the previous numerical example demonstrated that the choice between a first or second price auction has no impact on the expected revenue. However, this will only hold as long as the conditions of the revenue equivalence theorem hold. In light of the large number of auctions that will occur over time, risk neutrality of bidders seems to be a valid assumption. Advertisers will be less concerned with the outcome of individual auctions but will instead try to maximize their expected utility over many auctions, displaying risk neutrality. In comparison, bidders who are participating in a one-time auction for an expensive item are typically risk averse. If they are not able to secure the item, they do not have a second chance. This danger outweighs the drawback of a smaller profit which makes them bid higher than a risk neutral bidder would.

And, as mentioned in Chapter 2, advertisers have their own private values regarding how much an impression is worth. Without primary research, it is unclear whether advertisers draw their values from the same distribution; however, if we assume that they do and implement *any* auction model under which payments are dictated by bids alone, the expected revenue will always be the same.

The earlier numerical example also showed that given three bidders with [λ]-exponential value distributions, setting a list price will – on average – result in 60% more revenue than selling the item by auction ($0.93/\lambda$ vs. $0.58/\lambda$). However, Figure 3.5 shows that this superiority reverses as the number of bidders increases. Once more than five bidders exist in the market place, an auction will result in more revenue for the seller than a list price system. As mentioned before, an auction model that includes a reservation price will result in even higher revenue than a pure auction or list price system. This knowledge will have an impact on the recommendations given in Chapter 4.

Another observation that can be made from Figure 3.5 is the rapid increase in revenue when the number of bidders increases from three to about ten with decreasing rate of returns thereafter. This holds true regardless of λ; in particular, the expected revenue is always a multiple of $1/\lambda$. For example, if λ was 0.5 instead of 1, the seller could expect double the revenue (for each respective number of bidders). This sharp increase on the left hand side of the graph is a reflection of two negative effects that a small number of bidders have on revenue (from the seller's perspective): first, given a second price auction, there is a significant chance that there will only be a small number of bidders with high values who will merely have to pay the bids of low-value bidders. Congruently, in a first price auction, high-value bidders are aware of their high chance of winning the item even with a low bid. Second, a low number of bidders increases the chance that all bidders are of the same type and, for example, only one flotilla slot is sold.

Figure 3.5: Expected revenue of auction vs. list price (λ=1).

Because the actual distribution of advertisers' values for advertising slots is unknown, a number of distributions were tested and are displayed in Figure 3.6. The graph compares the [1]-exponential distribution to a [0.5,1.5]-uniform and a [1,0.1]-normal distribution in a second price auction setting.

The difference in expected revenue is particularly large when a small or large number of bidders exist in the market. When only a small number exists, there is a significant chance that winners will only have to pay amounts that are at the low end of the distribution. In the case of an exponential distribution, the left-hand tail end is very wide and bidders may only have to pay a very small amount. In comparison, with the [0.5, 1.5]-uniform distribution, bidders will never pay less than 0.5 and, even with only three bidders, the seller can expect to receive revenue of almost 1 (without types, this would be at least 1 but with types the probability of paying 0 exists).

Similarly, with many bidders in the market, the probability increases that bidders will have to pay prices at the high end of the distribution. This is particularly profitable for the seller if the distribution has a wide open tail on the right-hand; as a result, bidders whose values are derived from an exponential distribution will end up paying prices far beyond a scenario with bidders whose values are derived from a uniform distribution.

Following the reasoning above, distributions with higher variances should result in higher revenue as the number of bidders increases. Given a large number of bidders, revenue is increasing in variance. Since the variance of an exponential distribution is $1/\lambda^2$, the variance of a uniform distribution is $(b\text{-}a)^2/12$ (where a and b are the upper and lower bound of the distribution), and the variance of a normal distribution is σ^2 (where σ is the standard deviation of the distribution), the variance of the three distributions shown in Figure 3.6 are 1, 0.0833 and 0.01, respectively (Ross, 2002). However, all three distributions have a mean of 1.

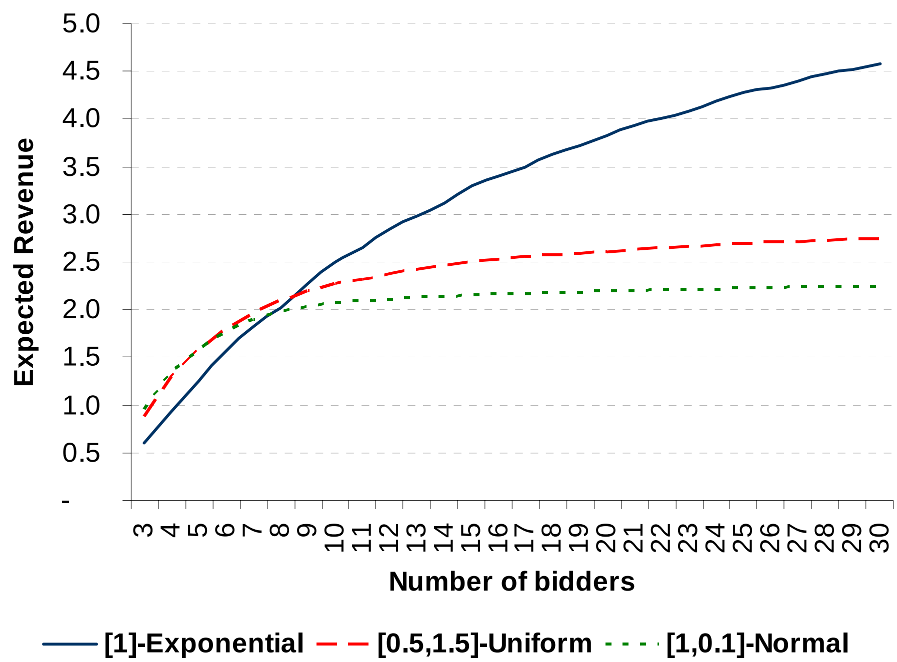

Figure 3.6: Expected revenue of second price auction for three value distributions.

To explore the relationship between revenue and variance of the value distributions further, Figure 3.7 compares three normal distributions, each with a mean of 1 but with standard deviation of 0.1, 0.2 and 0.4 respectively. The same effect as described above is clearly demonstrated: distributions with higher variances result in higher revenue as the number of bidders increases.

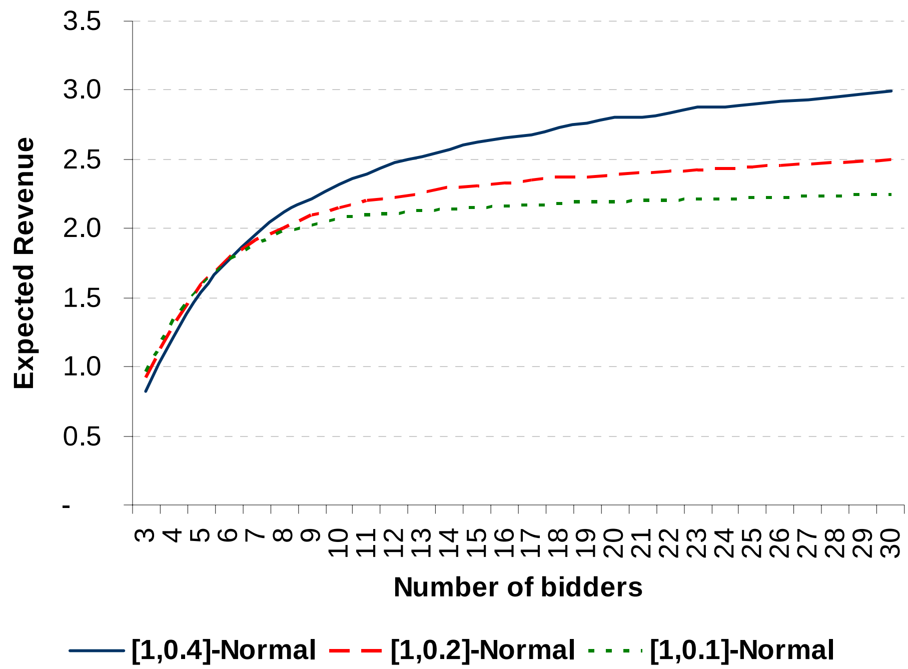

Figure 3.7: Expected revenue of second price auction for three normal value distributions.

3.8. Conclusions and Caveats

Even though the analytical model of this section does not capture all complexities of the targeted television advertising system, it provides a model that allows a mathematical analysis and several insights. The following sections highlight the main insights and how they should impact the seller's decision making.

3.8.1. Auction Model versus List Price System

The primary decision that the seller needs to make is whether to sell the flotilla slots using a list price system or whether to sell them through an auction. From a revenue perspective, an auction will lead to a better outcome as soon as the number of bidders exceeds a handful. This statement remains true regardless of the actual value distribution of the bidders, as shown in Appendix C. Furthermore, the list price system will perform even more poorly as the set list price deviates more from the optimal. This deviation may be intentional to achieve rounded prices. Or the deviation may be unintentional as a result of inaccurate approximations of the bidders' value distribution function. The seller should be aware that setting the list price too low decreases the revenue faster than setting the price too high. This can be concluded by comparing the steep slope before the optimal list price in Figure 3.2 with a flatter slope after the optimal list price. As a result, the seller is better off setting a higher price if they are uncertain about the optimal list price. Regardless of the cause and direction of the deviation from the optimal list price, any deviation will further decrease the expected revenue from the above results.

As mentioned before, there are other factors that will impact the choice between the two systems. A list price system will provide advertisers with the comfort of knowing that they have secured a certain number of impressions but burdens the seller with the obligation to fulfill this implied promise of service as well as to determine an appropriate price. The auction system, on the other hand, will guarantee an efficient allocation of flotilla slots – removing that burden from the seller – but it may result in unsatisfied advertisers who have become accustomed to the status quo. Ultimately, there are some strategy decisions that will have to be made in this regard by the seller but Chapter 4 will provide some further economic insight into this decision.

