TI 2007-041/3
Tinbergen Institute Discussion Paper

If Winning isn't Everything, Why do they keep Score? A Structural Empirical Analysis of Dutch Flower Auctions

Gerard J. van den Berg[1,2,3]
Bas van der Klaauw[1]

[1] Free University Amsterdam, Tinbergen Institute, and CEPR;
[2] IFS London;
[3] IZA.

Tinbergen Institute
The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at http://www.tinbergen.nl.

If Winning Isn't Everything, Why Do They Keep Score? A Structural Empirical Analysis of Dutch Flower Auctions

Gerard J. van den Berg *

Bas van der Klaauw †

May 14, 2007

Abstract

This paper provides a structural empirical analysis of Dutch auctions of houseplants at the flower auction in Aalsmeer, the Netherlands. The data set is unique for Dutch auctions in the sense that it includes observations of all losing bids in an interval adjacent to the winning bid. The size of this interval is determined by the speed of reaction of the auction participants, and as such these data are collectible due to neurological constraints on information processing. The data on losing bids are shown to be informative on the structural model determinants. The models are estimated using the Gibbs sampler with data augmentation. We take account of data limitations concerning the number of bidders. The estimation results are used to investigate whether actual reserve prices are optimal, and to determine the effects of reserve price changes.

*Free University Amsterdam, IFS London, IZA, and CEPR.

†Free University Amsterdam, Tinbergen Institute and CEPR.

Address: Department of Economics, Free University Amsterdam, De Boelelaan 1105, NL–1081 HV Amsterdam, The Netherlands.

Keywords: first-price auction, private value, speed of reaction, observing losing bids, data augmentation, Gibbs sampling, Markov Chain Monte Carlo, reserve price.

We are grateful to the Aalsmeer Flower Auction, and to Dirk Hogervorst, Chiel Post, Erik van 't Boveneind and André van Kruijssen in particular, for kindly providing the data and for useful discussions. Also thanks to Patrick Bajari, Douwe-Frits Broens, Jean-Pierre Florens, Harry Paarsch, Isabelle Perrigne, Martin Pesendorfer, Samita Sareen, Jos Stins, Eric van Damme, and participants in a conference at Yale University and in seminars at LSE and the University of Toulouse, for useful comments and suggestions.

1 Introduction

The Aalsmeer Flower Auction (AFA) is the largest flower auction and the largest commercial edifice in the world (Guinness World Records, 2001). The annual supply exceeds 4 billion flowers and 500 million plants, and the annual turnover of auctioned products amounts to 1.5 billion Euro. Daily, around 19 million flowers and 2 million plants are auctioned, resulting in around 50,000 daily transactions.[1] AFA is located at the city of Aalsmeer, the Netherlands, close to Amsterdam and Schiphol Airport. It is also the largest trade building in the world, with a surface size of 150 football fields and an average daily number of 10,000 working individuals, many of whom use a bicycle to get around.[2] AFA uses the Dutch auction mechanism, i.e., the price falls from an initial high price until a bidder stops the auction. This bidder then obtains the object for the price at which he stopped the auction.

We perform a full structural empirical analysis of auctions of potted houseplants at AFA. Structural empirical analysis aims at the estimation of the underlying distribution of valuations of the bidders. This in turn enables the evaluation of the imposition of counterfactual policies, notably concerning the minimum price (or reserve price or reservation price) set by AFA. The data are taken from electronically collected databases of AFA and cover all auctions of houseplants during 14 days at the end of August and the beginning of September, 1996.

In the paper we argue that the independent private value (IPV) paradigm is accurate for the AFA auctions, and we restrict attention to IPV models.[3] By now, there is an established literature on the structural estimation of IPV models (see for example Hendricks and Paarsch, 1995, Laffont and Vuong, 1996, and Laffont, 1997, for surveys). The structural empirical analysis of Dutch auctions suffers from a number of potential problems. First and most important, in general only a single bid is observed per auction. Obviously, observations of winning bids contain less information than observations of all bids (which may be available in a sealed-bid auction). Intuitively, it is difficult to obtain a reliable estimate of the whole distribution of private values from only winning bids, particularly when the

[1] See Van den Berg, Van Ours and Pradhan (2001), Kambil and Van Heck (1996), and Katok and Roth (2004), for descriptions of AFA in English.

[2] The magnitude of AFA reflects the importance of flowers and plants for the Dutch economy. The Netherlands is the leading producer and exporter of flowers in the world, and flowers are its most important export product. Around half of the products at AFA are imported and more than half are exported.

[3] See for example McAfee and McMillan (1987a) and Wolfstetter (1996) for surveys of auction theory.

number of bidders is large. The estimation results and policy evaluations may then be sensitive to the assumed functional form of the distribution of valuations.

We are in a unique position to deal with this, because our database does not only register the price at which the auctioning clock is stopped, but also all other bids of bidders who pushed the button to stop the clock up to one second after the moment of the winning bid. Such bidders are unable to inhibit their planned response within a very short time frame after the moment that the auction is over. The data on losing bids are therefore collectible due to neurological constraints in the speed of human information processing. As such, these data exploit that economic response behavior takes real time.

The data on the losing bids consist of observations of all bids in an interval adjacent to the winning bid. This constitutes an information gain compared to other studies of Dutch auctions. Of course, some bidders with a bid laying in the one-second time-frame are able to refrain from pushing because they realize on time that the auction is over. We can therefore only use the losing bids made within a fraction of a second after the winning bid. To quantify the size of the interval, we rely on relevant results in the psychological and behavioral science literature on stop signal reaction times (see e.g. Logan and Cowan, 1984, and the overview in Rieger, 2000). We infer that 0.2 seconds is an accurate and conservative approximation of the speed of reaction of bidders. It is important that this number does not exceed the smallest reaction time among the bidders, because otherwise the observed density of bids close to 0.2 seconds after the winning bid underestimates the true density at that point. The data provide over-identifying information on this. We use this to perform Hausman-type tests and we estimate models with upper bounds smaller than 0.2 seconds.

A second potential problem with structural empirical analysis of Dutch auctions arises with Maximum Likelihood (ML) estimation. As is well-known (e.g. see Donald and Paarsch, 1993, and Laffont, Ossard and Vuong, 1995), in most specifications of first-price auction models the support of the distribution of bids depends on all structural parameters. This has two implications for ML estimation, which are particularly problematic when the parameters of the distribution of valuations are expressed as functions of observed explanatory variables or covariates. First, finding the ML estimator involves maximizing a nonlinear function subject to nonlinear constraints, with heavy computational costs. This is exacerbated by the fact that the mapping from valuations to bids cannot be inverted analytically and depends on an integral that in general cannot be solved explicitly. Secondly, ML estimation is non-standard in the sense that standard asymptotic properties do not apply.

We deal with this by using recently developed Markov Chain Monte Carlo (MCMC) estimation methods. In particular, we use the Gibbs sampler with data augmentation to evaluate the marginal posterior distributions of the parameters (see Casella and George, 1992, and Gelfand and Smith, 1990). This estimation method is a special case of the novel collection of estimation methods called Markov-Chain Monte Carlo (MCMC) estimation methods. Because inference on auctions is simple and straightforward in case the private values of all bidders are observed, we use data-augmentation methods to construct these (latent) private values (see Tanner and Wong, 1987).[4] Estimation with the Gibbs sampler with data augmentation does not involve (numerical) optimization routines or (numerical) differentiation of objective functions. Instead, it only requires simulations of random variables. Because of the Bayesian nature of the estimation method, asymptotic rates of convergence are not an issue. Instead, we focus on posterior distributions based on conventional uninformative priors.[5]

We use the estimation results to investigate the effects of increases of the reserve prices on bids and on the expected revenue of the seller. The actual reserve price for houseplants does not vary across plants and is often less than 20% of the realized price. In fact, the price rarely falls below the reserve price. This may suggest that it is profitable to increase the reserve price and to differentiate it across plants.

The empirical analysis needs to address a number of complicating factors. First, we need to quantify the number of bidders. The problems we encounter here are similar to those in other applications with Dutch auctions, and the solutions we develop have therefore wider relevance. The data provide the number of individuals in an auction room who are logged on to a certain auction. Potential participants whose valuations are below the reserve price may not show up at the auction and are therefore not counted as a potential bidders. We therefore restrict attention to the sub-population of individuals whose valuation exceeds the reserve price. This implies that we can only examine counterfactual policies

[4]Sareen (1998) uses a similar approach to analyze a first-price auction. Her method differs from the method used in this paper, and we discuss the differences in more detail in Subsection 3.4. See also Bajari and Hortacsu (2003) for an application to internet auctions.

[5]Laffont, Ossard and Vuong (1995) developed and applied a simulated non-linear least squares estimation method for IPV models for Dutch auctions. This allows them to exploit standard asymptotic results and to avoid heavy computational burdens. However, the estimator focuses on data on the mean of the distribution of the winning bid, and as such it only uses limited information on observed bids. The method is unfeasible for estimation with our data on losing bids, since the number of observed losing bids is dispersed over the auctions in the data.

that result in a shrinking of the set of bidders, and not policies that result in expansion of this set. We demonstrate that the restriction to this sub-population is easily implementable and does not entail a loss of relevance.

Another problem with quantifying the number of bidders is that some of the individuals who are logged on may actually have a valuation below the reserve price and may be logged on just because they participate at the preceding and subsequent auction in the same room and do not bother to log off in between. We deal with this problem in two ways. First, we restrict attention to auctions at the beginning of the morning (or rather late at night), when the auction house opens and most participants arrive. Secondly, we estimate a model assuming that the true number of participating bidders is unknown to us (although known to the bidders). We demonstrate that the observations on losing bids are particularly informative on the distribution of the true number of bidders. From these data it follows that the number of individuals who are logged in (or, for that sake, the number of individuals seated in the auction rooms) grossly over-estimates the actual number of bidders, and as a result the actual auctioning market is much less close to a competitive market than may seem at first hand.

The second complicating factor concerns the sequential nature of certain auctions. At AFA, often, a certain number of homogeneous plants is auctioned sequentially with a buyer's option, i.e., the winner decides how many units he wants to buy in excess of the stipulated minimum purchase quantity, and if the desired number falls short of the total available number, the remaining units are auctioned in the same way. Theoretical models of sequential auctions are relatively complex, and therefore we only use data of auctions in which the minimum purchase quantity equals the total number of units available.

The outline of this paper is as follows. Section 2 provides a description of AFA and gives an overview of the data. We present some summary statistics and perform reduced-form analyses of the prices at which the plants are sold. Section 3 discusses our auction model, its parameterization, and our corresponding particular implementation of the Gibbs sampler with data augmentation. The estimation results of the structural analysis are presented in Section 4. This section also evaluates the reserve price of the seller. Subsection 4.3 presents results under alternative assumptions concerning the number of bidding participants. Section 5 concludes.

2 Institutions and data

2.1 The Aalsmeer Flower Auction

This subsection describes the organizational structure of AFA and the actual auctioning process. Next, in Subsection 2.2, we give an overview of the data we use in the empirical analysis, and we perform some reduced-form analyses of the winning bids. In Subsection 2.3 we discuss the losing bids that we observe.

AFA is a cooperative owned by about 3500 Dutch growers of the auctioned products.[6] The total number of growers participating in the auctions equals about 7100, of which almost 1500 are from abroad. The total number of buyers equals about 1500. The dispersion of their shares in total turnover is enormous. On the one hand, about 50 buyers each buy for more than 5 million Euro per year; together this amounts to about 50% of total turnover. On the other hand, about 725 buyers each buy for less than 50.000 Euro per year; together this amounts to about 1% of total turnover. These two extremes basically correspond to big exporting companies and small domestic retail shops, respectively.

There are 13 auctioning clocks in five separate auction rooms. These clocks are often used at the same time, so that simultaneous auctions take place within a room. A given individual can only participate at one auction, but a given buyer may of course delegate more than one individual to an auction room. The number of seats in an auction room is about 500. The average duration of a single auction (i.e., one transaction) equals between 3 and 4 seconds.

The products are auctioned as separate "lots", which are defined as the total supply of a given homogeneous product of a given grower on a given day. An "auction group" denotes a group of products with similar features. The daily ordering of auction groups over time is the same on every day. However, the daily ordering over time of different lots within a given auction group is randomized every day.

As explained in the introduction, we only consider non-sequential auctions, so our description below of the auctioning process ignores topics that are only relevant for sequential auctioning. The wall in front of the auctioning room contains a large board with a clock and an electronic display of the identity of the grower, name of the product, various quality indicators, size of the flower pot in case of plants, and monetary unit. The flowers or plants are transported through the room, and an employee takes a few items from the carriage to show them to the buyers (buyers also have the opportunity to closely examine the

[6] See the Annual Reports of AFA, e.g. Bloemenveiling Aalsmeer (1996a).

flowers some time before the actual auctioning). The auctioneer decides on a starting position for the clock which corresponds to an unreasonably high price for the product. He then sets the clock in motion. The value pointed at by the clock drops continuously until a buyer stops the clock by pushing the button in front of him.[7] The value pointed at by the clock at that moment is the price to be paid by that buyer for a single item. The identity of the buyer is then shown on the electronic display in front of the room. If the hand of the clock passes the minimum price then the lot is destroyed.

