Making Uncompetitive Auctions Competitive: A Survey of Experiments

Theo Offerman*and Sander Onderstal†

February 2, 2006

Abstract

In the past 15 years, governments around the world have often used auctions to sell scarce licenses to operate in markets. In many of these auctions, the number of interested competitors is relatively small compared to the number of available licenses. Auction design is crucial in such uncompetitive circumstances. Details of the design affect participants' decisions to compete seriously or not. Such decisions are important for the industry structure and the efficiency of the aftermarket as well as for the revenue raised in the auction. This paper provides a survey of emerging experimental work on the question how competition can be stimulated in uncompetitive license auctions. We consider papers that deal with the performance of standard auctions (such as the simultaneous ascending auction and the discriminative auction) in uncompetitive circumstances. We also discuss papers that investigate the performance of some less known auctions (such as the Anglo-Dutch auction, the Amsterdam auction, and Right-To-Choose auctions) that actively seek to foster competition among bidders who would not compete in standard auctions.

1 Introduction

In the years 2000 and 2001, several European governments auctioned licenses for third generation mobile telecommunication (UMTS). Governments raised an incredible amount of more than €100 billion in these auctions.[1] Although the auctions in the different countries may have

*University of Amsterdam, Economics Department, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands, +31-20-525-4294, T.J.S.Offerman@uva.nl, www1.fee.uva.nl/creed/people/offerman/index.shtml. The research of Offerman has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

†University of Amsterdam, Economics Department, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands, +31-20-525-7161, onderstal@uva.nl, www.fee.uva.nl/pp/onderstal. Onderstal gratefully acknowledges financial support from the Dutch National Science Foundation (NWO-VICI 453.03.606).

[1]The European UMTS auctions are not the only license auctions that took place in the recent past. In the past 15 years, governments all over the world have used auctions to sell scarce licenses to operate in markets.

been very similar to a layman's eye, they differed in small but crucial details. In uncompetitive circumstances, where the number of competitors is small compared to the number of licenses for sale, such details may have profound effect on the auction outcome. A good example is offered by the UMTS auction in the Netherlands and the UK. While both countries offered five licenses in a simultaneous ascending auction, the Netherlands' auction witnessed only one entrant besides the five incumbents and a revenue of €170 per capita was collected, while thirteen competitors participated in the UK auction which raised more than €650 per capita. What was the cause of this substantial difference in auction outcomes?

One difference between the Netherlands and the UK was that in the Netherlands, five firms were already active in the market for second generation mobile telecommunication, while in the UK this number was four. Incumbents have some important competitive advantages compared to newcomers, as they already have a network of base stations and a customer base. Furthermore, their willingness to pay may be larger for preemptive reasons as they prefer to keep out entrants to protect their profits in related markets. As both governments auctioned five licenses, it is quite possible that potential entrants in the UK perceived to have a much larger probability to win a license than in the Netherlands. Another difference was that the auction in the Netherlands was run after the auction in the UK. The sky-high prices in the UK may have deterred potential entrants in the Netherlands' auction. A third difference is that the Netherlands is a much smaller country than the UK, and that the perspectives to make money are smaller in the Netherlands. Finally, it can also not be excluded that part of the difference is just explained by luck.

Ideally, one would like to run a controlled field experiment to see which factor is responsible for differences in auction outcomes. In practice this is impossible. A good alternative option is to run a series of controlled laboratory experiments, in each of which a couple of auctions is compared that differs in one aspect only. Notice that laboratory experiments also offer the important possibility of replication, which clarifies a potential role of luck. With sufficiently

For instance, before the European governments decided to auction their spectrum, the Federal Communications Commission (FCC) had allocated licenses for second generation mobile telecommunication in a series of auctions. Other examples are auctions for commercial radio, petrol stations, wireless local loop, and passenger railway services.

repeated observations in identical circumstances, the researcher has the possibility to draw statistically meaningful conclusions. Notice that within the field of IO, the application of experiments to auctions is very fruitful, because it is straightforward to introduce the institutional details of the auction in an experiment.

An objection raised against laboratory experiments is that the subjects in the lab (usually undergraduate students) are not representative for the decision makers in the field. Notice, however, that the subject pool of the experiment is a choice of the experimenter, and the experiment can be run with decision makers of the field. The question whether students and decision makers in the field behave similarly is an empirical one that has been investigated. Most often the differences between students and professionals are small or absent (see Friedman and Sunder, 1994, p.39-43 and Abbink and Rockenbach, 2005 for references to this kind of work). For instance, Dyer, Kagel and Levin, (1989) observe in an auction experiment that experienced business executives in the construction contract industry perform qualitatively similarly as inexperienced undergraduate students. We would welcome more experiments of this kind, though, and think that this remains an important avenue for future research.

In this chapter, we offer an overview of experimental work that investigates the performance of auctions in uncompetitive circumstances. We consider papers that deal with the performance of standard auctions (such as the simultaneous ascending auction and the discriminatory auction) in uncompetitive circumstances. Questions that are addressed in those papers are: How vulnerable is an auction design to revenue reducing strategies? How likely is the auction to attract entry? Is the allocation of the auction efficient? How does the auction design affect competition in the market? How much revenue does the auction raise? We also discuss papers that investigate the performance of some less known auctions that actively seek to foster competition among bidders who would not compete in standard auctions: the Anglo-Dutch auction, the Dutch-Anglo auction, premium auctions, and right-to-choose auctions.[2]

Before we discuss these auctions, it is good the stress that the competition law alone is not sufficient to induce competition, neither in the auction, nor in the aftermarket. Of course, one

[2] We focus on situations where there are no strong synergies between licenses. Therefore we do not deal with combinatorial auctions in this chapter.

may also turn to competition law to prevent undesirable practices like collusion. For instance, 'bidding rings' in which bidders make agreements about who wins the license at which price, are considered illegal under article 81 of the EC Treaty. However, in both auctions and product markets, it is quite difficult to prove tacit collusion to be a violation of competition law. A well-designed auction may discourage (tacit) collusion, so that the government does not have to rely on competition law in the first place. Indeed, Motta (2004) argues in his book on competition policy that "it is better to try to create an environment that discourages collusion in the first place than trying to prove unlawful behavior afterwards. A clear advantage of auction markets is that the environment can be affected directly, since the rules of the game are specified at the beginning by the auctioneer." Motta's advice also applies to the market for which the licenses are auctioned: a well designed auction encourages entry into the market so that (tacit) collusion in the market becomes more difficult. One of the central questions of our survey is: which auction formats are more successful in encouraging competition, both in the auction and in the aftermarket?

The set-up of this chapter is as follows. In section 2, we discuss the goals that governments may pursue when they design license auctions. In Section 3, we briefly discuss some basic notions of auction theory. Section 4 contains an overview of laboratory experiments on standard auctions, while in section 5 and 6, we focus on hybrid auctions and right-to-choose auctions respectively. Section 7 concludes with the main lessons and some policy advice.

2 Objectives of license auctions

Usually governments pursue a hybrid of objectives, some of which relate to efficiency while others relate to revenue. We first deal with issues related to the efficiency of a license auction. The concept of efficiency is defined differently in the sale of licenses than in the sale of say a painting. In an auction of a painting, the result is efficient if the painting is won by the person who attaches the highest value to it. In a license auction, efficiency results if the stream of future surpluses in the aftermarket are as high as possible. The surplus in the aftermarket consists of a consumer part (the difference between consumers' willingness to pay for the product that

is traded in the aftermarket and the price that they pay) and a producers' part (roughly the difference between the price for the product and the costs of the product). Because it is quite impossible to directly measure the efficiency level of a real license auction, it is useful to identify a number of targets that correlate with efficiency:

1 *Entry by newcomers.* Entry is an important target, since according to mainstream theory, more competition leads to lower prices, better quality, and more innovation, and thus to a higher level of efficiency in the aftermarket. Notice also that newcomers may be useful to stir up a collusive product market.

2 *Make sure that the licenses end up in the hands of the firms that can serve the market most efficiently.* Cost efficiency is one of the contributors to producer surplus and welfare, so it is important that inefficient firms are unlikely to win licenses. Note that this target could be in conflict with the first, as an incumbent could be the more efficient producer in the market, while in some situations it may be better to have a somewhat less cost efficient firm enter the market to stir up competition.

3 *Make sure that firms use licenses to produce services that are valued by consumers.* Sometimes incumbent firms face preemptive reasons to bid high. They want to overbid potential entrants in order to maintain their attractive powerful position in the market (Gilbert and Newbery, 1982, and Hoppe et al., 2005). Once having obtained the license, they use it in an inferior way than an entrant would have used it. To pursue efficiency, licenses have to be assigned to firms that use them for the right reason, that is, to make consumers happy. Notice that this target may conflict with the second target, as it may be the case that an incumbent firm is the cost efficient but plans to use the license in a preemptive way only.

