Going, Going, Gone! A Swift Tour of Auction Theory and Its Applications*

Emiel Maasland† and Sander Onderstal‡

June 21, 2005

Abstract

This paper provides a swift tour of auction theory and its applications. Among the questions it considers are: How much do bidders bid in commonly studied single-object auctions? How efficient are these auctions? How much revenue do they generate? Which single-object auction maximizes the seller's expected revenue? What is the best way to auction incentive contracts? And, how efficient and complex are multi-object auctions?

Keywords: Auctioning incentive contracts; Auctions; Efficiency; Equilibrium bidding; Optimal auctions;

JEL classification: D44

Running title: Auction Theory

*We thank Bastiaan Overvest, Marta Stryszowska and two anonymous referees for their helpful comments.

†SEOR, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, +31-10-408-1513, fax +31-10-408-9650, emaasland@few.eur.nl, www.seor.nl/ecri/staff/emiel.html.

‡University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands, +31-20-525-7161, fax +31-20-525-5591, 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 INTRODUCTION

In the past few decades, auction theory has become one of the most active research areas in economic sciences. The focus on auctions is not surprising, as auctions have been widely used over thousands of years to sell a remarkable range of commodities. One of the earliest reports of an auction is by the old Greek historian Herodotus of Halicarnassus, who writes about men in Babylonia around 500 B.C bidding for women to become their wives.[1] Perhaps the most astonishing auction in history took place in 193 A.D. when the Praetorian Guard put the entire Roman Empire up for auction. Didius Julianus was the highest bidder. However, he fell prey to what, today, is known as the winner's curse: he was beheaded two months later when Septimus Severus conquered Rome.[2]

Nowadays, the use of auctions is widespread. There are auctions for perishable goods such as cattle, fish, and flowers; for durables including art, real estate, and wine; and for abstract objects like treasury bills, licenses for UMTS,[3] and electricity distribution contracts.[4] In some of these auctions, the amount of money raised is almost beyond imagination. In the 1990s, the US government collected tens of billions of dollars from auctions for licenses for second generation mobile telecommunications,[5] and in 2000, the British and German governments, together, raised almost 100 billion euros in UMTS auctions.[6]

The Dutch government is also becoming accustomed to auctions as allocation mechanisms. A beauty contest was used, as recently as 1996, to assign a GSM[7] license to Libertel. That year, however, seems to have been the turning point. A proposal to change the Telecommunication Law to allow auctions reached the Dutch parliament in 1996; the new law was implemented

[1] Some have called Herodotus the Father of History, while others have called him the Father of Lies (Pipes, 1998-1999). There may be some doubt, therefore, about whether auctions for women really took place.

[2] These, and other examples of remarkable auctions, can be found in Cassady (1967) and Shubik (1983).

[3] Universal Mobile Telecommunication Services: a third generation mobile telecommunications standard.

[4] Empirical investigations on these auctions include Zulehner (2005) (cattle), Pezanis-Christou (2001) (fish), van den Berg et al. (2001) (flowers), Ashenfelter and Graddy (2002) (art), Lusht (1994) (real estate), Ashenfelter (1989) (wine), Binmore and Swierzbinski (2000) (treasury bills), van Damme (2002) (UMTS licenses), and Littlechild (2002) (electricity distribution contracts).

[5] Cramton (1998).

[6] See, e.g., van Damme (2002) and Binmore and Klemperer (2002).

[7] Global System for Mobile communications: a second generation mobile telecommunications standard.

in 1997.[8] In 1998, GSM licenses were sold through an auction,[9] and in 2000, the UMTS auction took place (although this auction was not as successful as the English and German UMTS auctions in terms of money raised).[10] Moreover, since 2002, the Dutch government has auctioned licenses for petrol stations on a yearly basis.

In this paper, we present an overview of the theoretical literature on auctions. Auction theory is an important theory for two very different reasons. First, as mentioned, many commodities are being sold at auctions. Therefore, it is important to understand how auctions work, and which auctions perform best, for instance, in terms of generating revenues or in terms of efficient allocation. Second, auction theory is a fundamental tool in economic theory. It provides a price formation model, whereas the widely used Arrow-Debreu model, from general equilibrium theory, is not explicit about how prices form.[11] In addition, the insights generated by auction theory can be useful when studying several other phenomena which have structures that resemble auctions like: lobbying contests, queues, wars of attrition, and monopolists' market behavior.[12] For instance, the theory of monopoly pricing is mathematically the same as the theory of revenue maximizing auctions.[13] Reflecting its importance, auction theory has become a substantial field in economic theory.

Historically, the field of auction theory roughly developed along the following lines. William Vickrey's 1961 paper is usually recognized as the seminal work in auction theory. Vickrey studies auctions of a single indivisible object. In the symmetric independent private values (SIPV) model, Vickrey derives equilibrium bidding for the first-price and the second-price auction.[14] He finds that the outcome of both auctions is efficient in the sense that it is always the bidder that attaches the highest value to the object who wins. Moreover, he comes to the surprising conclusion that the two auctions yield the same expected revenue for the seller.

[8] Verberne (2000).

[9] van Damme (1999).

[10] van Damme (2002).

[11] Arrow and Debreu (1954).

[12] Klemperer (2003).

[13] Bulow and Roberts (1989).

[14] In the SIPV model, risk neutral bidders with unlimited budgets bid competitively for an object whose value each bidder independently draws from the same distribution function. If a bidder does not win the object, she is indifferent about who wins, and how much the winner pays. More details about this model can be found in the next section.

Vickrey's paper largely contribute to his 1996 Nobel prize in economics, which he shares with Sir James Mirrlees.[15]

It takes until the end of the 1970s before Vickrey's work is further developed. Roger Myerson, John Riley, and William Samuelson derive results with respect to auctions that maximize the seller's expected revenue. They show that Vickrey's 'revenue equivalence result' extends far beyond the revenue equivalence of the first-price auction and the second-price auction. In addition, they discover that the seller can increase his revenue by inserting a reserve price. In fact, in the SIPV model, both the first-price and the second-price auction maximize the seller's expected revenue if the seller implements the right reserve price.

From the early eighties onwards, the attention shifts to the effects of relaxing the assumptions underlying the SIPV model. Particularly, under which circumstances does the revenue equivalence between the first-price auction and the second-price auction ceases to hold? Important contributions, in this respect, are the affiliated signals[16] model of Paul Milgrom and Robert Weber, and the risk aversion model of Eric Maskin and John Riley. Under affiliated signals, the second-price auction turns out to dominate the first-price auction in terms of expected revenue, while with risk aversion, the opposite result holds true.

In the mid-eighties, Jean-Jacques Laffont, Jean Tirole, Preston McAfee, and John McMillan further develop auction theory by focusing on the auctioning of incentive contracts. In contrast to Vickrey's framework, the principal does not wish to establish a high revenue or an efficient allocation of an object, but aims at inducing effort from the winner after the auction. An example is the procurement for the construction of a road. The procurer hopes that the winner of the procurement will build the road at the lowest possible cost. The question that arises is then: What is the optimal procurement mechanism? One of the main results is that it is not sufficient to simply sell the project to the lowest bidder and make her the residual claimant of the social welfare that she generates. This is because the winner would put too much effort in the project relative to the optimal mechanism.[17]

[15] We will see that the techniques developed by Mirrlees to construct optimal taxation schemes (and other incentive schemes), turned out to be useful for auction theory as well.

[16] Affiliation roughly means that the signals of the bidders are strongly correlated.

[17] The methods used to derive the results are essentially all those of Mirrlees who developed them within the framework of optimal taxation. We thank an anomymous referee for pointing this out to us.

The most recent burst of auction theory follows in the mid-nineties and the first years of the new millennium, as a response to the FCC[18] auctions in the US, and the UMTS auctions in Europe. The main focus shifts from single-object auctions to auctions in which the seller offers several objects simultaneously. Rather simple efficient auctions can be constructed if all objects are the same, or if each bidder only demands a single object. In the general case, where multi-object demand and heterogeneous objects are concerned, the Vickrey-Clarke-Groves mechanism is efficient. Unfortunately, however, the mechanism has many practical drawbacks. Larry Ausubel, Peter Cramton, and Paul Milgrom have recently proposed the 'clock-proxy auction' to deal with these. However, more research is needed to determine the circumstances under which this auction generates desirable outcomes.

The aim of this paper is to give an easily accessible overview of the most important insights of auction theory. The paper adds the following to earlier surveys like Klemperer (1999) and Krishna (2002).[19] First, when discussing the results for single-object auctions, we try to find a compromise between the mainly non-technical treatment of Klemperer and the advanced treatment of Krishna by giving easily accessible proofs to the most elementary propositions. Second, we elaborate more on what happens if the assumptions of the SIPV model are relaxed. Third, we cover auctions of incentive contracts, which have been almost entirely ignored in earlier surveys, although the problem of auctioning incentive contracts is interesting from both a theoretical and practical point of view. Fourth, our treatment of multi-object auctions captures the progress of auction theory since Klemperer's and Krishna's work was published. The fact that most of the cited articles are very recently dated shows that the previous surveys are a little outdated with respect to multi-object auctions. Finally, we seek applications of the theory, looking both at auctions in practice and applications to other fields in economics, including lobbying, and 'standards battles'.

The setup of this paper follows the historical development of auction theory as above. We study single-object auctions in Section 2. We start this section by studying equilibrium

[18]FCC stands for Federal Communications Commission, the agency that organized the auctions for licenses for second generation mobile telecommunications.

[19]An overview of field studies on auctions can be found in Laffont (1997). Kagel (1995) presents a survey of laboratory experiments on auctions, while the books of Klemperer (2004) and Milgrom (2004) discuss the use of auction theory in the design of real-life auctions.

bidding in the SIPV model for standard auctions such as the English auction, and auctions that are important for modelling other economic phenomena such as the all-pay auction. Then we discuss the revenue equivalence theorem, and construct auctions that maximize the seller's expected revenue. We conclude this section by relaxing the assumptions of the SIPV model and discussing what happens to the revenue ranking of standard auctions. In Section 3, we solve the problem of auctioning incentive contracts. Section 4 moves our attention to multi-object auctions. Finally, Section 5 concludes with a short summary of the paper and an agenda for future research in the field of auction theory. The proofs of all propositions and lemmas are relegated to the Appendix.

2 SINGLE-OBJECT AUCTIONS

In this section, we study auctions of a single object. In Subsection 2.1, we introduce the symmetric independent private values (SIPV) model. In Subsection 2.2, we analyze equilibrium bidding for several auction types. Subsection 2.3 contains a treatment of the revenue equivalence theorem and optimal auctions. In Subsection 2.4, we relax the assumptions of the SIPV model, and discuss the effects on the revenue ranking of standard auctions. Subsection 2.5 contains a summary of the main findings.

2.1 The SIPV model

The SIPV model was introduced by Vickrey (1961). He models an auction game as a non-cooperative game with incomplete information. The SIPV model applies to any auction in which a seller offers one indivisible object to $n \geq 2$ bidders, and is built around the following set of assumptions.[20]

(A1) *Risk neutrality*: All bidders are risk neutral.

[20] For a more detailed discussion on this model, see for instance McAfee and McMillan (1987a).

(A2) *Private values:* Bidder i, $i = 1, ..., n$, has value v_i for the object. This number is private information to bidder i, and not known to the other bidders and the seller.

(A3) *Value independence*: The values v_i are independently drawn.

(A4) *No collusion among bidders*: Bidders do not make agreements among themselves in order to achieve the object cheaply. More generally, bidders play according to a Bayesian Nash equilibrium, i.e., each bidder employs a bidding strategy that tells her what to bid contingent on her value, and given the conditional bids of the other bidders, she has no incentive to deviate from this strategy.

(A5) *Symmetry*: The values v_i are drawn from the same smooth distribution function F on the interval $[0, \bar{v}]$ with density function $f \equiv F'$.

(A6) *No budget constraints*: Each bidder is able to fulfill the financial requirements that are induced by her bid.

(A7) *No allocative externalities*: Losers do not receive positive or negative externalities when the object is transferred to the winner of the auction.

(A8) *No financial externalities*: The utility of losing bidders is not affected by how much the winner pays.

2.2 Equilibrium bidding in the SIPV model

In this subsection we analyze equilibrium bidding in commonly studied auctions under the assumptions of the SIPV model. We start with the four 'standard' auctions that are used to allocate a single object: the first-price sealed-bid auction, the Dutch auction, the Vickrey auction, and the English auction. In addition, we examine two other auctions that are rarely used as allocation mechanisms, but that are useful in modeling other economic phenomena: the all-pay auction and the war of attrition. We focus on three types of questions. First, how much do bidders bid in equilibrium? Second, is the equilibrium outcome efficient?[21] And third, which of these auctions yields the highest expected revenue?

[21] By efficiency we mean that the auction outcome is always such that the bidder who wins the object is the one who attaches the highest value to it.

2.2.1 First-Price Sealed-Bid Auction

In the first-price sealed-bid auction (sealed high-bid auction), bidders independently submit sealed bids. The object is sold to the highest bidder at her own bid.[22] In the US, mineral rights are sold using this auction. In the Appendix, we consider two methods for deriving symmetric equilibrium bidding strategies, the 'direct' and the 'indirect' method. These methods turn out to be useful for determining equilibrium bidding, not only for the first-price sealed-bid auction, but for other auctions as well. The seller's expected revenue R_{FPSB} is the expectation of the bid of the highest bidder, which is equal to $E\{Y_2^n\}$, where Y_2^n is the second-order statistic of n draws from F. In other words, the expected revenue from the first-price sealed-bid auction is the expectation of the second highest value.

Proposition 1 *The n-tuple of strategies* $(B_{FPSB}, ..., B_{FPSB})$*, where*

$$B_{FPSB}(v) = v - \frac{\int_0^v F(x)^{n-1}dx}{F(v)^{n-1}},$$

constitutes a Bayesian-Nash equilibrium of the first-price sealed-bid auction. The equilibrium outcome is efficient. In equilibrium, the expected revenue is equal to

$$R_{FPSB} = E\{Y_2^n\}.$$

Observe that all bidders bid less than their value for the object, i.e., they shade their bids with an amount equal to

$$\frac{\int_0^v F(x)^{n-1}dx}{F(v)^{n-1}}.$$

This amount decreases when the number of bidders increases. In other words, more competition decreases a bidder's profit given that she wins.[23]

[22] Sometimes, a reserve price is used, below which the object will not be sold. Throughout the paper, when we do not explicitly specify a reserve price, we assume it to be zero.

[23] This result does not hold generally, though. In models with common values, increased competition may lead to lower bids. See, e.g., Goeree and Offerman (2003b).

2.2.2 Dutch Auction

In the Dutch auction (descending-bid auction), the auctioneer begins with a very high price, and successively lowers it, until one bidder bids, i.e., announces that she is willing to accept the current price. This bidder wins the object at that price, unless the price is below the reserve price. Flowers are sold this way in the Netherlands. The Dutch auction is strategically equivalent to the first-price sealed-bid auction because an n-tuple of bids $(b_1, ..., b_n)$ in both auctions yields the same outcome, i.e., the same bidder wins and she has to pay the same price.[24,25] This implies that the Bayesian-Nash equilibria of these two auctions must coincide, and that both are equally efficient and yield the same expected revenue.

Proposition 2 *The n-tuple of strategies $(B_{Dutch}, ..., B_{Dutch})$, where*

$$B_{Dutch}(v) = v - \frac{\int_0^v F(x)^{n-1} dx}{F(v)^{n-1}},$$

constitutes a Bayesian-Nash equilibrium of the Dutch auction. The equilibrium outcome is efficient. In equilibrium, the expected revenue is equal to

$$R_{Dutch} = E\{Y_2^n\}.$$

2.2.3 Vickrey Auction

In the Vickrey auction (second-price sealed-bid auction), bidders independently submit sealed bids. The object is sold to the highest bidder (given that her bid exceeds the reserve price). However, in contrast to the first-price sealed-bid auction, the price the winner pays is not her own bid, but the second highest bid (or the reserve price if it is higher than the second highest bid).