In 1996 (the year covered by our data), the reserve prices were fixed in co-operation with all (seven) major flower auction houses in the Netherlands. For a given product, the reserve price was the same at every auction house and was constant throughout the year. Variation across products was only very modest. The reserve prices are published in the publicly available annual AFA Codebook (see e.g. Bloemenveiling Aalsmeer, 1996b). For all the houseplants we consider in this paper, the reserve price in 1996 was 25 cents per single plant.

2.2 The data set

Our database describes auctions taking place in one of the auctioning halls of the AFA. This particular hall has four auctioning clocks on which the auctions of houseplants take place. The data consist of three different registers that were collected during the end of August and the beginning of September 1996. The first data register concerns transactions. These contain the moment and auctioning clock at which the auction takes place, the type of houseplants, a code for the grower, the price paid at the auction and a code for the winning bidder. The second data register concerns the registered bidders. This describes for each point in time the number of bidders that is registered at each of the four auctioning clocks. We do not observe the identities of the registered bidders. Combining the data on the transactions with the data on the number of bidders provides us with data on auctions during 14 different days. The third and last register concerns observed losing bids. We return to this in the next subsection.

Figure 1 shows, for auctions taking place during the first hour of the auctioning process, the number of bidders registered. Clearly, the average number of bidders increases relatively fast during the first 15 minutes and after that the increase slows down. Just before the auctioning starts, which is 6.30 a.m., bidders

[7] Actually, the clock is designed as a circle of small lamps each corresponding to a given monetary value, such that a clockwise movement corresponds to a decrease of this value. If the clock is set in motion then consecutive lamps light up sequentially.

start entering the auctioning hall and register as being a participant in the auctions taking place on one of the auctioning clocks in this hall. At any moment, a bidder can switch between the auctioning clocks, but a bidder is only allowed to participate at one auctioning clock at the time. When the bidder stops participating in auctions, he has to sign off as being a participant. However, some bidders do not sign off and remain registered as being participant. The registered number of bidders than exceeds the actual number of bidders. Note that this difference can only increase over the day, and that imperfect monitoring of the number of bidders may cause biases in the empirical analyses. To have a reliable measure of the number of participants in the auction, we restrict the data to auctions taking place during the first half hour after the auctioning starts. This reduces the data to 6633 lots that are auctioned.

We select sequences of observations of auctions of indivisible goods which are as homogenous as possible. In the ideal case, we would observe a particular type of houseplant which always has the same quality and is only supplied by a single grower. In practice, to achieve the objective of having data on indivisible goods, we restrict attention to lots for which the minimum purchase quantity equals the actual number of units in the lot. This implies that the "first" winner of the auction has to buy all units in the lot. By considering only these lots, we avoid that bidders anticipate on a sequence of auctions to sell the complete lot and there is no buyer's option. This reduces the data set to 2809 lots. This data set contains 274 different types of houseplants. Only 58 types of houseplants appear more than 10 times, and only 36 types of houseplants more than 20 times.

Some types of houseplants are supplied by many different growers. As will become clear below, the identity of the grower is a very important characteristic of a houseplant. Therefore, we select three types of houseplants that are supplied by a limited number of "large" growers. We thus create three subsamples, each containing the auctions of one of the types. The types of houseplants are: a single type of Begonia supplied in a mixture of three different colors and two particular types of Dieffenbachias, the "Camilla" and the "Compacta".

Some characteristics of the three subsamples are presented in Table 1. The prices are measured in Dutch Guilder cents per unit and the reserve prices for all three types are 25 cents per unit (we do not observe auctions where the price falls to this level and lots are destroyed). All lots of houseplants in all three subsamples are of the highest quality code. The subsample of Begonias is the largest with 178 observations, while the subsamples of Camillas contains 47 observations and the subsample of Compactas 34 observations. Auctions of the Begonias are observed on each day at which data is collected. Auctions of Camillas and Compactas are only observed during 11 and 9 days respectively. The Begonias and Compactas are both supplied by 3 growers, the Camillas only by 2 growers. Opposite to the subsample Begonias, the identity of the grower seems to have a large impact on the price per unit in both other subsamples. The standard deviation of the prices is large and the there is a large difference between the maximum and the minimum price observed. For both the Camillas and the Compactas we observe that the lots with the highest price per unit are all supplied by one grower. We return to this issue below. The average number of bidders lies around 50 for all subsamples. For the subsample of Camillas we do not observe any auctions taking place on Friday, for both other subsamples we observe auctions during any day of the week.

To get some insights into the set of covariates that determine the price of the winning bid at an auction, we perform some reduced-form analyses. We regress the logarithm of the price on the number of bidders participating in the auction, the identity of the grower and the day of the week at which the auction takes place. One may argue that the number of bidders participating is an endogenous regressor, as some lots of houseplants may have unobserved characteristics that bidders actually like. Lots that attract many bidders may thus not be sold for an higher price because there are many bidders present, but also because these lots may have some attractive unobserved characteristics. We use instrumental variable estimation to correct for possible endogeneity of the number of bidders.

First, consider again the sample of 2809 auctions that have the minimum purchase quantity equal to the number of units in the lot. As mentioned in the previous subsection, within the auctioning groups, the sequence at which different lots are auctioned is randomized early in the morning just before the auctioning starts. As can be seen from Figure 1 during the first 15 minutes after the actual auctioning started the number of bidders participating is increasing relatively fast. This suggests that during these first 15 minutes the time since the start of the auctioning has some effect on the number of bidders present, but due to the randomization in the order of appearance time does not have a direct effect on prices. Therefore, we use time since the start of the auctioning as an instrumental variable for the number of bidders. Since the increase in the number of bidders only occurs during the first 15 minutes, we restrict the sample to 1367 auctions that took place within the first 15 minutes. Furthermore, we exclude auctions of lots of which the type of the houseplants or the grower only appears once in the data. This restricts the subsample to 1257 auctions. In the regression we include 4 dummy variables for the day of the week (Tuesday until Friday), 109 dummy variables for the types of houseplants and 114 dummy variables for growers. Time has a significant effect on the number of bidders. The estimated coefficient equals 1.53 (the estimated standard error is 0.10), indicating that on average each minute around 1.5 additional bidders register to participate in auctions. The instrumental variable estimate of the number of bidders on the price equals 0.0035, with an standard error of 0.0020. The OLS estimate equals 0.0038, with an standard error of 0.0008. A Hausman test cannot reject the null hypothesis that the number of bidders is exogenous. In fact, the p-value of such a test equals 0.83, which is not even close to the usual significance levels. This provides rather strong evidence that the number of bidders is actually exogenous and that the price paid at the auction significantly increases with the number of bidders participating. This latter confirms the usual auction theory that the level of competition is higher if more bidders are present, which increases expected revenue. Furthermore, the type of houseplant and the grower are very relevant for the price. The coefficients for the different days of the week indicate that prices are constant from Monday until Thursday and are significantly lower on Friday.

Next, we use OLS to regress for the subsamples of Begonias, Camillas and Compactas, the logarithm of the price paid at an auction on the number of bidders participating, dummy variables for the growers and for the days of the week. The estimation results are presented in Table 2. The estimated covariate effect of the number of bidders is positive for the Begonias and the Camillas, but negative for the Compactas. The identity of the grower seems to be the most important covariate for the price paid at the auction. Recall that the houseplants in all lots are classified by the AFA as of the highest quality. This suggests that the reputation of the grower is very important at the AFA. The difference in prices are particularly striking for the Camillas and Compactas. Both growers supplying Camillas also supply Compactas. It is clear that Grower 4 generally has a better reputation than Grower 5. Even though the type of houseplants are very strictly defined there seem to be differences between similar houseplants supplied by different growers. Finally, the effect of the day of the week at which the auction takes place on the price is ambiguous. Whereas for Camillas the average prices are on Tuesday around 15% higher than on Wednesday, the opposite is true for Compactas.

2.3 The losing bids

The third data register contains bids that are made by losers of the auction. Such data have never been used before in the analysis of Dutch auctions. In the period covered by our data, the AFA recorded all bids made by the losers of an auction up until one second after the winner stopped the auctioning clock. Of the 2809 lots that have the minimum purchase quantity equal to the number of units in the lot, this information was collected in 803 cases. In the remainder of this subsection we focus on this subsample of 803 auctions.

It is important to stress that the auctioning clock stops at the moment the highest bidder pushes the button to reveal his bid. Once the other bidders note this, they know that bidding is useless. However, there is no penalty for pushing the button after the clock stopped. Hence, losing bidders may just push out of frustration or for fun. The is confirmed by the data, which show some losing bid far after the clock would have reached the reserve price. On the other hand, it takes some time before a bidder realizes that the auction actually stopped. The losing bids made in a very short time interval after the winning bid are thus most likely real bids.

In the subsample of 803 auctions that contain information on losing bids, in 92% of the auctions at least one losing bid is observed. This implies that in 8% of the auctions no losing bids was made within one second after the winning bid was registered. On average we observe 4 losing bids per auction, while the maximum number of observed losing bids in a single auction is 21.

In Figure 2 we have plotted a histogram of the time between the winning bid in an auction and all of the observed losing bids in the auction. In Figure 3 we have plotted univariate kernel estimates of the density of observed losing bids for some different values of the bandwidth (h). The observations are time intervals between two events and can therefore only take positive values. We use a reflection method to impose a boundary condition on the kernel density close to 0.[8] Note that both in the histogram and the kernel densities we observe a relatively large drop around 0.2 seconds.

There is a psychological literature on the time it takes individuals to inhibit performing some planned response or action (see for example Logan and Cowan, 1984, and the overview in Rieger, 2000). This is measured by so-called stop signal reaction times. The experiments in the literature do not consider the exact design of Dutch auctions, but some designs are quite similar, and it seems that they are

[8]The reflection methods implies that for every observation x_i an extra observation $-x_i$ is added. Standard methods can be used to compute the kernel density of the doubled sample. We use a normal density function. The resulting kernel density is truncated at 0 and multiplied by 2 for the positive values to ensure that the density integrates to 1. The implicit boundary condition imposed is that the right-derivative at 0 is equal to 0 (due to symmetry caused by the reflection). Other methods, like for example transformation to logarithms or truncation at 0 make similar type of arbitrary assumptions (see for an overview Silverman, 1986 p. 30–31).

sufficiently informative for our purposes. Most experiments indicate that the time needed to inhibit a planned action, like pushing a button at a planned time, lies in the interval from 0.20 to 0.25 seconds. One may claim that bidders at the AFA are well-trained, but the ability to inhibit actions does not seem to vary much across individuals or conditions (see the above literature). We therefore choose the cut-off time for what we consider as real losing bids at 0.20 seconds.

In fact, we can use the actual data to obtain some insight into the extent to which 0.20 is a reasonable choice. Define

$$y_k = \frac{b_{(1)} - b_{(k)}}{b_{(1)}} \qquad k = 2, \ldots$$

which denotes the difference between the winning bid and k^{th} highest bid as a fraction of the winning bid. Clearly, y_k is only observed if $b_{(1)} - b_{(k)}$ is less than some threshold $\delta = \kappa\bar{\tau}$ which depends on the bidders' speed of reaction $\bar{\tau}$ and the known speed of the auctioning clock κ. For a given choice of τ, we can estimate a censored regression model for a latent variable y_k^*,

$$\log(y_k^*) = x'\beta_k + \varepsilon_k$$

where ε_k is normally distributed with mean 0 and variance σ_k^2. The expected value of y_k equals

$$\mu_k = \exp\left(x'\beta_k + \frac{1}{2}\sigma_k^2\right)$$

Obviously, if we choose τ small, many observations for y_k become unobserved. But if y_k is observed, we can be sure that we observe the true value. A small τ thus yield consistent estimators for β_k and μ_k. However, the estimators are not as efficient as for larger τ, as long as the choice for τ is smaller than the true reaction time τ^*. Actually, if τ exceeds τ^*, the estimators may become inconsistent. We can use this to construct a Hausman type test in which we compare the estimator for μ_k for a small τ with the estimator for μ_k for a large τ. If the test rejects similarity of both estimators, the large τ obviously exceeded the true speed of reaction.

In the estimation, we can only use auctions which potentially have observed losing bids. For Begonias we only include 64 auctions in the estimation, for Camillas 29 auctions and 30 auctions of Compactas. In Table 3 we provide the results for $k = 2$. We take $\tau = 0.05$ as the initial small value for τ. However, for most auctions we already observe for this small value for τ the second highest bid. This implies that we cannot learn much about the bidders' reaction speed. However, μ_2 reflects the fraction that winning bidder could have lowered his bid without losing the auction, which can be interpreted as the a measure for information asymmetry in the auction. The estimates show that for auctions of Camillas the highest bidder could have lowered his bid with around 7% without losing the auction. The markets for Begonias and Compactas are slightly more competitive. The highest bidder in a auction of Begonias could have lowered his bid with almost 5% and in an auction of Compactas with slightly more than 5.5%. We repeated the empirical analyses for $k = 4$, which is reported in Table 4. The empirical analyses do not learn us much about the reaction speed of the bidders. The estimates values for μ are relatively constant for the different values of τ and the p-value for the test is never close to usual significance levels.