Other objectives of governments usually relate to revenue. Many sellers of commodities in the private market relentlessly pursue an as high expected revenue as possible. Quite often governments pursue a more sophisticated goal when it comes to revenue. In particular, many

governments refrain from setting the considerable reserve prices that would correspond to expected revenue maximizing auctions. Considerable reserve prices lead to the possibility that some licenses go unsold. Policy makers tend to dislike this outcome, probably both because it harms efficiency and because unsold licenses makes them look bad in the eyes of the public. Also, governments often explicitly state that they want to prevent the winner's curse, i.e., a situation in which a firm pays so much that it makes a loss. Although a high revenue may be served by a winner's curse, governments usually want to avoid it for the fear that winning firms go bankrupt or want to renegotiate the deal in court. Policy makers also dislike the opposite outcome where firms keep auction prices low and receive windfall profits. In other words, policy makers prefer to avoid auction formats with a high variance of revenue. Therefore, for policy makers it is usually important to know both about an auction format's expected revenue and its variance.

Revenue and efficiency are not independent targets. A high auction revenue implies that governments can levy less distortionary taxes. Indeed, Ballard et al. (1985) estimate deadweight losses of raising taxes to lie between 17 and 56 cents for every extra $1 raised. On the other hand, the experiment by Offerman and Potters (2006) suggests that auctioning licenses increases consumer prices in the aftermarket as firms become more willing to pursue collusion after they pay a high price to enter a market. A related theoretical argument is provided by Janssen (2005) who argues that the high revenue raised in an auction serves as a forward-induction device that helps firms coordinate on the most collusive equilibrium of the oligopoly game. High consumer prices in the aftermarket harm efficiency and counteract society's gain resulting from the decrease in distortionary taxes in other domains. Note also that realizing much revenue can easily conflict with pursuing efficiency. In some situations, the government could generate a very high revenue if it gives one firm the right to operate as a monopolist in the market rather than sell licenses to several firms, a solution that would clearly harm efficiency in the aftermarket.

Of course, governments have other means available besides varying the auction format. For instance, governments could give an advantage to weak bidders when they pursue a high revenue. One way to favor a weak bidder is to give her a bidding credit, i.e., give her a discount

on her bid. In one of the FCC auctions, companies owned by minorities or women only had to pay 80% of their bid when they won a license. Another possibility is to reserve licenses for weak bidders. For instance, in the UMTS auction in the UK, incumbent parties were not allowed to bid on one of the licenses. Still, favoring weak bidders has some disadvantages. First, it is not clear whether favoring one group of bidders over another is in conflict with anti-discrimination law, such as article 87(1) of the EC Treaty. Second, it is not so easy to get the details of the auction right. To which extent should weak bidders be favored? For instance, if the government wishes to maximize the revenue of the auction then the optimal bidding credit highly depends on the environment, such as the underlying distribution of signals or the utility functions of the bidders (Myerson, 1981). For these reasons, we only consider 'detail free' auctions, i.e., auctions that are context independent in the sense that 'the rules of the game' do not depend the parameters of the environment.[3]

3 Some theoretical notions

We first discuss the basics of auction theory, so that we have a framework in which the results of the experiments can be placed.[4] The most commonly studied auctions of a single object are the English, the Dutch, the first-price sealed-bid, and the second-price sealed-bid or Vickrey auction (the 'standard auctions'). In the English auction, the price is raised successively until one bidder remains. This bidder wins the object at the final price. The price can be raised by the auctioneer, or by having bidders call the bids themselves. In the Dutch auction, the auctioneer begins with a very high price, and successively lowers it, until one bidder announces that she is willing to accept the current price. This bidder wins the object at that price. In the first-price sealed-bid auction, bidders independently submit sealed bids. The object is sold to the highest bidder at her own price. The Vickrey auction has the same rules as the first-price sealed-bid auction, except that the price the winner pays is not her own bid, but the second highest bid.

[3] Wilson (1987) promotes the use of detail-fee mechanisms. See also Krishna (2002), p. 75.

[4] For excellent overviews of auction theory and its relation to practice see the books by Krishna (2002) and Klemperer (2004).

In his seminal theoretical work on auctions, Vickrey (1961) studies bidding behavior in standard auctions in the symmetric independent private values model. In this model, one indivisible object is sold to one out of a set of bidders who draw their value for the object independently from the same value distribution function. Vickrey proves two important 'equivalence' theorems: (1) all standard auctions are efficient, and (2) the (expected) revenue from all standard auctions is the same. This finding suggests that the seller does not have to worry much about the auction design, as many auctions generate the same outcome. Early laboratory experiments mainly tested the outcomes of Vickrey's model. Kagel (1995), in an excellent overview, observes that these experiments confirm quite a few of the predictions of Vickrey's theory. The main deviation is that first-price auctions raise more revenue than the other auctions.

Later theoretical research reveals that Vickrey's equivalence results rely on several crucial assumptions: (1) independent private values, (2) symmetric bidders, (3) no externalities, and (4) only one object is sold. The equivalence results may break down if at least one of these assumptions is violated. Let us shortly discuss alternative models in which deviations from the symmetric independent private values model are studied. More specifically, we compare first-price auctions (the first-price sealed-bid auction and the Dutch auction) and second-price auctions (the English auction and the Vickrey auction). In first-price auctions, the winner pays a price equal to her own bid, while in second-price auction, it is the runner-up's bid that determines how much the winner pays.[5]

- *Independent private values.* Vickrey assumes that all bidders draw their value for the object independent from the same distribution function. Milgrom and Weber (1982) show that in an environment in which values are correlated, the English auction and to a lesser extent the Vickrey auction yield more revenue than the first-price auction. The reason is that in the English auction, information is revealed while the price increases: active bidders observe who steps out at which price. This information may be useful to update a bidder's belief about the true value of the object, so that she is less prone to the

[5] In some formats of the English auction, strictly speaking, the winner pays her own bid. However, in equilibrium, this bid is just above the highest losing bid.

'winner's curse', and is hence willing to bid higher than in a first-price auction in which no information is released. This is one of the reasons why the simultaneous ascending auction was advocated to sell licenses, as this auction is a generalization of the English auction to multi-object environments.

- *Symmetric bidders.* The assumption that bidders are symmetric means that all bidders draw their value for the auctioned object from the same distribution function. Bulow et al. (1999) and Klemperer (1998) show that first-price auctions may yield (far) more revenue than second-price auctions if only a slight asymmetry is introduced. However, Plum (1992) and Maskin and Riley (2000) show in a more general model that the revenue ranking of first-price and second-price auctions is ambiguous. Moreover, the first-price auction turns out to be inefficient in the sense that the object does not always end up in the hands of the bidder that attaches the highest value to it, in contrast to second-price auctions. These models are very relevant, as they apply to situations in which 'strong' bidders (e.g. an incumbent) compete in the auction against 'weak' bidders (e.g. a potential entrant).

- *No externalities.* Vickrey assumes that the winning bidder imposes no externality on the losers. This assumption may be violated in license auctions. For instance, if an entrant wins a license, an incumbent may experience a negative externality as more competition in the market may induce the price, and hence his profit, to be lower. The incumbent may then bid more than his value for the object, i.e., bid preemptively in order to prevent entry (Gilbert and Newbery, 1982 and Hoppe et al., 2005). First-price and second-price auctions may perform quite differently in these environments (Jehiel et al., 1996, 1999, Jehiel and Moldovanu, 2000, Das Varma, 2002, and Goeree et al., 2004a).

- *Single unit.* Vickrey studies a model in which only one unit is sold. However, in many license auctions, the government simultaneously sells several licenses. In multi-object auctions, the discriminatory auction and the simultaneous ascending auction may be considered as multi-unit extensions of the first-price and second-price single-unit auctions respectively. In a discriminatory auction, the bidders submit sealed-bids on the objects,

and the winner of each object pays a price equal to the own bid. The simultaneous ascending auction is the auction format that has been used most frequently to assign licenses in markets.[6] The following rules apply to this auction. Multiple units are sold simultaneously and bidding occurs in a series of rounds. In each round, those bidders who are eligible to bid, make sealed bids for as many objects as they want or are allowed. At the end of each round, the auctioneer announces the standing high bid for each object along with the minimum bids for the next round, which is computed by adding a pre-determined bid increment such as 5% or 10% to the standing high bids. A standing high bid remains valid until it is overbid. The auction concludes when no new bids are submitted. The standing high bids are then declared winning bids, and each winner pays an amount equal to the standing high bid. One difference between a discriminatory auction and a simultaneous ascending auction is that the latter is more prone to tacit collusion, and may hence generate less revenue than a discriminatory auction. The following example illustrates why this is the case. Suppose two bidders, 1 and 2, bid for two licenses, A and B. Each bidder attaches value 100 to each license. The bidders may tacitly agree to divide the market such that 1 gets A and 2 gets B. The bidders may do so, by bidding 0 on 'their' license. In the simultaneous ascending auction, such agreement is stable. Imagine that bidder 1 deviates from the agreement, and overbids bidder 2 on license B. Then bidder 2 can immediately 'punish' bidder 1, by placing a bid on license A. These punishment strategies are not available in a discriminatory auction. In other words, the simultaneous ascending auction is prone to 'demand reduction', which may imply that revenues are low and that the allocation of the licenses is inefficient (Noussair, 1995, Engelbrecht-Wiggans and Kahn, 1998, Ausubel and Cramton, 2002, and Goeree et al., 2004a).