[24] Strictly speaking, in the Dutch auction, only one bidder submits a bid, namely the winner. However, each bidder has a price in her mind at which she wishes to announce that she is willing to buy the object. We consider this price as her bid.

[25] The strategic equivalence between the Dutch auction and the first-price sealed-bid auction is generally valid, i.e., not restricted to the SIPV model.

The Vickrey auction has an equilibrium in weakly dominant strategies in which each bidder bids her value. To see this, imagine that bidder i wishes to bid $b < v_i$. Let $\bar{b}$ be the highest bid of the other bidders. Bidding b instead of v_i only results in a different outcome if $b < \bar{b} < v_i$. If $\bar{b} > v_i$, bidder i does not win in either case. If $\bar{b} < b$, bidder i wins and pays $\bar{b}$ in both cases. However, in the case that $b < \bar{b} < v_i$, bidder i receives zero utility by bidding b, while she obtains $v_i - \bar{b} > 0$ when bidding v_i. Bidding $b > v_i$ only results in a different outcome if $v_i < \bar{b} < b$. A bid of v_i results in zero utility, whereas bidding b yields her a utility of $v_i - \bar{b} < 0$. Therefore, bidder i is always (weakly) better off by submitting a bid equal to her value. As all bidders bid their value and the winner pays the second highest value, the revenue from the Vickrey auction can be straightforwardly expressed as

$$R_{Vickrey} = E\{Y_2^n\}.$$

Proposition 3 *The n-tuple of strategies* $(B_{Vickrey}, ..., B_{Vickrey})$, *where*

$$B_{Vickrey}(v) = v,$$

constitutes a Bayesian-Nash equilibrium of the Vickrey auction. The equilibrium is in weakly dominant strategies and the equilibrium outcome is efficient. In equilibrium, the expected revenue is equal to

$$R_{Vickrey} = E\{Y_2^n\}.$$

Despite its useful theoretical properties, the Vickrey auction is seldom used in practice.[26] There may be several reasons why this is the case. First, bidding in the auction is not as straightforward as the theory suggests. At least in laboratory experiments, a substantial number of subjects deviates from the weakly dominant strategy, in contrast to the English auction.[27] Second, the Vickrey auction may cause political inconveniences. For instance, in a spectrum auction in New Zealand, the winner, who submitted a bid of NZ$ 7 million, paid only NZ$ 5,000, the bid of the runner-up.[28] Third, a reason why the auction may not be as efficient as

[26] Rothkopf et al. (1990).
[27] Kagel et al. (1987), Kagel and Levin (1993), Harstad (2000), and Englmaier et al. (2004).
[28] McMillan (1994).

the theory predicts is that bidders are reluctant to reveal their true value for the object, as the seller may use this information in later interactions. In the English auction, as we will see next, the highest bidder does not have to reveal how much she values the object, as the auction stops after the runner-up has left the auction.

Still, the Vickrey auction is closely related to the so-called proxy auction, which *is* frequently used in reality. For instance, Internet auction sites such as eBay.com, Amazon.com and ricardo.nl, use this auction format, and in the Netherlands, special telephone numbers, such as 0900-flowers, are also allocated via this auction.[29] In a proxy auction, a bidder indicates until which amount of money the auctioneer (commonly a computer) is allowed to increase her bid (in case she is outbid by another bidder). The proxy auction is strategically equivalent to the Vickrey auction if bidders are only allowed to submit a single bid, and no information about the bids of the other bidders is revealed.

2.2.4 English Auction

In the English auction (also known as English open outcry, oral, open, or ascending-bid auction), the price starts at the reserve price, and 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. We study here a version of the English auction called the Japanese auction, in which the price is raised continuously, and bidders announce to quit the auction at a certain price (e.g., by pressing or releasing a button). The English auction is the most famous and most commonly used auction type. Art and wine are sold using this type of auction.

In the SIPV model, the English auction is equivalent to the Vickrey auction in the following sense. In both auctions, bidders have a weakly dominant strategy to bid their own valuation. In the English auction, no bidder has a reason to step out at a price that is below or above her value. Therefore, the equilibrium outcome in terms of revenue and efficiency is the same for both auctions. However, unlike the first-price sealed-bid auction and the Dutch auction, these two

[29] Staatscourant (2004). The consultancy report underlying this decision is Janssen and Maasland (2002).

auctions are not strategically equivalent. In the English auction bidders can respond to rivals leaving the auction, which is not possible in the Vickrey auction. Therefore, the equilibrium outcomes are the same as long as the bidders' valuations are not affected by observing rivals' bidding behavior.

Proposition 4 *The n-tuple of strategies* $(B_{English}, ..., B_{English})$*, where*

$$B_{English}(v) = v,$$

constitutes a Bayesian-Nash equilibrium of the English auction. The equilibrium is in weakly dominant strategies and the equilibrium outcome is efficient. In equilibrium, the expected revenue is equal to

$$R_{English} = E\{Y_2^n\}.$$

2.2.5 All-Pay Auction

Now we turn to the all-pay auction and the war of attrition, mechanisms that are rarely used to allocate objects, but turn out to be useful in modeling other economic phenomena. The all-pay auction has the same rules as the first-price sealed-bid auction, with the difference that all bidders must pay their bid, even those who do not win the object. Although the all-pay auction is rarely used as a selling mechanism, there are at least three reasons why economists are interested in it. First, all-pay auctions are used to model several interesting economic phenomena, such as political lobbying, political campaigns, research tournaments, and sport tournaments.[30] Efforts of the agents in these models are viewed as their bids. Second, this auction has useful theoretical properties, as it maximizes the expected revenue for the auctioneer if bidders are risk averse or budget constrained.[31] Third, all-pay auctions are far better able to raise money for a public good than winner-pay auctions (such as the four auctions described above).[32] The reason is that in winner-pay auctions, in contrast to all-pay auctions, bidders

[30] See, e.g., Che and Gale (1998a) and Moldovanu and Sela (2001).
[31] See Matthews (1983) and Laffont and Robert (1996), respectively.
[32] See Goeree et al. (2005).

forgo a positive externality if they top another's high bid. The optimal fund-raising mechanism is an all-pay auction augmented with an entry fee and reserve price.

Most of the early literature on the all-pay auction and its applications focuses on the complete information setting.[33] This is somewhat surprising, as it seems to be more natural to assume incomplete information, i.e., the 'bidders' (e.g. interest groups) do not know each other's value for the 'object' (e.g. obtaining a favorable decision by a policy maker). In addition, in some situations, there is not less than a continuum of equilibria for the all-pay auction with complete information.[34] In contrast, there is a unique symmetric equilibrium for the all-pay auction with incompletely informed bidders, at least in the SIPV model.[35] The following proposition gives the equilibrium properties of the all-pay auction in the SIPV model.

Proposition 5 *The n-tuple of strategies* $(B_{All-Pay}, ..., B_{All-Pay})$*, where*

$$B_{All-Pay}(v) = (n-1)\int_0^v xF(x)^{n-2}dF(x),$$

constitutes a Bayesian-Nash equilibrium of the all-pay auction. The equilibrium outcome is efficient. In equilibrium, the expected revenue is equal to

$$R_{All-Pay} = E\{Y_2^n\}.$$

Note that, in contrast to the earlier models of the all-pay auction with complete information, there is no 'full rent dissipation': the total payments are below the value of the object to the winner. This finding suggests that Posner (1975) overestimates the welfare losses of rent-seeking when he assumes that firms' rent-seeking costs to obtain a monopoly position are equal to the monopoly profits.

2.2.6 War of Attrition

The war of attrition game was defined by biologist Maynard Smith (1974) in the context of animal conflicts.[36] For economists, this game has turned out to be useful to model certain

[33] See, e.g., Tullock (1967, 1980), and Baye et al. (1993).
[34] Baye et al. (1996).
[35] Moldovanu and Sela (2001).
[36] Maynard Smith speaks about 'contests' and 'displays' instead of 'wars of attrition'.

(economic) interactions between humans. An example is a battle between firms to control new technologies, for instance in mobile telecom the battle between the CDMA (code division multiple access), the TDMA (time division multiple access), and the GSM techniques to become the single surviving standard worldwide.[37]

Although at first sight, the war of attrition is not an auction, its rules could be used in the auction of a single object. In such an auction, the price is raised successively until one bidder remains. This bidder wins the object at the final price. Bidders who do not win the object pay the price at which they leave the auction. Observe that there are two differences between the war of attrition and the all-pay auction. First, the all-pay auction is a sealed-bid auction, whereas the war of attrition is an ascending auction. Second, in the war of attrition, the highest bidder only pays an amount equal to the second highest bid, and in the all-pay auction, the highest bidder pays her own bid.

For $n > 2$, it is not straightforward to construct a symmetric Bayesian Nash equilibrium of the war of attrition. Bulow and Klemperer (1999) show that in any efficient equilibrium all but the bidders with the highest two values should step out immediately. The remaining two bidders then submit bids according to a strictly increasing bid function. Strictly speaking, this cannot be an equilibrium, as there is no information available about whom of the bidders should step out immediately. Therefore, we restrict ourselves to the two-player case in the following proposition.

Proposition 6 *Let $n = 2$. The strategies (B_{WoA}, B_{WoA}), where*

$$B_{WoA}(v) = \int_0^v \frac{xf(x)}{1 - F(x)} dx,$$

constitutes a Bayesian-Nash equilibrium of the war-of-attrition. The equilibrium outcome is efficient. In equilibrium, the expected revenue is equal to

$$R_{WoA} = E\{Y_2^n\}.$$

[37] Bulow and Klemperer (1999).

2.3 Revenue equivalence and optimal auctions

Observe that in the SIPV model, all of the above auctions yield the same expected revenue. Is this result more general? Are there auctions that generate more revenue? And which auction yields the highest expected revenue? In his remarkable paper, published in 1981, Myerson answers these questions in a model that includes the SIPV model as a special case. In order to find the answers, Myerson derives two fundamental results, the *revelation principle*, and the *revenue-equivalence theorem.* In this subsection, we will discuss Myerson's results in the context of the SIPV model.[38] For simplicity, we assume that the seller does not attach any value to the object.[39]

A special class of auctions is the class of *direct revelation games.* In a direct revelation game, each bidder is asked to announce her value, and depending on the announcements, the object is allocated to one of the bidders, and one bidder, or several bidders, pay a certain amount to the seller. More specifically, let (p, x) denote a direct revelation game, where $p_i(\mathbf{v})$ is the probability that bidder i wins, and $x_i(\mathbf{v})$ is the expected payments by i to the seller when $\mathbf{v} \equiv (v_1, ..., v_n)$ is announced. There are two types of constraints that must be imposed on (p, x), an *individual rationality constraint* and an *incentive compatibility constraint.* The individual rationality constraint follows from the assumption that each bidder expects nonnegative utility. The incentive compatibility constraint is imposed as we demand that each bidder has an incentive to announce her value truthfully.

Lemma 1 (Revelation Principle) *For any auction there is an incentive compatible and individually rational direct revelation game that gives the seller the same expected equilibrium revenue as the auction.*

Lemma 1 implies that when solving the seller's problem, there is no loss of generality in only considering direct revelation games that are individually rational and incentive compatible.

[38] Independently, Riley and Samuelson (1981) derived similar results.

[39] Myerson (1981) assumes that the seller attaches some value to the object, which is commonly known among all bidders.

Now, consider the following definition of bidder i's marginal revenue:

$$MR(v_i) \equiv v_i - \frac{1 - F(v_i)}{f(v_i)}, \ \forall v_i, i. \tag{1}$$

We call the seller's problem *regular* if MR is an increasing function.

Lemma 2 *Let (p, x) be a feasible direct revelation mechanism. The seller's expected revenue from (p, x) is given by*

$$U_0(p, x) = E_{\mathbf{v}} \left\{ \sum_{i=1}^{n} MR(v_i) p_i(\mathbf{v}) \right\} - \sum_{i=1}^{n} U_i(p, x, \underline{v}_i), \tag{2}$$

where $U_i(p, x, \underline{v}_i)$ is the expected utility of the bidder with the lowest possible value.

Several remarkable results follow from Lemma 2. We start with the revenue equivalence theorem.

Proposition 7 (Revenue Equivalence Theorem) *The seller's expected revenue from an auction is completely determined by the allocation rule p related to its equivalent direct revelation game (p, x), and the expected utility of the bidder with the lowest possible value.*

From this proposition, it immediately follows that in the SIPV model, all standard auctions yield the same expected utility for the seller and the bidders, provided that all bidders play the efficient Bayesian Nash equilibrium. Efficiency implies that the allocation rule is such that it is always the bidder with the highest value who wins the object, so that the allocation rule is the same for all standard auctions. In addition, in the efficient equilibrium of all standard auctions, the expected utility of the bidder with the lowest possible value is zero.

Now, we use Lemma 2 to construct the revenue maximizing auction. Observe that in (2), apart from a constant, the seller's expected revenue is equal to the sum of each bidder's marginal revenue multiplied by her winning probability. If the sellers' problem is regular, then marginal revenues are increasing in v_i, so that the following result follows.[40]

[40] See Myerson (1981) for further discussion on the consequences of relaxing this restriction.

Proposition 8 *Suppose that the seller's problem is regular, and that there is an auction that in equilibrium, (1) assigns the object to the bidder with the highest marginal revenue, provided that the marginal revenue is nonnegative, (2) leaves the object in the hands of the seller if the highest marginal revenue is negative, and (3) gives the lowest types zero expected utility. Then this auction is optimal.*

This proposition has an interesting interpretation for the standard auctions:

Proposition 9 *When the seller's problem is regular, all standard auctions are optimal when the seller imposes a reserve price r with $MR(r) = 0$.*[41]

We only sketch the proof of this proposition. In the equilibrium of a standard auction with reserve price, bidders with a value below the reserve price abstain from bidding, and bidders with a value above the reserve price bid according to the same strictly increasing bid function. If the reserve price is chosen such that the marginal revenue at the reserve price is equal to zero, then all standard auction are optimal as (1) if the object is sold, it is always assigned to the bidder with the highest value and hence the highest nonnegative marginal revenue, (2) the object remains in the hands of the seller in the case that the highest marginal revenue is negative, and (3) the expected utility of the bidder with the lowest value is zero. Note that the reserve price does not depend on the number of bidders. In fact, it is the same as the optimal take-it-or-leave price when the seller faces a single potential buyer. The following example illustrates these findings in a simple setting with the uniform distribution.

Example 1 *Suppose all bidders draw their value from the uniform distribution on the interval $[0, 1]$. Then*

$$MR(v_i) = 2v_i - 1.$$

As MR is strictly increasing in v_i, the seller's problem is regular. Moreover, $MR(r) = 0$ implies $r = \frac{1}{2}$. So, from Proposition 9 it follows that all standard auctions are optimal when the seller choose reserve price $r = \frac{1}{2}$. Observe that the seller keeps the object with probability $\left(\frac{1}{2}\right)^n$.

[41] Myerson (1981) does not mention this result explicitly, but it follows from his study. Riley and Samuelson (1981) formally derive the result in an independent private values model.

2.4 Relaxing the SIPV model assumptions

In the previous subsection, we have observed, that in the SIPV model, all efficient auctions yield the same revenue to the seller as long as the bidder with the lowest possible value obtains zero expected utility. In this subsection, we relax the Assumptions (A1)-(A8) underlying the SIPV model and study the effect on the revenue ranking of the most commonly studied auctions, first-price auctions (like the first-price sealed-bid auction and the Dutch auction), and second-price auctions (like the Vickrey auction and the English auction).[42] We do so, by relaxing the assumptions one by one, while keeping the others satisfied, so that we obtain a clear view of the effect of each single assumption. Table 1 gives an overview.