Should the seller decide to proceed with an auction model, the next major decision will lie in the choice of auction model. This section has provided evidence that the choice of auction model will not impact the expected revenue given that the benchmark model assumptions hold. The one feature of the auction system that will impact the profits, however, is the inclusion of a reservation price. In particular when there are a low number of bidders, a reservation price can increase the expected revenue significantly as will be demonstrated in Chapter 4.

3.8.2. Number of Types

The analysis of this chapter is based on only two "types", whereas the actual bidders in targeted television advertising system will be of several more types. Depending on the permitted detail of specification, there could be 18 types (two gender and nine age cohorts) or even 180 types (two gender, nine age cohorts and ten income cohorts). This might lead the reader to conclude that with 180 types, the number of bidders may have to be as high as two or three times that amount (or 500) for an auction model to outperform a list price system. However, there is another variable that determines the optimal system – namely the number of flotilla slots.

In the analytical model studied above, two types of bidders competed for two flotilla slots which means that, in effect, each type was only competing with other members of its type. Since advertisers do not desire the same viewers as advertisers of the other type and there are flotilla slots for both types, they do not compete with each other. If we were to add a third type, the situation would change. An advertiser trying to secure a flotilla slot would not only have to outbid advertisers of the same type but would also have to ensure that his bid is higher than the highest bid of at least one of the two other types. With a Vickrey auction, this would not affect the bidding behaviour (always truth-telling) but it would have an effect on seller revenue. In fact, assuming an equal number of total bidders, the more types that exist, the more revenue the seller can expect. On the one end of the spectrum, if the number of types is equal to the number of bidders, then the seller will receive the same expected revenue as in a system without types, namely twice the value of the third highest bidder (assuming two flotilla slots). On the other end, with only one type of bidder, only one flotilla slot can be sold and the seller will only receive the value of the second highest bidder once. So if the targeted television advertising system consisted of four flotilla slots, the model predicts that if the number of bidders exceeds about ten, higher revenue can be expected in an auction system compared to a list price system. This question will be addressed in greater detail in Chapter 4.

3.8.3. Value Distributions

This chapter also illustrated the impact that the advertisers' value distributions have on expected revenue. If there are sufficient advertisers participating in the auction, the seller would benefit if these advertisers draw their values from a distribution with a high rather than a low standard deviation, and vice versa if there are less bidders participating. It is hard to predict with great accuracy what type of value distribution advertisers will have without undertaking primary market research. In addition, the additional complexities of the targeted television advertising system that were mentioned earlier may impact expected revenue so more detailed findings will be reported in Chapter 4.

3.8.4. Additional Complexities

The actual system under which the seller will operate encompasses three additional complexities which were not captured in this section; namely, each auction will be part of a series of auctions, there is a link between receiving a flotilla slot and actual impressions, and types may overlap. Overlapping types can occur as, for example, one advertiser targets females of ages 18-34 while another advertiser is targeting females of ages 25-49. With respect to impressions, a flotilla slot does not guarantee a certain static number of impressions in the actual targeted television advertising system. These complexities demand more detailed rules for any considered auction model and complicate the analysis significantly. This is the topic of the next chapter.

4. SMART MARKET AUCTION DESIGN

4.1. INTRODUCTION

The actual system under which the targeted television advertising will operate has some significant differences to the model described in Chapter 3 and this will add a number of complications. Because these complexities prohibit the mathematical derivation of closed form solutions, this section uses Monte Carlo simulations to derive insights to the problem. The complexities that need to be added to the previous model of Chapter 3 are the impact of multiple auctions in a row, the link between impressions and a flotilla slot and the existence of overlapping types.

4.1.1. Series of Auctions

One dimension of the problem omitted in the Chapter 3 model was that these auctions will actually occur repeatedly over time as avails occur in the cable operator's video stream. Advertisers typically engage in advertising campaigns that span weeks and will therefore participate in a large number of these auctions. Such a series of auctions can have a negative effect on a first price auction system. As became evident in the previous chapter, bidders in a first price auction shade their bids, bidding an amount less than their actual valuation. Since all bidders are shading their bids, a winning bid may actually be lower than the value of a losing bidder. In a one-time auction, the losing bidder cannot react to this knowledge because the outcome is revealed *after* the auction. In a series of auctions, however, the losing bidder can increase his bid to win an item in the following auction. But this will cause the previous winner to react by increasing his bid as well, which in return will trigger other reactions. Consequently, the auction system becomes very unstable with constantly changing bids.

4.1.2. Link between Impressions and Flotilla Slot

The previous model assumed that advertisers have a certain value for a flotilla slot. In truth, however, advertisers have a certain value for impressions and they can only acquire impressions by securing a flotilla slot. The reason this distinction is important is that the number of impressions derived from a flotilla slot is not constant over time and differs by type. For example, at a particular time, there may be five adult females, ten adult males and three viewers under the age of 18 in the system available to view an advertisement. As a result, an advertiser targeting adult females who secures a flotilla slot will receive five impressions, an advertiser targeting adult males who secures a flotilla slot will receive ten impressions, and an advertiser targeting viewers under the age of 18 who secures a flotilla slot will receive only three impressions. Assuming two flotilla slots, the number of total impressions available depends on which advertisers win slots in the flotilla.

4.1.3. Overlapping Types

In the previous chapter, only two types were assumed (males and females). As mentioned earlier, a more accurate model needs to be able to accommodate a much larger number of types. While this could be done quite easily using the model of chapter 3, the full model must also be able to accommodate overlapping types. For example, there may be three types: males, females, and a type encompassing both males and females. Under this scenario, let us assume that an advertiser targeting males secures the top flotilla slot. Clearly, all other advertisers targeting males have lost interest in the bottom flotilla slot and all advertisers targeting females are still as interested in the bottom flotilla. The difficult question to answer is how much value advertisers targeting males and females together will still have for the bottom flotilla slot. If they secure the slot, they will only reach females, or half their target.

The effect of this on the evaluation of the flotilla slot by these types of advertisers is uncertain: they may be willing to pay more than half the original bid if they value female impressions higher than male impressions; or, they may be willing to pay less than half the original bid if the value female impressions lower or are generally a balance between the two genders.

4.1.4. Motivating Example

To help understand the complexities of this model, a motivating example has been constructed and will be used to demonstrate the difficulty of the problem. In this example, there are four advertisers (represented in the columns of Table 4.1), one targeting males of ages 25 to 55, one targeting males between the age of 18 to 49, one targeting all females and one targeting all males. The rows of Table 4.1 represent viewers in the system by gender and age. The table totals the number of impressions available to each advertiser and multiplies this with the submitted bid, which is the amount advertisers are bidding to pay *per impression*, to calculate the total payment each advertiser is willing to make for a flotilla slot (assuming that he receives all appropriate viewers). Note that the amount of bid shading (difference between submitted bid and the value of the object) has been chosen arbitrarily. As an example, the advertiser targeting males of ages 18 to 49 faces three appropriate viewers (Male 18, Male 30 and Male 20) and values each possible impression at 0.55; since he has submitted a bid of 0.40 per impression, he is willing to pay 1.20 for a flotilla slot if he receives all three impressions.

		Advertisers targeting			
		Males 25-55	Males 18-49	Females	Males
Viewers	Male 18	0	1	0	1
	Male 50	1	0	0	1
	Male 30	1	1	0	1
	Male 55	1	0	0	1
	Male 20	0	1	0	1
	Femal 40	0	0	1	0
Total Impressions		3	3	1	5
Value of Object		0.65	0.55	0.60	0.05
Submitted Bid		0.50	0.40	0.45	0.01
Total Payment		1.50	1.20	0.45	0.05

Table 4.1: Example 1 – Full Information Matrix.

The challenge for the seller is to determine which two of the four advertisers receive flotilla slots (as well as their order) and how much to charge them for this slot. Since the goal is to promote truth-telling in the bidding behaviour (to create a stable system and increase revenue), advertisers must be charged a fair amount else they will revert to bid shading. To illustrate the difficulty of this problem, the following paragraphs describe various ways of conducting the auction and the resulting issues that could arise.

First price auction by submitted bid – Using the methodology of Chapter 3, the seller could declare the highest bidder, Males 25-55, to be the winner of the auction and have them pay their bid amounts of 0.50 per impression. Similarly, the next flotilla slot would go to advertiser Females for cost of 0.45 per impression. In a one-shot auction, this might be a valid outcome; however, the repetition of this auction over time will cause the system to become unstable. Even though a future avail will not have an identical composition of viewers available, advertisers can learn from the previous auction outcome and adjust their bids accordingly. For example, advertiser Males 18-49 will realize that if he increases his bid from 0.40 to 0.51, he can capture three impressions at a price that is still below his value of 0.55. Facing increased bids from his competitors, advertiser Males 25-55 will react by increasing his bid, too. However, if the other advertisers reduce their bids, or end their campaign, advertiser Males 25-55 will realize that he is paying more than necessary to win his flotilla slot, therefore reduce his bid and allow the cycle to start over again. The result of such an unstable system is an increasing amount of effort that advertisers put into gaming the system. This effort comes at a price which is typically passed on to the seller (Edelman and Ostrovsky, 2007) and may also cause advertiser frustration.

Second price auction by submitted bid – To solve this problem, the seller could declare the highest bid the winner of the top flotilla slot and make the winner pay the next highest bid of his type. Realizing the strategy dominance of truth-telling, advertisers' will submit bids equal to their respective values of the objects. Table 4.2 provides the updated bids and total payments for this example (the type and number of viewers has not changed).

	Advertisers targeting			
	Males 25-55	Males 18-49	Females	Males
Total Impressions	3	3	1	5
Submitted Bid	0.65	0.55	0.60	0.05
Total Payment	1.95	1.65	0.60	0.25

Table 4.2: Example 1 – Condensed matrix with updated bids.