In the remainder of this chapter we will see to what extent these theoretical considerations play a role in practical auctions.

[6] It was very often used in the American FCC auctions (see Plott, 1997, Roth, 2002) and in several UMTS auctions (Binmore and Klemperer, 2002, Abbink et al., 2005).

4 Standard auctions

In this section, we compare standard auctions (including the simultaneous ascending auction and the discriminatory auction) in uncompetitive circumstances. Klemperer (2002b) conjectures that in such situations, first-price auctions perform better than second-price auctions. The main question that is answered in this section is: do first-price auctions indeed outperform second-price auctions in thin markets? In subsection 4.1, this question is answered for auctions in which only a single object is sold, while in subsections 4.2 and 4.3 we focus on multi-object auctions, looking at homogeneous objects and heterogeneous objects respectively. Subsection 4.4 contains a summary of the main findings.

4.1 Single object auctions

In the previous section, we noted that in theory, the ranking of first-price and second-price auctions is ambiguous if bidders are asymmetric, for instance when an incumbent and a potential entrant in a market compete to get a license which gives them additional capacity to operate in that market. Does this theoretical result carry over to the lab? Pezanis-Christou (2002) and Güth, Ivanova-Stenzel and Wolfstetter (2004) conduct experiments in asymmetric environments in order to study the efficiency and the revenue generating properties of the first-price sealed-bid auction and the Vickrey auction. More precisely, they study an environment in which one object is auctioned to one out of two bidders. Let bidder 1 be the weak bidder and bidder 2 the strong one, and let v_i denote the value for the object to bidder i, $i = 1, 2$.

In both studies, it is assumed that the values for the object are independently drawn from uniform distributions. More precisely, Güth et al. (2004) considers an environment in which $v_1 \sim U[50, 150]$ and $v_2 \sim U[50, 200]$. Pezanis-Christou (2002) conducts two series of experiments. In his first [second] parametrization, he assumes $v_1 \sim U[-100, 100]$ [$v_1 \sim U[-300, 100]$] and $v_2 \sim U[0, 100]$.

The theoretical predictions are the following. According to Maskin and Riley (2000), the Vickrey auction is always efficient, in contrast to the first-price sealed-bid auction. Moreover the revenue ranking of the first-price sealed-bid auction and the Vickrey auction is ambiguous for

risk neutral bidders. Observe that in Güth et al.'s (2004) setting, the weak bidder's distribution function is 'stretched out' over a longer interval in order to obtain the distribution function of the strong bidder. Maskin and Riley show that in these type of environments, first-price auctions generate more revenue than second-price auctions. In contrast, in the model of Pezanis-Christou (2002), the weak bidder obtains his distribution function by shifting some mass of the strong bidder's distribution function to the point zero (as bidders with a negative valuation for the object behave as if their value is zero). Maskin and Riley prove that these types of distribution shifts imply that the Vickrey auction dominates the first-price auction in terms of revenue, in contrast to Güth et al.'s (2004) framework. The predicted revenue differences are the following. In Güth et al.'s (2004) set-up, the first-price sealed-bid auction generates slightly more revenue than the Vickrey auction (about 2%). In Pezanis-Christou's (2002) first [second] parametrization, the seller's expected revenue in the first-price sealed-bid auction is 10% [33%] less than in the Vickrey auction.

Güth et al. (2004) observe the following in their experiment. First, the Vickrey auction is slightly more efficient than the first-price sealed-bid auction (99% versus 98%).[7] The difference is statistically significant but small. The average revenue of the first-price [Vickrey] auction was more than 10% higher than [about the same as] what theory predicted, so that the first-price auction clearly dominated the Vickrey auction. Finally, the variance of the revenues was substantially smaller for the first-price auction compared to the Vickrey auction.

Pezanis-Christou (2002) finds qualitatively similar results as Güth et al. The efficiency levels for both the Vickrey auction and the first-price sealed-bid auction are very high (97% and 95% respectively for the first parametrization and 99% and 96% for the second), and are not statistically significantly different. Moreover, the first-price auction generates more revenue than the Vickrey auction as bidders bid substantially more than in the risk neutral equilibrium, while in the Vickrey auction, bids are very close to the theoretical prediction. The variance of

[7] Efficiency is defined as follows. Let the surplus of an allocation be the sum of the values the winners realize. The efficiency e of the auction outcome is the actual surplus (S^{act}) as a fraction of the maximal surplus ($S^{\max}$):

$$e = \frac{S^{\text{act}}}{S^{\max}}.$$

the revenues in the first-price auction is smaller than in the Vickrey auction.

To summarize, both Güth et al. (2004) and Pezanis-Christou (2002) observe that the first-price sealed-bid auction outperforms the Vickrey auction in asymmetric environments. The first-price auction yields substantially more revenue than the Vickrey auction, while the theory indicates that in Güth et al.'s environment this difference would be much less substantial, and that in Pezanis-Christou's environment the revenue ranking would be reversed. Moreover, the variance of the revenue is smaller for the first-price auction than for the Vickrey auction. Finally, the efficiency differences between the two auctions are negligible, although the theory indicates that the Vickrey auction is more efficient than the first-price auction. These observations support Klemperer's (2002b) conjecture that in thin markets, first-price auctions perform better than second-price auctions, in contrast to what the theory sometimes suggests.

4.2 Auctions of multiple homogeneous objects

Sometimes governments auction multiple homogeneous objects while they allow bidders to acquire more than one object. The UMTS auctions in Germany and Austria provide clear examples that fit this setup. Both countries made use of the simultaneous ascending auction to sell their UMTS licenses. In Germany as well as in Austria, twelve blocks of spectrum were sold from which bidders could buy two or three blocks. Thus the industry structure was determined by the auction itself. One salient outcome is that four bidders each buy three blocks, while another focal outcome is the one where six bidders each buy two blocks. Some observers noted that this design exploits preemptive motives of incumbent firms, as it allows them to coordinate and keep price-fighting entrants out of the aftermarket. In such a preemptive process incumbent firms win the available spectrum and drive up the revenue. Others commented that this design could easily lead to the opposite outcome where six bidders choose to reduce demand and be satisfied with two blocks each to ensure that auction prices stay low. In fact, the German outcome was in between these two extremes. Deutsche Telecom continued pushing up the price when the market could have been divided among six bidders, but then ended the auction before working one or two of the participating newcomers out of the market. In Austria the

auction outcome was in line with the strategic demand reduction prediction, as six bidders were satisfied with two blocks each after only a few rounds of bidding. These two outcomes suggest that a high variance may exist in the outcomes of a simultaneous ascending auction in this kind of setup, but two observations are of course way to few to draw statistically meaningful conclusions.

In a series of experiments, Goeree, Offerman and Sloof (2004a) investigate which outcome prevails in the simultaneous ascending auction in this kind of setup. Another goal of this study is to compare the performance of the simultaneous ascending auction with the discriminatory auction, that is easily defined in this environment. In the discriminatory auction, bidders may simultaneously submit multiple sealed bids. The seller awards the identical objects to the highest bids at prices equal to the bids. In theory, the discriminatory auction does neither support the strategic demand equilibrium nor the preemptive equilibrium of the simultaneous ascending auction. Thus, a comparison of the auction formats depends on the equilibrium that tends to be selected in the simultaneous ascending auction.

In particular, Goeree et al. (2004a) consider the following situation. Two incumbents compete with one entrant for six identical objects. To keep the theoretical model tractable and the experimental procedure simple, bidders have flat demand for the objects offered for sale, which means that each bidder draws one private value that is valid for each object acquired. Each bidder can only buy up to three objects. If the entrant acquires one or more objects, a negative external effect is inflicted upon the incumbents. The incumbents can only avoid the negative external effect if they manage to keep the entrant completely out of the market. Subjects either participate in the discriminatory auction or in the simultaneous ascending auction. Each auction format is run through three different regimes, one where the negative external effect is absent, another one where a mild negative external effect exists and a final one with a strong negative external effect.