These results imply that the choice for the reaction speed is not very important and we could take 0.3 seconds as well. However, since the earlier mentioned psychological literature indicates that this might be too large, we remain with our choice of 0.2 seconds as the maximum allowed speed of reaction. This means that we consider bids made within 0.2 seconds after the winner as actual losing bids and we ignore all other losing bids. After this restriction, we observed for Begonias for 58 auctions actually losing bids, which is around 36% of the sample. The subsample of Begonias contains 64 auctions that could potentially contain information about losing bids, implying that in 6 auctions no losing bids were made within 0.2 seconds after the winning bid. On average 5.6 losing bids were made within these 0.2 seconds and the maximum number of losing bids in a single auction is 13. For the Camillas and Compactas we observe respectively 29 and 30 auctions that could contain information on losing bids and 26 and 28 auctions respectively that actually contain losing bids. For the Camillas we observe on average 4.1 losing bid and the maximum number of losing bids in a single auction equals 11, while for Compactas the average is 3.2 and the maximum 9.

For the subsample of Begonias the speed of the auctioning clock is such that in 0.2 seconds the price drops 60 cents. On average the winning price is 202.25 cents, but the lowest winning price observed in the subsample is 130 cents. For this auction we thus observe all bids above 70 cents. For Camilla and Compactas the unit on the auctioning clock is not always the same. However, there are auctions for which we observe all bids higher than 35 cents.

3 The economic model and the econometric approach

3.1 The theoretical model for the Dutch flower auction of houseplants

In this section we present the theoretical model framework and we argue why it is appropriate for our auction setting. Furthermore, we discuss the identification, specification, and estimation method for the structural analyses.

The assignment rule of the Dutch auction is simple: the highest bidder obtains the object and pays his bid. Behavior of bidders is assumed to be governed by the standard symmetric IPV model in which both the seller and the potential buyers are risk neutral, and in which payments only depend on bids. This framework is covered by a number of lucid surveys, including McAfee and McMillan (1987a), Matthews (1995), and Wolfstetter (1996), so the present exposition can be very brief.

Consider the auction of a single indivisible object. Each bidder knows his valuation of the object, and knows that this valuation is strictly private. Moreover, each bidder knows that the values of the other bidders are random draws from a common known distribution function. The seller organizes the auction to maximize his return, and he only sells the object if the price is above some critical value or reserve price v_0. This reserve price is known beforehand to all bidders. In case the highest bid does not exceed v_0, the object is not sold and is not auctioned again. In fact, at AFA, these plants are destroyed. Therefore we take the seller's private valuation of the object to equal zero.

Suppose there are $n \geq 2$ potential buyers, denoted by $j = 1, \ldots, n$, who are identical ex-ante. Buyer j valuates the object at v_j, which is only privately observed. As the seller does not observe any of the buyers' valuations, the seller considers v_j, $j = 1, \ldots, n$ as independent draws from the same continuous distribution function F with support $[\underline{v}, \overline{v})$ with $0 \leq \underline{v} < \overline{v} \leq \infty$. The buyers consider the private values of the other buyers as random realizations from F. The equilibrium is a symmetric Bayesian-Nash equilibrium.

For given n and v_0, the optimal bid of a buyer with private value v equals

$$B(v|v_0, n) = v - \int_{v_0}^{v} \left[\frac{F(x)}{F(v)}\right]^{n-1} dx \tag{1}$$

Each bidder participating in the auction thus shades his private value with the amount $\int_{v_0}^{v}(F(x)/F(v))^{n-1}dx$.

The expected revenue of the seller is optimal if v_0 satisfies the first-order condition

$$v_0 f(v_0) = 1 - F(v_0) \tag{2}$$

and the second-order condition

$$2f(v_0) + v_0 \frac{\partial f(v_0)}{\partial v_0} \geq 0$$

Note that these two conditions are independent of the number of potential buyers. The reserve price can therefore be determined before the number of bidders is known, i.e. before the auctioning starts. In case there does not exist any v_0 satisfying these conditions, the optimal reserve price equals the lower bound of the support of the private values, $v_0 = \underline{v}$. The uniqueness of a reserve price in the support of v depends on the shape of F and is satisfied for most possible F (Van den Berg, 2007).

The model framework can be straightforwardly amended to account for the absence of potential bidders whose valuation falls short of v_0. Suppose that there are N potential bidders. The optimal value of v_0 does not depend on N. Only those n participants with $v \geq v_0$ decide to participate, leading to n actual bidders. These learn n, and they subsequently determine their bids. As a result, we have a game based on n bidders and on the distribution of valuations that equals the truncated original distribution of valuations, with truncation point v_0. In the sequel we simply denote the truncated distribution by F. There is an analogy to equilibrium search models of the labor market where the only labor market participants are those whose productivity exceeds the minimum wage (Van den Berg and Ridder, 1998).

A main advantage of the Dutch auction mechanism is that it works relatively fast. Therefore, it is often used to sell large amounts of perishable goods, such as flowers, vegetables and fruits. This is an advantage for all participants. The economic literature cites some additional reasons for why a Dutch auction could be attractive. For example, risk averse sellers prefer a Dutch auction, and collusive agreements among buyers are not self-enforcing (Wolfstetter, 1996).

We now discuss the accuracy of the above theoretical model for the AFA auctions of houseplants. First of all, consider the IPV framework. Bidders do not buy plants for other than commercial purposes. Many bidders are retailers with flower shops serving a local neighborhood. These act as monopolistic competitors on the consumption market for flowers and plants in their neighborhood. From their experience, they have an excellent knowledge of the demand functions of the products they sell to the consumers, and these functions may differ across different neighborhoods. In addition to these buyers, there are also large buyers who export flowers. These are typically active in a particular geographical region, where they have some market power. Plants are perishable and fragile goods, so there is no scope for extensive re-trading after the auction is held, at least not without substantial transaction costs. Milgrom and Weber (1982) also suggest that for nondurable consumer goods, like plants, the IPV framework suits better than the common value framework.

As noted above, the buyers include retailers as well as exporters. However, export of houseplants is much rarer than export of flowers, which lends some credence to the assumption of symmetry in the economic model. Similarly, by restricting attention to houseplants auctions where the mandatory minimum purchase quantity equals the actual total quantity, we restrict the data to non-sequential auctions.

Van den Berg, Van Ours and Pradhan (2001) argue that bidders at AFA are risk neutral. The main argument is that most bidders do not face strong binding financial constraints. The prices paid in single auctions are an extremely small fraction of the budget of a bidder. Furthermore, because other lots auctioned on the same day may be close substitutes, there is some insurance against losing an auction.

Most buyers have specific employees who perform the actual bidding. These individuals are present at AFA on many days of the year. It is therefore plausible that they know the exact number of actual bidders in the room. At the same time, our data only record the number of logged-on bidders, and this may over-estimate the number of actual bidders. We therefore also estimate models in which the number of actual bidders per auction is an unobserved drawing from a binomial distribution with the number of logged-on bidders as the highest outcome. This does not concern economic behavior, and it does not affect the market equilibrium for a given number of actual bidders. However, it implies that in the econometric analysis we need to consider a range of numbers of actual bidders for each given auction.

3.2 Non-parametric identification

We assume that all auctions are independent. With the number of actual bidders in each auction observed, F is non-parametrically identified from observations of the winning bid (this follows straightforwardly from Guerre, Perrigne and Vuong, 2000; see also Athey and Haile, 2002, for related identification results).

It is intuitively plausible that the observations of losing bids provide valuable additional information on F. This should reduce the dependence of the estimation results to arbitrary functional form assumptions on F. In the Appendix we investigate this by way of a simulation study.

As noted in the introduction and Subsection 3.1, we also estimate models where the actual number of bidders is unobserved to the econometrician and we assume that it is a drawing from a binomial distribution with the number of logged-on bidders as the highest outcome. Again, the observations of losing bids provide valuable information.[9] To see this, note that we observe the two highest bids for most separate auctions in the data. Song (2004) shows that the two highest bids identify F in eBay auctions, assuming an IPV framework and unknown numbers of bidders. In our setting, it is not difficult to demonstrate that with sufficient variation in the number of logged-on bidders, F is non-parametrically identified as well.[10]

3.3 Specification issues for the distribution of valuations

As explained in Subsection 3.1, we consider "truncated" models with absence of potential bidders whose valuation falls short of v_0. Consequently, the lower bound of the support of F equals the reserve price v_0. It is well known that for many specifications of F, the support of the distribution function of bids is bounded even if the support of F is unbounded. In particular, with n participating bidders the upper bound of the support depends on the structural parameters

$$\begin{aligned}\lim_{v\to\infty} B(v) &= \lim_{v\to\infty} v - \int_{v_0}^{v} \left[\frac{F(x)}{F(v)}\right]^{n-1} dx = \lim_{v\to\infty} v_0 \frac{F(v_0)^{n-1}}{F(v)^{n-1}} + \frac{\int_{v_0}^{v} x f(x) F(x)^{n-2} dx}{F(v)^{n-1}} \\ &= v_0 F(v_0)^{n-1} + \int_{v_0}^{\infty} x f(x) F(x)^{n-2} dx \\ &= \mathrm{E}\left[\max(V_{n-1}, v_0)\right]\end{aligned}$$

where V_{n-1} is the largest order statistic of $(n-1)$ draws from $F(\cdot)$ (see Laffont, Ossard and Vuong, 1995). Since v_0 is the lower bound of the support of F, it is

[9] First-price auctions are unidentified from only the winning bid if the number of bidders is unobserved.

[10] With two logged-on individuals, the actual number of bidders n is either 0 or 1 or 2. In the case $n = 0$ the item is not sold. In the case $n = 1$ it is sold at the observed reserve price v_0. In the case $n = 2$ the price exceeds the observed reserve price. Thus, the auction outcomes here identify the value of n. The fractions of the three possible outcomes $n = 0, 1, 2$ then identify the binomial parameter p_0. Subsequently, the bids in the case $n = 2$ non-parametrically identify F.

sufficient to impose

$$b_{(1)} \leq \mathrm{E}\left[V_{n-1}\right]$$

From an efficiency point of view, it is attractive to use the information captured in these bounds in the empirical analysis. Under some additional regularity conditions on the shape of the density function at its maximum, parameter estimates based on bounds are super-consistent, as the rate of convergence equals N (instead of the usual $\sqrt{N}$, where N is the number of observed auctions). However, estimation results using boundary conditions are generally not very robust against outliers. Possible causes for outliers are measurement errors or extreme behavior. In our setting, observing outliers is not very likely, the data are from an administrative database and all bidders are very experienced in the auctioning process, which excludes unusually high bids.

The distribution of private values F may differ between auctions. Suppose that all heterogeneity between auctions can be captured by a set of (exogenous) characteristics x of the houseplant and the auction. We assume that F can be uniquely characterized by a vector of unknown parameters θ and the set of known covariates x, $F_i(\cdot) = F(\cdot|x_i, \theta)$. Similar type of parametric assumptions are made in e.g. Donald and Paarsch (1996) and Laffont, Ossard and Vuong (1995). We take the distribution function of private values to be a transformed beta distribution, with density function

$$f(v|\alpha, \beta, \gamma) \sim \left(\frac{v - v_0}{\exp(\gamma) - v_0}\right)^{e^{\alpha}-1} \left(\frac{\exp(\gamma) - v}{\exp(\gamma) - v_0}\right)^{e^{\beta}-1} \qquad v_0 \leq v \leq \exp(\gamma) \quad (3)$$

where α, β and γ are unknown parameters, possibly depending on characteristics x of the auction. In Section 3.4 we discuss in detail how these parameters depend on explanatory variables. The density function is transformed such that it has support from the current minimum price v_0 up till an unknown finite upper bound $\exp(\gamma)$.[11]

For our purposes, the beta distribution is relatively flexible. The density is only symmetric if $\alpha = \beta$. The uniform distribution is a special case ($\alpha = \beta = 0$). The shape at the lower bound of the support of the density function is determined by α and at the upper bound by β. Close to the lower bound the density increases (decreases) if $\alpha > 0$ ($\alpha < 0$). Similar $\beta < 0$ ($\beta > 0$) implies that the density increases (decreases) close to the upper bound. Subsequently,

[11]This specification does not guarantee the existence of an optimal reserve price within the support of the distribution function of private values, which satisfies equation (2). To illustrate, let $\exp(\gamma) < 2v_0$ and $\alpha = \beta = 0$. This implies that $v_0^* f(v_0^*) = 1 - F(v_0^*)$ is solved for $v_0^* = \exp(\gamma)/2 < v_0$, which is below the lower bound of the support.

if both α and β are negative the density is U-shaped, and hump-shaped if these both parameters are positive. The density is strictly increasing (decreasing) if $\alpha > 0$ and $\beta < 0$ ($\alpha < 0$ and $\beta > 0$). The expectation (of v in equation (3)) equals $(\exp(\beta)v_0 + \exp(\alpha)\exp(\gamma))/(\exp(\alpha) + \exp(\beta))$ and the variance $(\exp(\gamma) - v_0)^2 \exp(\alpha)\exp(\beta)/((\exp(\alpha) + \exp(\beta))^2(\exp(\alpha) + \exp(\beta) + 1))$. Not only the uniform distribution is a special case of the beta distribution. Other special cases are for example the gamma distribution ($\beta \to \infty$ and $\exp(\gamma) = \exp(\gamma^*)(\exp(\alpha) + \exp(\beta)))$ and the exponential distribution ($\beta \to \infty$, $\alpha = 0$ and $\exp(\gamma) = \exp(\gamma^*)(\exp(\alpha) + \exp(\beta)))$ (see McDonald, 1984).