When trying to apply this methodology to the example shown in Table 4.1, we face the earlier described problem of overlapping types: advertiser Males 25-55 is the highest bidder (0.65), will therefore receive the top flotilla slot, and must pay the bid of the second highest bidder of his type; however, there is no bidder of *precisely* his type. A solution to this problem will be presented later but for now let us assume that for billing purposes any multiple types with overlapping demographics will be considered to be of the same type. As a result, advertiser Males 25-55 is charged the bid of advertiser Males 18-49, namely 0.55 for each impression for a total of 1.65. Similarly, the bottom slot of the flotilla goes to the second highest bid, advertiser Females with a bid of 0.60, for the price of zero since there is no other advertiser of this type. Clearly, this is not a good outcome since there are two other advertisers that are willing to pay more than zero, namely 0.55 and 0.05 per impression. As mentioned in an earlier discussion, if more types than flotilla slots exist, advertisers will start competing not only against other advertisers of their type but against *all* other advertisers. Therefore, advertiser Females should be charged the next highest bid of 0.55 per impression for a total of 0.55.

The problem with applying the methodology of Chapter 3 in this more complex setting is the fact that there now exists a link between a flotilla slot and the resulting impressions – ignoring this link can cause inefficiencies in raising revenue. For example, advertiser Males 18-49 is willing to pay more than advertiser Females for the second flotilla slot because although he submitted a slightly smaller bid per impression, there are many more impressions available to him. Ignoring this fact leaves "money on the table" because an advertiser with a large number of appropriate viewers may be left out of the flotilla because is bid was slightly below another bidder who has very viewer appropriate viewers in the system. While the seller may receive more revenue per impression if the latter wins, the total revenue will be less than if the former advertiser had won.

Second price auction by next highest total payment – An obvious solution to this problem is to determine winners by total payment instead of submitted bid. In the above example, advertiser Males 25-55 still wins the top flotilla slot (he has the highest total payment of 1.95) and is charged 1.65 (the next highest total payment of his type). Before determining the winner of the bottom flotilla slot, we need to update Table 4.1 since some of the viewers have been captured by the winning advertiser and are no longer available. Table 4.2 shows the updated table with changes highlighted in bold.

		Advertisers targeting			
		Males 25-55	Males 18-49	Females	Males
Viewers	Male 18	0	1	0	1
	Male 50	1	0	0	**0**
	Male 30	1	**0**	0	**0**
	Male 55	1	0	0	**0**
	Male 20	0	1	0	1
	Femal 40	0	0	1	0
Total Impressions		3	**2**	1	**2**
Submitted Bid		0.65	0.55	0.60	0.05
Total Payment		1.95	**1.10**	0.60	**0.10**

Table 4.3: Example 1 – Revised full information matrix.

The table shows that the bottom flotilla slot will go to advertiser Males 18-49 for the price of 0.60 (again, this advertiser is competing against all other bidders since the number of types exceeds the number of flotilla slots).

Unfortunately, our assumption that, for billing purposes, overlapping types will be treated as the same type creates unfair results. While advertiser Males 25-55 will end up paying 0.55 (1.65/3) per impression, advertiser Males 18-49 will only pay 0.30 (0.60/2) per impression even though they are targeting similar demographics and their bid differ by only 0.10. Knowing that these kinds of outcomes are possible, advertiser Males 25-55 may have an incentive to shade his bid; for example, had he reduced his bid to 0.50, advertiser Males 18-49 would have won the top flotilla slot, paid 1.50 for three impressions, and advertiser Males 25-55 would have received the bottom flotilla slot and paid only 0.60 for two impressions. In other words, this type of bid shading would have left advertiser Males 25-55 with 33% less impressions but with a price reduction of over 45% per impression.

Second price auction by third highest total payment – One possible solution to this problem would be to charge both winning advertisers the amount of the third highest total payment. In this case, the winner is still advertiser Males 25-55 and the second flotilla slot goes to advertiser Males 18-49. Both pay 0.60, which implies a cost per impression of 0.20 and 0.30 respectively. However, this is a poor result since advertiser Males 18-49 is willing to pay 0.55 for each impression, including the viewer Male 30 who is captured by advertiser Males 25-55 for only 0.20.

In addition to the above issues, it may not be appropriate to use the submitted bid of advertiser Males 18-49 once advertiser Males 25-55 has secured the top flotilla slot. The original bid reflected the desire to reach all males in the appropriate age group. However, the winning advertiser captured all males above 25 years of age. Consequently, advertiser Males 18-49 can only capture males between the ages of 18 to 24 which may impact the price he is willing to pay per impression. Ignoring this issue may lead to undesirable consequences: advertisers targeting broad demographic groups, such as all males, may reduce their bid in anticipation of receiving impressions only from less sought-after demographic groups. Alternatively, these advertisers may break down their bid into multiple bids for more specific demographics. This segmentation can lead to inefficiencies as it is possible that the same advertiser receives multiple flotilla slots when one slot would have been sufficient.

One solution to this problem would be to allow the advertiser to indicate whether his bid reflects his expectation to receive a representative share of each subgroup within his target demographic (i.e., no advertiser of an overlapping type may be in any of the flotilla slots above him) or whether he values all demographic subgroups equally and is indifferent toward the proportion of demographic subgroups he receives within his target group.

Another issue illustrated above is the fact that bidders may be able to achieve substantial cost savings (or increase in profit) in return for a slightly lower number of total impressions. In the above example, advertiser Males 25-55 could game the system to pay 0.60 for two impressions, rather than 1.65 for three impressions. Since he values each impression at 0.65, his total profit would increase to 0.70 (0.65*2 - 0.60), from 0.30 (0.65*3 - 1.65). Alternatively, a scenario could be constructed in which a bidder may be able to achieve cost savings *per impression* (receiving fewer, but cheaper impressions) while his overall profit (value minus cost) decreases. Given that a series of auctions will occur over time, the argument could be made that a bidder may be willing to sacrifice overall profit in a single auction because over time he can still achieve all his desired impressions, but at a lower cost. To deal with this problem, the following assumption was made for the remainder of this chapter: within the range of the possible number of impressions, advertisers value the number of impressions uniformly. In other words, the first impression they receive is as valuable to them as the second impression, which is as valuable to them as the last impression they receive. Therefore, within any given auction, their goal is to maximize *total* profit, not profit per impression. However, the seller still needs an auction system that charges the advertisers appropriately.

Vickrey-Clark-Groves mechanism – The mechanism that is known in the literature to induce truth-telling and allocate items efficiently in a multi-item situation is the VCG mechanism. Under this system, bidders are charged the opportunity cost that they impose on other bidders. In the context of this motivating example, the VCG mechanism would work as follows: first, the combination of bidders yielding the highest revenue would be selected, in this case advertisers Males 25-55 and Males 18-49. To determine the payment of advertiser Males 25-55 (received the top flotilla slot), we find the additional payments that other advertisers would have made without the former being present in the system. In this case, advertiser Males 18-49 and advertiser Females would have won the auction, which means that advertiser Males 18-49 would have gained one more impression (the 30 year old male) at 0.55 and advertiser Females would have gained a flotilla slot to capture one impression at 0.60. Therefore, advertiser Males 25-55 would be charged 1.15 (0.55 - 0.60). Similarly, advertiser Males 18-49 would be charged at 0.60 (without him, advertiser Females could capture a flotilla slot and one impression).

Unfortunately, the VCG mechanism has some downsides as well. In particular, it is vulnerable to *shill bidding* in which advertisers submit artificial bids to influence the outcome positively for them. The following example demonstrates this possibility. In Table 4.4 there are three advertisers and each has indicated that they are not willing to accept overlapping types in the flotilla slots above them. According to the VCG mechanism, advertiser Males 25-55 wins the top flotilla slots and advertiser Females wins the bottom slot. Advertiser Males 25-55 must pay 1.25 since he is depriving advertiser Males of a flotilla slot and advertiser Females pay zero since his presence does not deprive any other advertiser of impressions.

		Advertisers targeting		
		Males 25-55	Females	Males
Viewers	Male 18	0	0	1
	Male 50	1	0	1
	Male 30	1	0	1
	Male 55	1	0	1
	Male 20	0	0	1
	Femal 40	0	1	0
Total Impressions		3	1	5
Submitted Bid		0.65	0.05	0.25
Total Payment		1.95	0.05	1.25

Table 4.4: Example 2 – Full information matrix.

In this situation, advertiser Males 25-55 has an incentive to submit a shill bid targeting Males 18-24 as shown in Table 4.5.

		Advertisers targeting			
		Males 25-55	Males 18-24	Females	Males
Viewers	Male 18	0	1	0	1
	Male 50	1	0	0	1
	Male 30	1	0	0	1
	Male 55	1	0	0	1
	Male 20	0	1	0	1
	Femal 40	0	0	1	0
Total Impressions		3	2	1	5
Submitted Bid		0.65	0.65	0.05	0.25
Total Payment		1.95	1.30	0.05	1.25

Table 4.5: Example 2 – Full information matrix with shill bid.

As a result, advertiser Males 25-55 wins his real bid and his shill bid targeting Males 18-49. For his real bid, he must now pay only 0.05 (without him, the flotilla slots would go to his shill bid and advertiser Females), and for his shill bid he must also pay 0.05 (without his shill bid, the slots go to his real bid and advertiser Females). As a result, he now only pays 0.10

instead of 1.25 and he gets more impressions. Obviously, this would be a very negative consequence for the seller.

Some of these concerns have led other companies such as Google and Yahoo! to invent their own auction rules (Edelman, Ostrovsky and Schwarz, 2007). The following section describes a number of proposed auction models that are aimed to solve the described challenges of implementing an auction into this system while maximizing revenue to the seller.

4.2. PROPOSED SMART MARKET AUCTIONS

4.2.1. Introduction and Notation

To enable a clear description of the auction models, the following notation will be used:

- H Number of households.
- A Number of advertisers.
- S Number of slots in the flotilla.
- V_a Value per impression of advertiser a.
- $T_{a,h}$ Is 1 if household h is in target audience of advertiser a, else 0.
- Z_a Is the sum of the target audience multiplied with the respective value for each advertiser a; $Z_a = V_a \sum_h T_{a,h} \quad \forall a$.