Assumption	**Alternative model**
(A1) Risk neutrality	Risk aversion
(A2) Private values	Almost common values
(A3) Value independence	Affiliation
(A4) No collusion	Collusion
(A5) Symmetric bidders	Asymmetry
(A6) No budget constraints	Budget constraints
(A7) No allocative externalities	Allocative externalities
(A8) No financial externalities	Financial externalities

Table 1: Models that relax assumptions (A1)-(A8).

2.4.1 Risk aversion

The first assumption in the SIPV model is risk neutrality. Under risk aversion, the expected revenue in the first-price auction is higher than in the second-price auction. A model of risk aversion is the following. The winning bidder receives utility $u(v-p)$ if her value of the object is v and she pays p, where u is a concave increasing function with $u(0) = 0$. Note that risk aversion may play a role as a bidder has uncertainty about the values, and hence the bids, of the other bidders. In a second-price auction, bidding one's value remains a dominant strategy.

[42] From this section on, when writing 'standard auctions', we only refer to the first four auction types dealt with in the subsection 2.2: the first-price sealed-bid auction, the Dutch auction, the Vickrey auction, and the English auction.

In a first-price auction, a risk-averse bidder will bid higher than a risk-neutral bidder as she prefers a smaller gain with a higher probability. By bidding higher she insures herself against ending up with zero. This result still holds true if the value of the object is ex ante unknown to the bidders.[43]

2.4.2 Almost common values

Assumption (A2) states that each bidder knows her own value for the object, which may be different than the values of the other bidders. In almost common value auctions, the actual value of the object being auctioned is almost the same to all bidders - but the actual value is not known to anyone. For instance, in the case of two bidders, bidder 1 attaches value $v_1 = t_1 + t_2$ to the object, and bidder 2 has value $v_2 = t_1 + t_2 + \epsilon$, where t_i is bidder i's signal, and ϵ is strictly positive but small. In the equilibrium of the second-price auction, bidder 1 bids 0 and bidder 2 a strictly positive amount, so that the revenue to the seller will be zero. The intuition is as follows. Suppose that bidder 1 intends to continue bidding until B. If the high-valuation bidder goes beyond B, the low-valuation bidder's profit is zero. If the high-valuation bidder stops bidding before B, she obviously is of the opinion that the object is worth less than B to her. But in that case, it is certain that it is worth less than B to the low-valuation bidder. For each positive B for the low-valuation bidder, there is an expected loss. Therefore, bidder 1 bids zero in equilibrium. In the first-price auction, in contrast, the auction proceeds are strictly positive. Bidder 1 bids more than zero as she knows that she has a chance of winning as bidder 2 does not know exactly how much she should shade her bid in order to still win the auction.[44]

2.4.3 Affiliation

The third assumption is value independence, i.e., the values are independently drawn. If these values are 'affiliated', the second-price auction yields more expected revenue than the first-price auction. Affiliation roughly means that there is a strong positive correlation between the signals

[43] Maskin and Riley (1984).
[44] Klemperer (1998).

of the bidders. In other words, if one bidder receives a high signal about the value of the good, she expects the other to receive a high signal as well. Let us consider a situation with pure common values, i.e., all bidders have the same value for the object, for instance the right to drill oil in a certain area. As the actual value of the oil field is not known to the bidders before the auction, they run the risk of bidding too high, and fall prey to what is called the winner's curse: for a bidder, winning is bad news as she is the one who has the most optimistic estimate for the true value of the object. Taking the winner's curse into account, a bidder is inclined to shade her bid substantially. However, if bidders can base their final bid on other bidders' information, then they feel more confident about bidding - and will hence bid less conservatively. An auction generates more revenue if the payment of the winning bidder has greater linkage to the value estimates of other bidders. In a first-price sealed-bid auction, there is no such linkage (the winner pays her own bid). A second-price auction has more linkage since the winner pays the second highest bid - which is linked to the value estimate of the second highest bidder.[45]

2.4.4 Collusion

According to Assumption (A4), bidders do not collude, i.e., they play according to a Bayesian Nash equilibrium. Collusive agreements are easier to sustain in a second-price auction than in a first-price auction, so that the expected revenue is higher in the latter. The optimal agreement in a second-price auction is for the bidder with the highest value to bid her true value and for all other bidders to abstain from bidding. This agreement is stable as the bidders with the lower values cannot improve their situation by bidding differently. In a first-price auction, the optimal agreement for the highest value bidder is to bid a very small amount and for all other bidders to abstain from bidding. This agreement is not stable as the bidders with the lower values have an incentive to cheat on the agreement by bidding just a little bit higher than the bid of the highest value bidder.[46]

[45] Milgrom and Weber (1982).
[46] Robinson (1985).

2.4.5 Asymmetry

The fifth assumption is symmetry, which means that the values are drawn from the same distribution function. Asymmetry in bidders' value distributions has an ambiguous effect on the revenue ranking of the first-price and second-price auctions. In some situations, the expected revenue from a first-price auction is higher. Imagine, for instance, that the strong bidder's distribution is such that, with high probability, her valuation is a great deal higher than that of a weak bidder. In a first-price auction the strong bidder has an incentive to outbid the weak bidder (to enter a bid slightly higher than the maximum valuation in the weak bidder's support) in order to be sure that she will win. In a second-price auction the expected payment will only be the expected value of the weak bidder's valuation, as for both bidders it is a weakly dominant strategy to bid their own value. In other situations, however, the expected revenue from the first-price auction may be lower. Suppose, for instance, that across bidders, distributions have different shapes but approximately the same support. A strong bidder, with most mass in the upper range of the distribution, has not much reason to bid high in the first-price auction as she has a substantial probability to beat the weak bidder by submitting a low bid. This incentive to 'low ball' is absent in a second-price auction, so that the expected revenue from the latter may be higher.[47]

2.4.6 Budget constraints

Under Assumption (A6), bidders face no budget constraints. If this assumption is violated, the first-price auction yields more revenue than the second-price auction. This is trivially true when all bidders face a budget constraint $\bar{b}$ such that $B_{FPSB}(\bar{v}) < \bar{b} < \bar{v}$. Clearly, the expected revenue of the first-price auction is not affected, as no bidder wishes to submit a bid above $\bar{b}$ in equilibrium. In contrast, in the second-price auction, bidders with a value in the range $[\bar{b}, \bar{v}]$ cannot bid higher than $\bar{b}$, so that the expected revenue from the second-price auction decreases relative to the situation that there are no budget constraints. This finding turns out to hold

[47] Maskin and Riley (2000).

more generally.[48]

2.4.7 Allocative externalities

According to Assumption (A7), losers face no allocative externalities when the object is transferred to the winner. If allocative externalities are present, the second-price auction and the first-price auction are only revenue equivalent under specific circumstances. Allocative externalities arise when losing bidders receive positive or negative utility when the auctioned object is allocated to the winner. As an example, think about a monopolist suffering a negative externality when a competitor wins a licence to operate in 'his' market.[49] Jehiel et al. (1999) show that the Vickrey auction (weakly) dominates other sealed-bid formats, such as the first-price sealed-bid auction. Das Varma (2002) derives circumstances under which first-price auctions and second-price auctions are revenue equivalent, namely when externalities are 'reciprocal', i.e., for each pair of bidders, the externality imposed on each other is the same. However, when externalities are nonreciprocal, the revenue ranking becomes ambiguous.

The following example shows why this is the case. Imagine that two bidders bid for a single object in an auction. We assume that all conditions of the SIPV model hold, except that the utility of bidder i when bidder j wins at a price of p is given by

$$U_i(j,p) = \begin{cases} v_i - p & \text{if } i = j \\ -a_i & \text{if } i \neq j \end{cases},$$

$i, j \in \{1, 2\}$, where a_i is the negative externality imposed on bidder i when the other bidder wins. Assume also that a_i is private information to bidder i. Note that in equilibrium, each bidder submits a bid as if her value for the object were $v_i + a_i$. Now, if a_i is drawn from different distribution functions, this model is isomorphic to a model with asymmetry in bidders' value distributions. Recall from subsection 2.4.5 that in such a model, the revenue ranking between the two auctions is ambiguous.

[48] Che and Gale (1998b).
[49] Gilbert and Newbery (1982).

2.4.8 Financial externalities

Finally, we relax the assumption that the bidders face no financial externalities. The seller generates more revenue in the second-price auction than in the first-price auction in situations with financial externalities. A losing bidder enjoys financial externalities when she obtains a positive externality from the fact that the winning bidder has to pay some money from winning the object. In soccer, the Spanish team FC Barcelona may obtain positive utility when the Italian club AC Milan spends a lot of money when buying a new striker. Assuming that AC Milan faces a budget constraint, AC Milan becomes a weaker competitor to FC Barcelona in future battles for other soccer players.

Formally, financial externalities can be described as follows. The utility of bidder i when bidder j wins at a price of p is given by

$$u_i(j,p) = \begin{cases} v_i - p & \text{if } j = i \\ \varphi p & \text{if } j \neq i, \end{cases},$$

where $\varphi > 0$ is the parameter indicating the financial externality. In this model, given that the other assumptions of the SIPV model hold, the expected revenue from a second-price auction is higher than from a first-price auction for reasonable values of φ. The intuition is that, in contrast to the first-price auction, a bidder in a second-price auction can directly influence the level of payments made by the winner by increasing her bid.[50]

2.5 Summary

In the SIPV model, a remarkable result arises with respect to the seller's expected revenue: it is the same for the four standard auctions! Vickrey (1961) was the first to show this result for the simple case of a uniform value distribution function on the interval $[0,1]$. Also the all-pay auction and the two-player war of attrition turn out to yield the same revenue to the seller.[51] Observe that the seller does not always realize all gains from trade, although he has some market power as he can determine the rules of the auction. In expectation, he obtains the expected

[50] Maasland and Onderstal (2005) and Goeree et al. (2005).

[51] In the next section, we will give an alternative proof of this 'revenue equivalence result', and argue that it holds more generally.

value of the second highest value, whereas under complete information, his revenue could be equal to the highest value. The seller can exploit his market power a bit more by inserting a reserve price in any standard auction, which is indeed a way to implement an optimal auction.

Table 2 summarizes how the ranking of the standard auctions changes when one of the Assumptions (A1)-(A8) is relaxed while the other assumptions remain valid. In this table, we compare first-price auctions (F), like the first-price sealed-bid auction and the Dutch auction, and second-price auctions (S), like the Vickrey auction and the English auction. S $\prec$ F [S $\succ$ F] means that a second-price auction yields strictly lower [strictly higher] expected revenue than a first-price auction. S ? F implies that the revenue ranking is ambiguous, that is, in some circumstances S $\prec$ F holds, and in other S $\succ$ F.

Assumption	**Alternative model**	**Ranking**
(A1) Risk neutrality	Risk aversion	S $\prec$ F
(A2) Private values	Almost common values	S $\prec$ F
(A3) Value independence	Affiliation	S $\succ$ F
(A4) No collusion	Collusion	S $\prec$ F
(A5) Symmetric bidders	Asymmetry	S ? F
(A6) No budget constraints	Budget constraints	S $\prec$ F
(A7) No allocative externalities	Allocative externalities	S ? F
(A8) No financial externalities	Financial externalities	S $\succ$ F

Table 2: Revenue ranking of standard auctions when the Assumptions (A1)-(A8) are relaxed.

3 AUCTIONING INCENTIVE CONTRACTS

In this section, we turn to the problem of auctioning incentive contracts. In a large range of countries, governments use procurements to select a firm to establish a certain project, e.g., constructing a road. These procurements give flesh and blood to Demsetz' (1968) idea of competition 'for' the market. McAfee and McMillan (1986, 1987b) and Laffont and Tirole (1987, 1993) study these types of situations, thus building a bridge between auction theory and incentive theory. Auction theory applies as the buyer wishes to select a bidder out of a set of several bidders, and incentive theory is relevant as the buyer may wish to stimulate the winning bidder to put effort in the project. The question that arises is then: what is the optimal procurement mechanism? Is it optimal, in the example of road construction, to simply select the cheapest bidder and make her the residual claimant of all cost savings, or are there more advanced mechanisms that increase the buyer's utility? McAfee and McMillan and Laffont and Tirole have answered these questions using the techniques that were first developed by Mirrlees (1971, 1976, 1999 (first draft 1975)).

The problem of auctioning incentive contracts is not only of theoretical interest. For instance, in several countries, the government procures welfare-to-work programs as a part of their active labor market policy.[52] In these procurements, the government allocates welfare-to-work projects to employment service providers. A welfare-to-work project typically consists of a number of unemployed people, and the winning provider is rewarded on the basis of the number of these people that find a job within a specified period of time. According to OECD (2001) procurements for welfare-to-work projects should be organized as follows. The government defines an incentive contract that guarantees an employment service provider a fixed reward for each person that finds a job. This reward is equal to the increase in social welfare if this person does find a job. The government sells the contract to the highest bidder in an auction, who has to pay her bid. Onderstal (2005) shows that the mechanism proposed by OECD indeed performs almost as well as the optimal mechanism.

In the next subsection, we present a simple model. In Subsection 3.2, we construct the optimal mechanism, and in Subsection 3.3, we summarize the main findings.

3.1 The model

Let us describe a simple setting, in which a risk neutral buyer wishes to procure a project. We assume that n risk neutral firms participate in the procurement. Each firm i, $i = 1, ..., n$, when winning the project, is able to exert observable effort e_i at the cost

$$C_i(e_i, \alpha_i) = \frac{1}{2}e_i^2 + e_i - \alpha_i e_i.$$

[52] SEO and TNO (2004) provide a comparison of welfare-to-work procurements in Australia, Denmark, the Netherlands, Sweden, the UK, and the US.

In the road construction example, the effort level e_i may be interpreted as a decrease in the cost to build the road, while in procurements of welfare-to-work programs, effort is related to the number of people that find a job. We choose this specific cost function so that by construction, in the first-best optimum, i.e., the optimum under complete information, the winning firm's effort is equal to α_i. In addition, note that $C_i''(e_i) = 1 > 0$. In other words, the marginal costs of effort is strictly increasing in effort. In road construction, this seems to make sense: the first euro in cost savings is easier to obtain than the second euro, and so forth. We assume that diseconomies of scale do not play a role, as otherwise the government would have a good reason to split up the project in smaller projects, and have several firms do the job.

The firms differ with respect to their efficiency level $\alpha_i \in [0,1]$, which is only observable to firm i. Note that the costs per unit of effort are increasing in effort. The firms draw the α_i's independently from the same distribution with a cumulative distribution function F on the interval $[0,1]$ and a differentiable density function f. F is common knowledge. We assume that

$$\alpha_i - \frac{1 - F(\alpha_i)}{f(\alpha_i)} \text{ is strictly increasing in } \alpha_i, \tag{3}$$

which is the same as the regularity condition we imposed in the problem of revenue maximizing auctions.

Firm i has the utility function

$$U_i = t_i - C_i,$$

where t_i is the monetary transfer that it receives from the buyer. Let S denote the buyer's utility from the project. We assume that

$$\begin{aligned} S &= e_i - t_i \\ &= e_i - U_i - C_i(e_i, \alpha_i), \end{aligned} \tag{4}$$

where i is the firm the buyer has selected for the project. In the road construction example, S can be viewed as the net cost savings for the government.[53] An optimal mechanism maximizes S under the restriction that the firms play a Bayesian Nash equilibrium, and that the mechanism

[53] When the government does not select one of the bidders in the procurement, a public firm builds the road. We assume that in that case $S = 0$.

satisfies a participation constraint (in equilibrium, each participating firm should at least receive zero expected utility).