The optimal reserve prices only depends on the distribution function of private values. Given that we imposed a beta distribution, we can investigate the optimal reserve price and the percentage increase in expected revenue when moving from an auction without reserve price to an auction with an optimal reserve price. Because the expected revenue depends on the number of bidders participating in the auction, the percentage increase in revenue is a function of the number of bidders as well.

For the moment we choose a beta distribution function with support on 25 cents to 200 cents. Figures 4 to 6 show for α equal to -1, 0 and 1, respectively, what the percentage increase in expected revenue is if β is also -1, 0 and 1. The effect of imposing an optimal reserve price depends strongly on the value of the parameters (shape of the density). Note that if $\alpha = -1$ and β equals 1, the optimal reserve price remains equal to 25 cents. The percentage increase in expected revenue never exceeds 30%. In general, the effect of imposing a reserve price is higher for smaller values of β, i.e. when the density function is increasing close to the upper bound of the support. Reserve prices are particularly effective if the density is U-shaped (both α and β are smaller than 0). However, the impact of the reserve price on the expected revenue diminishes quickly if the number of bidders becomes large. If the number of bidders exceeds 5, reserve prices are not a very efficient instrument for generating additional revenue.

3.4 Estimation

We use a sampling Bayes approach to evaluate the vector of parameters θ. Inference reduces to evaluating the posterior density of the vector of parameters, $p(\theta|(b, x, n)_i, i = 1, \ldots, N)$, where N is the number of auctions observed in the data set and $(b, x, n)_i$ represent the bids observed in the i^{th} auction, the observed covariates and the number of bidders, respectively. We use Gibbs sampling to create Markov chains of values for vector of parameters θ. We use noninformative prior distributions for the vector of parameters θ.

When sampling the new vector of $\theta = \{\alpha, \beta, \gamma\}$, we have to ensure that it lies in the feasible parameter space bounded by the observed (winning) bids. In particular, after sampling a new set of values for θ, all observed bids must be within the support of the bid distribution function. Since we compute exact posterior distributions of the parameters, we do not have to rely on asymptotics.

Inference on auctions is straightforward if private values are observed. However, we observe bids instead of private values. Furthermore, only the winning bids and sometimes bids close to the winning bid are observed. For those bids, that are not observed, we know however that these are less than $b_{(1)} - \delta$, where $\delta \geq 0$ is the earlier mentioned interval in which we can observe losing bids.

We use data augmentation to construct the private values of all bidders in all auctions. For the observed bids we use the one-to-one relation given in equation (1) between bids and private values to compute the private values, i.e. $v = B^{-1}(b)$, where B^{-1} is the inverse of the equilibrium bid function. Data augmentation is used to sample all remaining private values from the distribution function $F(v|v \leq B^{-1}(b_{(1)} - \delta))$ (see Tanner and Wong, 1987). We use Gibbs sampling to construct a Gibbs sequence of values for the vector of parameters θ. Below we discuss the sampling method more formally.

Let n_i denote the actual number of bidders in auction i, which lies between the number of observed bids m_i and the registered number of bidders r_i. The outline of our sampling Bayes approach is as follows:

Initialization Choose an initial value p_0 in the interval between 0 and 1 and set $n_i = r_i$. Choose a set of parameters θ^0 within the possible parameter space, i.e. for each auctions the boundary condition on the support of the bid distribution function is satisfied, $b_{i,(1)} \leq \mathrm{E}\left[V_{n_i-1}|x_i, \theta^0\right]$. Since n_i and x_i are known we can compute the expectation for any particular choice of θ^0. Because also $b_{i,(1)}$ is observed, it is easy to check if the inequality is satisfied.

Augment number of bidders For each auction i we generate $n_i \in \{m_i, m_i + 1, \ldots, r_i\}$ from

$$f(n_i|(b,x,r,m)_i, p_0, \theta) \sim \begin{pmatrix} r_i \\ n_i \end{pmatrix} p_0^{n_i}(1-p)^{r_i-n_i} \left(\prod_{j=1}^{m_i} f(b_{i,j}|n_i, x_i, \theta)\right) \times F(b_{i,(1)} - \delta_i|x_i, n_i, \theta)^{n_i - m_i}$$

where $f(\cdot|n_i, x_i, \theta)$ and $F(\cdot|n_i, x_i, \theta)$ are the density and distribution function of observed bids conditional on the true number of bidders, covariates and the structural parameters. Once a new value for n_i is sampled it is only accepted if it satisfies the boundary condition $b_{i,(1)} \leq \mathrm{E}\left[V_{n_i-1}|x_i, \theta\right]$. If the boundary condition is not satisfied a new value for n_i is drawn.

Construct the 'observed' private values Use the equilibrium bid rule to compute the private values corresponding to all observed bids conditional on the set of parameters θ_t, i.e. for all auctions $i = 1, \ldots, N$, we compute for the observed bids $j = 1, \ldots m_i$ the private values $v_{i,(j)} = B^{-1}(b_{i,(j)}|n_i, x_i, \theta_t)$. Solving the equilibrium bid rule requires the use of a numerical procedure. However, in most cases 3 iterations are sufficient to find the value of $v_{i,(j)}$.[12]

Augment the 'latent' private values We know that the remaining latent bids are below $b_{(1)} - \delta$. This implies that the corresponding private values are below $B^{-1}(b_{(1)} - \delta)$. Therefore for each auction $i = 1, \ldots, N$ we compute $v_i^* = B^{-1}(b_{i,(1)}|n_i, x_i, \theta_t)$ and we sample for the latent bids $j = m_i+1, \ldots, n_i$ private values from the distribution function $F(v|v \leq v_i^*, x_i, \theta_t)$. This gives a complete set of private values v_{ij}, $j = 1, \ldots, n_i$, $i = 1, \ldots, N$.

Generate parameters We use the set of private values v to sample a new vector of parameters θ_{t+1}. Ignoring auction specific characteristics, we sample α_{t+1} from the density function

$$f(\alpha|\beta_t, \gamma_t, v) \sim \prod_{i=1}^{N} \prod_{j=1}^{n_i} \frac{\Gamma(\exp(\alpha) + \exp(\beta_t))}{\Gamma(\exp(\alpha))} \left(\frac{v_{ij} - v_0}{\exp(\gamma_t) - v_0} \right)^{e^{\alpha}-1}$$

where $\Gamma(\cdot)$ is the gamma function. A draw of α is only accepted if it satisfied the boundary condition $b_{i,(1)} \leq \mathrm{E}\left[V_{n_i-1}|x_i, \alpha, \beta_t, \gamma_t\right]$ for all auctions $i = 1, \ldots, N$. If the boundary condition is satisfied $\alpha_{t+1} = \alpha$, otherwise we draw a new value for α. Similarly, β_{t+1} is sampled from the density function

$$f(\beta|\alpha_{t+1}, \gamma_t, v) \sim \prod_{i=1}^{N} \prod_{j=1}^{n_i} \frac{\Gamma(\exp(\alpha_{t+1}) + \exp(\beta))}{\Gamma(\beta)} \left(\frac{\exp(\gamma_t) - v_{ij}}{\exp(\gamma_t) - v_0} \right)^{e^{\beta}-1}$$

again taking account of the boundary condition. And finally γ_{t+1} is drawn from the density function

$$f(\gamma|\alpha_{t+1}, \beta_{t+1}, v) \sim \prod_{i=1}^{N} \prod_{j=1}^{n_i} \frac{(\exp(\gamma) - v_{ij})^{e^{\beta}-1}}{(\exp(\gamma) - v_0)^{e^{\alpha}+e^{\beta}-1}}$$

[12]Sareen (1998) proposes a different approach. She avoids solving the equilibrium bid rule numerically by direct sampling of private values from $F(v|x, \theta_t)$. A sampled private value is only accepted if it corresponds to the observed bid.

where again an acceptable draw of γ should satisfy the boundary condition.

We also have to generate a new value for p_0. The new value of p_0 is sampled from the density function

$$f(p_0|(r,n)_i) \sim p_0^{\sum_{i=1}^N n_i}(1-p_0)^{\sum_{i=1}^N (r_i-n_i)}$$

which is a beta density function with parameters $\sum_{i=1}^N n_i$ and $\sum_{i=1}^N (r_i - n_i)$.

This provides us we a new vector of parameter θ_{t+1} and we can return to augmenting the number of bidders.

The Gibbs sequence only approximates the stationary distribution of the parameters if it is ergodic. At first sight ergodicity may seem problematic, as the support of the conditional distribution of a specific element of θ depends on the current values of the other elements in θ. This may cause the Gibbs sequence not to be irreducible (see Robert and Casella, 1995). Let $\mathcal{S}$ denote the support of the stationary distribution of θ and $K(\theta_t, \theta_{t+1})$ denotes the transition density resulting from the Gibbs sampler. Indeed there exist $\theta_t \in \mathcal{S}$ and $\theta_{t+1} \in \mathcal{S}$ for which $K(\theta_t, \theta_{t+1}) = 0$. However, it can be shown that

$$\int_{\mathcal{S}} K(\theta_{t+1}, \theta_{t+2}) K(\theta_t, \theta_{t+1}) d\theta_{t+1} > 0$$

for all $\theta_{t+2} \in \mathcal{S}$ and $\theta_t \in \mathcal{S}$.[13] Thus within two steps from each particular value of $\theta_t \in \mathcal{S}$ any $\theta_{t+2} \in \mathcal{S}$ can be reached, which implies that the Gibbs sequence is irreducible. Because the Gibbs sequence is also aperiodic, it is ergodic. With simulated data the Gibbs sampler performed well, i.e. once the Gibbs sequence converged the mean values of the parameters are close to the true values. However, it should be noted that convergence of the Gibbs sequence is slowed down by the limited support of the conditional distributions.

Our main interests are the means and the standard deviations of the posterior distribution. To approximate these we need realizations from the Gibbs sequence when it converged to the stationary distribution. To monitor convergence to the stationary distribution, we generate simultaneously 5 Gibbs sequences each with a

[13]The conditional density for α_{t+1} has support from some lower bound until infinity, where the lower bound decreases if β_t decreases or γ_t increases. The conditional density for β_{t+1} has support from minus infinity until an upper bound that increases if α increases or γ_t increases. And the conditional density for γ_{t+1} has support from some lower bound until infinity, with the lower bound decreasing if α_{t+1} increases or β_{t+1} decreases. It is intuitively clear that from some $\theta_t \in \mathcal{S}$ any $\theta_{t+2} \in \mathcal{S}$ is possible if γ_{t+1} is large, which always lies in the support of the conditional density of γ_{t+1}.

different initial value in the support $\mathcal{S}$. Convergence to the stationary distribution implies that the within variance of the Gibbs sequences is approximately equal to the between variance (see Gelman, Carlin, Stern and Rubin, 1995). Let $\theta_{r,t}$ denote the t^{th} value from the r^{th} Gibbs sequence (recall that we only keep each second draw). The mean of the posterior distribution is approximated by

$$\bar{\theta} = \frac{1}{5}\frac{1}{T}\sum_{r=1}^{5}\sum_{t=1}^{T}\theta_{r,t}$$

and the between and within variance are respectively computed by

$$V_B = \frac{T}{4}\sum_{r=1}^{5}\left(\bar{\theta}_r - \bar{\theta}\right)^2 \qquad \text{with} \qquad \bar{\theta}_r = \frac{1}{T}\sum_{t=1}^{T}\theta_{r,t}$$

$$V_W = \frac{1}{5}\sum_{r=1}^{5}\frac{1}{T-1}\sum_{t=1}^{T}\left(\theta_{r,t} - \bar{\theta}_r\right)^2$$

The variance of the marginal posterior distribution equals

$$\text{var}\left(\theta\right) = \frac{T-1}{T}V_W + \frac{1}{T}V_B$$

If all Gibbs sequences are drawn from the stationary distribution than

$$R = \sqrt{\frac{\text{var}\left(\theta\right)}{V_W}}$$

approach to 1 if T becomes sufficiently large.