These models will assume that advertisers bid truthfully so that their bids are equal to their values for impressions. As mentioned before, the new targeted television advertising system will run as a smart market in which the outcomes of the auctions are determined automatically through a computer program. Advertisers are not participating directly in these auctions (there will be hundreds of these automatic auctions a day) but instead advertisers provide bids that will be taken into account automatically in multiple small auctions. Therefore, interactive auction models such as ascending auctions are not part of this analysis and only sealed bid auctions are considered. The proposed auction models are first introduced and then compared computationally in a later section.

4.2.2. Second Largest Ad

A simple second price auction implementation could emulate the Google system described by Edelman, Ostrovsky and Schwarz (2007). In such a system, the advertiser that will create the highest revenue (the highest Z) will receive the top flotilla slot and pay the second highest revenue raised by any advertiser (the second highest Z). After the first winner has been determined, all households that have been captured by this advertiser are removed from the system. Because they are no longer accessible to other advertisers, all Z values will have to be recalculated. At this point, the winner of the next flotilla slot can be determined using the same method as before (highest Z wins and pays second highest Z). This cycle continues until all flotilla slots have been filled. The algorithm Second Largest Ad shows these steps in a more precise mathematical form.

Algorithm Second Largest Ad

---

1. Initialize
   - 1.a. Sort all Z from highest to lowest so Z_1 is the highest
   - 1.b. Sort all V and T to be in congruence with the ordered Z's; i.e. $V_a \sum_h T_{a,h} = Z_a \quad \forall a$
2. Determine winner and payment
   - 2.a. Z_1 receives first free flotilla slot
   - 2.b. Z_1 must pay Z_2
3. Remove previous winner
   - 3.a. Set $Z_1 = 0$
   - 3.b. Set $T_{2,\mathrm{h}}$ .. $T_{\mathrm{A,h}}$ to zero for all h where $T_{1,\mathrm{h}} = 1$
   - 3.c. Recalculate $Z_2 \ldots Z_{\mathrm{A}}$
   - 3.d. Resort $Z_1 \ldots Z_{\mathrm{A}}$ (and all V and T accordingly)
4. Terminate algorithm
   - 4.a. If another free flotilla slot exists
     - 4.a.1. Return to step 2
   - 4.b. Else
     - 4.b.1. Terminate algorithm

The problem with this algorithm is its failure to address the effect of types; i.e., the advertiser placed in the top flotilla slot may capture most of the next highest advertiser's target viewers or he may capture none of them. A small example will demonstrate why this matters.

Imagine a situation with two flotilla slots and three advertisers, two of which have very high values (advertiser A and B, with A slightly higher than B) and one of very low value (advertiser C). If advertiser A and B are targeting the same demographic, advertiser A will capture the first flotilla slot and pay the Z of advertiser B. Since the target demographic of advertiser B has now been captured, advertiser B will no longer have interest in the auction and advertiser C will receive the second flotilla slot for a price of zero. In this scenario, the algorithm does promote truth-telling as the advertisers cannot gain from shading their bid. However, a slight change in the previous scenario reveals the weakness of the algorithm.

Let us assume the same situation as before, except that advertiser A and B are now targeting *different* demographics. In this case, advertiser A still captures the first flotilla slot and pays the Z of advertiser B. However, advertiser B remains interested in the auction and secures the second flotilla slot for the Z of advertiser C. Since advertiser B had a very high value and advertiser C a very low value, this outcome implies that advertiser A will pay a high price and advertiser B a low price. In this situation, advertiser A would rather be in advertiser B's position because he would still get the identical number of impressions but for a much lower price. Therefore, advertiser A will shade his bid in an attempt to make a bid below advertiser B and secure the second flotilla slot. As proven by example, this algorithm is no longer a truth revealing mechanism.

4.2.3. Reimburse

One method of solving the deficiency of the previous algorithm is to charge the winning advertisers an amount congruent with the number of viewers he is "taking away" from other advertisers. The winner is still determined by the advertiser with the highest Z but his payment is now calculated as follows: for each non-winning advertiser the sum of their target viewers captured by the winning ad is calculated and multiplied with their value to derive Z^*. The winner (highest Z) must pay the highest of these Z^*s. Once the winner has been determined, he is removed from the system together with all the viewers he captured. Then the process is started over to determine the next winner.

In the previous example, if the top two advertisers are targeting the same demographic, this algorithm will force the winner to make a high payment (since he is taking away many viewers from the second highest advertiser). On the other hand, if the top two advertisers are targeting different demographics, then the winner will only have to pay a smaller amount (depending on how many viewers he is taking away from other advertisers). Since his payment now only depends on other bidders, this is a truth promoting mechanism. Below is a detailed description of algorithm Reimburse.

Algorithm Reimburse

---

1. Initialize
   1.a. Sort all Z from highest to lowest so Z_1 is the highest
   1.b. Sort all V and T to be in congruence with the ordered Z's; i.e. $V_a \sum_h T_{a,h} = Z_a \quad \forall a$
2. Determine winner and payment
   2.a. Z_1 receives first free flotilla slot
   2.b. Set $T^*_{a,h} = T_{a,h} * T_{1,h}$ for all $a \neq 1$ and all h
   2.c. Set $Z^*_a = X_a \sum_h T^*_{a,h} \quad \forall a$
   2.d. Sort all Z^* from highest to lowest so Z^*_1 is the highest
   2.e. Z_1 must pay Z^*_1
3. Remove previous winner
   3.a. Set $Z_1 = 0$
   3.b. Set $T_{2,h}$ .. $T_{A,h}$ to zero for all h where $T_{1,h} = 1$
   3.c. Recalculate $Z_2 \ldots Z_A$
   3.d. Resort $Z_1 \ldots Z_A$ (and all V and T accordingly)
4. Terminate algorithm
   4.a. If another free flotilla slot exists
      4.a.1. Return to step 2
   4.b. Else
      4.b.1. Terminate algorithm

While not immediately apparent, later calculations will show that this algorithm will frequently result in low revenue for the seller. The revised algorithm, described below, leads to higher revenue while maintaining the truth-telling property.

4.2.4. Reimburse Revised

The idea behind this revision is that the sale of the last flotilla slot has special implications. The ad that captures the last flotilla slot does not only seize the advertiser's particular demographic but, in effect, takes away the chance to capture *any* demographic from all other advertisers. Therefore, the winning ad of the final slot should pay the next highest Z (instead of the highest Z^*), as shown in algorithm Reimburse Revised.

Algorithm Reimburse Revised

---

1. Initialize
    - 1.a. Sort all Z from highest to lowest so Z_1 is the highest
    - 1.b. Sort all V and T to be in congruence with the ordered Zs; i.e. $V_a \sum_h T_{a,h} = Z_a \quad \forall a$
2. Determine winner and payment
    - 2.a. Z_1 receives first free flotilla slot
    - 2.b. If a free slot remains
        - 2.b.1. Set $T^*_{a,h} = T_{a,h} * T_{1,h}$ for all $a \neq 1$ and all h
        - 2.b.2. Set $Z^*_a = V_a \sum_h T^*_{a,h} \quad \forall a$
        - 2.b.3. Sort all Z^* from highest to lowest so Z^*_1 is the highest
        - 2.b.4. Z_1 must pay Z^*_1
    - 2.c. Else
        - 2.c.1. Z_1 must pay Z_2
3. Remove previous winner
    - 3.a. Set $Z_1 = 0$
    - 3.b. Set $T_{2,h}$ .. $T_{A,h}$ to zero for all h where $T_{1,h} = 1$
    - 3.c. Recalculate $Z_2 \ldots Z_A$
    - 3.d. Resort $Z_1 \ldots Z_A$ (and all V and T accordingly)
4. Terminate algorithm

4.a. If another free flotilla slot exists

4.a.1. Return to step 2

4.b. Else

4.b.1. Terminate algorithm

4.2.5. Vickrey-Clarke-Groves Mechanism

The next smart market auction that will be tested is the previously described VCG mechanism. Through the literature we know that the VCG mechanism is a truth revealing mechanism and it will be useful as a benchmark for the other proposed smart market auctions. An implementation of the VCG mechanism in the targeted television advertising system would consist of the following steps:

1. Find optimal combination of ads for the flotilla and their positions within the flotilla
   a. Save the resulting revenue R, given that each ad must pay his bid multiplied with his received impressions
2. For each of the winning ads W:
   a. Remove $Advertiser_w$ from the system and find the new optimal combination of ads
   b. Save the resulting revenue R_w
   c. $Advertiser_w$ must pay the amount that he contributed to R (his bid * his impressions) minus the difference between R and R_w

In words, the payment that advertisers will have to make is their bid reduced by the additional revenue that they bring to the system.

As Step 1 of the VCG mechanism shows, we need an optimization model to determine the winning ads. Below is a non-linear integer programming optimization model that accomplishes this task.

Non-linear Integer Programming Optimization Model – VCG

Model parameters

$H, A, S, V_a, T_{a,h}$ are defined as before.

Decision variables

$$x_{s,a} = \begin{cases} 1 & \text{slot } s \text{ is filled by ad } a \\ 0 & \text{else} \end{cases}$$

Objective function

$$\max_x R = \sum_a \left[V_a \sum_h \left(\sum_{s=1}^{S} \left[T_{a,h} x_{s,a} \sum_{i=1}^{s-1} \left(1 - \sum_{k \neq a} (T_{k,h} x_{i,k}) \right) \right] \right) \right] \qquad \text{(revenue for the seller)}$$

Constraints

$$\sum_s x_{s,a} \leq 1 \quad \forall a \qquad \text{(each ad can only receive one flotilla slot or none)}$$

$$\sum_a x_{s,a} \leq 1 \quad \forall s \qquad \text{(each slot can only be filled by one ad or none)}$$

$$x_{s,a} \in \{0,1\} \quad \forall s, a$$

The reason that the objective function is non-linear lies in the way that the flotilla advertisement ordering influences which ad will be seen. Each ad a contributes its value per impression V_a for each household h, $\sum_a \left[V_a \sum_h (...) \right]$, if it has received a flotilla slot s and the household is in its target demographic, $\sum_{s=1}^{S} [T_{a,h} x_{s,a} ...]$, and no ad that is higher in the flotilla has already captured that household, $\sum_{i=1}^{s-1} \left(1 - \sum_{k \neq a} (T_{k,h} x_{i,k}) \right)$. Note that if any higher ranked ad has captured the household, the last term becomes 0 and else it becomes 1.

