Structural Empirical Analysis of Dutch Flower Auctions

Gerard J. van den Berg *

Bas van der Klaauw †

April 2000

PRELIMINARY

Abstract

This paper analyzes Dutch auctions of houseplants at the flower auction in Aalsmeer, The Netherlands. We perform a structural empirical analysis of the independent private values model. Our dataset is unique for Dutch auctions in the sense that it includes observations of all losing bids in a time interval next to the winning bid. The length of this interval is determined by the speed of reaction of the auction participants. To investigate the amount of information in the observations of the losing bids, a simulation study is performed. The model is analyzed using the Gibbs sampler with data-augmentation. The results of the structural empirical analysis are used to investigate whether the actual reservation prices set by the auctioneer are optimal.

*Free University Amsterdam, Tinbergen Institute and CEPR.

†Free University Amsterdam and Tinbergen Institute.

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

Keywords: first-price auctions, private values, speed of reaction, observing losing bids, data-augmentation, Gibbs sampling.

We are grateful to the Aalsmeer Flower Auction and to Dirk Hogervorst, Chiel Post, Eric van 't Boveneind and Andre van Kruijsse in particular for kindly providing the data and for useful discussion. We thank Douwe-Frits Broens, Jos Stins of the Faculty of Human Movement Sciences of the Free University and Eric van Damme for useful comments.

1 Introduction

At Dutch auctions, the price falls from an initial high price until a bidder stops the auction. This bidder then obtains the object against the price at which he stopped the auction. An advantage of the Dutch auction mechanism is that it works relatively fast. It is therefore often used to sell large amounts of perishable goods, such as flowers, vegetables and fruits. In this paper we study the Aalsmeer Flower Auction (AFA), at which almost 1200 transaction per auctioning clock can take place within one hour.

In The Netherlands the flower industry is an important sector of economic activity. The Dutch flower auction in Aalsmeer is the largest auction of ornamental plant products (henceforth called 'houseplants') in the world. Approximately 45% of the international export of flowers and plants takes place at the AFA. Daily, 18 million flowers and 2 million plants are auctioned, resulting in around 50,000 transactions. There are 13 auctioning clocks in 5 separate halls. The largest of these halls contains 650 seats for bidders. The products are auctioned in lots, which are defined as the total supply of a given homogeneous article of a given grower on a given day. A lot consists of a number of units, which are a fixed number of plants. The lots are auctioned sequentially, implying that the winner decides how many units he wants to buy. If the number of units in the lot exceeds the number of units bought by the winner, the remaining units are auctioned in the same way. Again starting at the initial high price. This process stops if either the complete lot is sold or the price decreases the fixed reservation price. In the latter case all remaining units are destroyed. In The Netherlands, reservation prices for agricultural product already exist since 1933 (see Broens and Meulenberg, 1999). At the beginning of the year, AFA announces the reservation price for each plant type. In 1996, only two levels were used (25 cents and 50 cents).

In this paper we perform a full structural analysis of the independent private values (IPV) model of Dutch auctions of houseplants. The data used in this paper are from an administrative database of the AFA. This database covers all auctions of houseplants during 14 days at the end of August and the beginning of September 1996. From this database we extracted all auctions in which the minimum quantity to buy equals the total number of units in the lot. We impose this requirement to avoid that bidders anticipate on a sequence of auctions necessary to sell the complete lot.

By now, there is a substantial literature on the structural estimation of IPV models. For example Elyakime, Laffont, Loisel and Vuong (1994) choose this approach to investigate the auctions of timber, Laffont, Ossard and Vuong (1995) to study a market of agricultural goods and Paarsch (1997), who also uses data from timber sales. Laffont (1997) gives for a survey on structural estimation of such models and Wolfstetter (1996) provides a recent survey on the theory of auction models.

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. Intuitively, it may be clear that it is difficult to identify the whole distribution of private values from only the winning bid. Moreover, in most specifications of first-price auction models the support of the distribution of bids depends on all structural parameters. This implies that in the classical structural empirical analysis not all standard regularity conditions are fulfilled to be able to use the standard asymptotic properties of Maximum Likelihood estimation (e.g. Donald and Paarsch, 1993a, and Laffont, Ossard and Vuong, 1995).[1]

From an efficiency point of view, it is generally unattractive to ignore such information (see Donald and Paarsch, 1993a, 1993b, Hong, 1998, and Paarsch 1994). Within the classical framework, it is most straightforward to optimize the loglikelihood function subject to the constraints on the support of the parameter space as is proposed by Donald and Paarsch (1993b). However, this implies maximizing a nonlinear function subject to nonlinear constraints, which generally raises computational difficulties. The asymptotical properties require the evaluation of the loglikelihood function at all possible cross-points of the constraints in the parameter space. In general, when the number of parameters is not very low, for example because some explanatory variables are observed, the computational requirements rise very fast with the number of observations.