We generate Gibbs sequences with an initial length of 100 draws and compute for each element in θ the value of R. If one of the values of R exceeds 1.1 we continue drawing elements of the Gibbs sequences. After each draw we compute R again, based on the last 100 draws of each Gibbs sequence until all values of R are below 1.1. If this is the case, we assume that all 5 Gibbs sequences converged to the stationary distribution. Then we draw for each Gibbs sequences 400 new realizations for θ, which we use to compute $\bar{\theta}$ and $\text{var}\left(\theta\right)$.

4 Structural analysis

In the empirical literature on Dutch auctions, it is common to observe only the winning bid and the number of participants. To obtain insight in the importance of actually observing losing bids, we first analyze our model ignoring the losing bids, and next we take into account that we also observe losing bids. In both of these analyses we assume that all registered bidders are actual participants in the auctions. Next, we extend our model such that we allow for flexibility in the number of bidders participating in the auctions.

We investigate the fit of the different specifications by computing the expected winning prices at the auction and the expected difference between the winning bid and the highest losing bid. (Table 5 provides these statistics for each of the models we estimate.) The expected winning price in an auction equals

$$\mathrm{E}[b_{(1)}|n, v_0] = n \int_{v_0}^{\overline{v}} \left(vf(v) - (1 - F(v))\right) F(v)^{n-1} dv$$

and the expected highest losing bid equals

$$\mathrm{E}[b_{(2)}|n, v_0] = n(n-1) \int_{v_0}^{\bar{v}} \left(vf(v) \frac{1 - F(v)}{F(v)} - \log(F(v)) + (1 - F(v)) \right) F(v)^{n-1} dv$$

We use the empirical results to investigate the importance of reserve prices. In particular, we compute the optimal reserve price for each auction and how much the expected revenue would increase if reserve prices would have set optimally.

We allow for heterogeneity between auctions of different houseplants by allowing the upper bound of the support of the distribution of private values to depend on observed characteristics x of an auction ($\gamma_i = x_i\gamma$). The parameters determining the shape of the density function of private values α and β are similar in all auctions.

4.1 Standard Dutch auction without losing bids

Table 6 presents the means and standard deviations of the marginal posterior distributions of the parameters when ignoring the losing bids. For all three subsample α is less than 0 and β is slightly larger than 0. This implies that the density functions of private values are downward sloping and many bidders have private values close to the reserve price (the lower bound of the support). The estimated model turns out to overestimate average winning prices for all three subsamples. For Begonias, the model predicts the expected average winning price equal to 219 cents, while in the data it is only 202 cents. For Camillas and Compactas the estimated expected average winning prices are 183 and 188, while the observed average winning prices are only 167 and 168 respectively. Also the difference between the winning bid and the highest losing bid is overestimated. For the subsamples of Begonias, Camillas and Compactas the model predicts these differences to be equal to 31, 22 and 24, while the observed values are 9, 9 and 8 respectively.

Overestimation of the average winning bid might indicate that our parametric specification is too restrictive. We have tried specifications where also β (and α) depends on observed characteristics x. However, the data on only the winning bid are not sufficiently informative to analyze such specifications. The winning bid alone cannot distinguish between covariate effects in the upper bound of the support and in β, which describes the shape of the density function close to the upper bound.

4.2 Dutch auction with losing bids

Observing losing bids allows to analyze more flexible parametric specifications of the structural model. Because losing bids are observed close to the winning bid, losing bids are mainly informative on the right-tail of the density function of private values. Therefore, we extended the model by allowing β to be different for each grower and we specify $\alpha = \alpha_0 + \alpha_1\beta$, so that also α can differ between growers. However, we only managed to use this specification for the subsample of Begonias, both other subsamples are too small. Therefore, for these subsamples we have set $\alpha_1 = 0$. Furthermore, for the subsample of Compactas the number of observations per grower is too small to let β be different between the growers, so that for Compactas we restricted the β to be similar for all growers. So, for Compactas we analyze the same specification as in Subsection 4.1

Table 7 shows the means and the standard deviations of the marginal posterior distributions for these specifications. The information on the losing bids did not improve the model predictions on the winning prices. Because the specification for auctions of Begonias is more flexible than the earlier specification, the prediction of the winning price is slightly better. However, for Compactas the prediction of the winning price is worst than in the previous specification. This is obviously caused by the fact that the same specification should now explain not only the winning bid but also losing bids. The fit of the highest losing bid improved for all three subsamples compared to the models in Subsection 4.1.

The shape of the density function of private values shows that for all three subsamples most participants in the auctions have private values close to the reserve prices. This might indicate that many bidders registered at an auction are actually not interested in buying the houseplants. The model implies that all winning prices fall within the support of the bid distribution function. If not all registered bidders are actual participants at the auction, the predicted variance in the winning prices is smaller than the observed variation in transaction prices. It is indeed the case that the observed variation is transaction prices (of auctions with similar characteristics) is large given the large number of registered participants. So, overestimation of the number of participants might explain why the model performs poorly in explaining winning prices.

4.3 Unknown number of bidders

Next, we relax the restriction that all registered bidders are true participants in an auction. So we allow for the possibility that buyers registered at a particular auction do not have a private value above the reserve price. However, we retain the assumption that bidders have an exact measure of the number of bidders participating at an auction. Recall that the parameter p_0 denotes the probability that a registered bidder is an actual participant at an auction. As a sensitivity analysis we also try a specification where we let the probability of being a true bidder be dependent on the moment in time at which the auction takes place.

The means and standard deviations of the marginal posterior distributions are provided in Table 8. The mean values for p_0 are between 0.10 and 0.14, implying that a large fraction of the registered bidders is not really interested in buying the houseplants. Most likely these bidders are just waiting for one of the next scheduled auctions. Recall from Subsection 2.2 that the average number of bidders registered during an auction was around 50. This implies that in every auction on average only 5 to 7 buyers are actually participating. The marginal posterior distributions of p_0 provided in Figure 7 have hardly any support above 0.20, which indicates that the models investigated in the previous subsections where all registered bidders are actual participants in the auction are clearly rejected, i.e. this model is a special case with $p_0 = 1$.

Bidders typically enter the auctioning hall at the moment flowers or plants are auctioned in which they are interested. This has two implications. First, there are less registered bidders early in the morning (see Figure 1). Secondly, early in the morning, the fraction of the registered bidders that actually intend to participate can be expected to be higher than later during the day. Let τ^* denote the elapsed time in minutes since the start of the auctioning process (at 6.30am). We extend our specification to

$$p = \begin{cases} p_0 + p_1 \frac{\tau *}{15} & \text{if} \quad \tau * < 15 \\ p_0 + p_1 & \text{if} \quad \tau * \geq 15 \end{cases}$$

Table 9 provides the results of this specification. As expected, in all cases p_1 is less than 0, indicating that the fraction of interested bidders decreases shortly after the auctioning starts. Immediately, after the auctioning starts around 20% of the registered bidders are true bidders, 15 minutes later this decreases to only around 6% to 8%. The expected number of interested bidders varies between 3 and 11 in all subsamples.

Compared to the previous subsection, the fit of the models improves. The second specification, where the fraction of the bidders interested in buying the houseplants depends on time does slightly better than the first specification. Therefore, we only discuss this second specification. For the Begonias the expected winning bid equals 208 and the expected difference between the highest and the second bid is 12. The expected winning bid is thus still overestimated, but the difference is much smaller than in the earlier specification. Also the difference between the highest and the second highest bid is slightly overestimated by the model (mainly due to the overestimation of the highest bid). We see the same for the subsamples of Camillas and Compactas, the fit improves but still the winning bid is slightly overestimated. For the Camillas the expected winning bid equals 173 and for the Compactas 174.

The grower seems to be the most informative covariate. Although the houseplants in our samples all have the highest quality codes, the distribution function of private values differs between growers. For both Camillas and Compactas the upper bound of the support is smaller for houseplants supplied by Grower 5 (as compared to houseplants supplied by Growers 4 and 6). Obviously Grower 4 and Grower 6 have better reputations among potential buyers of Camillas and Compactas than Grower 5. For Begonias both β and γ are larger for Grower 3 than for both other Growers or Begonias. The expected revenue does not differ much between the growers, but the variation in revenue between lots is much larger for Grower 3 than for Grower 1 and Grower 2.

The second covariate effect concerns the day of the week at which the auction takes place. For Begonias expected revenue is lowest on Thursday. For Camillas this is the case on Wednesday and the prices of Compactas are lower on Wednesday and Thursday compared to the other days. The variation in private values might have to do with the day at the week when transport of houseplants to foreign countries takes place. However, the pattern is not similar for the three types of flowers. Therefore, transportation can only explain these differences if different countries have preferences for different types of houseplants.

As expected allowing not all registered bidders to be true participants at the auctions changes the shape of the distribution function of private values. For all subsamples α exceeds 0 and thus the density function of private values is increasing close to the lower bound (reserve price). For the subsamples of Begonias and Camillas the density function is decreasing close to the upper bound ($\beta > 0$). For the subsample of Compactas the density function of private values is not hump-shaped, but increasing ($\beta < 0$). From Subsection 3.3 we know that reserve prices are not very effective when the density function of private values is increasing close to the lower bound. However, since the actual number of bidders is not very large, reserve prices might have some impact on the expected revenue.

4.4 Optimal reserve prices

The structural model allows us to evaluate the impact of reserve prices on the expected revenue at the auction. At the moment the data were collected, there was almost no differentiation of reserve prices between different types of plants. AFA believes that the reserve prices are low. This is confirmed by the data, which do not show any winning bids close to the reserve price. Recall that without ad hoc functional form assumptions we can not study the consequences of decreasing reserve prices.

For each of the parameter values obtained in the Gibbs sequence, we can compute the optimal reserve price. We compare the current situation in which there is only a single reserve prices of 25 cents with a situation where there is full differentiation of reserve prices. This means that we allow for different reserve prices for each grower on each day of the week.

It turns out that from the point of view of revenue maximization, the actual reserve prices in the auctions of Begonias are sub-optimally low, for all plants of all growers on all days of the week. The optimal reserve prices are around 141 cents, and these do not vary much between growers and over the days of the week (except that on Thursdays, reserve prices should be slightly lower than on other days). Increasing the reserve prices to the optimal values increases the expected revenue with only 0.2%. In the auctions of Camillas and Compactas, the actual reserve prices are also always sub-optimally low. For Grower 4, the reserve prices should be around 190 cents when supplying Camilla and 200 cents when supplying Compacta. The reserve prices for Camillas supplied by Grower 5 should only increase to 45 cents and to 50 cents if this grower supplies Compactas. The optimal reserve price for Compactas supplied by Grower 6 is about 150 cents. When choosing the reserve prices optimally (and thus differentiated between growers and over the days of the week) the expected revenue increases with 1.3% for both Camillas and Compactas.

We conclude from this that AFA can generate more revenue if it increases reserve prices. Although the increase in revenue would be small, one may wonder whether we overlooked reasons which caused actual reserve prices to be set at such a low level in the first place. One explanation focuses on the bidder entry decision. Levin and Smith (1994), who endogenize the number of bidders, prove that if bidders enter only if their expected marginal profit exceeds their expected marginal costs, reserve prices are an instrument to discourage bidder entry. Imposing a reserve price exceeding the seller's valuation is therefore not necessarily beneficial to the seller. McAfee and McMillan (1987b) show that in the presence of entry costs, equation (2) no longer holds as the solution to the optimal reserve price. In this case the auctioneer should impose a reserve price equal to the valuation of the seller. At AFA, bidders do not make costs to participate in a single auction. Instead, bidders pay an annual fee. However, they also have to be present in the auctioning halls during the auctioning process, and they make transportation costs. The costs associated to preparing bids can also be considered as entry costs. It is therefore possible that bidder entry costs that are not explicitly included in our model negatively affect the optimal reserve price.

Broens and Meulenberg (1999) conduct a survey among growers, bidders, and auctioneers at AFA to gauge their opinion on aspects of the auctioning process. The growers of houseplants believe that higher reserve prices increase the bids without affecting the percentage of houseplants that will be destroyed. However, the bidders believe that this percentage will increase, and that in particular low-quality houseplants will not be sold. The auctioneers like the idea of "crowding-out" low-quality houseplants, as it improves the average quality of plants sold at auction. But they fear that higher (reserve) prices will drive buyers to competing auction houses and/or may induce production in other countries of cheaper plant varieties to be sold locally in those countries. Recall that even though AFA is the largest auction house of plant products in the world, it is not a monopoly. In the Netherlands, there are 6 other medium-sized and large flower auction houses, and other countries have internationally oriented flower auction houses as well. As we saw in Section 2, the reserve prices were homogeneous across the main flower auction houses in the Netherlands. This can be seen as an outcome of a game played by the joint Dutch auction houses against auction houses abroad. A substantial increase of AFA's reserve price may well induce a loss of buyers. Such effects are not incorporated in our optimal revenue analysis. Finally, with higher reserve prices, the costs of destroying the unsold goods is higher.