While this formulation is relatively compact, non-linear integer optimization models are challenging to solve (Hromkovic, 2001). Therefore, this formulation is not practical for the seller's purposes as solving any iteration of this problem can easily exceed the time period that the seller has between receiving the inputs and having to construct the flotilla. The actual computational times required will be explored in a later section; however, first a linearization of the problem was attempted. One way of making non-linear formulations linear is to add additional decision variables. The next formulation is now a linear integer programming optimization model.

Linear Integer Programming Optimization Model – VCG

Model parameters

H, A, S, V_a, $T_{a,h}$ are defined as before.

Decision variables

$$x_{s,a} = \begin{cases} 1 & \text{slot } s \text{ is filled by ad } a \\ 0 & \text{else} \end{cases}$$

$$y_{a,h} = \begin{cases} 1 & \text{ad } a \text{ is seen by household } h \\ 0 & \text{else} \end{cases}$$

Objective function

$$\max_{x,y} R = \sum_a \left[V_a \sum_h y_{a,h} \right] \qquad \text{(revenue for the seller)}$$

Constraints

$$\sum_s x_{s,a} \le 1 \quad \forall a \qquad \text{(each ad can only receive one flotilla slot or none)}$$

$$\sum_a x_{s,a} \le 1 \quad \forall s \qquad \text{(each slot can only be filled by one ad or none)}$$

$$y_{a,h} \le T_{a,h} \sum_s x_{s,a} \quad \forall a,h$$ (each ad can only be seen by any household if it has received a slot and the household is interested in the ad)

$$y_{a,h} \le \sum_{i=1}^{s-1} x_{i,a} + 1 - \sum_{j \ne a} T_{j,h} x_{s,j} \quad \forall a,h,s < S$$ (each ad can only be seen by any household if the household has not been captured by a higher ranked flotilla slot)

$$x_{s,a}, y_{a,h} \in \{0,1\} \quad \forall s,a,h$$

Due to its linear nature, this optimization model will solve much faster; however, integer programming problems are still challenging problems to solve as the size of the problem increases.

4.2.6. Reservation Prices

As Chapter 3 showed, revenue can be increased further through an appropriate reservation price. In this case, a reservation price would imply that all advertisers with values below this price will not participate in the auction. The winner determination would remain the same as in the previously described algorithms; however, the payment of the winners would now be the maximum between the previously calculated payment and the required payment to satisfy the reservation price per impression.

4.3. LIST PRICE ALTERNATIVE

To allow for a revenue comparison between the smart market auctions and a list price system, the following list price model was developed. In this model, advertisers must only decide whether or not they want to purchase impressions at the list price and must not place any bids. If their value V_a is greater than the list price L_a, they will buy and else they will not. Since different demographic groups can have differing list prices, some advertisers will need to pay more than others.

Once the purchasing decision has been made by the advertisers, the seller must decide how to fill the flotilla. Again, they could use an optimization model that maximizes their revenue. The required formulation is very similar to the previous optimization model with two main differences: the objective function uses the list prices (not the individual values) and an additional constraint is needed to ensure that only advertisers with values greater than the list price can participate in the auction.

Linear Integer Programming Optimization Model – List Price

Model parameters

H, A, S, V_a, $T_{a,h}$ are defined as before.

L_a List price charged to advertiser *a*.

Decision variables

$$x_{s,a} = \begin{cases} 1 & \text{slot } s \text{ is filled by ad } a \\ 0 & \text{else} \end{cases}$$

$$y_{a,h} = \begin{cases} 1 & \text{ad } a \text{ is seen by household } h \\ 0 & \text{else} \end{cases}$$

Objective function

$$\max_{x,y} R = \sum_a \left[L_a \sum_h y_{a,h} \right]$$ (revenue for the seller)

Constraints

$$L_a \sum_s x_{s,a} \le V_a \quad \forall a$$ (each ad can only receive a flotilla slot if its value is higher than the list price)

$$\sum_s x_{s,a} \le 1 \quad \forall a$$ (each ad can only receive one flotilla slot or none)

$$\sum_a x_{s,a} \le 1 \quad \forall s$$ (each slot can only be filled by one ad or none)

$$y_{a,h} \le T_{a,h} \sum_s x_{s,a} \quad \forall a,h$$ (each ad can only be seen by any household if it has received a slot and the household is interested in the ad)

$$y_{a,h} \le \sum_{i=1}^{s-1} x_{i,a} + 1 - \sum_{j \ne a} T_{j,h} x_{s,j} \quad \forall a,h,s < S$$ (each ad can only be seen by any household if the household has not been captured by a higher ranked flotilla slot)

$$x_{s,a}, y_{a,h} \in \{0,1\} \quad \forall s,a,h$$

4.4. COMPUTATIONAL ANALYSIS

4.4.1. Introduction and Parameters

This section analyzes the new auction models to determine which one will result in the highest revenue and what impact the market structure (in terms of number of advertisers, their value distribution, etc.) will have on the suitability of the models. Furthermore, the expected revenue of the auction models is compared to that of a list price system. The results are obtained using Monte Carlo simulation. For an excellent summary of this method, see Rubinstein and Kroese (2008).

To allow for a computational analysis and Monte Carlo simulation, this work is completed under the following framework: it is assumed that the flotilla consists of two slots. Further, it is assumed that the population of television viewers consists of four distinct demographic groups, females and males either under or over the age of 30. Table 4.6 shows the percentage of the total population that each of these demographic groups encompasses.

Abbreviation	Demographic Group	Percent of Total Population
1F	Females under the age of 30	20%
2F	Females over the age of 30	20%
1M	Males under the age of 30	30%
2M	Males over the age of 30	30%

Table 4.6: Sample data for demographic composition of viewers.

Advertisers can target any combination of these demographics which leads to nine possible target groups which are listed in Table 4.7. The table also shows the percent of advertisers that fall into those respective groups and the average value per impression that these demographic groups are worth to advertisers. These values have been chosen arbitrarily and do not reflect real market prices; however, as would be expected in the market place, target demographic groups with higher demand to supply ratios are valued higher by advertisers.

Abbreviation	Target Demographic Group	Percent of Advertisers	Average Value of Demographic Group (G_D)
1F	Females under the age of 30	20%	1.4
2F	Females over the age of 30	5%	0.6
1M	Males under the age of 30	15%	1.2
2M	Males over the age of 30	5%	0.8
F	All Females	10%	1.0
M	All Males	10%	1.0
1	All people under the age of 30	20%	1.3
2	All people over the age of 30	10%	0.7
FM	All people	5%	1.0

Table 4.7: Sample data for advertiser composition in terms of target demographics.

The random inputs to the Monte Carlo simulation are as follows:

1. Each viewer is randomly chosen to be one of the four demographic types in accordance with the probabilities given in Table 4.6
2. Each advertiser is randomly chosen to be targeting one of the nine demographic groups in accordance with the probabilities listed in Table 4.7
3. Each advertiser receives a value (per impression of his target group) based on an exponential distribution with an average value in accordance with G_D (see Table 4.7)

4.4.2. Base Case Scenario

In the base case scenario, there are 10 viewers and 5 advertisers. Figure 4.1 shows the expected revenue for each of the previously described auction models and the list price system. These results are based on 200 runs, and the 90% confidence interval is indicated by black bars. For this base case scenario, the list price for each demographic target group was set at 20% above the average value. For example, advertisers targeting females under the age of 30 were charged 1.68 (=1.4+20%*1.4, using Table 4.7). The perfect (price) discrimination bar in Figure 4.1 shows how much revenue could be extracted if it was possible to charge the precise value of each advertiser. Of course, this is not possible because the seller does not know the precise values of individual advertisers. If the seller were to charge the advertisers their full bids, then this would represent a first price auction system which would result in bid shading by the advertisers. Even though the perfect discrimination revenue is not attainable, it shows the maximum amount of value that is available.

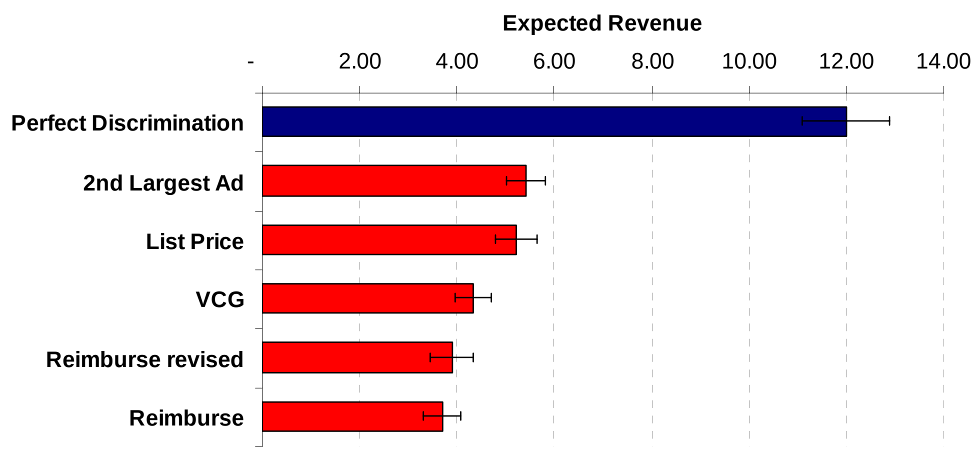

Figure 4.1: Expected revenue of auctions and list price in base case.