The first-best optimum has the following properties. First, the buyer selects the most efficient firm, i.e., the firm with the highest type α_i, as this firm has the lowest C_i for a given effort level. Second, the buyer induces this firm to exert effort α_i. Finally, the buyer exactly covers the costs C_i. We will see that this first-best optimum cannot be reached in our setting with incomplete information: the buyer has to pay informational rents to the firm.

3.2 The optimal mechanism

What is the optimal mechanism, i.e., the mechanism that maximizes (4)? As in the problem of finding a revenue maximizing auction, we apply the revelation principle: without loss of generality we restrict our attention to incentive compatible and individually rational direct revelation games. Let

$$\tilde{\boldsymbol{\alpha}} = (\tilde{\alpha}_1, ..., \tilde{\alpha}_n)$$

be the vector of announcements by firm $1, ..., n$ respectively. We consider mechanisms $(x_i, e_i, t_i)_{i=1,...,n}$ that induce a truthtelling Bayesian Nash equilibrium, where, given $\tilde{\boldsymbol{\alpha}}$, $x_i(\tilde{\boldsymbol{\alpha}})$ is the probability that firm i wins the contract, and, given that firm i wins the contract, $e_i(\tilde{\boldsymbol{\alpha}})$ is its effort and $t_i(\tilde{\boldsymbol{\alpha}})$ is the monetary transfer it receives from the buyer.

Proposition 10 *The optimal mechanism $(x_i^*, e_i^*, t_i^*)_{i=1,...,n}$ has the following properties:*

$$\begin{aligned} x_i^*(\boldsymbol{\alpha}) &= \begin{cases} 1 & \text{if } \alpha_i > \max_{j \neq i} \alpha_j \text{ and } \alpha_i \geq \underline{\alpha} \\ 0 & \text{otherwise} \end{cases}, \\ e_i^*(\boldsymbol{\alpha}) &= \alpha_i - \frac{1 - F(\alpha_i)}{f(\alpha_i)}, \text{ and} \\ t_i^*(\boldsymbol{\alpha}) &= C_i(e_i^*(\boldsymbol{\alpha}), \alpha_i) + \int_{\underline{\alpha}}^{\alpha_i} e_i^*(y) F(y)^n dy, \end{aligned}$$

where $\underline{\alpha}$ is the unique solution to y in $y = \frac{1-F(y)}{f(y)}$.

The optimal mechanism has the property that the buyer optimally selects the most efficient firm, provided that its efficiency level exceeds $\underline{\alpha} > 0$. This firm exerts effort according to e_i^*, and t_i^* determines the payments it receives from the buyer. Observe that the desired effort level $e_i^*(\boldsymbol{\alpha})$ and $\underline{\alpha}$ do not depend on the number of bidding firms.

Finally, let us go back to the example of the road construction project. Is it optimal to simply auction the project to the lowest bidder and gives her a compensation equal to the cost savings e that she realizes? The answer turns out to be 'no'. This can be seen as follows. The winner i of the auction maximizes her utility, which is equal to

$$t_i + e_i - C(e_i, \alpha_i) = t_i - \frac{1}{2}e_i^2 + \alpha_i e_i \tag{5}$$

where $t_i + e_i$ is the transfer the government makes to the winner. It is routine to derive that $\hat{e}_i = \alpha_i$ maximizes (5). In other words, the winner puts too much effort in the project relative to the optimal mechanism, as $\hat{e}_i > e_i^*(\boldsymbol{\alpha})$. It can be checked that the optimal mechanism can be implemented by selling a non-linear contract to the lowest bidder. For instance, if the efficiency levels are drawn from the uniform distribution on the interval $[0, 1]$, the government optimally pays the winner

$$b + \frac{1}{4}e^2 + \frac{1}{2}e$$

if her winning bid is b and she puts effort e in the project.

3.3 Summary

In this section, we have studied auctions of incentive contracts in a stylized model. Table 3 summarizes the main results of this model. Note that three types of inefficiency arise from the optimal mechanism under incomplete information relative to a situation with complete information. First, since $e_i^*(\boldsymbol{\alpha}) < \alpha_i$ for all $\alpha_i < 1$, the firm's effort is lower than in the full-information optimum. Second, the buyer will not contract with any firm whose efficiency level is below $\underline{\alpha}$, whereas in the full-information world, the buyer would contract with any provider. The latter is analogous to a reserve price in an optimal auction. Third, as $t_i^*(\boldsymbol{\alpha}) > C_i(e_i^*(\boldsymbol{\alpha}))$ for all $\alpha_i > 0$, the government covers more than the costs that are actually born by the winning provider. These types of inefficiency give the buyer the opportunity to reduce the informational rents that he has to pay to the winner because of incomplete information.

	First-best mechanism	Optimal mechanism
Winner	$x_i^{**}(\boldsymbol{\alpha}) = \begin{cases} 1 & \text{if } \alpha_i > \max_{j \neq i} \alpha_j \\ 0 & \text{otherwise} \end{cases}$	$x_i^{*}(\boldsymbol{\alpha}) = \begin{cases} 1 & \text{if } \alpha_i > \max\{\max_{j \neq i} \alpha_j, \underline{\alpha}\} \\ 0 & \text{otherwise} \end{cases}$
Effort	$e_i^{**}(\boldsymbol{\alpha}) = \alpha_i$	$e_i^{*}(\boldsymbol{\alpha}) = \alpha_i - \frac{1-F(\alpha_i)}{f(\alpha_i)}$
Payment	$t_i^{**}(\boldsymbol{\alpha}) = C_i(e_i^{**}(\boldsymbol{\alpha}), \alpha_i)$	$t_i^{*}(\boldsymbol{\alpha}) = C_i(e_i^{*}(\boldsymbol{\alpha}), \alpha_i) + \int_{\underline{\alpha}}^{\alpha_i} e_i^{*}(y) F(y)^n dy$

Table 3: Properties of the first-best mechanism and the optimal mechanism under incomplete information.

4 MULTI-OBJECT AUCTIONS

In the previous sections, we have observed that the seller faces a trade-off between efficiency and revenue. When selling a single object, the seller maximizes his revenue by imposing a reserve price. This causes inefficiency as the object remains unsold when none of the bidders turns out to be willing to pay the reserve price, while they may assign a positive value to it. Equivalently, in auctions of incentive contracts, the revenue maximizing buyer only assigns the incentive contract if a sufficiently efficient firm enters the auction.

In multi-object auctions, a new trade-off enters the picture: the trade-off between efficiency and complexity. We will see that if each bidder in the auction only demands one object, and if the seller offers homogeneous objects, the main results from the single-object case carry over: straightforward generalizations of the standard auctions are efficient (and revenue equivalent). However, as soon as objects are heterogeneous, or when bidders demand more than one object, an efficient outcome is no longer guaranteed. Luckily, rather simple efficient auctions can be constructed with multi-object demand if objects are homogeneous and with heterogeneous objects if there is single-object demand. In the general case, with multi-object demand and heterogeneous objects, the Vickrey-Clarke-Groves mechanism is efficient. However, this auction has several practical drawbacks, for instance that it is complex as bids are needed on a large range of packages. These disadvantages are only partially mitigated in innovative new designs that have been recently proposed in the literature, such as Ausubel, Cramton, and Milgrom's clock-proxy auction.

From the above it is clear that the results are highly dependent on whether the objects are identical or not and whether the bidders are allowed to win several objects or only one. We have therefore decided to build up this section along these two crucial points. In the 2x2 matrix in Table 4 it is shown which part is covered in which subsection.

	Identical Objects	Non-Identical Objects
Single-Object Demand	Subsection 4.1	Subsection 4.2
Multi-Object Demand	Subsection 4.3	Subsection 4.4

Table 4: Set-up of this section.

In Subsection 4.1, we deal with auctions of multiple identical objects when bidders are allowed to win only one object/desire at most one object. A real-life example of such an auction is the Danish UMTS auction (by which licenses for third generation mobile telecommunication were sold). Four identical licenses were put up for sale and firms were only allowed to win one license. Subsection 4.2 introduces auctions of multiple non-identical objects when bidders are allowed to win only one object/desire at most one object. A good example of such an auction is the Dutch UMTS-auction (and most of the other European UMTS-auctions). In the Netherlands, five non-identical licenses (differing with respect to the amount of spectrum) were put up for sale and firms were only allowed to win one license. In Subsection 4.3, we analyze auctions of multiple identical objects when bidders are allowed to win multiple objects. Examples are treasury bond auctions, electricity auctions, and initial public offerings (IPOs) of companies shares (e.g. Google's IPO). In Subsection 4.4, we discuss auctions of multiple non-identical objects when bidders are allowed to win multiple objects. The Dutch GSM auction is an example of such an auction. Subsection 4.5 contains a conclusion with the main findings of this section.

4.1 Auctions of multiple identical objects with single-object demand

In this subsection, we present the multi-unit extensions of the four standard auctions dealt with in Section 2 when bidders are allowed to win only one unit/each bidder wants at most one unit.[54] These four extensions are the pay-your-bid auction, the multi-unit Dutch auction, the uniform-price auction, and the multi-unit English auction. We assume that $2 \leq k < n$ units are put up for sale.

4.1.1 Pay-Your-Bid Auction

In the pay-your-bid auction, bidders independently submit sealed bids (each bidder submits one bid). The k units are sold to the k highest bidders at their own bid. This auction is sometimes called a discriminatory auction as it involves price discrimination (bidders pay different prices for an identical object), and can be seen as the generalization of the first-price sealed-bid auction. In the SIPV model, the equilibrium outcome is efficient.

4.1.2 Multi-Unit Dutch Auction

In the multi-unit Dutch auction, the auctioneer begins with a very high price, and successively lowers it, until one bidder bids. This bidder wins the first unit at that price. The price then goes further down until a second bidder bids. This bidder wins the second unit for the price she bid. The auction goes on until all k units are sold (or until the auction has reached a zero price). Note that also this auction involves price discrimination. In contrast to the single-unit case in Section 2, the multi-unit Dutch auction is not strategically equivalent to the pay-your-bid auction, as bidders may update their bid when bidders leave the auction (after winning one of the k units). In the SIPV model, the Bayesian-Nash equilibria of the multi-unit Dutch auction and the pay-your-bid auction still coincide though, so that also the multi-unit Dutch auction is efficient.

[54] Early articles on multiple identical object auctions with single-object demand are Vickrey (1962) and Weber (1983).

4.1.3 Uniform-Price Auction

In the uniform-price auction with single unit demand, bidders independently submit sealed-bids (each bidder submits one bid). The k units are sold to the k highest bidders (given that these bids exceed the reserve price). The winners pay the $(k+1)$-th highest bid, i.e. the highest rejected bid. The uniform-price auction has an efficient equilibrium, as each bidder has a weakly dominant strategy to bid her value. The intuition is the same as in the Vickrey auction.

4.1.4 Multi-Unit English Auction

In the multi-unit English auction, the price starts at the reserve price, and is successively raised until k bidders remain. These bidders each win one unit at the final price. As in the SIPV model, the multi-unit English auction is equivalent to the uniform-price auction, the equilibrium outcome in terms of revenue and efficiency is the same for both auctions.

4.1.5 Results

When objects are identical and bidders desire at most one unit, several results from the single-unit case generalizes to the multi-unit case. In the SIPV model, all standard auctions remain efficient. Moreover, the revenue equivalence theorem continues to hold, which implies that the above four auction types yield the same revenue in expectation.[55] Another result is that all four auctions are revenue maximizing, provided that the seller imposes the optimal reserve price. What changes with respect to the single-unit case is that the multi-unit Dutch auction (open first-price format) is not strategically equivalent to the pay-your-bid auction (sealed-bid first-price format) anymore.

[55] See e.g., Harris and Raviv (1981), Maskin and Riley (1989).

4.2 Auctions of multiple non-identical objects with single-object demand

In this subsection, we keep the assumption that each bidder only desires one object, but now we assume that the seller auctions non-identical objects.

4.2.1 The Simultaneous Ascending Auction

The best-understood auction format in this environment is the simultaneous ascending auction (SAA).[56] The rules of this auction are the following. Multiple objects 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. 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 he computes 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 or withdrawn. The auction concludes when no new bids are submitted. The standing high bids are then deemed to be winning bids, and the winners pay an amount equal to the standing high bid.

In the simple case in which each bidder can buy at most one object, the SAA is reasonably well understood in the SIPV context. An equilibrium is established when bidders bid 'straightforwardly', i.e., in each round, each bidder that currently does not have a standing high bid, bids for that object which currently offers the highest surplus (the highest difference between value and price), and they drop out once the highest available surplus becomes negative.[57] This equilibrium is efficient. Indeed, this is one important reason why auction experts have convincingly advocated the use of the SAA to sell license for mobile telecommunication, both in the US (second generation mobile telecommunication) and Europe (UMTS).

Although the SAA has nice theoretical properties, there are some practical disadvantages. For instance, the SAA may perform poorly with respect to revenue in uncompetitive situations, i.e. when the number of objects available exactly equals the number of 'advantaged' bidders.

[56] See e.g. Milgrom (2004).

[57] Leonard (1983) and Demange et al. (1986).

Weaker bidders are reluctant to participate in the auction, and those that are present bid especially cautiously because of the enhanced 'winner's curse' they face. Klemperer (2002) suggests incorporating a first-price element to bolster competition in this case. Indeed, Goeree and Offerman (2003a) show in a laboratory experiment that the seller's revenues are the highest among a number of first-price formats when the licenses are sold sequentially, in decreasing order of quality.

Another practical drawback of the SAA is related to the time it takes for the auction to complete. For instance, in the UMTS auctions in Europe, it sometimes took several weeks for the auction to finish. The 'proxy auction' is much faster. This auction format implements straightforward bidding: a computer bids on behalf of the bidders, who indicate for each object which amount of money they are maximally willing to pay. The computer takes this maximal willingness to pay as a bidder's value, and then bids straightforwardly for all bidders until the auction ends. Suppose for simplicity that bidders have fixed valuations which are possibly different for different licenses, but which were not affected by information held by other companies, or by information released during the auction. Then it is a dominant strategy for each bidder to reveal her true willingness to pay for each object, so that the outcome of the proxy auction is efficient. When valuations are not private, a disadvantage of the proxy auction relative to the SAA is that the bidders cannot adjust bids when during the bidding process information is revealed which affects their estimation of the objects' values.

4.3 Auctions of multiple identical objects with multi-object demand

In this subsection, we return to the situation in which all objects are identical, now assuming that bidders are allowed to win more than one unit. We present multi-unit extensions of the four standard auctions dealt with in Section 2: the pay-your-bid auction, the uniform-price auction, the multi-unit Vickrey auction, and the Ausubel auction. We will see that several nice properties of single-object auctions which still remain valid under single-object demand, seize to hold under multi-object demand. Finally, we return to the SAA, and discuss some practical drawbacks of this auction under multi-object demand.

4.3.1 Pay-Your-Bid Auction

As said, the pay-your-bid auction can be viewed as the multi-unit extension of the first-price sealed-bid auction. Bidders now simultaneously submit several sealed bids. These bids should comprise weakly decreasing inverse demand curves $p_i(q)$, for $q \in \{1, ..., k\}$ and $i \in \{1, ..., n\}$. $p_i(q)$ represents the price offered by bidder i for the q-th unit. Each bidder wins the quantity demanded at the clearing price, and pays the amount that she bid for each unit won.