5 Conclusions

In this paper we have used a Bayesian approach to structurally analyze Dutch flowers auctions. In particular, we focused on the IPV model for the flower auction. This model is analyzed with Gibbs sampling methods using data augmentation to sample the (latent) private values of all participants in the auction. This method appeared to perform well.

In most cases the Dutch auction reveals only the highest bid, as the auction stops when this bid has been made. However, at AFA also all losing bids in a time interval next to the winning bid are observed. The length of the time interval is determined by the speed of reaction of the auction participants.

The empirical results show that reputation is very important at AFA. Growers get different prices for their plants, even if these plants have the same quality code. Although the current reserve prices at AFA are too low to optimize the expected revenue, increasing the reserve prices does not generate substantial additional revenue. Still, current reserve prices at AFA are higher and closer to our predicted optimum values than they were in the period covered by our data.

Some topic for future research emerge. First, the data registers at our disposal include observations of sequential auctions of multiple homogeneous units. We can use these data to estimate structural models of such auctions, incorporating the so-called buyer's option that bidders at AFA have. Secondly, one may empirically analyze competition and collusion among different flower auction houses. Flower auction houses may use their reserve prices to compete with each other. During the period covered by our data there was no variation in reserve prices across Dutch auction houses and no variation over time for given auction houses. Currently, reserve prices are more differentiated across auction houses and across products than before. At the same time, market concentration has been increasing as auction houses have been involved in mergers and takeovers since decades (see e.g. Elshof, 2000). It would be interesting to examine to what extent certain changes in the variation of reserve prices antedate mergers between auction houses.

References

Athey, S. and P.A. Haile (2002), Identification of standard auction models, *Econometrica* 70, 2107–2140.

Bajari, P. and A. Hortacsu (2003), The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions, *Rand Journal of Economics* 34, 329–355.

Bloemenveiling Aalsmeer (1996a), *Annual report* (In Dutch).

Bloemenveiling Aalsmeer (1996b) *Code-booklet 1996* (in Dutch).

Broens, D.F. and M.T.G. Meulenberg (1999), Minimumprijzen voor potplanten (in Dutch), Report, Wageningen University.

Casella, G. and E.I. George (1992), Explaining the Gibbs sampler, *American Statistician* 46, 167–174.

Donald, S.G. and H.J. Paarsch (1993), Piecewise pseudo-maximum likelihood estimation in empirical models of auctions, *International Economic Review* 34, 121–148.

Donald, S.G. and H.J. Paarsch (1996), Identification, estimation and testing in parametric empirical models of auctions within the independent private values paradigm, *Econometric Theory* 12, 517–567.

Elshof, P. (2000), The Dutch flower sector: structure, trends and employment, Working paper, ILO, Geneva.

Gelfand A.E. and A.F.M. Smith (1990), Sampling-based approaches to calculating marginal densities, *Journal of the American Statistical Association* 85, 398–409.

Gelman, A., J.B. Carlin, H.S. Stern and D.B. Rubin (1995), *Bayesian Data Analysis*, Chapman and Hall, London.

Guerre, E., I. Perrigne and Q. Vuong (2000), Optimal nonparametric estimation of first-price auctions, *Econometrica* 68, 525–574.

Guinness World Records (2001), *Guinness World Records 2002 (Guinness Book of Records)*, Guinness, Enfield Middlesex.

Hendricks, K. and H.J. Paarsch (1995), A survey of recent empirical work concerning auctions, *Canadian Journal of Economics* 28, 403–426.

Kambil, A. and E. van Heck (1996), Competition in the Dutch flower markets, Working paper, New York University, New York.

Katok, E. and A.E. Roth (2004), Auctions of homogeneous goods with increasing returns: experimental comparison of alternative "Dutch" auctions, *Management Science* 50, 1044–1063.

Laffont, J.-J. (1997), Game theory and empirical economics: The case of auction data, *European Economic Review* 41, 1–35.

Laffont, J.-J., H. Ossard and Q. Vuong (1995), Econometrics of first-price auctions, *Econometrica* 63, 953–980.

Laffont, J.-J. and Q. Vuong (1996), Structural analysis of auction data, *American Economic Review, Papers and Proceedings* 80, 414–420.

Levin, D. and J.L. Smith (1994), Equilibrium in auctions with entry, *American Economic Review* 84, 585–599.

Logan, G.D. and W.B. Cowan (1984), On the ability to inhibit thought and action: A theory of an act of control, *Psychological Review* 91, 295–327.

Matthews, S.A. (1995), A technical primer on auction theory I: independent private values. Working paper, Northwestern University, Evanston.

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

McAfee, R.P. and J. McMillan (1987b), Auctions with entry, *Economics Letters* 23, 343–347.

McDonald, J.B. (1984), Some generalized functions for the size distribution of income, *Econometrica* 52, 647–663.

Milgrom, P.R. and R.J. Weber (1982), A theory of auctions and competitive bidding, *Econometrica* 50, 1089–1122.

Rieger, M. (2000), *Stop! And Go? – Neuroanatomical correlates and consequences of the inhibition of ongoing responses*, Phillips-University Marburg, Marburg/Lahn.

Robert, C.P. and G. Casella (1999), *Monte Carlo Statistical Methods*, Springer-Verlag, New York.

Sareen, S. (1998), Likelihood based estimation of symmetric parametric structural auction models with independent private values, Working paper, University of Toronto, Toronto.

Silverman, B.W. (1986), *Density Estimation*, Chapman and Hall, London.

Song, U. (2004), Nonparametric estimation of an eBay auction model with an unknown number of bidders, Working paper, University of British Columbia, Vancouver.

Tanner, M.A. and W.H. Wong (1987), The calculation of posterior distributions by data augmentation, *Journal of the American Statistical Association* 82, 528–540.

Van den Berg, G.J. and G. Ridder (1998), An empirical equilibrium search model of the labor market, *Econometrica* 66, 1183–1221.

Van den Berg, G.J., J.C. van Ours and M.P. Pradhan (2001), Declining prices in the sequential Dutch flower auction of roses, *American Economic Review* 91, 1055–1062.

Van den Berg, G.J. (2007), On the uniqueness of optimal prices set by monopolistic sellers, *Journal of Econometrics*, forthcoming.

Wolfstetter, E. (1996), Auctions: An introduction, *Journal of Economic Surveys* 10, 367–420.

Appendix: Simulation study

We perform a small simulation study to investigate to what extent the losing bids data contribute to the precision of the estimates. To keep matters manageable, we assume that the number of bidders is observed, and we ignore the economic model in the simulation study and simply focus on the distribution function of bids instead of the distribution function of private values. We draw samples of bids from 3 different distribution functions. To remain as close as possible to the IPV first-price auction model, we consider distribution functions for which the right-tail is truncated. In particular, we sample from (*i*) a uniform distribution function with support from 0 until 2, (*ii*) an exponential distribution function ($\lambda = 2$), which is truncated at 2, and (*iii*) a log-normal distribution function ($\mu = 0$ and $\sigma = 1$), which is also truncated at 2.

We generate two samples more or less similar to the data sets discussed in the previous section. Moreover, the first sample contains 64 auctions with the number of bidders similar to the number of bidders in the data set of Begonias and the second sample has 29 auctions with the number of bidders from the data set of Camillas (the data set of Compactas is almost of the same size as that of Camillas). For each bidders in each auction we draw a bid. We use these bids to construct three samples, the first sample contains only the winning bid, the second sample contains also the losing bids close to the winning bid and the third sample contains all bids. We choose the threshold point for observing losing bids (the length of the interval just below the winning bid) such that we observe approximately 5.5% of all losing bids.

We perform six simulation experiments. For both sample sizes we sample from the three distribution functions mentioned above. Once we have constructed new samples, we estimate the distribution function of bids using the beta distribution function (see equation (3)) only taking into account the winning bid, also considering some losing bids and using all bids. We repeat each experiment 100 times. In Figures 8, 9 and 10 we show the average estimated densities of the simulation studies for the sample size of 64 auctions and the Figures 11, 12 and 13 show these graphs for the sample size of 29 auctions. Each figure presents four lines, the true density and the estimated densities based on (*i*) all bids, (*ii*) the winning bid and some losing bids and (*iii*) only the winning bid. As expected, the true density is best estimated by the samples containing all bids. But the estimated density based on the winning bids and some of the losing bids lies relatively close to the estimated density based on all bids. Obviously, the fit improves already enormously if only a small proportion of the losing bids is observed. The information contained in the losing bids is thus fairly large. According to the results it seems hard to approximate an underlying density if one observes only a sequence with the highest observations. Finally, the assumed numbers of participants in the auctions is rather large in the light of the final estimation results in the paper, and one may expect that the simulation conclusions on the usefulness of losing bids are even stronger in the case of smaller numbers of participants.

	Begonia	Camilla	Compacta
Number of observation	178	47	34
Number of days	14	11	9
Number of growers	3	2	3
Average price (in cents)	202	167	168
	(23)	(94)	(95)
Minimum price	130	48	53
Maximum price	250	320	320
Average number of bidders	56	54	50
	(12)	(9.1)	(12)

Explanation: Standard deviations are given in parentheses.

Table 1: Some characteristics of the datasets.

	Begonia		Camilla		Compacta	
Number of bidders	0.0012	(0.0008)	0.0098	(0.0030)	−0.0051	(0.0034)
Grower 1	5.20	(0.046)				
Grower 2	5.31	(0.052)				
Grower 3	5.23	(0.048)				
Grower 4			4.88	(0.18)	5.90	(0.20)
Grower 5			3.66	(0.16)	4.31	(0.14)
Grower 6					5.40	(0.19)
Monday	0		0		0	
Tuesday	0.029	(0.021)	0.13	(0.067)	−0.074	(0.086)
Wednesday	0.014	(0.024)	−0.026	(0.060)	0.074	(0.082)
Thursday	−0.10	(0.025)	−0.051	(0.085)	−0.017	(0.084)
Friday	0.0040	(0.037)	−0.0005	(0.12)		
R^2	0.29		0.95		0.95	

Explanation: Estimated standard errors are given in parentheses.

Table 2: OLS regression results on the logarithm of the winning price.

	Begonia			**Camilla**			**Compacta**		
τ	μ_2	obs	p	μ_2	obs	p	μ_2	obs	p
0.05	0.049 (0.0073)	0.80		0.070 (0.030)	0.62		0.058 (0.022)	0.57	
0.10	0.049 (0.0068)	0.94	0.68	0.071 (0.030)	0.76	0.73	0.055 (0.014)	0.80	0.70
0.15	0.049 (0.0067)	1.00	0.42	0.068 (0.024)	0.90	0.59	0.064 (0.017)	0.83	0.50
0.20				0.068 (0.023)	0.93	0.59	0.056 (0.014)	0.97	0.73
0.25				0.071 (0.024)	0.93	0.51	0.057 (0.014)	0.97	0.83
0.30				0.074 (0.024)	0.93	0.50	0.058 (0.014)	0.97	0.90

Explanation: τ is the speed of reaction, μ_2 the fraction that the highest losing bid is lower than the winning bid. 'obs' is the fraction of the auctions in which the highest losing bid is observed. p is the p-value of a test for a significant difference with the estimate under $\tau = 0.05$.

Table 3: Estimated difference between the winning bid and the highest losing bid (as a fraction of the winning bid).

	Begonia			**Camilla**			**Compacta**		
τ	μ_4	obs	p	μ_4	obs	p	μ_4	obs	p
0.10	0.14	0.67					0.20	0.30	
	(0.021)						(0.10)		
0.15	0.13	0.84	0.42	0.31	0.52		0.19	0.50	0.83
	(0.016)			(0.18)			(0.069)		
0.20	0.14	0.89	0.62	0.25	0.69	0.45	0.19	0.60	0.81
	(0.016)			(0.10)			(0.055)		
0.25	0.14	0.94	0.53	0.27	0.72	0.51	0.22	0.63	0.55
	(0.016)			(0.11)			(0.076)		
0.30	0.14	0.94	0.75	0.28	0.76	0.53	0.21	0.70	0.66
	(0.016)			(0.12)			(0.063)		

Explanation: τ is the speed of reaction, μ_4 the fraction that the fourth highest bid is lower than the winning bid. 'obs' is the fraction of the auctions in which the fourth highest bid is observed. p is the p-value of a test for a significant difference with the estimate under $\tau = 0.10$ for Begonias and Compactas and $\tau = 0.15$ for Camillas.

Table 4: Estimated difference between the winning bid and the third highest losing bid (as a fraction of the winning bid).