As Figure 4.1 shows, none of the described auction models can capture even half of the available value (the perfect price discrimination case). However, this is a reflection of the low ratio of numbers of advertisers to available flotilla slots, as will become apparent later. Furthermore, the domination of algorithm 2nd Largest Ad is misleading because these results were based on truth-telling bidding behaviour. As previously mentioned, this is not an accurate assumption for the algorithm 2nd Largest Ad. For this reason, this algorithm will no longer be included in the analysis. This does not necessarily imply that this algorithm is not a viable auction model; however, the expected revenue from this auction must be calculated under the assumption of bid shading. This may be an interesting area for future research.

However, truth-telling is valid assumption for the VCG mechanism and algorithm Reimburse, and Figure 4.1 shows that algorithm Reimburse is not far behind the VCG mechanism in its ability to raise revenue (expected revenue is 15% and 10% below VCG for algorithm Reimburse and Reimburse Revised, respectively). In addition, the algorithm Reimburse can execute in a fraction of time of the VCG mechanism. A later section will investigate the implications of the execution time.

In this base case, the list price leads to more revenue than the auction models. However, using the results of Chapter 3, we expect that the list price only dominates the auction systems when the ratio of number of advertisers to available flotilla slots is very low. The next section will explore the breakeven point at which auction models result in higher expected revenue than the list price system.

4.4.3. Number of Advertisers

Figure 4.2 shows how the list price mechanism initially dominates the auction models but raises less revenue than the auction models as the number of advertisers increases. Since the revised version of algorithm Reimburse always outperforms the former, only the revised version is represented in the figure.

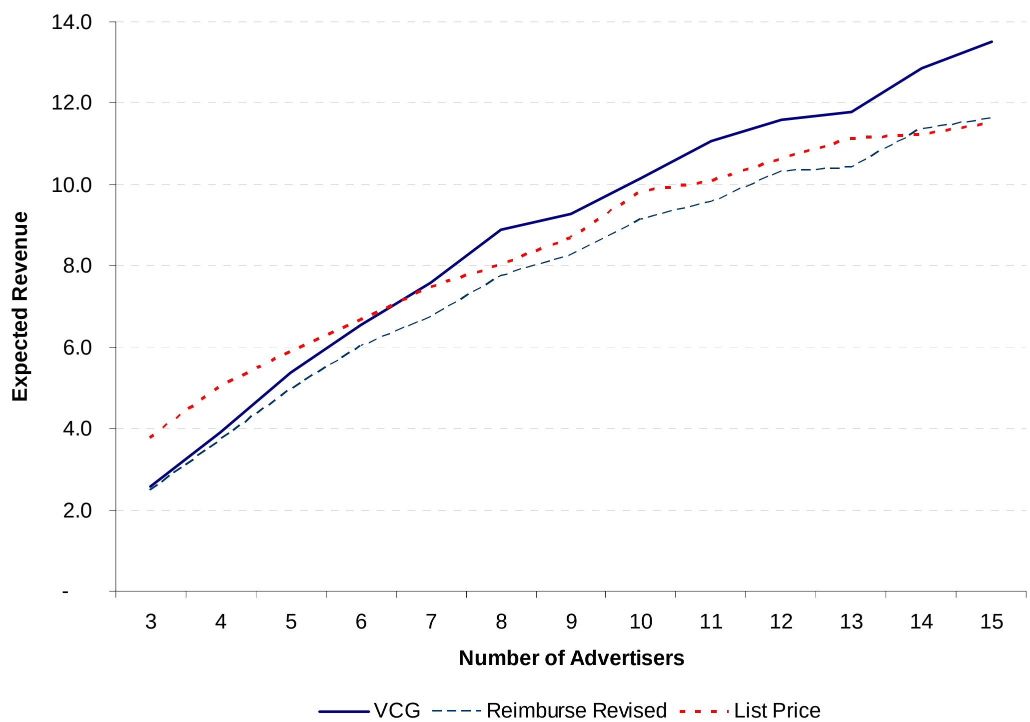

Figure 4.2: Expected revenue of auctions and list price by number of advertisers.

Figure 4.2 provides two main insights: first, the algorithm Reimburse Revised raises a similar amount of revenue as the VCG mechanism when only a small number of advertisers are in the system. However, the higher the number of bidders, the wider the gap between the expected revenue from the VCG mechanism and the algorithm Reimburse Revised. Second, the list price only raises more revenue than the VCG mechanism if the number of advertisers is less than seven. In addition, using the list price system implies the additional challenge of finding the optimal list price. Therefore, the next section will investigate how sensitive the expected revenue is to changes in the list price.

4.4.4. Setting the Correct List Price

This section analyzes the case of 10 viewers with 4 advertisers. Figure 4.3 compares expected revenue of the list price system to the VCG auction as the list price increases. Since the number of advertisers is low, the list price is expected to dominate the auction. More specifically, Figure 4.3 shows that the list price will dominate as long as it is set between 0.7 and 1.9 of the average value of the respective demographic group (roughly 40% below or 60% above the optimal list price factor of 1.2).

Figure 4.3: Expected revenue of list price dependent on list price set.

As before, the seller should rather overestimate the list price than underestimate it if they are uncertain of the optimal value. In the case of four advertisers there is a relatively broad range of list prices that will still lead to higher revenues than the VCG mechanism. Chapter 3 illustrated how the benefits of the auction model and the list price system can be combined by using an auction with a reservation price. The following section investigates how such a model would perform in the targeted television advertising system.

4.4.5. Setting a Reservation Price

Using the base case scenario, Figure 4.4 shows the substantial gains in expected revenue that can be made by using a reservation price with the auction models.

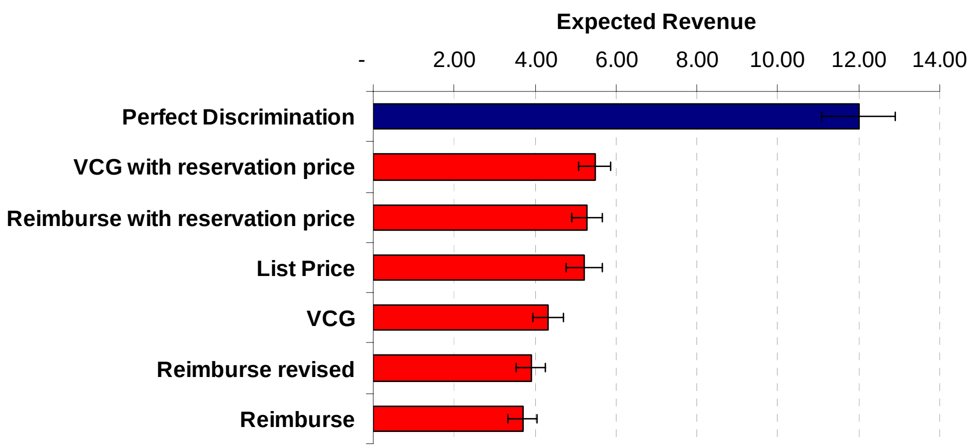

Figure 4.4: Expected revenue of auctions with reservation price and list price in base case.

Figure 4.4 shows that the seller can expect the highest revenue by using an auction model with a reservation price. For example, the expected revenue increases by 28% (from 4.3 to 5.5) when a reservation price is added to the VCG mechanism. Even with only a small number of advertisers, the auction models with reservation price will outperform the list price system. However, as with the list price, the performance of the auction model with a reservation price depends on how the reservation price is set. A reservation price that is set too low will fail to make an impact while a reservation price that is set too high will cause too many advertisers to stay away from the auction. Assuming 10 viewers and 4 advertisers, Figure 4.5 demonstrates this trade-off and shows how the reservation price should be set at roughly the expected value of the advertisers.

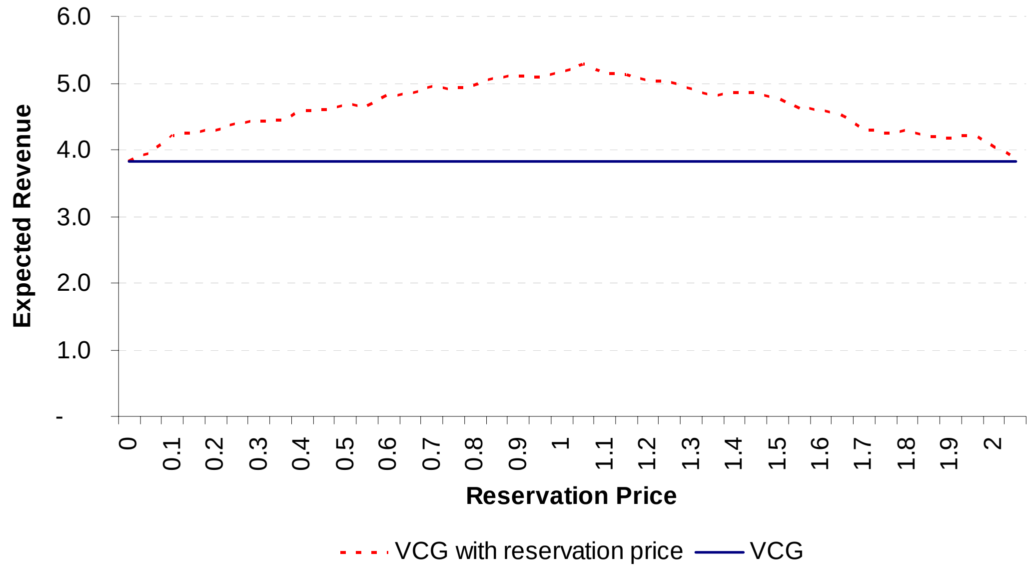

Figure 4.5: Expected revenue of VCG auction with reservation price.

The range of reservation prices that lead to higher expected revenues is even broader than for the list price. A reservation price that is set very low will not have a worse effect than its complete omission but it will fail to increase revenue. This effect can be seen in Figure 4.5 where the VCG with reservation price leads to higher expected revenue as soon as the reservation price exceeds zero. And the two methods lead to identical expected revenues when the reservation price is exactly zero. Furthermore, the reservation price can exceed double its optimal value and still lead to higher expected revenues than the equivalent auction model without the reservation price.