In a model where q can be any positive real number below a certain threshold value, Ausubel and Cramton (2002) show that under certain circumstances the pay-your-bid auction results in an efficient allocation of the object. To be more precise, the pay-your-bid auction is efficient if (1) bidders are symmetric, in the sense that the joint distribution governing the bidders' valuations is symmetric with respect to the bidders, (2) each bidder i has a constant marginal valuation for every quantity $q \in [0, \lambda_i]$, where λ_i is a capacity limitation on the quantity of units that bidder i can consume, and (3) the bidders are symmetric in their capacity limitations: $\lambda_i = \lambda$, for all bidders i. Otherwise, the auction may be inefficient. One reason for this is that bidders submit a higher bid on the first unit than on the second if they have the same value for each unit.[58]

However, Swinkels (1999) shows in a general setting that if the number of bidders gets arbitrarily large, the inefficiency in the pay-your-bid auction goes to zero. This result is interesting, as it shows that the pay-your-bid auction does rather well when the seller is able to attract a large number of bidders. This may partly explain why the pay-your-bid auction is popular in practice, for instance to sell treasury bills.[59] In addition, this result also formalizes the more general idea that competition in a market leads to an efficient allocation of resources when none of the agents has market power.

[58] Engelbrecht-Wiggans and Kahn (1998b).

[59] Binmore and Swierzbinski (2000).

4.3.2 Uniform-Price Auction

In the uniform-price auction, bidders simultaneously submit sealed-bids comprising inverse demand curves. Each bidder wins the quantity demanded at the clearing price, and pays the clearing price for each unit she wins. This auction format raised quite some confusion among economists. For instance, Nobel Prize winners Milton Friedman and Merton Miller thought incorrectly that the uniform-price auction was the multi-unit extension of the single unit Vickrey auction. Both argued that, like in the Vickrey auction, bidders would truthfully reveal their values in the auction, so that the auction outcome is efficient.[60] This is actually only true under similarly strong symmetry conditions as for the pay-your-bid auction.[61] Usually, however, bidders have an incentive to shade their bidding on some units, so that the uniform-price auction is inefficient.

The reason is simple: if a bidder can demand more than one unit, she can influence the market price and she will profit from a lower market price on all the units that she gets. Suppose that a bidder is bidding for two units to which she attaches the same value. Imagine that she bids her true value on both units. Then the other bidders may happen to bid in such a way that her bid on the second unit is the highest rejected bid. In that case, she wins one unit for which she pays a price equal to her own bid on the second unit. Therefore, she strictly prefers to bid lower on the second unit. As a consequence, bidding the same for both units cannot happen in equilibrium.[62] This phenomenon has become known as 'demand reduction': bidders understate their true value for some units.

The above example might suggest that demand reduction is limited to a small number of bidders, or is mainly a theoretical concept. However, this is not the case as the following stylized example shows. Imagine that n bidders compete for n units, which are each worth \$1 for every bidder. Suppose that each bidder bids \$1 for one unit, and \$0 for a second, third, etc. unit. The resulting price is \$0 (the highest rejected bid is equal to \$0).

[60] Quotations from press articles can be found in Ausubel and Cramton (2002).

[61] Ausubel and Cramton (2002).

[62] Noussair (1995) and Engelbrecht-Wiggans and Kahn (1998a) examine uniform-price auctions where each bidder desires up to two identical units. They find that a bidder generally has an incentive to bid sincerely on her first unit but to shade her bid on the second unit.

While demand reduction implies that a bidder will not win some units that she would have liked to win, it is advantageous because it reduces the price which the bidder has to pay for all those units which she will win. Demand reduction may imply that the auction outcome is inefficient, and may also result in lower revenues than an efficient auction.[63]

Like in the pay-your-bid auction, if the number of bidders becomes very large, the outcome of the uniform-price auction is efficient. The intuition is simple. The probability that a bidder is the highest losing bidder is zero, and so is the probability that she determines the price. This may explain why the uniform-price auction is sometimes used in practice, for instance in auctions for treasury bills.[64]

The ranking of the uniform-price auction and the pay-your-bid auction in terms of revenue is ambiguous: depending on the circumstances, the expected revenue of one auction may be higher than the other.[65] Also in practice, like in treasury bill auctions, one auction does not dominate the other in terms of expected revenue.[66]

4.3.3 Multi-Unit Vickrey Auction

We have just observed that both the pay-your-bid auction and the uniform-price auction are inefficient in many circumstances. The multi-unit Vickrey auction (proposed by Vickrey, 1961) solves this problem. In fact, the multi-unit Vickrey auction is the correct generalization of the single unit Vickrey auction. The rules are the following. Bidders submit demand schedules. The auctioneer orders the bids from highest to lowest and awards the k highest bids. Each bidder pays for the j-th unit that she wins an amount equal to the j-th highest rejected bid of her opponents. In other words, each bidder should pay for the externality that she imposes on the other bidders by winning.

Let us discuss a simple example to clarify the rules of the multi-unit Vickrey auction.

[63] Ausubel and Cramton (2002) and Engelbrecht-Wiggans and Kahn (1998a).
[64] Binmore and Swierzbinski (2000).
[65] Ausubel and Cramton (2002).
[66] Binmore and Swierzbinski (2000).

Example 2 *Imagine that three bidders bid for three units. Their bids on the j-th unit are given in Table 5. Bidder 1 wins two units, as she submits the highest two bids, and bidder 2 wins a single unit, as her bid on the first unit is the third highest bid. Bidder 1 pays 10, as the highest two rejected bids from the other bidders are 6 and 4. Similarly, bidder 2's payment is 7.*

Unit\Bidder	1	2	3
1st	10	8	6
2nd	9	4	3
3rd	7	3	3

Table 5: Bids on the j-th unit for each bidder.

The multi-unit Vickrey auction has an equilibrium in weakly dominant strategies, in which each bidder bids her value for each unit. In such an equilibrium, the outcome is efficient. Still, the multi-unit Vickrey auction is rarely used in practice, presumably for the very same reasons why the Vickrey auction is hardly ever applied: it may not be very obvious to bidders that it is optimal for them to reveal their demand, items may be sold far below the winner's willingness-to-pay, and bidders may not be willing to reveal information in the auction as the seller may use this information in later occasions. In the multi-unit case, there is another reason why Vickrey auctions are not very popular. As the above example shows, the per unit price may differ substantially, usually at the advantage of 'large' bidders, i.e., bidders who win many units in the auction. Politically, this may be hard to sell.

4.3.4 Ausubel Auction

The Ausubel auction (named after Lawrence Ausubel) is the dynamic version of the multi-unit Vickrey auction. In this sense, the Ausubel auction is the correct generalization of the single unit English auction. The rules are the following. The auction price starts at zero and then increases continuously. At any price, each bidder indicates how many units she wants. At a certain price, it will happen that one bidder is guaranteed to win one or more units. This is the case when the aggregate demand of the other bidders is smaller than the available number of units. The bidder then wins the residual supply at the current price. In Ausubel's words: the bidder 'clinches' these units. The auction then continues from this price with the remaining units, until the next unit is clinched. This process continues until all units are allocated. The Ausubel auction has roughly the same advantages over the multi-unit Vickrey auction as the English auction has over the Vickrey auction under single-unit supply.[67]

4.3.5 The Simultaneous Ascending Auction

The SAA may also be implemented to allocate identical objects to bidders who demand more than one unit. However, the SAA loses many of the nice properties we discussed in the previous subsection. In particular, the SAA is sensitive to collusion, demand reduction, and the 'exposure problem'.

The German GSM auction shows how bidders may be able to collude under the SAA. In 1999, the German government auctioned ten GSM licenses using the SAA with a 10% minimum bid increment. In the first round, Mannesmann made a jump-bid of 18.18 million Deutsche Mark (DM) per megahertz (MHz) on the first five licenses, and 20 million DM on the last five. Doing so, Mannesmann signalled to its main competitor T-Mobile that it would be happy to share the ten licenses equally. To see this, observe that if T-Mobile overbids the standing high bid on licenses 1-5 with 10%, the final price is almost exactly 20 million DM per MHz. T-Mobile indeed understood Mannesmann's signal, and the auction ended after just two rounds.[68]

This example also shows that the SAA is sensitive to demand reduction. If Mannesmann and T-Mobile had bid up to their willingness-to-pay for the ten licenses, the auction probably would have ended with a much higher price per MHz. Even ignoring the possibility to collude, bidders have a strong incentive to reduce their demand in order to avoid winning items at a high price. Demand reduction is also reported in US spectrum auctions.[69]

Finally, the SAA is sensitive to the so-called exposure problem. The exposure problem occurs if bidders face the risk of winning too few objects. This may result in low revenue

[67] Ausubel (2004).
[68] Jehiel and Moldovanu (2001) and Grimm et al. (2003).
[69] Weber (1997).

and an inefficient allocation of the objects.[70] The 1998 GSM auction in the Netherlands is an example of a situation in which the exposure problem was present. The Dutch government had split up the spectrum in 18 licenses. Two of these licenses contained sufficient spectrum to operate a GSM network. The other 16 licenses, however, were so small that an entrant needed at least four of them to be able to operate a network. Bidders were discouraged to submit high bids on the small licenses, as they faced the risk of being overbid on a fraction of them, and winning insufficient spectrum. In fact, the per MHz price on the large licenses turned out to be about two and a half times as high as for the small licenses. This violation of 'the law of one price' may indicate that the allocation of the licenses was inefficient.[71]

4.4 Auctions of multiple non-identical objects with multi-object demand

In the previous subsections, we have observed that with single-object demand or identical objects, rather straightforward auctions generate an efficient allocation of the objects. In this subsection, we will see that in the case of multi-object demand and non-identical objects, there are still mechanisms that result in an efficient allocation, but that these mechanisms have serious practical drawbacks. We finish by discussing the clock-proxy auction, an innovative design that may partly mitigate these disadvantages.

4.4.1 The Vickrey-Clarke-Groves mechanism

We start by describing the Vickrey-Clarke-Groves (VCG) mechanism, which is efficient under a large range of circumstances. The VCG mechanism is developed by Clarke (1971) and Groves (1973), and generalizes the (multi-unit) Vickrey auction.[72] The most important property of VCG mechanisms is that these are able to allocate objects efficiently under fairly general conditions. Let us study the following model to make this claim more precise.

[70] Onderstal (2002b) and van Damme (1999).

[71] Onderstal (2002b) and van Damme (1999).

[72] Clarke and Groves constructed this mechanism for a class of problems that is far more general than the allocation of objects: their mechanism applies to any public choice problem.

Assume that a seller wishes to allocate k objects among n bidders, labeled $i = 1, ..., n$. Let Γ be the set of possible allocations of the k objects over the bidders, and t_i the monetary transfer by bidder i (where a negative number indicates a payment to bidder i). Bidder i has the following quasi-linear utility function:

$$u_i(\gamma, t_i, \theta_i) = v_i(\gamma, \theta_i) - t_i,$$

where $\gamma \in \Gamma$ is a feasible allocation, v_i a valuation function, and θ_i bidder i's type, which is private information to bidder i. Note that this model is very general: no assumptions are made with respect to the risk attitude of the bidders and the distribution of the types, no structure is assumed on the complementarity or substitutability of the objects, even allocative externalities are included in this model. The main restrictions are (1) that utility is additively separable in money and the allocation of the objects, (2) a bidder's utility does not depend on other bidders' types, and (3) the exclusion of financial externalities, i.e., a bidder's utility does not depend on how much other bidders pay. Note that in this model, an allocation γ^* of the objects over the bidders is ex post efficient if and only if

$$\sum_i v_i(\gamma^*, \theta_i) \geq \sum_i v_i(\gamma, \theta_i) \text{ for all } \gamma \in \Gamma. \tag{6}$$

A VCG mechanism has the following properties. All bidders are asked to announce a type $\tilde{\theta}_i$. Let $\tilde{\theta}$ be the vector of announcements, i.e., $\tilde{\theta} \equiv (\tilde{\theta}_1, ..., \tilde{\theta}_n)$. The objects are allocated efficiently under the assumption that the $\tilde{\theta}_i$'s are the true types. Let $\gamma^*(\tilde{\theta})$ denote such an allocation. Moreover, bidder i pays an amount $t_i(\tilde{\theta})$ equal to

$$t_i(\tilde{\theta}) = \sum_{j \neq i} v_j(\gamma^*_{-i}(\tilde{\theta}), \tilde{\theta}_j) - \sum_{j \neq i} v_j(\gamma^*(\tilde{\theta}), \tilde{\theta}_j), \tag{7}$$

with $\gamma^*_{-i}(\tilde{\theta})$ an allocation that would be efficient if there were only $n-1$ bidders $j \neq i$ and the announced types were the true types. In words, bidder i pays an amount equal to the externality that she imposes on the other bidders.

Lemma 3 *For each bidder, it is a weakly dominant strategy to announce her true type in a VCG mechanism.*

The following result then immediately follows from Lemma 3, as in equilibrium the allocation of the objects is efficient by (6).

Proposition 11 *The VCG mechanism has an efficient equilibrium in weakly dominant strategies.*

Unfortunately, the VCG mechanism has many practical disadvantages. We have elaborated on several of these disadvantages while discussing the Vickrey auction and the multi-unit Vickrey auction: (1) it is not straightforward for bidders to understand how to play the auction, (2) the outcome may be politically problematic as items may be sold far below the willingness-to-pay of the winner, (3) bidders may not be willing to reveal information in the auction as the seller may use this information on later occasions, and (4) the per unit price may differ substantially, usually at the advantage of 'large' bidders. In the case of heterogeneous objects, there are several additional drawbacks. First, bidding is complex as bidders have to specify bids on all packages they desire to win. If 10 objects are for sale, a bidder may specify a bid on $2^{10} - 1 = 1023$ packages. Second, more competition may lead to lower prices, which may be as low as zero even if competition is fierce. Third, the VCG mechanism is sensitive to collusion. Finally, the VCG mechanism is not robust against 'shill bidding'.

Milgrom (2004) constructs a very nice example to illustrate the last three problems. Suppose that two spectrum licenses (A and B) are put up for sale. Assume first that there are only two entrants interested in buying the licenses. For each of them the value of a single license is \$0. The pair of licenses is worth \$1 billion for bidder 1 and \$900 million for bidder 2. If these bidders are the only bidders in the auction, the VCG mechanism boils down to a Vickrey auction where the pair of licenses is for sale. Bidder 1 will win both licenses for a price of \$900 million.

Now, suppose that two incumbents also participate in the auction. Bidder 3 [bidder 4] is already active in region B [region A] and is only interested in license A [license B], for which she is willing to pay \$1 billion. If all four bidders play their weakly dominant strategy in the VCG mechanism, then the licenses will be allocated to the incumbents (as this is the most efficient allocation). Surprisingly, both bidders get the licenses for free. Why? If one of the incumbents does not show up, the value to the other bidders is \$1 billion, which happens to be exactly the same value to the other bidders if she is present. In other words the externality she imposes on the other bidders is \$0. To summarize, more intense competition may lead to lower revenue to the seller.

Next, imagine that bidders 3 and 4 only have a value for a license equal to \$400. In the case that they play their weakly dominant strategies, they will not win a license. However, if they coordinate in such a way that both bid \$1 billion on their license, they do win for a price of \$0. In other words, the VCG mechanism is sensitive to collusion.

Third, consider another situation in which bidders 1 and 2 participate, together with a third entrant who values the two licenses at \$800 million. If all players play their weakly dominant strategies, bidder 3 will not win. However, bidder 3 can hire a 'shill' bidder. If she and the shill bidder bid \$1 billion on license A and B respectively, then they will win both licenses for a price equal to zero. Hence, the VCG mechanism is not robust against shill bidding.

4.4.2 The Clock-Proxy Auction

In the previous subsection, we have seen that the VCG mechanism, though efficient in theory, has some serious practical drawbacks. Ausubel, Cramton, and Milgrom (2005) propose the clock-proxy auction (CPA) as an alternative to mitigate these disadvantages. The CPA consists of two stages: a clock phase and a proxy round.