In the experiments, demand reduction in the simultaneous ascending auction is widespread. For each level of external effect, demand reduction is observed more frequently (60.4%, 25.5% and 30.3% for the no, low and high negative external effect respectively) than the preemptive bidding outcome (7.2%, 14.8% and 20.4% for the no, low and high negative external effect

respectively).[8] As a consequence, the ascending auction raises substantially lower revenue at a higher variance than the discriminatory auction: with no (mild) [strong] external effect, it raises 26.0% (51.0%) [51.9%] of the revenue of the discriminatory auction. Although equilibrium predictions track the observed revenues of the discriminatory auctions quite well, subjects' bidding behavior departs from predicted bidding in several ways. Perhaps the most striking departure is that subjects, unlike theory predicts, submit different bids for identical objects. This feature of actual bidding is the main reason why the discriminatory auction leads to higher entry than predicted in the equilibrium. Similar high levels of entry characterize the simultaneous ascending auction, because there incumbents and entrants tend to peacefully divide the supplied objects. From this perspective the two auction formats lead to similar levels of efficiency. However, the discriminatory auction does a better job in assigning the objects to the bidders with the higher private values, as in the demand reduction of equilibrium of the ascending auction each bidder receives two objects independently of her private value. In this sense, the discriminatory auction yields higher efficiency levels than the simultaneous ascending auction. This result contrasts strongly with the rhetorics of some popular texts on this topic, that advocate the simultaneous ascending auction as a useful tool to "put the licenses in the hands of the firms who value them most".

The bidding data of the simultaneous ascending auction reveal an empirical selection device. It turns out that incumbents only pursue the preemptive equilibrium when they have high private values. With low private values incumbents figure out that there is not much chance to keep the entrant out of the market. If one incumbent settles for demand reduction, the other bidders often reduce their demand quickly. An intuitive empirical finding is that with a stronger negative external effect the threshold value above which incumbents tend to pursue preemption decreases.

A sizable minority of the outcomes of the simultaneous ascending auction is characterized as a cheap preemptive attempt (8.1%, 22.1% and 18.9% for the no, low and high negative external effect respectively). Like in the German auction, in quite a few auctions the incumbents try

[8] There are also cases in practice where firms successfully 'divided the market' (see Cramton and Schwartz, 2000 for the FCC auctions and Jehiel and Moldovanu, 2001, and Grimm et al., 2003 for the auction for second generation mobile telecommunication licenses in Germany) .

to work the entrant out of the market, but if this attempt is not successful at moderate prices one incumbent chooses to reduce demand and the other incumbent and the entrant follow swiftly. In the regime with a mild negative external effect, the experimental auctions that are characterized as cheap preemptive attempts even outnumber the auctions characterized as true preemptive attempts (where incumbents are willing to bid their value or more on all three objects).

The study by Goeree et al. (2004a) complements previous experimental studies that investigate the phenomenon of demand reduction in situations without negative external effects. Alsemgeest, Noussair and Olson (1998) compare the ascending auction and a sealed bid auction where the price equals the lowest accepted bid. They consider two regimes, one where bidders only demand one unit and another where bidders demand two units. The sealed bid mechanism raises higher revenue than the ascending auction in both regimes. Alsemgeest et al. observe some demand reduction in the ascending auctions where bidders have demand for two units. In a field experiment, List and Lucking-Reiley (2000) observe some demand reduction when they sell sports cards in a uniform-price auction (the auction where all winning bidders pay a price equal to the highest rejected bid). Nevertheless, the uniform-price auction does not raise less revenue than the Vickrey auction because bidders tend to bid too high on their first units.[9] Kagel and Levin (2001) have a human bidder with two units demand participate against a computer in a uniform-price auction, an ascending auction and the Vickrey/Ausubel auction. The ascending auction produces a higher level of demand reduction than the uniform-price auction, even though a substantial level of demand reduction is observed in the latter auction. Engelmann and Grimm (2003) run five auction formats: the uniform-price, the ascending, the discriminatory, the Vickrey and the Ausubel auction. They observe a higher level of demand reduction in the ascending auction than in the uniform-price auction. The discriminatory auction outperforms the other auctions in terms of raising revenue.

Pooling across studies, the conclusion is that demand reduction is a real danger in uniform-

[9] In an m-object Vickrey auction bidders submit as many individual unit bids as they like. The top m bids are declared winning bids. For the k-th unit won by a bidder, she pays a price equal to the k-th highest of the rejected bids submitted by others. Demand reduction is not supported as an equilibrium phenomenon in a Vickrey auction. The Ausubel auction is a continuous version of the Vickrey auction.

price auctions and even more so in simultaneous ascending auctions. Surprisingly, demand reduction is even more likely than preemptive bidding in situations where entrants inflict negative externalities upon incumbents. In the setting where bidders may buy multiple homogeneous objects, the discriminatory auction performs better than the simultaneous ascending auction.

4.3 Auctions of multiple heterogeneous objects

Quite often governments consider auctions to assign heterogeneous licenses with the restriction that each bidder can obtain at most one license. The results reported in the previous section cannot directly be extrapolated to this case. Most importantly, the rules of the discriminatory auction need to be adapted when the licenses are heterogeneous and bidders are restricted to one license each. In a discriminatory auction with heterogenous licenses, it may occur that one and the same bidder has the highest bid on more than one license. In that case it may be true that the revenue for the seller is larger when this bidder is awarded her second highest instead of her highest bid. For instance, assume that there are two bidders and two licenses. Bidder 1 bids 20 on license A and 18 on license B, while bidder 2 bids 10 on license A and 1 on licence B. If license A (B) is awarded to bidder 1 (2), the revenue is 21, whereas the other allocation where bidder 1 wins license B raises a revenue of 28. Still, even if it is not straightforward to incorporate a first-price element in the auction, the evidence obtained for other cases suggests that it may be worthwhile to search for ways to introduce one.

Goeree, Offerman and Schram (2005) compare the standard simultaneous ascending auction with three versions of a discriminatory auction. The first version of the discriminatory auction is the so-called "simultaneous first-price auction". In this format, each bidder simultaneously submits (at most) one bid for each license for sale. The seller collects all the bids and assigns the licenses such that the revenue is as high as possible and each bidder is assigned at most one license. In the sequential first-price version the licenses are sold sequentially one after the other. In a sequential auction the seller has to decide whether she want to put up the best license first ("best foot forward") or whether she keeps the best licenses for last ("best for last"). In the simultaneous descending auction, for each license a clock starts moving down from a very

high price. The first bidder who stops the clock of a license buys the license and pays a price equal to the level of the clock. The clocks of the other licenses move further down until one is stopped by a second bidder. This bidder buys the license at a price equal to the level where she stopped the clock. This process continues until all available licenses are sold.[10]

The sale of licenses in the thin Dutch FM-radio market was taken as the motivating example for the study of Goeree et al. (2005), and the experiments were designed accordingly. That is, in each auction four bidders competed for three licenses. Subjects participated in one auction format only. They were run through different environments with increasing complexity. In the first environment, subjects drew a separate private value for each of the licenses. In this setting, there were differences in the quality of the licenses but bidders were symmetric (i.e., distributions that were used to draw the values varied across licenses but not across bidders). In the second environment asymmetries between the bidders were introduced. In these auctions, three "strong" bidders competed against one "weak" bidder (the support of the distributions of the private values of the weak bidder started below the support of the strong bidders and partly overlapped the support of the strong bidders). The third environment differed from the second in the sense that the private values of the bidders were multiplied by an unknown common value component (that reflected unknown developments in the demand of the product market). Each bidder received a private estimate of this common value component. In the final environment, an additional layer of realism was added to the experiment. Here, prior to the auction subjects sequentially made the decision whether or not to enter the auction at an opportunity cost.

Although this type of auction is intractable from a theoretical point of view, it is a useful test-bed for some of the intuitions of auction experts. As said, Klemperer (2002b) pointed out that in uncompetitive circumstances, first-price auctions may be more attractive to entrants since they offer the possibility of surprise. In a first-price auction, an incumbent firm may be too optimistic about the profit-margin in its bid such that there is scope for a weak entrant to submit a higher bid. Another intuition of auction experts is that in a sequential auction with

[10] The authors do not consider a right-to-choose auction, which could be an interesting alternative in environments with heterogeneous objects (see section 6).

common value uncertainty, it is better to save the best licenses for last. The idea is that in the sale of the first licenses information about the common value component is released. With less uncertainty bidders can bid more aggressively later in the auction. From this perspective, it is better to sell the good licenses last, because then the most aggressive bidding is observed on the superior licenses. In agreement with this intuition, in March 2000 the Swiss auctioned three nation-wide wireless-local-loop licenses in increasing order of value. Surprisingly, the first two identical licenses were sold for 121 and 134 million francs respectively, while the third superior license sold for only 55 million francs. This result suggests that it may be wiser to sell the good licenses first when all interested bidders are present to drive up the price, although again it is hard to generalize from one observation only.