4.4.6. Value Distribution Functions

All the previous figures were based on [λ]-exponential value distributions of the advertisers. This section explores how the auction models perform based on a [a, b]-uniform distribution and a [μ, σ]-normal distribution, where $a=G_D*0.5$, $b=G_D*1.5$, $\mu=G_D$, $\sigma=G_D*0.2$ depending on the demographic group (see Table 4.7). Figure 4.6 compares the three different value distributions in a scenario with 5 and with 10 advertisers (10 viewers in both cases).

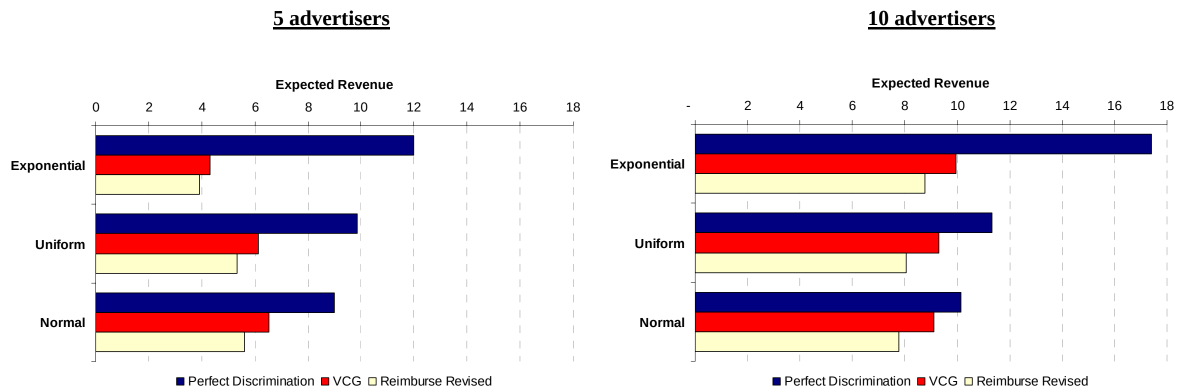

Figure 4.6: Expected revenue of auctions depending on value distributions.

Again, the results are consistent with the findings of Chapter 3. With a small number of bidders, the auction models perform well if the variance of the value distributions is low and vice versa.

4.5. LIMITATIONS

4.5.1. Large-scale Optimization

One of the downsides of the VCG mechanism is the required computational time to execute the optimization. In a scenario with 15 advertisers and 10 viewers and 3 flotilla slots, the Excel Solver requires 8.6 seconds (average, based on 100 runs) to find the optimal solution using the optimization model described in Section 4.2.5. To complete the VCG method, the optimization must run once with all advertisers included and then once again for each of the winning advertisers. Therefore, the system must optimize multiple times to determine the outcome of the VCG auction which increases the total run time even further (in this case to 8.5*3=25.5 seconds). As a comparison, the average execution time of the algorithm Reimburse is 0.0003 seconds assuming the same number of advertisers, viewers and flotilla slots.

The execution time can be drastically shortened by using a more powerful solver. The website NEOS (http://www-neos.mcs.anl.gov/) offers state-of-the-art optimization software which can execute much faster than the standard Excel Solver. One of their featured solvers, MINTO, solves mixed-integer linear programs through a branch-and-bound algorithm with linear programming relaxations. In the above case with 15 advertisers and 10 viewers and 3 flotilla slots, MINTO decreases the average execution time of a single optimization by 90% (from 8.5 to 1 seconds).

However, the seller may face optimization problems of a size that challenges even very powerful solvers such as MINTO. For example, increasing the problem to 30 viewers and 30 advertisers (still 3 flotilla slots) extends the average execution time of a single optimization to roughly 100 seconds using MINTO. For even larger problems, the required optimization time may not be feasible as the seller only has limited time available to construct the flotilla.

4.5.2. Main Assumptions

Aside from the parameters of the simulation model (such as the value distributions and number of advertisers) and the framework of a private value auction, there are a few other assumptions that drive the previous results.

First, the expected revenue is derived from these auction models assume that advertisers *bid truthfully*. An example given earlier showed that this is definitely not valid for the algorithm 2nd Largest Ad and that advertisers can gain from shading their bid. In comparison, in both the VCG mechanism and the algorithm Reimburse, the advertisers' bids only influence the probability of winning and not the amount to be paid, which makes them both truth-telling mechanisms. However, as mentioned earlier, there are other considerations that may influence how advertisers bid such as privacy concerns (revealing their true value) or a perceived lack of fairness. Both can lead to bid shading.

Second, the results assume that there is *no collusion* between advertisers. Given the limited information of advertisers regarding the other bids that are in the system at any given time, this is likely a valid assumption. The motivating example of Section 4.1.1 showed that there are cases in which advertisers can successfully use shill bids in a VCG auction; however, they would need to know the full information matrix (current households and advertisers in the marketplace). Although they could try to guess, this information is not accessible to them and a shill bid could easily backfire on the advertiser and lead to higher costs. Furthermore, within the context of a smart market in which there is no direct bidding, it is nearly impossible to collude successfully. Too many factors are outside of the advertisers' control such as when the individual avails occur and which households will be watching at those particular moments.

4.6. CONCLUSIONS

The analysis showed that many of the findings of Chapter 3 still hold in this full model. However, some of the standard auctions are not directly applicable to this problem. Instead, new auction models must be developed to deal with the complexity of the targeted television advertising system.

4.6.1. Algorithm Reimburse vs. VCG Mechanism

An optimization model was developed that allows for a VCG implementation but the required optimization may require excessive execution time. The auction models Reimburse and Reimburse Revised were developed that also lead to truth-telling bidding behaviour without requiring a time-consuming optimization.

The revised version of the algorithm Reimburse always outperforms the former version. In addition, the revised version performs well against the VCG mechanism (in terms of revenue) if the number of advertisers in the system is small. However, the more advertisers are in the system, the larger the gap between the expected revenue of the VCG mechanism and the algorithm Reimburse Revised.

The type of value distribution has a lesser impact on the relative performance of the algorithm Reimburse Revised compared to the VCG mechanism. Nevertheless, a lower variance in the bidders' value distribution leads to an increasing performance margin between the two algorithms in terms of their expected revenue. Consequently, the choice of auction model depends on the preferences of the seller and the market structure he faces.

4.6.2. List Price System vs. VCG Mechanism

As is the model of Chapter 3, a list price system only leads to higher revenue if the ratio of advertisers to flotilla slots is low. This result is not price sensitive and even if the list price deviates significantly from the optimal price. However, the VCG mechanism outperforms the list price system as soon as more than seven advertisers are in the system (assuming two flotilla slots).

With the introduction of reservation prices the auction models become superior to the list price system regardless of the parameters. Similar to the list price, the range of reservation prices that lead to higher expected revenue is broad. Consequently, slight over- or underestimations will not lead to negative results.

5. CONCLUSIONS AND FUTURE RESEARCH

5.1. CONCLUSIONS

This thesis has analyzed how a seller with new technology to deliver targeted television advertising can successfully sell its avails through auctions and has explored when it is beneficial to do so. There are many aspects to the way that the seller offers television advertising impressions that are new and original; in particular, the information structure under which the seller must decide which advertisers receive slots in the flotilla and how much they should be charged has not been previously studied. This thesis suggests specific auction rules that will lead to desirable results such as revenue maximization and limitation of computational complexity. Furthermore, this thesis provides the seller with guidance on the decision between a list price system and an auction model.

5.1.1. List Price versus Auction Model

If there are relatively few advertisers in the marketplace, the list price system will generally outperform an auction. However, auction models with reservation prices will lead to higher expected revenues compared to a list price system regardless of the number of advertisers. The expected revenue is relatively insensitive to the precise value of the list price but under uncertainty the seller is better off overestimating the price.

The seller may choose to implement a list price system in some local markets in which they expect a smaller number of advertisers. They may also use some of these local markets to experiment with how they should set the list price.

5.1.2. Type of Auction Model

First price auction models create undesirable instability in the system as bidders will try to game the system and will keep updating their bids. Some auctions, such as the algorithm 2nd Largest Ad, may seem like a truth revealing auction but will actually lead to bid shading and attempts to game the system. However, the seller now has a method of applying the VCG mechanism to their system. This mechanism truly promotes truth-telling in the bidding behaviour but it may require long execution times. As an alternative, the seller can use the algorithm Reimburse which also promotes truth-telling but executes much faster.

5.2. FUTURE RESEARCH

The new targeted television advertising system provides many avenues for future research. Companies such as Google, who have also implemented smart market auctions, have spent many years of manpower to develop the successful system that they now operate. For example, in Google's system advertisers can choose to bid a maximum amount that they are willing to pay (which was also assumed in this thesis) or an *average amount*. The latter forces Google to ensure that the advertiser pays this average price by inserting his advertisements into appropriate auctions. Risk-averse advertisers may be very interested in such an option in the new targeted television advertising system.

In addition, there will have to be some feedback to the bidders after the completion of the auctions. Because there will be many auctions in a day, as well as future advertising campaigns by the same advertisers, information provided to the bidders may influence future bidding behaviour. A crucial question to address is therefore *when* to provide information to the bidders and *which* information to provide. The depth of information provided to the bidder could range from only one number (the number of viewers reached) or a more detailed report indicating statistics such as average price paid for each viewer or the average completion of campaign targets by other bidders.

Furthermore, it seems intuitive to allow bidders to increase their bids throughout their campaign because it should lead to revenue increases for the seller. The question remains whether the bidders should be allowed to increase *and* decrease their bids and how frequent they should be allowed to update their bids. This question should be researched in conjunction with the information sharing aspect because they will likely influence each other.