In the clock phase, the auctioneer announces a price for each object put up for sale. The bidders express the quantities of each object desired at the specified prices. Then the prices are increased for objects in excess demand, while other prices remain unchanged. Next, the bidders express quantities at the new prices. This process is repeated until there is no object with excess demand. The market-clearing prices serve as a lower bound on the prices in the proxy phase.

The proxy phase consists of a single round in which each bidder reports her values to a proxy agent for all packages she is interested in. The proxy agent iteratively submits package bids in an ascending package auction on behalf of the real bidder, selecting these packages that would maximize the real bidder's profit given the bidder's reported values. After each round, the auctioneer selects the provisionally-winning bids that maximize revenues, also considering the bids submitted in the clock phase. This process continues until the proxy agents have no new bids to submit.[73] The winners pay an amount equal to the standing high bids.

There are several reasons why the CPA may be expected to result in desirable outcomes. First of all, the CPA is efficient, just as the VCG mechanism.[74] Second, the auction ends at a 'core' allocation for the reported preferences when bidding is straightforward, implying seller revenues are competitive.[75] In other words, the seller will always generate sufficient revenue if competition is fierce. This is in contrast to the VCG mechanism, in which the revenue may decrease all the way down to zero if additional bidders enter the auction, as the example in the previous subsection showed. Third, the CPA is expected to handle pretty well other complications of the VCG mechanism (such as collusion), and of the SAA (collusion, demand reduction, and the exposure problem).[76] Finally, the CPA has the advantage that in the clock phase, valuable information about the prices can be revealed, in contrast to sealed-bid formats such as the VCG mechanism.

Still, the CPA has at least three disadvantages. First, bidding in the proxy phase is as complex as bidding in the VCG mechanism, as bidders may wish to specify bids on a large range of packages. Second, the proxy phase consists of a single round, so that there is no way for bidders to learn from each others' bid about the values of the packages. The practical reason for this is that the allocation problem is 'NP-hard': there are no general ways for solving the problem in reasonable time, so that it may take a long time for subsequent rounds to finish.[77] Third, if competition is strong and objects are mostly substitutes, then a clock auction without a proxy round may be the best approach, since it offers the greatest simplicity and transparency, while being highly efficient. In fact, clock auctions have been implemented in the field for

[73] It would take very long for the process to complete if the auctioneer ran the proxy auction with negligible bid increments. However, the process can be accelerated by using various methods, see Day and Raghavan (2005), Hoffman et al. (2005), and Wurman et al. (2005).

[74] Ausubel and Milgrom (2002).

[75] Ausubel and Milgrom (2002).

[76] Ausubel et al. (2005).

[77] See e.g. de Vries and Vohra (2003).

products like electricity in recent years with considerable success.[78] In this simple setting, the SAA also performs well. However, a clock auction is to be preferred as it has a couple of advantages over the SAA: (1) complex bid signalling and collusive strategies are eliminated, since the bidders cannot see individual bids, but only aggregate information, (2) the exposure problem is eliminated: bidders are free to reduce quantities on any object (as long as at least one price increases), and (3) clock auctions are faster, as the SAA is in particular slow near the end when there is little excess demand.

4.5 Summary

In this section, we have studied multi-object auctions, focussing on the trade-off between efficiency and complexity. Table 6 gives an overview of efficient auctions under different assumptions with respect to the bidders' demands. These auctions are 'simple', apart from the VCG mechanism when bidders have demand for several non-identical objects. We have argued that the CPA is a reasonable substitute for this mechanism.

	Identical Objects	Non-Identical Objects
Single-Object Demand	Pay-Your-Bid Auction Multi-Unit Dutch Auction Uniform-Price Auction Multi-Unit English Auction	SAA
Multi-Object Demand	Pay-Your-Bid Auction (for many bidders) Uniform-Price Auction (for many bidders) Multi-Unit Vickrey Auction Ausubel Auction	VCG mechanism

Table 6: Efficient auction formats.

[78] Ausubel and Cramton (2004).

5 CONCLUSIONS

In this paper, we have made a swift tour of auction theory and its applications. We have roughly followed the historical development of the field which started, in the 1960s, with the work of William Vickrey. For single-object auctions, we have observed how auction theorists have developed the celebrated revenue equivalence theorem, how they have constructed revenue maximizing auctions, and how they have shown that under various circumstances the revenue equivalence between commonly used auctions breaks down. One of the main findings has been the trade-off between efficiency and revenue. In addition, we have discussed how to optimally auction incentive contracts. Finally, we have seen that in the case of auctions for more than one object a new trade-off enters the picture, namely the trade-off between efficiency and complexity.

5.1 Further research

Despite the great achievements of the past decades, numerous issues, especially with regards to applied work, remain open for further exploration. Let us distinguish between two types of applications: (1) auctions in practice and (2) applications to other economic situations.

5.1.1 Auctions in practice

An important contribution of auction theory is to help auction designers understand under which circumstances what type of auction works best. This is not straightforward as, according to Klemperer (2002):

> "[i]n auction design, the devil is in the details."

Some details are not, yet, well understood and may require more attention in future research. We mention three of them.

One detail that requires further attention is the design of multi-object auctions under multi-object demand. We have observed that when the auctioneer sells identical objects such as government bonds or when there is single-object demand like for UMTS-licenses, rather straightforward auctions work well both in theory and in practice. However, things may become very complex when bidders demand several heterogeneous objects. Theoretically, a VCG mechanism generates an efficient allocation of the objects, but this mechanism is probably too complex to be successfully implemented in practice. The CPA proposed by Ausubel, Cramton, and Milgrom may solve some of the practical problems arising from the VCG mechanism. However, it is not yet understood under what circumstances we may expect the CPA to work well. In the near future, this design may be implemented in practice, which might generate valuable lessons. It may be a good idea to thoroughly test the CPA in laboratory experiments.

A second topic that deserves further research is related to firms going bankrupt after overbidding in auctions. This issue plays an important role in practice. For instance, some telecommunication firms were on the edge of bankruptcy because of the high amounts bid in UMTS auctions. Bankruptcy may be undesirable from the consumers' point of view as it may increase the market power of the remaining providers in the market, and consumers may need to spend time and money searching for a new provider. Assuming that bankruptcy is undesirable from a social point of view, it would be interesting to investigate what governments can do to prevent bidders from overbidding.[79] In addition, several critics have claimed that firms will have to raise prices for UMTS services[80] and that they will invest less in new technology compared to a situation in which the firms had not overbid on the licenses. It would be interesting to analyze whether the possibility of bankruptcy affects firms' bidding behavior and their pricing and investment decisions.

A third topic that seems to be worth investigating is the choice between auctions and beauty contests. A beauty contest differs from an auction in the sense that not all bidding criteria

[79] First steps in these directions have been made by Zheng (2001) and Haan and Toolsema-Veldman (2003).

[80] This argument is in contrast with the theory of sunk costs. Still, theorists and experimental economists have found some support (see Janssen (2005), Janssen and Karamychev (2005), and Offerman and Potters (2003)). Haan and Toolsema-Veldman (2003), in contrast, show that prices may *decrease* when firms pay high prices in auctions if firms are limitedly liable.

are objective.[81] For the allocation of UMTS licenses, some of the European governments used auctions, while others used beauty contests.[82] Several auction theorists argued that, in this particular context, auctions perform better.[83] Perhaps this makes sense with regards to UMTS licenses, as the main public goals are ex ante contractible (such as nation-wide coverage). However, the latter is not always true. Think about a procedure to select an architect for a new city hall. Several architects present a proposal and the council of the city chooses the most appealing one. The selection is based on objective criteria such as the price, but also on subjective criteria including the 'beauty' of the design. It seems to be virtually impossible to ex ante specify how these subjective criteria will be evaluated. New research may develop formal theory to explain under which circumstances auctions are more appropriate versus beauty contests.

5.1.2 Applications to other economic situations

Some economists argue that the insights and techniques gained from auction theory are not only valuable for the specialized field of auction design, but also for mainstream economics. For instance, Klemperer (2003) claims that:

> "[a]uction theory refreshes the parts other economics cannot reach."

In this paper, we have observed how auction-like games such as the all-pay auction and the war of attrition have been applied to a range of economic situations, such as lobbying and technology battles. In fact, any situation in which several 'agents' 'compete' for a 'prize' might be fruitfully modeled as an auction. Examples are tax competition (several jurisdictions offer advantages to attract new factories), the labor market (several firms make a job offer to a potential employee), and advertising (several firms spend money and effort in advertising to

[81] Dykstra and van der Windt (2004).
[82] Maasland and Moldovanu (2004).
[83] See, e.g., Binmore and Klemperer (2002).

attract customers).[84] New insights may emerge when auction theory is applied to these fields, and established ones may be refreshed.

[84]First steps in this direction have been made by Menezes (2003) (tax competition), Julien et al. (2000) (labor market), and Onderstal (2002a) (advertising).

6 APPENDIX

Proofs of Propositions 1, 5, 6, and 10, and Lemmas 1-3 follow.

PROOF OF PROPOSITION 1

Two different techniques can be used to prove this proposition. The first is the 'direct' method, in which we assume that all bidders but bidder 1 use the same bidding strategy $b : [0, \bar{v}] \to [0, \infty)$. Then we construct the best response b_1 for bidder 1, and we derive conditions under which this best response is equal to $b(v_1)$. This yields us an educated guess of what may be a symmetric Bayesian-Nash equilibrium. Finally, we need to check whether this strategy indeed constitutes an equilibrium.

For the moment, assume that b is strictly increasing and differentiable. Let b^{-1} be the inverse function of b. Bidder 1's expected payoff from bidding b_1 is given by the difference between her value and her bid, multiplied by the probability that she wins:

$$\begin{aligned} U_1(v_1, b_1, b) &= (v_1 - b_1) \Pr(b_1 > b(v_2), ..., b_1 > b(v_n)) \\ &= (v_1 - b_1) F(b^{-1}(b_1))^{n-1}. \end{aligned}$$

When deriving bidder 1's best response, the first order condition is

$$\frac{\partial U_1}{\partial b_1} = (v_1 - b_1) \frac{\partial F(b^{-1}(b_1))^{n-1}}{\partial b_1} - F(b^{-1}(b_1))^{n-1} = 0.$$

We wish to construct a symmetric equilibrium, so that $b_1 = b$, or $b^{-1}(b_1) = v_1$ for all v_1. Solving for the resulting differential equation we find

$$b(v) F(v)^{n-1} = \int_0^v x(n-1) F(x)^{n-2} dF(x) + c,$$

where c is the constant of integration. As a bidder with value 0 will always bid 0 in equilibrium, the boundary condition is $b(0) = 0$, so that $c = 0$. Hence, a natural candidate for a symmetric Bayesian-Nash equilibrium is

$$b(v) = v - \frac{\int_0^v F(x)^{n-1} dx}{F(v)^{n-1}}. \tag{8}$$

It can be shown that b is indeed strictly increasing and differentiable, the assumptions we started with. An implication is that the equilibrium outcome is efficient, i.e., it is always the highest bidder who obtains the object.

Finally, we have to check that b is indeed an equilibrium. Assume that all bidders but bidder 1 submit a bid according to b. Is it optimal for bidder 1 to follow this strategy as well? As b is strictly increasing, the bidder with the highest value submits the highest bid and wins the auction. Obviously, bidder 1 does not wish to submit a bid $b_1 > b(\bar{v})$. As b is strictly increasing and continuous, a bid $0 \leq b_1 \leq b(\bar{v})$ corresponds to a unique value w for which $b(w) = b_1$. We can write bidder 1's expected profit from bidding b_1 as

$$\begin{aligned} U(w, v_1) &= F^{n-1}(w)[v_1 - b(w)] \\ &= F^{n-1}(w)(v_1 - w) + \int_0^w F(x)^{n-1} dx \\ &\leq \int_0^{v_1} F(x)^{n-1} dx \\ &= U(v_1, v_1). \end{aligned}$$

Therefore, bidder 1 optimally chooses $b_1 = b(v_1)$ so that indeed b constitutes a symmetric Bayesian-Nash equilibrium.

The 'indirect' method is an alternative way to derive a symmetric Bayesian-Nash equilibrium. Let $b_i : [0, \bar{v}] \to [0, \infty)$ be the equilibrium bid function for bidder i. We assume that all bidders reveal a value to the auctioneer, and that the auctioneer calculates the equilibrium bids as if the signals were the true signal. To prove that the bidding function b_i constitutes an equilibrium, we must show that all bidders have an incentive to reveal their true value. Again, we start by assuming that all bidders but bidder 1 use the same bidding strategy $b : [0, \bar{v}] \to [0, \infty)$. Next, we define the utility $U(v, w)$ for bidder 1 having value v who misrepresents herself as having value w, whereas the other bidders report truthfully. Then we derive conditions under which bidder 1 wishes to honestly reveal her type. From these conditions, we are able to derive the equilibrium bidding strategy.

For the moment, assume that b is strictly increasing and differentiable. Then

$$U(v,w) = (v - b(w))F(w)^{n-1},$$

so that

$$\frac{\partial U(v,w)}{\partial w} = (v - b(w))\frac{\partial F(w)^{n-1}}{\partial w} - b'(w)F(w)^{n-1}. \tag{9}$$

For bidder 1, it should be optimal to reveal her true value, so that $U(v,w)$ is maximized at $w = v$. Hence, the first order condition of the equilibrium is

$$\left.\frac{\partial U(v,w)}{\partial w}\right|_{w=v} = 0.$$

The resulting differential equation turns out to be the same as under the 'direct' method, so that (8) is a solution.

We still need to verify whether the second order condition $sign\left(\frac{\partial U(v,w)}{\partial w}\right) = sign(v-w)$ holds true. Observe that (9) can be rewritten as:

$$\begin{aligned}\frac{\partial U(v,w)}{\partial w} &= \left.\frac{\partial U(y,w)}{\partial w}\right|_{y=w} + (v-w)(n-1)F(w)^{n-2}f(w) \\ &= (v-w)(n-1)F(w)^{n-2}f(w). \end{aligned} \tag{10}$$

The second equality follows from the observation that

$$\left.\frac{\partial U(x,w)}{\partial w}\right|_{x=w} = 0.$$

From (10), it immediately follows that the second order condition is satisfied.

As it is always the bidder with the highest value who submits the highest bid, the expected revenue R_{FPSB} equals the expected bid of the bidder with the highest value:

$$
\begin{aligned}
R_{FPSB} &= \int_0^{\overline{v}} B_{FPSB}(v)dF(v)^n \\
&= \int_0^{\overline{v}} vdF(v)^n - \int_0^{\overline{v}} nf(v) \int_0^{v} F(x)^{n-1}dxdv \\
&= \int_0^{\overline{v}} vdF(v)^n - \int_0^{\overline{v}} \int_x^{\overline{v}} nf(v)dvF(x)^{n-1}dx \\
&= \int_0^{\overline{v}} vdF(v)^n - n\int_0^{\overline{v}} (1-F(v))F(v)^{n-1}dx \\
&= \int_0^{\overline{v}} vd\left[F(v)^n + n(1-F(v))F(v)^{n-1}\right] \\
&= E\{Y_2^n\}.
\end{aligned}
$$

Proof of Proposition 5

We use the 'indirect' method to solve for the symmetric equilibrium bidding strategies. Let b be a strictly increasing and differentiable bidding function, and assume that all bidders but 1 use this strategy. Imagine that bidder 1 with value v wishes to act as if having value w. Let $U(v,w)$ be bidder 1's expected utility. Then,

$$U(v,w) = vF(w)^{n-1} - b(w).$$

Maximizing $U(v,w)$ with respect to w yields the first order condition of the equilibrium:

$$(n-1)vF(v)^{n-2}f(v) - b'(v) = 0.$$

Integrating over v and applying the boundary condition $b(0) = 0$ yields

$$b(v) = (n-1)\int_0^{v} xF(x)^{n-2}dF(x).$$

It is readily verified that b is indeed strictly increasing and continuous, and that the second order condition holds. As all bidders pay their bid, the seller expects to collect $nE[b(v)]$ which can be shown to be equal to $E\{Y_2^n\}$.