The experimental results of Goeree et al. (2005) underline the strength of first-price auctions when it comes to raising high revenue at low variance. From the revenue perspective all first price auctions perform better than the simultaneous ascending auction. In line with the Swiss experience, the sequential first-price auction where the seller uses the best foot forward strategy provides the most promising way to incorporate a first-price element when the goal is to generate high revenue: pooled across all environments, it raises 15% more revenue than the simultaneous ascending auction. The high revenues of the first-price auctions come at a price for efficiency though. The simultaneous ascending auction is the superior mechanism when the goal is to assign the licenses to the bidders with the highest private values. Notice that here the simultaneous ascending auction is not hindered by the possibility of demand reduction, because each bidder is restricted to one license.

The experimental results suggest a refinement of the idea that first-price auctions are conducive to entry of weak bidders. In fact, with an exogenous number of bidders, the first-price auctions promise less profits to the winners of the auction. As expected, first-price auctions increase the prospects for weak bidders to win a license relative to second-price auctions. At the same time however, the amount that winners win is less than in the simultaneous ascending since first-price auctions elicit more aggressive bidding. As a result, Goeree et al. do not observe more entry by weak bidders in first-price auctions than in the simultaneous ascending auction. In fact, with endogenous entry the differences in the auctions' performances tend to

diminish.

Finally, the simultaneous ascending auction shows the highest incidence of winner's curse outcomes and the highest variance in per license profits. These facts may cause feelings of injustice and may result in bankruptcies or costly lawsuits after the auction.

4.4 Summary

In this section, we have observed that laboratory experiments confirm Klemperer's (2002b) conjecture that first-price auctions perform better than second-price auctions in uncompetitive environments. First-price auctions raise more revenue at a lower variance than second-price auctions. The results regarding efficiency depend on the setting. If the possibility of demand reduction exists, the simultaneous ascending auction may achieve a lower level of efficiency. If bidders are restricted to one license each, second-price auctions and the simultaneous ascending auction are more efficient than first-price auctions, but often the differences are small, and sometimes not even statistically significant.

5 Hybrid Auctions

In the previous section, we concluded that in thin markets, first-price auctions often outperform second-price auctions. One disadvantage of first-price auctions relative to second-price auctions could be that in second-price auctions, bidders can take into account information that is released by the bidding strategies of other bidders. This feature is even more prominent in the English auction, as bidders observe when others leave the auction. In order to combine the best of both worlds, Klemperer (1998) proposes the Anglo-Dutch auction as an alternative auction format that could work well in uncompetitive markets. The Anglo-Dutch auction is a hybrid auction that combines features of first-price and second-price auctions. We discuss this auction in subsection 5.1. Other hybrid auctions that are tested in the lab are the Dutch-Anglo auction and premium auctions, which are the topics of subsections 5.2 and 5.3. Subsection 5.4 summarizes the main findings in these experiments.

5.1 The Anglo-Dutch auction

The Anglo-Dutch auction has the following rules for k identical licenses and the restriction of one license per bidder. The auction consists of two stages:

1. The auctioneer raises the price round-by-round. In each round, each bidder can decide to leave or stay in the auction. If a bidder leaves the auction, she cannot re-enter. This round ends as soon as no more than $k+1$ bidders are left.

2. The $k+1$ remaining bidders submit a sealed-bid, which should at least be equal to their highest first-stage bid. The highest k bidders win a license. The amount they pay depends on the format. In the *discriminatory* Anglo-Dutch auction, each winner pays her bid. In the *uniform* Anglo-Dutch auction, each winner pays the kth highest bid.

Note that this auction is a hybrid between two types of standard auctions. Klemperer (1998) conjectures that the Anglo-Dutch format combines the best of both worlds. The first-price element (stage 2) makes the auction attractive for weak bidders, as they are better able to compete against the strong ones. Moreover, in the ascending phase (stage 1), bidders can deduce valuable information from the other bidders' behavior, which could increase the efficiency of the auction and the seller's revenue, as bidders are less likely to be prone to the winner's curse. To our knowledge, this claim has not been proved yet in a formal model.

In an experiment, Abbink et al. (2005) explore a design that resembles the UK UMTS auction as closely as possible. They study the situation that was prevalent two years before the actual auction, when the British Radiocommunication Agency (BRA) intended to sell four almost identical licenses (instead of five as in the actual auction). As the number of incumbents was also four, BRA's advisors feared that a simultaneous ascending auction would hardly attract entry, as they expected entrants to have little chance to win a license in this format (Binmore and Klemperer, 2002). Abbink and his coauthors compare the performance of the discriminatory Anglo-Dutch auction, the uniform Anglo-Dutch auction and the English auction.

They study a situation with eight bidders: four incumbents (denoted by INC) and four potential new entrants (NEW). The value of a license to bidder i is the sum of a private value component v_i (which may differ between bidders) and a common value component c (which is the same for all bidders). Each bidder i draws v_i from the set $\{-100, -99, ..., +100\}$. INC-type [NEW-type] bidders draw their value with 80% [20%] probability independently and uniformly from the set $\{0, 1, ..., 100\}$ and with the remaining 20% [80%] probability from $\{-100, -99, ..., 0\}$. The common value c is randomly drawn from the set $\boldsymbol{C} \equiv \{1000, 1001, ..., 1500\}$. The bidders were not informed about c or $\boldsymbol{C}$, but each bidder i received a signal s_i which was commonly known to be independently and uniformly drawn from the set $\boldsymbol{D}(c) \equiv \{c - 200, c - 199, ..., c + 200\}$.

In a closely related setting, Goeree and Offerman (2002) show both theoretically and experimentally that an efficient outcome is not always reached. The reason is that an efficient allocation can only be realized if bidders condition their bids only on their private information. But, naturally, the equilibrium bids are strictly increasing in a bidder's signal on the common value. Inefficient allocations arise as in equilibrium, bidders with a high signal on the common value and a low private value could outbid bidders with a low signal on the common value and a high private value.

Abbink et al. compare the auctions on several dimensions. They define the surplus of an allocation as the sum of the values the winners realize. The efficiency e of the auction outcome is the difference between the actual surplus (S^{act}) and the minimal surplus (S^{min}) as a fraction of the difference between the maximal surplus (S^{max}) and the minimal surplus:

$$e = \frac{S^{\text{act}} - S^{\text{min}}}{S^{\text{max}} - S^{\text{min}}}$$

They also look at entry, i.e., the number of NEW-type bidders that win a license. Moreover, they measure revenue as standardized revenue, i.e., total revenue as fraction of S^{max}. The following table summarizes the outcomes of the three auction formats relative to the efficient outcome.

The experiment of Abbink et al. shows that the nice properties of first-price auctions in thin

Auction	Efficiency	Entry	Revenue
Discriminatory Anglo-Dutch	74.9%	1.5	93.8%
Uniform Anglo-Dutch	69.8%	1.7	95.1%
English	69.4%	1.9	94.7%
Efficient	100%	1.1	

Table 1: Efficiency, entry, and revenue in Abbink et al.'s (2005) experiment.

markets relative to second-price auctions, do not carry over to the Anglo-Dutch auction, even though this auction is expected to combine the best of first-price and second-price auctions. The Anglo-Dutch formats that are considered do not outperform the English auction in any of three dimensions: efficiency, entry, and revenue.

Binmore and Klemperer (2002) criticize Abbink et al.'s design in the sense that potential newcomers were relatively strong: on average, one entrant has a higher value than at least one incumbent. It may be more realistic that all incumbents have higher values than all entrants, and that this is common knowledge to all bidders. Moreover, entry into the auction was costless, while in practice, bidders do have to make some costs to enter. Although these costs may be negligible relative to the total value of a license, if a bidder expects not to win a license, she will not enter. Therefore, it remains interesting for future research to see how Anglo-Dutch auctions perform in other circumstances.

5.2 The Dutch-Anglo auction

The 'Dutch-Anglo auction' is another hybrid auction that is sometimes used in practice to sell a single indivisible object. The auction consists of the following two stages:

1. All bidders submit a sealed bid. The highest bidder wins and pays her bid if the difference between her bid and the second highest bid is at least equal to some $z \in (0, \infty)$. Otherwise, all bidders for whom the difference between the highest bid and their own bid is smaller than z, enter stage 2.

2. The remaining bidders then play the English auction with a reserve price equal to the

highest bid in the first stage.

This auction format was used by the Brazilian government in the partial privatization of the telecommunications firm Telebras (Dutra and Menezes, 2001). Dutra and Menezes (2001) study an auction in the lab that is closely related to the Dutch-Anglo auction: the Dutch-Vickrey auction. The Dutch-Vickrey auction has the same rules as the Dutch-Anglo auction with the exception that in the second stage, the Vickrey auction is played instead of the English auction. Dutra and Menezes make the following additional assumptions:

- Three bidders compete for the object.
- Each bidder draws her private value for the object independently from the set $\{0, 3, 6\}$, where the first [second, third] element has probability 0.4 [0.3].