	Begonia	Camilla	Compacta
Data			
Reserve price	25.00	25.00	25.00
Average revenue	202.25	166.70	167.97
Average difference winning and second bid	9.38	8.82	8.43
Model ignoring losing bids			
Average revenue	219.17	182.95	187.57
	(0.96)	(2.60)	(5.19)
Average difference winning and second bid	31.42	22.46	24.26
	(1.95)	(1.80)	(3.32)
Average optimal reserve price	97.13	80.83	80.82
Average revenue at optimal reserve price	219.44	183.11	187.71
	(0.94)	(2.62)	(5.16)
Model with losing bids			
Average revenue	218.88	184.44	192.98
	(0.85)	(2.37)	(2.44)
Average difference winning and second bid	25.29	3.55	7.70
	(0.89)	(3.18)	(2.70)
Average optimal reserve price	75.06	99.96	103.45
Average revenue at optimal reserve price	218.93	187.56	195.39
	(0.87)	(1.67)	(1.99)
Model unknown bidders			
Average revenue	209.74	173.02	173.29
	(0.93)	(2.56)	(2.81)
Average difference winning and second bid	12.03	10.03	7.24
	(0.42)	(3.21)	(2.87)
Average optimal reserve price	139.06	133.50	133.07
Average revenue at optimal reserve price	209.96	175.29	176.36
	(0.92)	(2.05)	(2.01)
Model unknown bidders and p dependent on time			
Average revenue	208.44	173.22	174.40
	(1.03)	(2.27)	(2.77)
Average difference winning and second bid	11.90	10.20	8.48
	(0.45)	(2.98)	(2.64)
Average optimal reserve price	141.84	133.40	132.07
Average revenue at optimal reserve price	208.83	175.50	176.75
	(1.01)	(1.88)	(2.21)

Explanation: Estimated standard errors are given in parentheses.

Table 5: Model fit and policy simulations.

	Begonia		Camilla		Compacta	
α	−1.97	(0.088)	−2.05	(0.17)	−2.12	(0.19)
β	0.046	(0.034)	0.094	(0.14)	0.24	(0.34)
Grower 1	5.61	(0.012)				
Grower 2	5.65	(0.0096)				
Grower 3	5.62	(0.020)				
Grower 4			5.86	(0.042)	6.06	(0.22)
Grower 5			4.41	(0.035)	4.56	(0.20)
Grower 6					5.83	(0.16)
Monday	0		0		0	
Tuesday	0.0003	(0.0038)	0.10	(0.024)	0.011	(0.098)
Wednesday	0.16	(0.019)	−0.069	(0.017)	−0.083	(0.097)
Thursday	−0.055	(0.017)	0.064	(0.034)	−0.087	(0.11)
Friday	−0.051	(0.011)	−0.035	(0.064)		

Explanatory note: Standard deviations in parentheses.

Table 6: The mean and the standard deviation of the marginal posterior densities of the structural model ignoring the losing bids.

	Begonia		Camilla		Compacta	
α_0	−1.15	(0.14)	−2.11	(0.10)	−2.11	(0.12)
α_1	0.27	(0.17)	0		0	
β					−0.32	(0.070)
β_1	0.41	(0.11)				
β_2	0.26	(0.092)				
β_3	1.65	(0.44)				
β_4			−0.38	(0.062)		
β_5			−0.27	(0.082)		
Grower 1	5.61	(0.026)				
Grower 2	5.63	(0.017)				
Grower 3	6.30	(0.27)				
Grower 4			5.74	(0.0045)	5.76	(0.010)
Grower 5			4.32	(0.012)	4.14	(0.0080)
Grower 6					5.61	(0.012)
Monday	0		0		0	
Tuesday	−0.025	(0.012)	0.086	(0.0037)	0.085	(0.012)
Wednesday	0.046	(0.011)	−0.061	(0.0050)	0.12	(0.0093)
Thursday	−0.14	(0.023)	0.063	(0.017)	0.11	(0.0086)
Friday	−0.055	(0.017)	−0.046	(0.026)		

Explanatory note: the standard deviations are given in parentheses.

Table 7: The mean and the standard deviation of the marginal posterior densities of the structural model taking the losing bids into account.

	Begonia		Camilla		Compacta	
α_0	3.53	(0.12)	0.95	(0.18)	0.59	(0.15)
α_1	−0.86	(0.065)	0		0	
β					−0.19	(0.11)
β_1	1.49	(0.071)				
β_2	1.38	(0.079)				
β_3	2.16	(0.13)				
β_4			0.066	(0.28)		
β_5			0.46	(0.33)		
Grower 1	5.62	(0.016)				
Grower 2	5.66	(0.021)				
Grower 3	6.11	(0.099)				
Grower 4			5.77	(0.031)	5.78	(0.041)
Grower 5			4.38	(0.055)	4.31	(0.047)
Grower 6					5.58	(0.014)
Monday	0		0		0	
Tuesday	−0.019	(0.012)	0.045	(0.019)	0.029	(0.039)
Wednesday	−0.034	(0.014)	−0.084	(0.018)	−0.064	(0.041)
Thursday	−0.16	(0.014)	0.019	(0.037)	−0.093	(0.046)
Friday	−0.028	(0.023)	−0.030	(0.082)		
p_0	0.10	(0.0043)	0.13	(0.011)	0.14	(0.013)

Explanatory note: the standard deviations are given in parentheses.

Table 8: The mean and the standard deviation of the marginal posterior densities of the structural model with an unknown number of bidders (first specification).

	Begonia		**Camilla**		**Compacta**	
α_0	3.02	(0.28)	0.91	(0.16)	0.61	(0.18)
α_1	−0.61	(0.15)	0		0	
β					−0.11	(0.15)
β_1	1.35	(0.14)				
β_2	1.18	(0.16)				
β_3	2.27	(0.22)				
β_4			0.046	(0.22)		
β_5			0.47	(0.35)		
Grower 1	5.60	(0.019)				
Grower 2	5.64	(0.023)				
Grower 3	6.20	(0.15)				
Grower 4			5.78	(0.023)	5.85	(0.055)
Grower 5			4.40	(0.075)	4.39	(0.061)
Grower 6					5.59	(0.019)
Monday	0		0		0	
Tuesday	−0.025	(0.012)	0.036	(0.017)	−0.037	(0.054)
Wednesday	0.016	(0.020)	−0.096	(0.016)	−0.13	(0.054)
Thursday	−0.18	(0.014)	0.013	(0.040)	−0.16	(0.061)
Friday	0.021	(0.034)	0.070	(0.13)		
p_0	0.21	(0.012)	0.18	(0.026)	0.21	(0.041)
p_1	−0.15	(0.013)	−0.10	(0.042)	−0.15	(0.076)

Explanatory note: the standard deviations are given in parentheses.

Table 9: The mean and the standard deviation of the marginal posterior densities of the structural model with an unknown number of bidders (second specification).

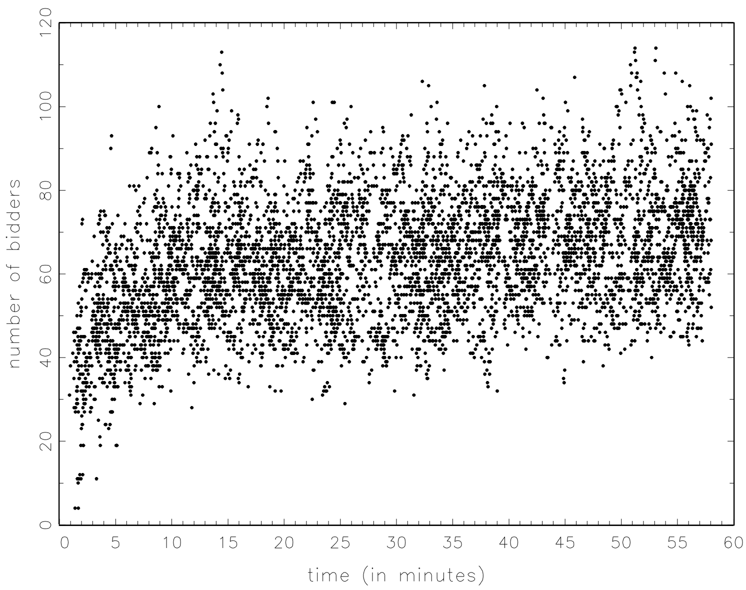

Figure 1: The number of bidders participating in auctions shortly after the early-morning opening of the auction house.

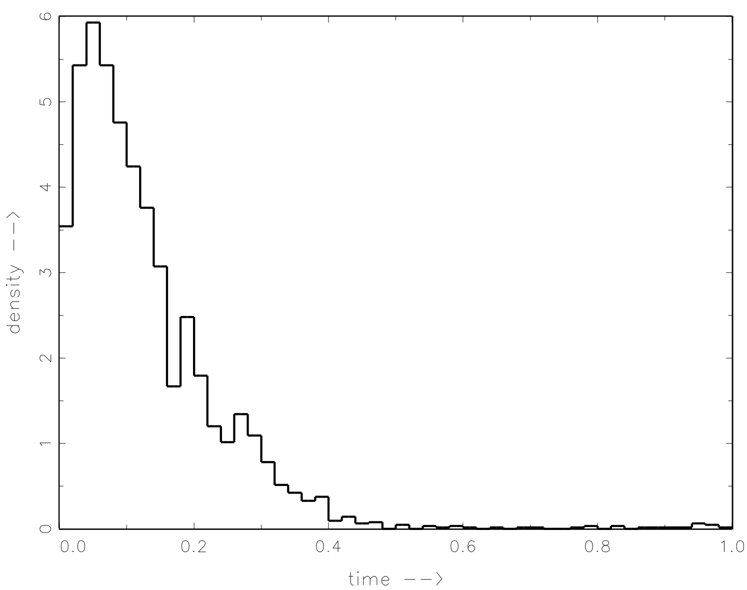

Figure 2: Scaled histogram of the time (in seconds) between the moment that the highest bidder bids and the moment the next bidder is observed to bid.

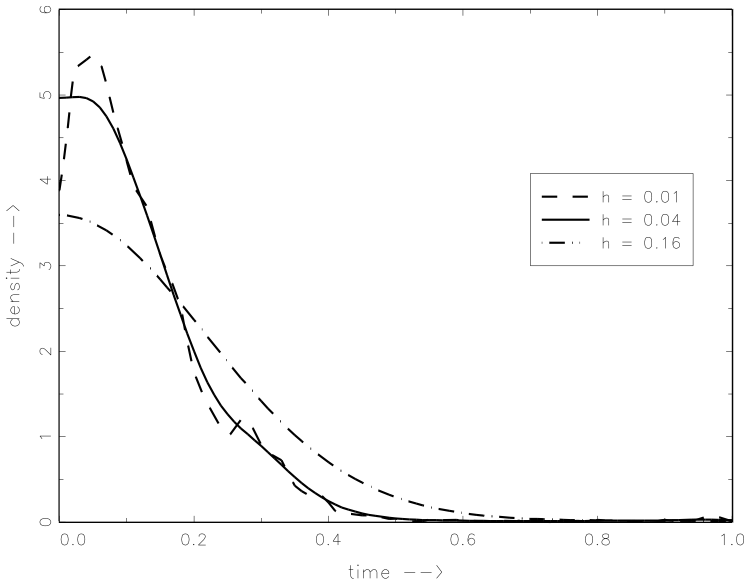

Figure 3: Kernel estimates of the time (in seconds) between the moment that the highest bidder bids and the moment the next bidder is observed to bid (for different values of the bandwidth (h)).

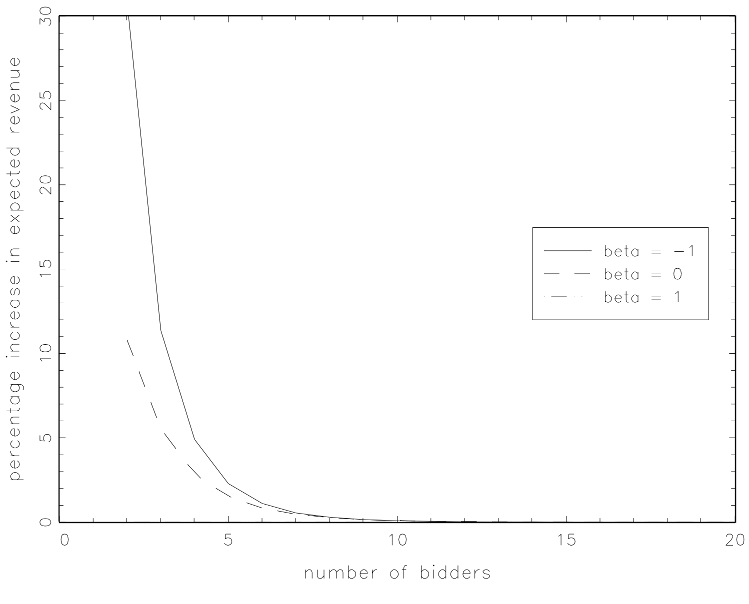

Figure 4: The percentage increase in expected revenue if the auction sets an optimal reserve price instead of no reserve price, as a function of the number of bidders ($\alpha = -1$).