Proof of Proposition 6

We use the 'direct' method to derive the equilibrium. Assume that bidder 2 employs the strictly increasing and differentiable bidding function b. The expected utility for bidder 1 when bidding b_1 is equal to

$$U(v_1, b_1, b(\cdot)) = \int_0^{b^{-1}(b_1)} (v_1 - b(x))dF(x) - b_1(1 - F(b^{-1}(b_1)))$$

where the first [second] term on the right hand side indicates bidder 1's utility when she wins [loses]. The first order condition can be expressed as

$$\frac{\partial U}{\partial b_1} = (v_1 - b_1)f(b^{-1}(b_1))(b^{-1}(b_1))' - (1 - F(b^{-1}(b_1))) + b_1 f(b^{-1}(b_1))(b^{-1}(b_1))' = 0.$$

Substituting $b^{-1}(b_1) = v_1$ and some routine calculations yield

$$b'(v_1) = \frac{v_1 f(v_1)}{1 - F(v_1)}.$$

Solving for b with boundary condition $b(0) = 0$, we obtain

$$b(v) = \int_0^v \frac{xf(x)}{1 - F(x)}dx \tag{11}$$

as a candidate for the symmetric Bayesian Nash equilibrium. It is then straightforwardly checked that b indeed constitutes an equilibrium. The expected equilibrium revenue R_{WoA} equals twice the bid of the lowest bidder, which, once again, is equal to $E\{Y_2^n\}$ (with $n = 2$).

Proof of Lemmas 1 and 2

Let us first introduce some additional notation and concepts. Define the sets $V \equiv [0, \bar{v}]^n$ and $V_{-i} \equiv [0, \bar{v}]^{n-1}$ with typical elements $\mathbf{v} \equiv (v_1, ..., v_n)$ and $\mathbf{v}_{-i} \equiv (v_1, ..., v_{i-1}, v_{i+1}, ..., v_n)$ respectively. Let $g(\mathbf{v}) \equiv f(v_1) \cdots f(v_n)$ be the joint density of $\mathbf{v}$, and let $g_{-i}(\mathbf{v}_{-i}) \equiv f(v_1) \cdots f(v_{i-1})f(v_{i+1}) \cdots f(v_n$ be the joint density of $\mathbf{v}_{-i}$.

We define an auction as follows. In an auction, bidders are asked to simultaneously and independently choose a bid. Bidder i chooses a bid $b_i \in B_i$, where B_i is the set of possible bids for bidder i, $i = 1, ..., n$. Let $\mathbf{b} = (b_1, ..., b_n)$ be the vector of bids. An auction is characterized by its outcome functions $(\widehat{p}_i, \widehat{x}_i)_{i=1,...,n}$, where $\widehat{p}_i(\mathbf{b})$ is the probability that bidder i wins the object, and $\widehat{x}_i(\mathbf{b})$ is the expected payment of bidder i to the seller.

Lemma 1 follows from the following considerations. Consider an auction and the following direct revelation game. First, the seller asks each bidder to announce her value. Then, he determines the bid that each bidder would have chosen in the equilibrium of the auction given her announced value. Next, he implements the outcomes that would result in the auction from these bids. As the strategies form an equilibrium of the auction, it is an equilibrium for each bidder to announce her value truthfully in the direct revelation game. Therefore, the revelation game has the same outcome as the auction, so that both the seller and the bidders obtain the same expected utility as in the equilibrium of the auction.[85]

Bidder i's utility of direct revelation mechanism (p, x) given $\mathbf{v}$ equals $v_i p_i(\mathbf{v}) - x_i(\mathbf{v})$, so that if bidder i knows her value v_i, her expected utility from (p, x) can be written as

$$U_i(p, x, v_i) \equiv \int_{V_{-i}} [v_i p_i(\mathbf{v}) - x_i(\mathbf{v})] g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i}, \tag{12}$$

with $d\mathbf{v}_{-i} \equiv dv_1...dv_{i-1}dv_{i+1}...dv_n$.

The individual rationality constraint follows from the assumption that each bidder expects nonnegative expected utility, so that

$$U_i(p, x, v_i) \geq 0, \; \forall v_i, i. \tag{13}$$

The incentive compatibility constraint is imposed as we demand that each bidder has an incentive to announce her value truthfully. Thus,

$$U_i(p, x, v_i) \geq \int_{V_{-i}} [v_i p_i(\mathbf{v}_{-i}, w_i) - x_i(\mathbf{v}_{-i}, w_i)] g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i}, \; \forall v_i, w_i, i,$$

where $(\mathbf{v}_{-i}, w_i) \equiv (v_1, ..., v_{i-1}, w_i, v_{i+1}, ..., v_n)$.

The seller aims at finding an auction which gives him the highest possible expected revenue (*the seller's problem*). The seller's expected revenue of (p, x) is

[85] Note that we have already applied the revelation principle in the 'indirect method' for deriving an equilibrium for an auction (see the proof of Proposition 1).

$$U_0(p, x) \equiv \int_V \sum_{i=1}^{n} x_i(\mathbf{v}) g(\mathbf{v}) d\mathbf{v}, \tag{14}$$

with $d\mathbf{v} \equiv dv_1...dv_n$.

Now, let

$$Q_i(p, v_i) \equiv E_{v_{-i}}\{p_i(\mathbf{v})\}$$

be the conditional probability that bidder i wins the object given her value v_i. Lemma A1 gives a characterization of direct revelation games that are individually rational and incentive compatible.

Lemma A1 *The direct revelation game (p, x) is individually rational and incentive compatible if and only if*

$$Q_i(p, w_i) \geq Q_i(p, v_i) \text{ if } w_i \geq v_i, \forall w_i, v_i, i, \tag{15}$$

$$U_i(p, x, v_i) = U_i(p, x, \underline{v}_i) + \int_{\underline{v}_i}^{v_i} Q_i(p, y_i) dy_i, \forall v_i, i, \text{ and} \tag{16}$$

$$U_i(p, x, \underline{v}_i) \geq 0, \forall i. \tag{17}$$

Proof. Incentive compatibility implies

$$U_i(p, x, w_i) \geq U_i(p, x, v_i) + (w_i - v_i) Q_i(p, v_i), \tag{18}$$

so that (p, x) is individually rational and incentive compatible if and only if (13) and (18) hold. With (18),

$$(w_i - v_i) Q_i(p, w_i) \geq U_i(p, x, w_i) - U_i(p, x, v_i) \geq (w_i - v_i) Q_i(p, v_i),$$

from which (15) follows. Moreover, these inequalities imply

$$\frac{\partial U_i(p, x, v_i)}{\partial v_i} = Q_i(p, v_i), \tag{19}$$

at all points where p_i is differentiable in v_i. By integration of (19), (16) is obtained. Finally, with (13) and (16), individual rationality is equivalent to (17). ∎

Now, with (12), the seller's expected revenue (14) can be rewritten as

$$U_0(p,x) = \sum_{i=1}^{n} \int_V v_i p_i(\mathbf{v}) g(\mathbf{v}) d\mathbf{v} - \sum_{i=1}^{n} \int_{\underline{v}_i}^{\bar{v}_i} U_i(p,x,v_i) f_i(v_i) dv_i. \tag{20}$$

Taking the expectation of (16) over v_i and using integration by parts, we obtain

$$E_{v_i}\{U_i(p,x,v_i)\} = U_i(p,x,\underline{v}_i) + E_{v_i}\left\{\frac{1-F_i(v_i)}{f_i(v_i)} p_i(\mathbf{v})\right\}. \tag{21}$$

Lemma 2 follows from (20) and (21).

Proof of Proposition 10

The proof consists of two parts: first we consider which mechanisms are incentive compatible by looking at firms' bidding behavior, after which we derive which incentive compatible mechanism maximizes the buyer's utility.

Firms' bidding behavior

If all firms bid truthfully, firm i's expected utility given its efficiency parameter α_i is equal to

$$U_i(\alpha_i) = E_{\alpha_{-i}}[t_i(\boldsymbol{\alpha}) - x_i(\boldsymbol{\alpha})\{\varphi(e_i(\boldsymbol{\alpha})) - \alpha_i e_i(\boldsymbol{\alpha})\}], \tag{22}$$

where

$$\varphi(e) = \frac{1}{2}e^2 + e.$$

Let $U_i(\alpha_i, \tilde{\alpha}_i)$ be firm i's utility if it has efficiency parameter α_i, it announces $\tilde{\alpha}_i$, and all other firms truthfully reveal their type. Then

$$U_i(\alpha_i,\tilde{\alpha}_i) = E_{\alpha_{-i}}[t_i(\alpha_{-i},\tilde{\alpha}_i) - x_i(\alpha_{-i},\tilde{\alpha}_i)\varphi(e_i(\alpha_{-i},\tilde{\alpha}_i))] + \alpha_i E_{\alpha_{-i}}[e_i(\alpha_{-i},\tilde{\alpha}_i)x_i(\alpha_{-i},\tilde{\alpha}_i)]. \tag{23}$$

Incentive compatibility requires that firm i optimally announces its own type, so that

$$\frac{\partial U_i(\alpha_i,\tilde{\alpha}_i)}{\partial \tilde{\alpha}_i} = 0 \tag{24}$$

at $\tilde{\alpha}_i = \alpha_i$. From (22), (23), and (24), it follows that

$$\begin{aligned} \frac{dU_i(\alpha_i)}{d\alpha_i} &= \left.\frac{\partial U_i(\alpha_i, \tilde{\alpha}_i)}{\partial \tilde{\alpha}_i}\right|_{\tilde{\alpha}_i=\alpha_i} + E_{\alpha_{-i}}[e_i(\boldsymbol{\alpha})x_i(\boldsymbol{\alpha})] \\ &= E_{\alpha_{-i}}[e_i(\boldsymbol{\alpha})x_i(\boldsymbol{\alpha})]. \end{aligned} \tag{25}$$

The participation constraint then immediately reduces to

$$U_i(0) \geq 0. \tag{26}$$

The buyer's problem

To maximize the buyer's utility given (25) and (26), we apply the Pontryagin principle. The buyer solves

$$\begin{array}{ll} \max\limits_{(x_i(\cdot),e_i(\cdot),U_i(\cdot))} & E_{\boldsymbol{\alpha}} \sum\limits_i \left\{x_i(\boldsymbol{\alpha})\left[e_i(\boldsymbol{\alpha}) - C_i(e_i(\boldsymbol{\alpha}), \alpha_i)\right] - U_i(\alpha_i)\right\} \\ \text{s.t.} & \begin{array}{c} \dot{U}_i(\alpha_i) = E_{\alpha_{-i}}[e_i(\boldsymbol{\alpha})x_i(\boldsymbol{\alpha})] \\ U_i(0) \geq 0 \end{array} \quad . \end{array}$$

This problem looks horrendously complicated, but we can rely on the following three tricks to make it solvable. First, it can be shown that $e_i(\boldsymbol{\alpha})$ only depends on α_i. We do not prove this formally, but the intuition is that e_i is a stochastic scheme when it depends on announcements other than α_i. As the firm's cost function is convex, there is a deterministic scheme that only depends on α_i which strictly improves the objective function of the buyer.

The second trick is to keep the x_i's fixed, and solve the problem. Let

$$X_i(\alpha_i) \equiv E_{\boldsymbol{\alpha}_{-i}}\left\{x_i(\boldsymbol{\alpha})\right\}.$$

For given $X_i(\alpha_i)$, the buyer's problem can be decomposed into the following n independent maximization problems:

$$\begin{array}{ll} \max\limits_{(e_i(\cdot),U_i(\cdot))} & \int_0^1 \left\{X_i(\alpha_i)\left[e_i(\alpha_i) - C_i(e_i(\boldsymbol{\alpha}), \alpha_i)\right] - U_i(\alpha_i)\right\} dF(\alpha_i) \\ \text{s.t.} & \begin{array}{c} \dot{U}_i(\alpha_i) = e_i(\alpha_i)X_i(\alpha_i) \\ U_i(0) \geq 0 \end{array} \quad . \end{array}$$

The third trick in solving the buyer's problem is to realize that these problems amount to dynamic optimization programs where U_i is the state variable and e_i the control variable. We can now apply the Pontryagin principle to find a solution. The Hamiltonian H_i of each program is given by

$$H_i(\alpha_i, e_i, U_i, \mu_i) = \{X_i(\alpha_i)\left[e_i - C_i(e_i, \alpha_i)\right] - U_i\} f(\alpha_i) + \mu_i e_i X_i(\alpha_i).$$

Using the Pontryagin principle, we obtain the first-order condition of the programs:[86]

$$\begin{aligned}\dot{\mu}_i(\alpha_i) &= f(\alpha_i) \\ \mu_i(\alpha_i) &= \left[\frac{\partial C_i(e_i^*(\boldsymbol{\alpha}), \alpha_i)}{\partial e_i} - 1\right] f(\alpha_i) \\ \mu_i(1) &= 0.\end{aligned}$$

Substituting $C_i(e_i, \alpha_i) = \frac{1}{2}e_i^2 + e_i - \alpha_i e_i$ together with some straightforward calculations yields the optimal effort levels:

$$e_i^*(\boldsymbol{\alpha}) = \begin{cases} \alpha_i - \frac{1-F(\alpha_i)}{f(\alpha_i)} & \text{if } \alpha_i \geq \underline{\alpha} \\ 0 & \text{if } \alpha_i < \underline{\alpha} \end{cases}, \tag{27}$$

with $\underline{\alpha}$ the unique solution to $y = \frac{1-F(y)}{f(y)}$ with respect to y. Under these effort levels, the buyer's problem is reduced to

$$\max_{X_i(\cdot)} \sum_i E_{\alpha_i} \left\{ X_i(\alpha_i)\left[e_i^*(\alpha_i) - C_i(e_i^*(\boldsymbol{\alpha}), \alpha_i)\right] - \int_0^{\alpha_i} e_i^*(y) X_i(y) dy \right\},$$

which is equivalent to

$$\max_{X_i(\cdot)} \frac{1}{2} \sum_i E_{\alpha_i} \left\{ X_i(\alpha_i) e_i^*(\alpha_i)^2 \right\}. \tag{28}$$

From (28), it is straightforward to see to which firm i the project should be allocated. The buyer's expected utility is proportional to the sum of the firms' winning probability and the square of the optimal effort. By (3) and (27), $(e_i^*(\alpha_i))^2$ is strictly increasing in α_i for $\alpha_i \geq \underline{\alpha}$. Therefore, it is optimal for the buyer to maximize the winning probability of the firm with the highest α_i, i.e., to always allocate the project to the most efficient firm.

Proof of Lemma 3

[86] The second order conditions can be shown to hold as well.

The proof is by contradiction. Suppose that telling the truth is not a weakly dominant strategy for all bidders. Then for some bidder i there exist a θ_i, $\tilde{\theta}_i$, and $\theta_{-i} \equiv (\theta_1,\theta_{i-1},\theta_{i+1},\theta_n)$ such that

$$v_i(\gamma^*(\tilde{\theta}_i,\theta_{-i}),\theta_i) - t_i(\tilde{\theta}_i,\theta_{-i}) > v_i(\gamma^*(\theta_i,\theta_{-i}),\theta_i) - t_i(\theta_i,\theta_{-i}).$$

Substituting t_i for (7) implies that

$$\sum_j v_j(\gamma^*(\tilde{\theta}_i,\theta_{-i}),\theta_j) > \sum_j v_j(\gamma^*(\theta_i,\theta_{-i}),\theta_j),$$

which contradicts (6). Thus it must be a weakly dominant strategy for each bidder to announce her true type.