APPENDIX

APPENDIX A – EXPECTED VALUES

$$E[X_3] = \int_0^\infty 1 - F(x)^3 dx$$
$$= \int_0^\infty 1 - \left(1 - 3e^{-\lambda x} + 3e^{-2\lambda x} - e^{-3\lambda x}\right) dx$$
$$= \left(\frac{3}{\lambda} - \frac{3}{2\lambda} + \frac{1}{3\lambda}\right) = \frac{11}{6\lambda}$$

$$E[X_1] = \int_0^\infty R(x)^3 dx$$
$$= \int_0^\infty e^{-3\lambda x} dx$$
$$= \frac{1}{3\lambda}$$

$$E[X_2] = \left(3 - E[X_3] - E[X_1]\right)$$
$$= \left(\frac{3}{\lambda} - \frac{11}{6\lambda} - \frac{1}{3\lambda}\right)$$
$$= \frac{5}{6\lambda}$$

APPENDIX B – INTEGRATION BY PARTS

$$\int_a^\infty x\lambda e^{-\lambda x} dx$$
$$= \left[-xe^{-\lambda x}\right]_a^\infty - \int_a^\infty -e^{-\lambda x} dx$$
$$= \left[-xe^{-\lambda x}\right]_a^\infty - \left[\frac{1}{\lambda} e^{-\lambda x}\right]_a^\infty$$
$$= ae^{-\lambda a} + \frac{1}{\lambda} e^{-\lambda a}$$
$$= \left(a + \frac{1}{\lambda}\right) e^{-\lambda a}$$

$$\int_a^\infty x\lambda e^{-2\lambda x} dx$$
$$= \left[-\frac{1}{2} xe^{-2\lambda x}\right]_a^\infty - \int_a^\infty -\frac{1}{2} e^{-2\lambda x} dx$$
$$= \left[-\frac{1}{2} xe^{-2\lambda x}\right]_a^\infty - \left[\frac{1}{4\lambda} e^{-2\lambda x}\right]_a^\infty$$
$$= \frac{1}{2} ae^{-2\lambda a} + \frac{1}{4\lambda} e^{-2\lambda a}$$
$$= \left(\frac{1}{2} a + \frac{1}{4\lambda}\right) e^{-2\lambda a}$$

$$\int_a^\infty x\lambda e^{-3\lambda x} dx$$
$$= \left[-\frac{1}{3} xe^{-3\lambda x}\right]_a^\infty - \int_a^\infty -\frac{1}{3} e^{-3\lambda x} dx$$
$$= \left[-\frac{1}{3} xe^{-3\lambda x}\right]_a^\infty - \left[\frac{1}{9\lambda} e^{-3\lambda x}\right]_a^\infty$$
$$= \frac{1}{3} ae^{-3\lambda a} + \frac{1}{9\lambda} e^{-3\lambda a}$$
$$= \left(\frac{1}{3} a + \frac{1}{9\lambda}\right) e^{-3\lambda a}$$

Calculation use integration by parts (Stewart, 1999).

APPENDIX C – AUCTION VS LIST PRICE BY VALUE DISTRIBUTION

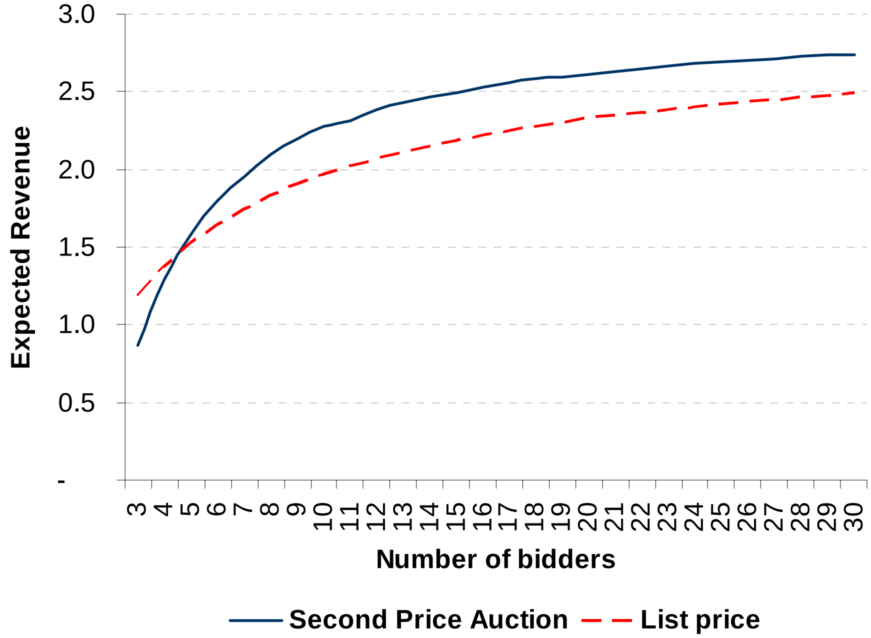

Figure 6.1: Second price auction vs. list price assuming [0.5,1.5]-Uniform distribution.

Figure 6.2: Second price auction vs. list price assuming [1,0.1]-Normal distribution.

REFERENCES

Bollapragada, S., H. Cheng, M. Phillips, M. Garbiras, M. Scholes, T. Gibbs, M. Humphreville. 2002. NBC's Optimization System Increase Revenues and Productivity, *Interfaces*, Vol. 32, No. 1, pp. 47-60.

Caldentey, R., G. Vulcano. 2007. Online Auctions and List Price Revenue Management, *Management Science*, Vol. 53. No. 5, pp. 795-813.

Clarke, E. 1971. Multipart pricing of public goods. *Public Choice*. Vol. 8, pp. 19-33.

Cramton, P.C. 1995. Money Out of Thin Air: The Nationwide Narrowband PCS Auction, *Journal of Economics & Management Strategy*, Vol. 4, No. 2, pp. 267-343.

Edelman B., M. Ostrovsky. 2007. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*. Vol. 43, pp. 192-198.

Edelman B., M. Ostrovsky, M. Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*. Vol. 97, No. 1, pp. 242-259.

Gallien, J., L.M. Wein. 2005. A Smart Market for Industrial Procurement with Capacity Constraints, *Management Science*, Vol. 51, No. 1, pp. 76-91.

Groves, T. 1973. Incentives in Teams. *Econometrica*. Vol. 41, pp. 671-631.

Hohner, G., J. Rich, E. Ng, G. Reid, A.J. Davenport, J.R. Kalagnanam, H.S. Lee, C. An. 2003. Combinatorial and Quantity-Discount Procurement Auction Benefit Mars, Incorporated and Its Suppliers, *Interfaces*, Vol. 33, No. 1, pp. 23-35.

Hromkovic, J. 2001. Algorithms for Hard Problems, Heidelberg: Springer-Verlag.

Jones, J.L. 2000. Incompletely Specified Combinatorial Auctions: an alternative allocation mechanism for business-to-business negotiations. Doctoral Dissertation, University of Florida.

Jones, J.L., G.J. Koehler. 2002. Combinatorial auctions using rule-based bids, *Decision Support Systems*, Vol. 34, pp. 59-74.

Katz, H. 1995. The Media Handbook, Lincolnwood: NTC Business Books.

Klemperer, P. 1999. Auction Theory: A Guide to the Literature, *Journal of Economic Surveys*, Vol. 13, No. 3, pp. 227-286.

Klemperer, P. 2002. What Really Matters in Auction Design, *The Journal of Economic Perspectives*, Vol. 16, No. 1, pp. 169-189.

Krishna, V. 2002. Auction Theory, San Diego: Academic Press.

Menezes, F.M., P.K. Monteiro. 2005. An Introduction to Auction Theory, New York: Oxford University Press Inc.

Milgrom, P.R., R.J. Weber. 1982. A Theory of Auctions and Competitive Bidding, *Econometrica*, Vol. 50, No. 5, pp. 1089-1122.

Milgrom, P. 2004. Putting Auction Theory to Work, Cambridge: Cambridge University Press.

Millet, I., D.H. Parente, J.L. Fizel, R.R. Venkataraman. 2004. Metrics for Managing Online Procurement Auctions, *Interfaces*, Vol. 34, No. 3, pp. 171-179.

Pekec, A., M.H. Rothkopf. 2003. Combinatorial Auction Design, *Management Science*, Vol. 49, No. 11, pp. 1485-1503.

Rassenti, S.J., V.L. Smith, R.L. Bulfin. 1982. A Combinatorial Auction Mechanism for Airport Time Slot Allocation, The Bell Journal of Economics, Vol. 13, No. 2, pp. 402-417.

Riley, J.G., W.F. Samuelson. 1981. Optimal Auctions, *The American Economic Review*, Vol. 71, No. 3, pp. 381-392.

Ross, S.M. 2002. A First Course in Probability, 6th Ed., Upper Saddle River: Prentice-Hall Inc.

Ross, S.M. 2003. Introduction to Probability Models, 8th Ed., San Diego: Academic Press.

Rothkopf, M.H., S. Park. 2001. An Elementary Introduction to Auctions, *Interfaces*, Vol. 31, No. 6, pp. 83-97.

Rothkopf, M.H. 2007. Thirteen Reasons Why the Vickrey-Clarke-Groves Process Is Not Practical, *Operations Research*, Vol. 55, No. 2, pp. 191-197.

Rubinstein, R.Y., D.P. Kroese. 2008. Simulation and the Monte Carlo Method, 2nd Ed., Hoboken, New Jersey: John Wiley & Sons, Inc.

Sissors, J.Z., R.B. Baron. 2002. Advertising Media Planning, New York: McGraw-Hill.

Stewart, J. 1999. Calculus, 4th Ed., Pacific Grove: Brooks/Cole Publishing Company.

TNS Media Intelligence, 2007. TNS Media Intelligence Forecasts 2.6 Percent Increase in U.S. Advertising Spending for 2007 [Online]. Available from: http://www.tns-mi.com/news/01082007.htm [Accessed: 22 January 2008]

Varian, H.R. 2007. Position Auctions, *International Journal of Industrial Organization*, Vol. 25, pp. 1163-1178.

Vickrey, W. 1961. Counterspeculation, Auctions, and Competitive Sealed Tenders, *The Journal of Finance*, Vol. 16, No. 1, pp. 8-37.