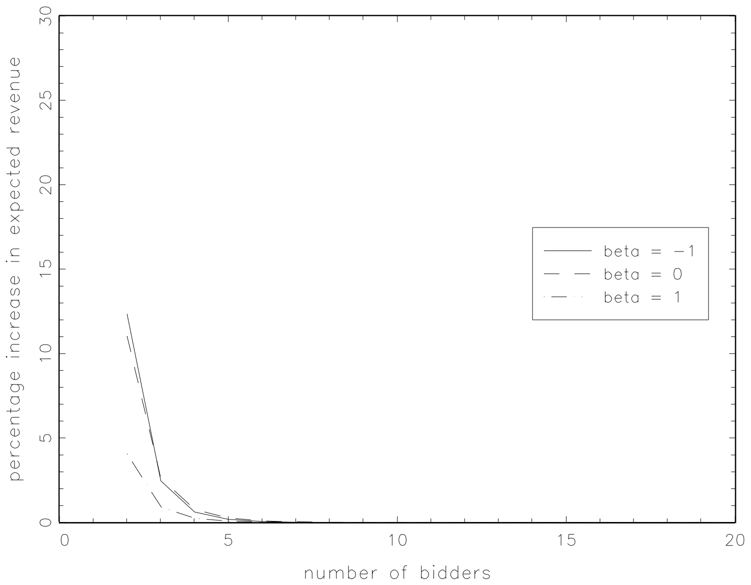

Figure 5: The percentage increase in expected revenue if the auction sets an optimal reserve price instead of no reserve price, as a function of the number of bidders ($\alpha = 0$).

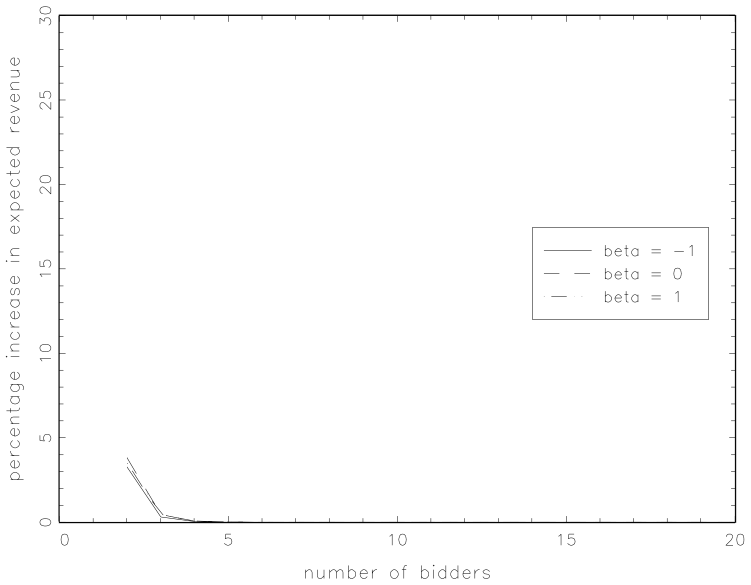

Figure 6: The percentage increase in expected revenue if the auction sets an optimal reserve price instead of no reserve price, as a function of the number of bidders ($\alpha = 1$).

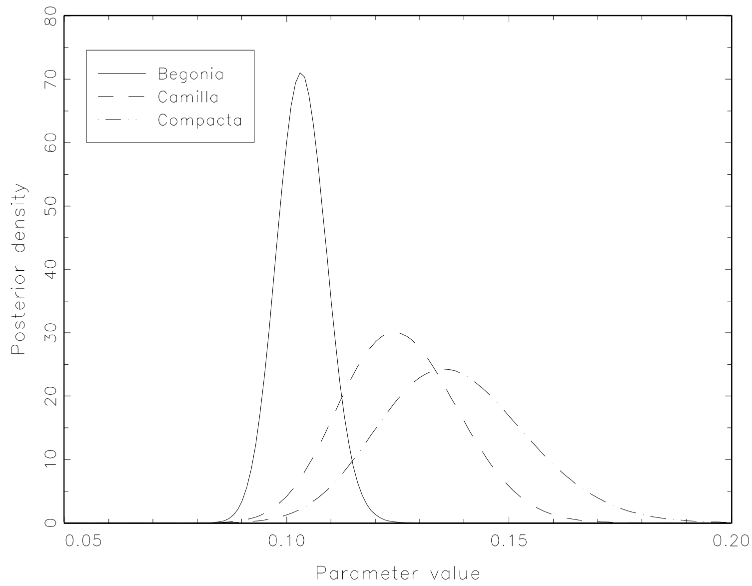

Figure 7: Marginal posterior density for the fraction of the registered bidders that is actually participating.

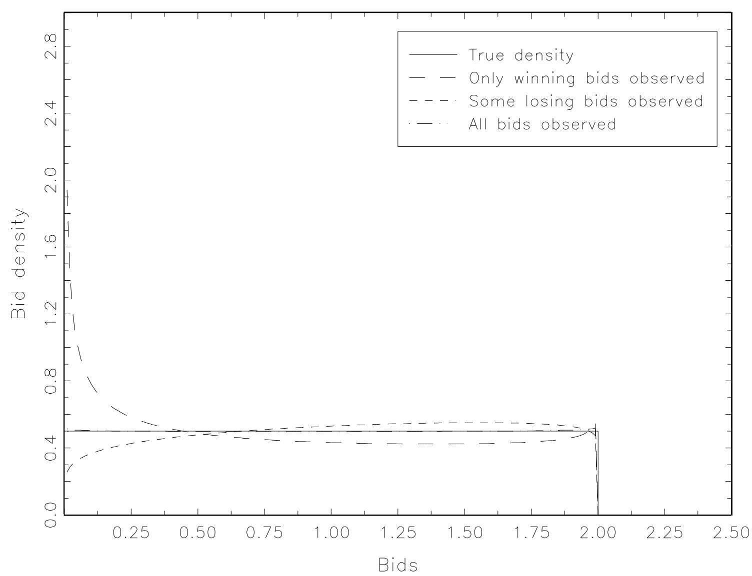

Figure 8: The true density and estimated densities of the simulation study with a uniform distribution.

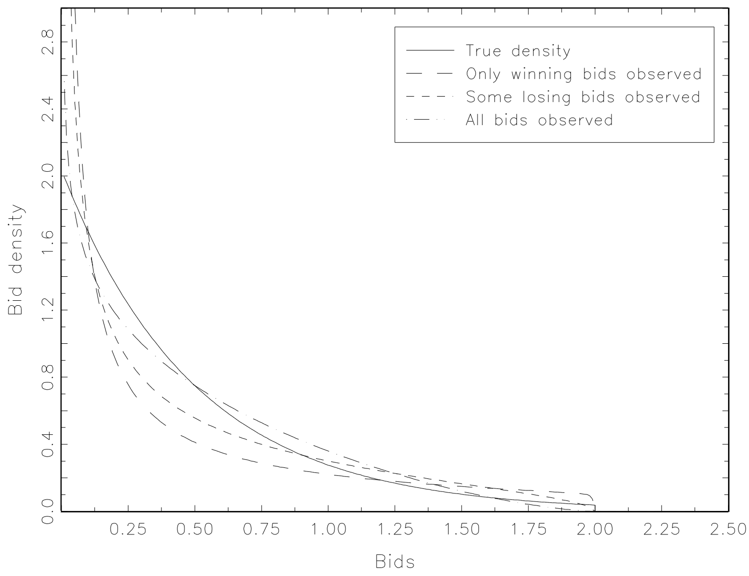

Figure 9: The true density and estimated densities of the simulation study with a truncated exponential distribution.

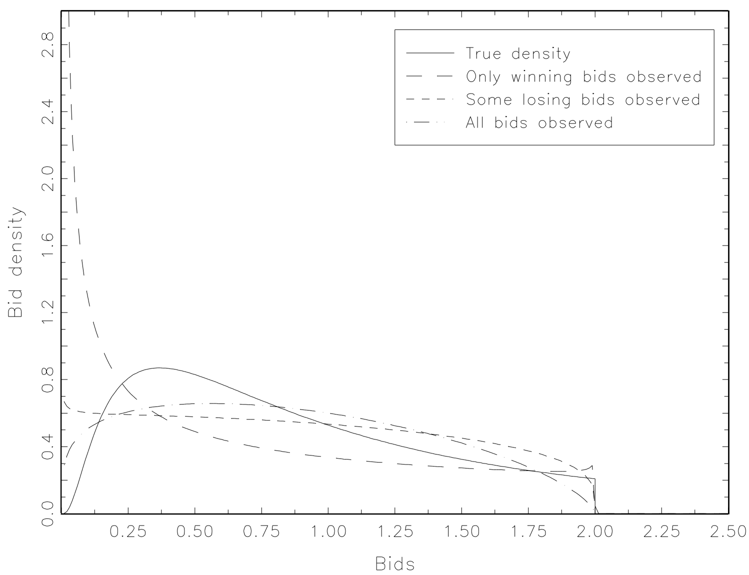

Figure 10: The true density and estimated densities of the simulation study with a truncated exponential distribution.

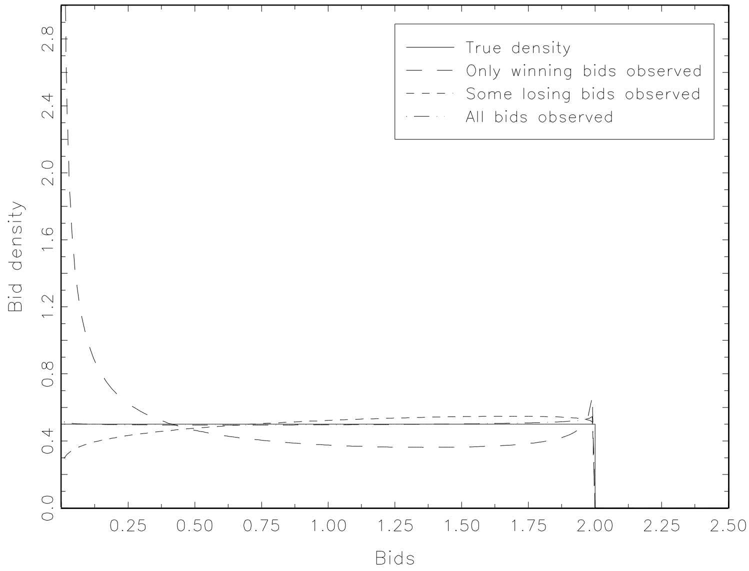

Figure 11: The true density and estimated densities of the simulation study with a uniform distribution.

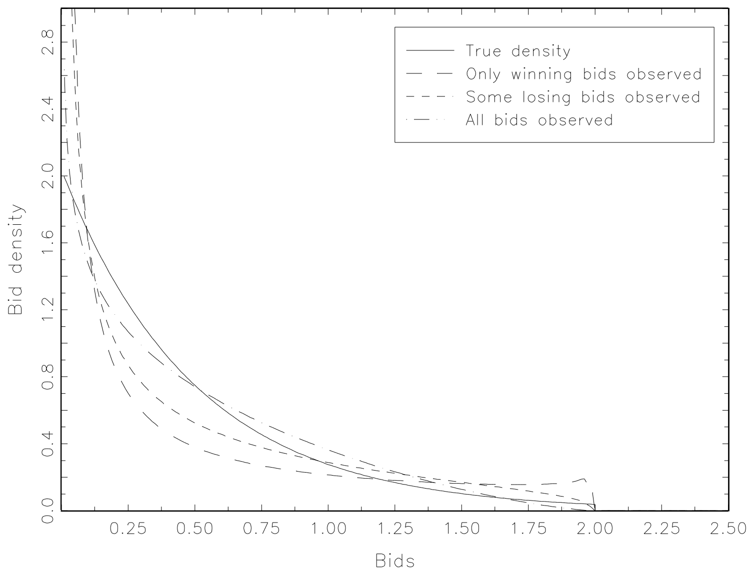

Figure 12: The true density and estimated densities of the simulation study with a truncated exponential distribution.

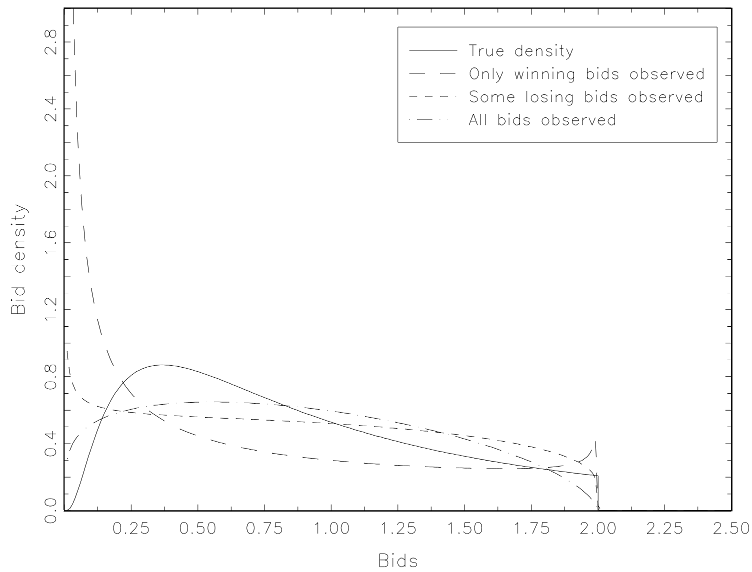

Figure 13: The true density and estimated densities of the simulation study with a truncated exponential distribution.