Dutra and Menezes (2002) solve for the equilibrium of this auction, and show that for risk neutral bidders, the Dutch-Vickrey auctions revenue dominates the first-price sealed-bid auction.[11] Moreover, both auction formats are efficient in the sense that the winner of the auction is the bidder who attaches the highest value to the object.

In their experiment, Dutra and Menezes (2001) compare the first-price sealed-bid auction with two instances of the Dutch-Vickrey auction, with $z = 1$ and $z = 1\frac{1}{2}$. They do so using a within-subject design. In the first six rounds, the object was sold in the first-price sealed-bid auction. In rounds 7-12 [13-18], the Dutch-Vickrey auction with $z = 1$ [$z = 1\frac{1}{2}$] was used to sell the object.

Dutra and Menezes' (2001) experiment partly confirms the theory. They find that both auction formats are highly efficient. However, they cannot reject the hypothesis that the three auctions yield the same expected revenue. Most observations are higher than the risk-neutral equilibrium, but none of the auction formats revenue dominates the others.

[11] The theoretical revenue result may be an artefact of the chosen design. When bidders draw their private values from the same continuous distribution, the revenue equivalence theorem seems to apply.

Still, Dutra and Menezes' experiment leaves several questions open for further research: Would a between-subject treatment induce the same results as the within-subjects treatment that the researchers choose? How does the Dutch-Vickrey auction perform relative to the English auction? How does the Dutch-Vickrey auction perform in thin markets?

5.3 Premium auctions

In practice, hybrid auctions often offer a premium to the highest losing bidder. Premium auctions are employed in Belgium and the Netherlands since the Middle Ages (Sikkel, 2001). These auctions are used to sell houses, land, boats, machinery and equipment. Many Belgium and Dutch cities claim that their own variant is unique in the world. Although actual premium auctions differ in the institutional details, they all share the feature of offering a premium to the highest losing bidder in order to stir up competition. Goeree and Offerman (2004) consider a stylized premium auction that captures the essential features of a premium auction. They refer to this format as the Amsterdam auction, since Amsterdam has a particular prominent history of premium auctions.

Goeree and Offerman investigate the situation where a seller sells a single object in an Amsterdam auction. Like the Anglo-Dutch and the Dutch-Anglo auction, the Amsterdam auction consist of two phases.

1. The price level rises until all but two bidders have dropped out. Both bidders enter phase 2.

2. The level at which the last bidder dropped out of the first phase is the reserve price for the second phase. The two remaining bidders submit a sealed bid in the second phase that must be at least as high as the reserve price of the first phase. The highest bidder wins the object for sale and pays a price equal to the own bid in the first-price Amsterdam auction and a price equal to the other finalist's sealed bid in the second-price Amsterdam auction. Both finalists receive a premium proportional to the difference of the lowest

sealed bid in the second phase and the reserve price.[12]

In a symmetric private value model, the Amsterdam auction is revenue equivalent with standard auctions like the English auction or first-price auction (the outcome is efficient and with more than two bidders, the bidder with the lowest possible value expects to earn a zero profit like in standard auctions). However, actual real-estate auctions tend to be characterized by asymmetries between bidders. Usually genuinely interested buyers compete with speculators out for a bargain. Premium auctions may provide an effective tool to exploit asymmetries between bidders. The premium stimulates weak bidders (speculators) to set an endogenous reserve price for the strong bidders (the genuinely interested buyers). Competition between weak bidders dissipates the premium that they can earn. For the single unit case, the optimal auction is well known (Myerson, 1981) and can easily be used as a benchmark.

In a series of experiments, Goeree and Offerman compare the performance of the two Amsterdam auctions, the first-price auction, the English auction and the optimal auction under varying degrees of asymmetry. They vary the auction format across sessions and the degree of asymmetry within sessions. First, four subjects compete in a symmetric environment where all draw an independent private value from a uniform U[0,60] distribution. Then mild asymmetries are introduced as one strong bidder who receives a private value from a U[40,100] distribution competes against three weak bidders who continue to draw their values independently from a U[0,60] distribution. The final part makes use of strong asymmetries where the distributions of the strong buyer U[70,100] and the weak buyers U[0,60] do not overlap anymore. In the premium auctions the two finalists receive a premium of 30% of the difference of the lowest sealed bid and the reserve price.

In the symmetric environment, the two versions of the Amsterdam auction raise a roughly equal revenue as the English auction, as theory predicts. As usual, the English auction is out-

[12]Notice that it is quite straightforward to design a k-identical objects Amsterdam auction. One possibility is that in the first phase the price rises until k+1 bidders are left. These bidders submit a sealed bid in the second phase, the highest k bids win an object and bidders in the second phase receive a premium proportional to the difference of the lowest sealed bid and the reserve price. Other implementations also seem plausible and future work should identify the best way to generalize the Amsterdam auction to the multi-unit case.

performed by the first-price auction when bidders are symmetric. The first price auction raises about 10% more revenue at a lower variance. With mild asymmetries, the first-price auction, the first-price Amsterdam auction and the optimal auction all generate roughly 20% more than the English auction. With strong asymmetries, the second-price Amsterdam auction and the optimal auction raise about 15% more revenue than the first-price auction, and about 45% more revenue than the English auction. Interestingly, the optimal auction hardly outperforms the Amsterdam auction.

The English auction performs better in assigning the object to the bidder with the highest private value (although the efficiency differences with the first-price auction are surprisingly small (98.5% for the English auction versus 96.7% for the first-price auction, pooled across all environments). Raising revenue through an Amsterdam auction comes at a cost of reducing efficiency (efficiency levels drop to 94.0% in the first-price Amsterdam auction, 90.3% in the second-price Amsterdam auction and 91.2% in the optimal auction). On the other hand, the Amsterdam auctions stimulate serious participation by bidders with low private values. In contrast to the English auction, weak bidders are more inclined to participate in the Amsterdam auctions as asymmetries grow. In fact, with strong asymmetries 39.1% of the bids in the English auction are almost zero, while in the Amsterdam auctions only 4.2% (first-price Amsterdam) and 6.5% (second-price Amsterdam) of the bids are almost zero. As such, the Amsterdam auction may be an attractive format to pursue efficiency if a government wants to stimulate entry into a collusive industry.

5.4 Summary

Among the hybrid auctions (Anglo-Dutch, Dutch-Anglo, Amsterdam), the results for the Amsterdam auction seem most promising. In the environments that are investigated, both the Anglo-Dutch auction and the Dutch-Anglo auction do not outperform standard auctions. This is in sharp contrast to the Amsterdam auction in which the seller stirs up competition by offering a premium to the highest losing bidder. In environments with asymmetries between bidders, the Amsterdam auction raises substantially more revenue than a standard ascending

auction. Moreover, it occurs relatively more often that a newcomer wins a license even if it is less cost-effective than an incumbent firm.

6 Selling the right to choose

In this section, we consider the case in which a seller sells multiple heterogeneous objects of comparable value. When auctioneers fear that there is little demand for each object separately, they often switch to using so-called "right-to-choose auctions" or "pooled auction". In these auctions, the auctioneer gives the winners the opportunity to choose their most favorite object. In subsections 6.1 and 6.2, we focus on right-to-choose auctions and pooled auctions respectively. Section 6.3 includes a summary.

6.1 Right-to-choose auctions

In a right-to-choose auction (also known as "bidder's choice auctions" or "choice selling"), bidding proceeds in multiple rounds. The high bidder of the first round selects the object that she likes best. This item is deleted from the list and the remaining objects are offered in the second round. Again the high bidder of this round selects her preferred object which is then deleted from the list of objects, and so on, until all objects are sold. This is a common method to sell real-estate and condominiums in the US (Ashenfelter and Genesove, 1992). It is also frequently employed to sell jewelry, movie posters, lamps, watches, pottery, glassware and militaria. In 2004, the Dutch government used a right-to-choose auction to assign scarce phone numbers to businesses.

Auctioneers sometimes switch to using this auction format because they believe that there are situations where right-to-choose auctions are the superior revenue raising mechanism. Auction law attorney Steve Proffitt describes this belief as follows: "Choice selling sometimes has the ability to generate higher selling prices than straight item-by-item offerings. Auctioneers often don't know who has the most interest in what piece and how much that bidder might be willing to pay to own it. The auctioneer wants to push that bidder as far as he can (and he should), so long as it's done legally and ethically. One way to do this is

to force bidders to compete who would not otherwise do so, and choice selling does that" (http://www.maineantiquedigest.com/articles/feb04/ethi0204.htm).