7 REFERENCES

Arrow, K.J. and G. Debreu (1954), 'Existence of an Equilibrium for a Competitive Economy,' *Econometrica*, 22, pp. 265-290.

Ashenfelter, O. (1989), 'How Auctions Work for Wine and Art,' *Journal of Economic Perspectives*, 3, pp. 23-36.

Ashenfelter, O. and K. Graddy (2002), Art Auctions: A Survey of Empirical Studies, NBER working paper nr. 8997, Cambridge, MA.

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

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

Ausubel, L.M. and P. Cramton (2004), 'Auctioning Many Divisible Goods,' *Journal of the European Economic Association*, 2, pp. 480-493.

Ausubel, L.M., P. Cramton and P. Milgrom (2005), 'The Clock-Proxy Auction: A Practical Combinatorial Auction Design,' in: P. Cramton, Y. Shoham and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA, MIT Press, forthcoming.

Ausubel, L.M. and P.R. Milgrom (2002), 'Ascending Auctions with Package Bidding,' *Frontiers of Theoretical Economics*, 1(1), Article 1, pp. 1-42.

Baye, M.R., D. Kovenock and C.G. de Vries (1993), 'Rigging the Lobbying Process: An Application of the All-Pay Auction,' *American Economic Review*, 83, pp. 289-294.

Baye, M.R., D. Kovenock and C.G. de Vries (1996), 'The All-Pay Auction with Complete Information,' *Economic Theory*, 8, pp. 291-305.

Berg, G.J. van den, J.C. van Ours and M.P. Pradhan (2001), 'The Declining Price Anomaly in Dutch Rose Auctions,' *American Economic Review*, 91, pp. 1055-1062.

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

Binmore, K. and J. Swierzbinski (2000), 'Treasury Auctions: Uniform or Discriminatory?,' *Review of Economic Design*, 5, pp. 387-410.

Bulow, J.I. and P.D. Klemperer (1999), 'The Generalized War of Attrition,' *American Economic Review*, 89, pp. 175-189.

Bulow, J.I. and J. Roberts (1989), 'The Simple Economics of Optimal Auctions,' *Journal of Political Economy*, 97, pp. 1060-1090.

Cassady, R. (1967), *Auctions and Auctioneering*, Berkeley/Los Angeles, University of California Press.

Che, Y-K. and I.L. Gale (1998a), 'Caps on Political Lobbying,' *American Economic Review*, 88, pp. 643-651.

Che, Y-K. and I.L. Gale (1998b), 'Standard Auctions with Financially Constrained Bidders,' *Review of Economic Studies*, 65, pp. 1-21.

Clarke, E.H. (1971), 'Multipart Pricing of Public Goods,' *Public Choice*, 11, pp. 17-33.

Cramton, P. (1998), 'The Efficiency of the FCC Spectrum Auctions,' *Journal of Law and Economics*, 41, pp. 727-736.

Damme, E. van (1999), 'The Dutch DCS-1800 Auction,' in: F. Patrone, I. García-Jurado and S. Tijs (eds.), *Game Practice: Contributions from Applied Game Theory*, Boston, Kluwer Academic Publishers, pp. 53-73.

Damme, E. van (2002), 'The European UMTS-Auctions,' *European Economic Review*, 46, pp. 846-869.

Das Varma, G. (2002), 'Standard Auctions with Identity Dependent Externalities,' *RAND*

Journal of Economics, 33, pp. 689-708.

Day, R.W. and S. Raghavan (2005), Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions, mimeo, University of Maryland.

Demange, G., D. Gale and M. Sotomayor (1986), 'Multi-Item Auctions,' *Journal of Political Economy*, 94, pp. 863-872.

Demsetz, H. (1968), 'Why Regulate Utilities?,' *Journal of Law and Economics*, 11, pp. 55-66.

Dykstra, M.J. and N. van der Windt (2004), 'Beauty Contest Design,' in: M.C.W. Janssen (ed.), *Auctioning Public Assets: Analysis and Alternatives*. Cambridge, UK, Cambridge University Press.

Engelbrecht-Wiggans, R. and C.M. Kahn (1998a), 'Multi-Unit Auctions with Uniform Prices,' *Economic Theory*, 12, pp. 227-258.

Engelbrecht-Wiggans, R. and C.M. Kahn (1998b), 'Multi-Unit Pay-Your-Bid Auctions with Variable Rewards,' *Games and Economic Behavior*, 23, pp. 25-42.

Englmaier, F., P. Guillén, L. Llorente, S. Onderstal and R. Sausgruber (2004), The Chopstick Auction: A Study of the Exposure Problem in Multi-Unit Auctions, FEEM working paper 10.2004, Milan.

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

Goeree, J.K. and T. Offerman (2003a), Using First-Price Auctions to Sell Heterogeneous Licenses, mimeo, University of Amsterdam.

Goeree, J.K. and T. Offerman (2003b), 'Competitive Bidding in Auctions with Private and Common Values,' *Economic Journal*, 113, pp. 598-614.

Goeree, J.K., E. Maasland, S. Onderstal and J. Turner (2005), 'How (Not) to Raise Money,' *Journal of Political Economy*, forthcoming.

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, pp. 1557-1569.

Groves, T. (1973), 'Incentives in Teams,' *Econometrica*, 41, pp. 617-631.

Haan, M. and L. Toolsema-Veldman (2003), License Auctions when Winning Bids are Financed through Debt, mimeo, University of Groningen.

Harris, M. and A. Raviv (1981), 'Allocation Mechanisms and the Design of Auctions,' *Econometrica*, 49, pp. 1477-1499.

Harstad, R.M. (2000), 'Dominant Strategy Adoption, Efficiency, and Bidders' Experience with Pricing Rules,' *Experimental Economics*, 3, pp. 261-280.

Hoffman, K., D. Menon, S. van den Heever and T. Wilson (2005), 'Observations and Near-Direct Implementations of the Ascending Proxy Auction,' in: P. Cramton, Y. Shoham and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA, MIT Press, forthcoming.

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

Janssen, M.C.W. and V.A. Karamychev (2005), Auctions, Market Prices and the Risk Attitude Effect, Tinbergen Institute discussion paper 2005-025/1.

Janssen, M.C.W. and E. Maasland (2002), Second Opinion Onderzoek naar de Veiling van Telefoonnummers (Second Opinion Research on the Auction of Telephone Numbers), consultancy report, SEOR-ECRi, available at www.ecri.nl.

Jehiel, P. and B. Moldovanu (2001), The European UMTS/IMT-2000 License Auctions, CEPR discussion paper 2810.

Jehiel, P., B. Moldovanu and E. Stacchetti (1999), 'Multidimensional Mechanism Design for Auctions with Externalities,' *Journal of Economic Theory*, 85, pp. 258-294.

Julien, B., J. Kennes and I. King (2000), 'Bidding for Labor,' *Review of Economic Dynamics*, 3, pp. 619-649.

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, pp. 501-585.

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

Kagel, J. H. and D. Levin (1993), 'Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders,' *The Economic Journal*, 103, pp. 868-879.

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

Klemperer, P. (1999), 'Auction Theory: A Guide to the Literature,' *Journal of Economic Surveys*, 13, pp. 227-286.

Klemperer, P. (2002), 'What Really Matters in Auction Design,' *Journal of Economic Perspectives*, 16, pp. 169-190.

Klemperer, P. (2003), 'Why Every Economist Should Learn Some Auction Theory,' in: M. Dewatripont, L. Hansen and S. Turnovksy (eds.), *Advances in Economics and Econometrics: Invited Lectures to Eighth World Congress of the Econometric Society (2000)*, Cambridge, UK, Cambridge University Press, pp. 25-55.

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

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

Laffont, J.-J. (1997), 'Game Theory and Empirical Economics: The Case of Auction Data,' *European Economic Review*, 41, pp. 1-35.

Laffont, J.-J. and J. Robert (1996), 'Optimal Auction with Financially Constrained Buyers,' *Economics Letters*, 52, pp. 181-186.

Laffont, J.-J. and J. Tirole (1987), 'Auctioning Incentive Contracts,' *Journal of Political Economy*, 95, pp. 921-937.

Laffont, J.-J. and J. Tirole (1993), *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA, MIT Press.

Leonard, H. (1983), 'Elicitation of Honest Preferences for the Assignment of Individuals to Positions,' *Journal of Political Economy*, 91, pp. 461-479.

Littlechild, S. (2002), 'Competitive Bidding for a Long-term Electricity Distribution Contract,' *Review of Network Economics*, 1, pp. 1-39.

Lusht, K. (1994), 'Order and Price in a Sequential Auction,' *Journal of Real Estate Finance and Economics*, 8, pp. 259-266.

Maasland, E. and B. Moldovanu (2004), 'An Analysis of the European 3G Licensing Process,' in: M.C.W. Janssen (ed.), *Auctioning Public Assets: Analysis and Alternatives*. Cambridge, UK, Cambridge University Press.

Maasland, E. and S. Onderstal (2005), Auctions with Financial Externalities, mimeo, Erasmus University Rotterdam and University of Amsterdam.

Maskin, E.S. and J.G. Riley (1984), 'Optimal Auctions with Risk Averse Buyers,' *Econometrica*, 52, pp. 1473-1518.

Maskin, E.S. and J.G. Riley (1989), 'Optimal Multi-Unit Auctions,' in: F. Hahn (ed.), *The Economics of Missing Markets, Information, and Games*, New York, Oxford University Press, pp. 312-335.

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

Matthews, S. (1983), 'Selling to Risk Averse Buyers with Unobservable Tastes,' *Journal of Economic Theory*, 30, pp. 370-400.

Maynard Smith, J. (1974), 'The Theory of Games and the Evolution of Animal Conflicts,' *Journal of Theoretical Biology*, 47, pp. 209-221.

McAfee, R.P. and J. McMillan (1986), 'Bidding for Contracts,' *RAND Journal of Economics*, 17, pp. 326-338.

McAfee, R.P. and J. McMillan (1987a), 'Auctions and Bidding,' *Journal of Economic Literature*, 25, pp. 699-738.

McAfee, R.P. and J. McMillan (1987b), 'Competition for Agency Contracts,' *RAND Journal of Economics*, 18, pp. 296-307.

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

Menezes, F. (2003), 'An Auction Theoretical Approach to Fiscal Wars,' *Social Choice and Welfare*, 20, pp. 155-166.

Milgrom, P.R. (2004), *Putting Auction Theory to Work*, Cambridge, UK, Cambridge University Press.

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

Mirrlees, J.A. (1971), 'An Exploration in the Theory of Optimal Income Taxation,' *Review of Economic Studies*, 38, pp. 175-208.

Mirrlees, J.A. (1975), The Theory of Moral Hazard and Unobservable Behaviour: Part I, mimeo, Oxford University.

Mirrlees, J.A. (1976), 'The Optimal Structure of Incentives and Authority within an Organization,' *Bell Journal of Economics*, 7, pp. 105-131.

Mirrlees, J.A. (1999), 'The Theory of Moral Hazard and Unobservable Behaviour: Part I,' *Review of Economic Studies*, 66, pp. 3-21.

Moldovanu, B. and A. Sela (2001), 'The Optimal Allocation of Prizes in Contests,' *American Economic Review*, 91, pp. 542-558.

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

Noussair, C. (1995), 'Equilibria in a Multi-Object Uniform Price Sealed Bid Auction with Multi-Unit Demands,' *Economic Theory*, 5, pp. 337-351.

OECD (2001), 'Innovations in Labour Market Policies: The Australian Way,' Policy report, Paris, France.

Offerman, T. and J. Potters (2003), Does Auctioning of Entry Licenses Induce Collusion? An Experimental Study, mimeo, University of Amsterdam and Tilburg University.

Onderstal, S. (2002a), Socially Optimal Mechanisms, CentER discussion paper nr. 34, Tilburg University.

Onderstal, S. (2002b), The Chopstick Auction, CentER discussion paper nr. 35, Tilburg University.

Onderstal, S. (2005), Bidding for the Unemployed: An Application of Mechanism Design to Welfare-to-Work Programs, mimeo, University of Amsterdam.

Pezanis-Christou, P. (2001), Sequential Auctions with Asymmetric Buyers: Evidence from a Fish Market, IAE-UAB working paper nr. 558.03, Bonn University.

Pipes, D. (1998-1999), 'Herodotus: Father of History, Father of Lies,' *Loyola University New Orleans: The Student Historical Journal*, 30.

Posner, R.A. (1975), 'The Social Costs of Monopoly and Regulation,' *Journal of Political Economy*, 83, pp. 807-828.

Riley, J.G. and W.F. Samuelson (1981), 'Optimal Auctions,' *American Economic Review*, 71, pp. 381-392.

Robinson, M. (1985), 'Collusion and the Choice of Auction,' *RAND Journal of Economics*, 16, pp. 141-145.

Rothkopf, M.H., T.J. Teisberg and E.P. Kahn (1990), 'Why Are Vickrey Auctions Rare?,' *Journal of Political Economy*, 98, pp. 94-109.

SEO and TNO (2004), 'Buitenlandse Ervaringen met de Aanbesteding van Reïntegratiediensten' ('Foreign Experiences with Tendering Reintegration Services'), consultancy report, Amsterdam, the Netherlands.

Shubik, M. (1983), 'Auctions, Bidding, and Markets: An Historic Sketch,' in: R. Engelbrecht-Wiggans, M. Shubik, and J. Stark (eds.), *Auctions, Bidding, and Contracting*, New York, New York University Press, pp. 33-52.

Staatscourant (2004), 'Regeling Veilingprocedure en Lotingprocedure Nummers,' ('Ruling on the Rules for Conduct of the Auction and Lottery for the Award of Phone Numbers') *Staatscourant,* 172, 8 September 2004, p. 21.

Swinkels, J. (1999), 'Asymptotic Efficiency for Discriminatory Private Value Auctions,' *Review of Economic Studies*, 66, pp. 509-528.

Tullock, G. (1967), 'The Welfare Costs of Tariffs, Monopolies, and Theft,' *Western Economic Journal*, 5, pp. 224–232.

Tullock, G. (1980), 'Efficient Rent-Seeking,' in: J.M. Buchanan, R.D. Tollison and G. Tullock (eds.), *Towards a Theory of the Rent-Seeking Society*, College Station, Texas A&M University Press, pp. 97-112.

Verberne, M.L. (2000), *Verdeling van het Spectrum* (*Allocation of Spectrum*), Ph.D. thesis, University of Amsterdam.

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

Vickrey, W. (1962), 'Auction and Bidding Games,' in: *Recent Advances in Game Theory*, Princeton University Conference, pp. 15-29.

Vries, S. de and R.V. Vohra (2003), 'Combinatorial Auctions: A Survey,' *INFORMS Journal on Computing*, 15(3), 284-309.

Weber, R.J. (1983), 'Multiple-Object Auctions,' in: R. Engelbrecht-Wiggans, M. Shubik, and R.M. Stark (eds.), *Auctions, Bidding, and Contracting: Uses and Theory*, New York, New York University Press, pp. 165-191.

Weber, R.J. (1997), 'Making More from Less: Strategic Demand Reduction in the FCC Spectrum Auctions,' *Journal of Economics and Management Strategy*, 6, pp. 529-548.

Wurman, P.R., J. Zhong, and G. Cai (2005), 'Computing Price Trajectories in Combinatorial Auctions with Proxy Bidding,' *Electronic Commerce Research and Applications*, forthcoming.

Zheng, C. (2001), 'High Bids and Broke Winners,' *Journal of Economic Theory*, 100, pp. 129-171.

Zulehner, C. (2005), 'Strategic Behavior of Asymmetric Bidders in Sequential Auctions: Evidence from Austrian Cattle Auctions', mimeo, University of Vienna.