Burguet (1999, 2004) provides a theoretical investigation of right-to-choose auctions when a seller sells two items and each bidder's private values for the two items can be summarized by a single parameter. When bidders are risk neutral, the right-to-choose auction is revenue equivalent to standard good by good auctions (like a sequential Vickrey auction and a simultaneous ascending auction). With risk averse bidders, right-to-choose auctions generate higher revenues than the sequential Vickrey auction or the simultaneous ascending auction: the intuition is that risk-averse bidders dislike the uncertainty arising from not winning the first round -the winner of the first round might pick their most preferred good- and as a consequence they are willing to raise their bids in the first round.[13] Eliaz, Offerman and Schotter (2004) extend the theoretical analysis to the case where the seller sells $k > 2$ objects and where there are $n > k$ interested bidders for each object. Again, revenue equivalence results for the case where bidders are risk-neutral. Eliaz et al. identify an interesting trick that a seller may use in a right-to-choose auction. By selling rights to choose instead of heterogenous objects, the seller artificially creates monopoly power. As a consequence, the seller may raise more revenue by restricting the number of rights to choose offered for sale. Selling fewer rights than available goods enhances the expected revenue of the seller.

Experimental tests of right-to-choose auctions are provided by Goeree, Plott and Wooders (2003) and Eliaz et al. (2004). Goeree et al. (2003) consider the case where a seller sells two heterogenous goods to four buyers through either an ascending right-to-choose auction or a simultaneous ascending auction. In both auctions, bidders who do not have the current high bid have the possibility to submit a higher bid than the current high bid. Bidding stops if in a given time interval no new bid is submitted. Bidders can buy at most one good. The experimental results confirm the revenue raising virtues of the ascending right-to-choose auction: the right-to-choose auction raises almost 20% more revenue than the simultaneous ascending auction. Maximum likelihood estimation shows that a utility function of $u(x) = x^{0.39}$ matches

[13] Burguet also explains why prices decline in successive rounds of a righ-to-choose auction, a phenomenon observed by Ashenfelter and Genesove (1992).

the data well. Thus, consistent with Burguet, the authors are able to explain their data with risk aversion. Interestingly, the increase in revenue hardly affects the efficiency level of the auctions: the right-to-choose auction achieves an efficiency level of 98.4%, only slightly below the efficiency level of the simultaneous ascending auction (100.0%).

In their experiment, Eliaz et al. (2004) use a tougher benchmark to judge the performance of the second-price right-to choose auction. They compare how this auction performs in comparison to the theoretically optimal auction. For completeness, they also consider a standard Vickrey good by good auction. All auctions are implemented via the second-price format, which means that the highest bidder wins and pays a price equal to the second highest bid. In particular, Eliaz et al. focus on the setup where a seller has four heterogeneous goods for sale while each good only attracts two interested buyers. They also investigate the possibility of reducing quantity via a right-to-choose auction, by having a treatment where the seller offers three rights to choose from the four available goods. The right-to-choose auction strongly outperforms the standard good-by-good auction in terms of raising revenue: the seller collects 40% more revenue in the right-to-choose auction. Like in Goeree et al., the right-to-choose auction achieves an efficiency level close to the one of the good-by-good auction (98.2% for the right-to-choose auction versus 98.3% for the good-by-good auction).

Surprisingly, the right-to-choose auction even raises more revenue than the theoretically optimal auction. Indeed, the bids in the right-to-choose auction are so competitive that it is not possible to explain them with reasonable degrees of risk aversion. Instead, the authors argue that the right-to-choose auction induces a competitive mindset such that only bidders with very high private values dare to submit bids below value in early phases as all are supposed to do. Given the staggering amounts raised in the right-to-choose auction, there is hardy room for improvement when the seller reduces quantity. Therefore, it is not surprising that with quantity reduction the seller raises a roughly similar revenue as without quantity reduction. Nevertheless, a right-to-choose auction with quantity reduction is still an interesting option because it is the format with the smallest variance in revenue.

6.2 Pooled auctions

A related type of auction is the so-called "pooled auction". In this auction, $n > 2$ bidders simultaneously submit one sealed bid for k heterogeneous objects, $n > k > 1$. The highest bidder chooses her preferred object and pays a price equal to the own bid. The second highest bidder then chooses from the list of remaining objects and pays a price equal to the own bid. This procedure continues until all objects are sold. Menezes and Monteiro (1998) characterize the risk neutral equilibrium of the pooled auction. Salmon and Iachini (2003) test the pooled auction in a experiment, and compare its performance with that of a simultaneous ascending auction. Like the right-to-choose auctions, pooled auctions elicit competitive bidding and thereby raise higher revenue than the simultaneous ascending auction. The disadvantage of a pooled auction is that bidders often experience losses, when they have to pay a high bid while their preferred objects have already been chosen by others. In practice, these losses may lead to bankruptcy or costly renegotiation by bidders who prefer to get out of the deal.

6.3 Summary

Both right-to-choose auctions and pooled auctions induce more competitive bidding in thin markets with heterogeneous objects relative to standard auctions. In the lab, both auctions are found to raise considerably more revenue than standard auctions, while achieving an almost equal level of efficiency as the simultaneous ascending auction. For pooled auctions, these advantages may come at the cost of overbidding.

7 Conclusion

Auction design matters in thin markets. Experiments demonstrate that if the number of bidders is small compared to the number of objects for sale, details of the auction may have tremendous effect on the outcome. For standard auctions, the most robust results pertain to the revenue raising qualities of the mechanisms. First-price auctions consistently raise higher revenue than the simultaneous ascending auction and they do so at a lower variance. This is even true for cases where standard theory predicts the opposite result.

The results for efficiency are more delicate and depend in more detail on the specific situation of the auction. For instance, when a government sells multiple licenses and each bidder can buy more than one license, experiments show that the theoretical possibility of demand reduction in the simultaneous ascending auction is indeed a real practical danger. With demand reduction, each bidder gets part of the cake at a low price irrespective of its cost-effectiveness. Therefore, against the popular wisdom, the simultaneous ascending auction may not succeed in putting the licenses in the hands of the firms that value them most. In fact, in such a situation a discriminatory auction does a much better job in fulfilling that goal. Interestingly, demand reduction also leads to the rejection of another popular wisdom, which says that the simultaneous ascending auction discourages entry because it facilitates the process by which incumbent firms can keep newcoming firms out of the market. If newcomers anticipate a demand reduction outcome, the ascending auction may be very attractive to them because it gives them a foothold in the industry for a very low price. In fact, the experimental study by Goeree et al. (2004a) demonstrates that the two opposite forces approximately balance each other, and that newcomers enter at a similar pace in a discriminatory auction as they do in a simultaneous ascending auction.

If, on the other hand, each bidder is restricted to one license only, demand reduction is excluded by design. In such situations, simultaneous ascending auctions usually do the better job in assigning the licenses to the cost effective firms, although differences in efficiency levels between auction formats are often not that large. However, Goeree et al. (2005), in an experiment that endogenizes the entry decision, do not find support that first-price auctions attract more entry than the simultaneous ascending auction. It is indeed the case that in first-price auctions, a newcomer has a higher chance to win a license, but it is also true that the amount that they win tends to be less due to aggressive bidding. In the aggregate, the two opposite effects tend to cancel.

The experimental literature has also identified some promising "exotic" auctions. Among the hybrid auctions that are held in two phases (Anglo-Dutch, Amsterdam, Dutch-Anglo) the results for the Amsterdam auction seem most promising. By offering a premium to the highest losing bidder, the seller stirs up competitive bidding by weak bidders. In environments with

asymmetries between bidders, the Amsterdam auction raises substantially more revenue than a standard ascending auction. In an Amsterdam auction, it occurs relatively more often that a newcomer wins a license even if it is less cost-effective than an incumbent firm. Such outcomes may be desirable if governments want to shake up a collusive industry by selling a license to a newcomer.

Another interesting format that is often used in practice is a right-to-choose auction. In this format bidders bid for the right to choose between heterogeneous objects. This mechanism makes bidders compete who are interested in different objects and who would not compete otherwise. In thin markets where there is little interest per license, a right-to-choose auction may raise considerably more revenue than a standard ascending auction, while it achieves an almost equal level of efficiency as the simultaneous ascending auction does. As such, this auction, like the Amsterdam auction, is a prime candidate for practical mechanism design.

There is still a lot to be done.[14] A shortcoming of all studies mentioned in this chapter is that they aim to say something about efficiency in the aftermarket without actually modelling the aftermarket. In current experimental work, the aftermarket is reduced to a set of private values. This is not optimal for a couple of reasons. First, it blurs issues that relate to consumer surplus in the aftermarket. Second, Offerman and Potters (2006), the only experimental study that explicitly models the auction together with the aftermarket, shows that auctioning licenses may affect how firms behave in the subsequent market. In particular, in an oligopoly setting, the prices that firms pay to obtain a license enhance the prices that firms charge to consumers in the aftermarket. The explanation offered by Offerman and Potters is that auctions enhance the firms' willingness to embark on a risky but profitable collusive price path. In agreement with this hypothesis, they do not find a price-enhancing effect of auctioning a monopoly license, where collusion is excluded by design. The message of their study is that new insights can be gained if the interaction of firms in license auctions and aftermarkets is not reduced to a simple auction game. In our view, this opens an interesting avenue for future experimental research.

[14] If you plan to do an experiment yourself it is interesting to surf to Jacob Goeree's webpage, who is currently developing software that can be used for laboratory experiments on auctions. See www.hss.caltech.edu/~jkg/jAuctions.html.

The challenge is to design richer games that explicitly model the license auction together with the important processes before and after the auction.

8 References

Abbink, K., B. Irlenbusch, P. Pezanis-Christou, B. Rockenbach, A. Sadrieh, and R. Selten (2005). "An experimental test of design alternatives for the British 3G/UMTS auction." European Economic Review 49, 505-530.

Abbink, K. and B. Rockenbach (2005). "Option Pricing by Students and Professional Traders: A Behavioural Investigation." Managerial and Decision Economics, forthcoming.

Alsemgeest, P., C. Noussair, and M. Olson (1998). "Experimental Comparisons of Auctions under Single- and Multi-unit Demand." Economic Inquiry 36, 87–97.

Ashenfelter, O. and D. Genesove (1992). "Testing for Price Anomalies in Real-Estate Auctions." American Economic Review 60, 501-505.

Ausubel, L.M. and P. Cramton (2002). "Demand Reduction and Inefficiency in Multi-Unit Auctions." Working paper, University of Maryland.

Ballard, C. L., J. B. Shoven, and J. Whalley (1985). General Equilibrium Computations of the Marginal Welfare Costs of Taxes in the United States, American Economic Review 75, 128-38.

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

Bulow, J., M. Huang, and P. Klemperer (1999). "Toeholds and Takeovers." Journal of Political Economy 107, 427-454.

Burguet, R. (1999). "Right-to-choose in Oral Auctions." Working Paper, Institute for Economic Analysis, Barcelona, Spain.

Burguet, R. (2004). "The Condominium Problem; Auctions for Substitutes." Review of Economic Design, forthcoming.

Cramton, P. and J. Schwartz (2000). "Collusive Bidding: Lessons from the FCC Spectrum Auctions." Journal of Regulatory Economics 17, 229–252.

Das Varma, G. (2002). "Standard Auctions with Identity-dependent Externalities." Rand Journal of Economics 33, 689–708.

Dutra, J. C. and F. Menezes (2001). "Hybrid Auctions II: Experimental Evidence." Working Paper, Australian National University.

Dutra, J. C. and F. Menezes (2002). "Hybrid Auctions." Economics Letters 77, 301-307.

Dyer, D., J. H. Kagel, and D. Levin (1989). "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis." Economic Journal 99, 108-115.

Eliaz, K., T. Offerman, and A. Schotter (2004). "Creating Competition Out of Thin Air: Market Thickening and Right-to-Choose Auctions." Working paper, University of Amsterdam.

Engelbrecht-Wiggans, R. and C. Kahn (1998). "Multi-unit Auctions with Uniform Prices." Economic Theory 12, 227–258.

Engelmann, D. and V. Grimm (2003). "Bidding Behavior in Multi-unit Auctions: An Experimental Investigation and Some Theoretical Insights." Working paper, Humboldt University of Berlin.

Friedman, D. and S. Sunder (1994). *Experimental Methods - A Primer for Economists* (Cambridge University Press).

Gilbert, R. and D. Newbery (1982). "Preemptive Patenting and the Persistence of Monopoly." American Economic Review 72, 514-526.

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

Goeree, J. K. and T. Offerman (2004). "The Amsterdam Auction." Econometrica 72, 281-294

Goeree, J. K., T. Offerman, and R. Sloof (2004a). "Demand Reduction and Preemptive Bidding in Multi-Unit License Auctions." Working paper, University of Amsterdam.

Goeree, J., C. Plott, and J. Wooders (2004b). "Bidders' Choice Auctions." Journal of the European Economic Association 2, 504-515.

Goeree, J. K., T. Offerman, and A. Schram (2005). "Using First-Price Auctions to Sell Heterogeneous Licenses." Working paper, University of Amsterdam.

Grimm, V., F. Riedel and E. Wolfstetter (2003). "Low Price Equilibrium in Multi-Unit Auctions: The GSM Spectrum Auction in Germany." International Journal of Industrial Organization 21, 1557-1569.

Güth, W., R. Ivanova-Stenzel, and E. Wolfstetter (2004). "Bidding Behavior in Asymmetric Auctions." Working paper, Humboldt University.

Hoppe, H., P. Jehiel, and B. Moldovanu (2005), "License Auctions and Market Structure." Journal of Economics and Management Strategy, forthcoming.

Janssen, M.C.W. (2005). "Auctions as Coordination Devices." European Economic Review, forthcoming.

Jehiel, P. and B. Moldovanu (2000). "Auctions with Downstream Interaction among Buyers." Rand Journal of Economics 31, 768–791.

Jehiel, P. and B. Moldovanu (2001). "The European UMTS/IMT-2000 License Auctions." Working paper, University College London.

Jehiel, P., B. Moldovanu, and E. Stacchetti (1996). How (Not) to Sell Nuclear Weapons. American Economic Review 86, 814–829.

Kagel, J. and D. Levin (2001). "Behavior in Multi-unit Demand Auctions: Experiments with Uniform Price and Dynamic Auctions. Econometrica 69, 413–454.

Kagel, J.H. (1995). "Auctions: A Survey of Experimental Research." In: J.H. Kagel and A.E. Roth (eds.), The Handbook of Experimental Economics, Princeton, NJ, Princeton University Press, 501-585.

Klemperer, P. (1998). "Auctions with Almost Common Values: The "Wallet Game" and Its Applications." European Economic Review 42, 757-769.

Klemperer, P. (2002a). "How (Not) to Run Auctions: the European 3G Telecom Auctions." European Economic Review 46, 829-845.

Klemperer, P. (2002b). "What Really Matters in Auction Design." Journal of Economic Perspectives 16, 169-189.

Klemperer, P. (2004). *Auctions: Theory and Practice* (Princeton University Press).

Krishna, V., (2002). *Auction Theory* (Academic Press, London).

List, J. and D. Lucking-Reiley (2000). "Demand Reduction in Multiunit Auctions: Evidence

from a Sportscard Field Experiment." American Economic Review 90, 961–972.

Maskin, E.S. and J.G. Riley (2000). "Asymmetric Auctions." Review of Economic Studies 67, 413-438.

McMillan, J. (1994). "Selling Spectrum Rights." Journal of Economic Perspectives 8, 145-162.

Menezes, F. and P. Monteiro (1998). "Simultaneous Pooled Auctions." Journal of Real Estate Finance and Economics 17, 219-232.

Milgrom, P.R. and R.J. Weber (1982). "A Theory of Auctions and Competitive Bidding." Econometrica 50, 1089-1122.

Motta, M. (2004). *Competition Policy* (Cambridge University Press, Cambridge).

Myerson, R.B. (1981). "Optimal Auction Design." Mathematics of Operations Research 6, 58-73.

Noussair, C. (1995). "Equilibria in a Multi-object Uniform-price Sealed Bid Auction with Multi-unit Demands." Economic Theory 5, 337–351.

Offerman, T. and J. Potters (2006). "Does Auctioning of Entry Licenses Induce Collusion? An Experimental Study." Review of Economic Studies, forthcoming.

Pezanis-Christou, P. (2002). "On the Impact of Low-balling: Experimental Results in Asymmetric Auctions." International Journal of Game Theory 31, 69-89.

Plott, C.R. (1997). "Laboratory Experimental Testbeds: Applications to the PCS Auction." Journal of Economics and Management Strategy 6, 605-638.

Plum, M. (1992). "Characterization and Computation of Nash-equilibria for Auctions with Incomplete Information." International Journal of Game Theory 20, 393-418.

Roth, A. E. (2002). "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." Econometrica 70, 1341-1378.

Salmon, T. and M. Iachini (2003). "Risk Aversion vs. Loss Aversion in Pooled Auctions." Working paper, Florida State University.

Sikkel, M. (2001). "Plokken en Mijnen," Parool (main Dutch newspaper), April 3, 2.

Van Damme, E. (2002). "The European UMTS Auctions." European Economic Review 46, 846-858.

Vickrey, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." Journal of Finance 16, 8-37.

Wilson, R. B. (1987). "Game-Theoretic Approaches to Trading Processess." In T. Bewley (ed.), *Advances in Economic Theory: Fifth World Congress* (Cambridge University Press, Cambridge).