In this paper, we deal with these problems. First, the database is unique in a sense that it does not only register the price at which the auctioning clock is stopped, but also all bids of bidders who pushed the button to stop the clock in a short time interval after the winner of the auction. All bids made in a time interval

[1]To overcome this difficulty, a number of alternative estimation schemes are available. For example, Laffont, Ossard and Vuong (1995) use simulated nonlinear least squares estimation and Paarsch (1992) uses ordinary nonlinear regression. An alternative approach is to introduce measurement errors to avoid the support problems. Also estimation methods specific to auctions and related models have been derived. Elyakime, Laffont, Loisel and Vuong (1994) use a two step procedure. They first estimate the bid density function, which is used to construct a set of pseudo valuations. The underlying valuation function is estimated using the set of pseudo valuations. Also based on this principle is the nonparametric estimation method of Guerre, Perrigne and Vuong (1998). However, these methods can only be applied if all bids are observed, this makes nonparametric estimation nonfeasible in the context of Dutch auctions.

until one second after the winning bid are registered. To distinguish between actual losing bids and 'fake' losing bids, we use in the empirical analysis only the losing bids made within the first 0.2 seconds after the winning bid, which seems to be a reasonable approximation of the speed of reaction of bidders. Compared to other studies of Dutch auctions our data contain additional information on a small part of the distribution function of the bids. Second, we use a sampling Bayes approach. In particular, we use the Gibbs sampler to generate the posterior distribution of the parameters (see Casella and George, 1992, and Gelfand and Smith, 1990). Because inference on auctions is simple and straightforward in case we observe the private values of all bidders, we use data-augmentation methods to construct these (latent) private values (see Tanner and Wong, 1987). Lancaster (1997) shows that this method of empirical analysis performs well in case of the structural evaluation of job search models. Bontemps, Robin and Van den Berg (2000) show a strong similarity between job search models and IPV models of first-price auctions, and their empirical inference.

Finally, we use the results from the structural analysis to investigate the effect of the reservation prices on bids and the expected revenue of the seller. At the AFA, it is commonly felt that the reservation prices are still too low and that more differentiation can increase profits. In particular, the prices of the larger and more expensive plants are still far above the reservation price. It is therefore not very likely that the reservation price affects the bids.

The outline of this paper is as follows. In Section 2 we give a description of the AFA. Section 3 discusses first-price auctions and method for the structural empirical analysis. In Section 4 we give an overview of the data. We present some summary statistics and perform reduced-form analyses on the price at which the flowers are sold. Section 5 shows the results of a simulation study performed to clarify the issue of identification of the model. The results of the structural analysis are presented in Section 6. Section 7 concludes.

2 The Aalsmeer Flower Auction

2.1 General statistics

In this section we give some general statistics concerning the AFA, and we provide details on the actual auctioning process. Most of the information on the general statistics is from the Annual Reports of the AFA recent years (see e.g. Bloemenveiling Aalsmeer (1994)).

The AFA is located in Aalsmeer, close to Amsterdam, The Netherlands. It is the largest auction of ornamental plant products (cut flowers, houseplants, gardenplants, etcetera) in the world. The current annual turnover exceeds 3 billion Dutch guilders (approximately 1.5 billion US Dollars). The AFA is a cooperative owned by about 4000 Dutch growers of the auctioned products. The magnitude of the AFA reflects the importance of the market for ornamental plant products for the Dutch economy. Indeed, The Netherlands is the world's leading producer and distributer of cut flowers, and they are The Netherlands' most important export product. The AFA itself employs around 1800 workers, but on a given day almost 10,000 individuals do their work in the auction buildings (the latter number includes suppliers and buyers).

To give some further indication of the size of the AFA, the current total annual supply consists of approximately 4.3 billion single flowers, 330 million houseplants and 150 million gardenplants. The current annual import includes 1.8 billion single flowers. Of these, the largest shares are supplied by Israel, Kenya and Spain. The value of the current annual export of flowers and houseplants equals 4.4 and 1.8 billion guilders, respectively. For flowers, Germany, France and the United Kingdom are the most important markets, while for houseplants these are Germany, France and Italy.

Of the total number of 7100 growers participating in the auctions, almost 1500 are from abroad. The total number of buyers equals 1700. The dispersion of their shares in total turnover is enormous. On the one hand, about 50 buyers each buy for more than 10 million guilders a year; together this amounts to around 50% of total turnover. On the other hand, about 725 buyers each buy for less than 0.1 million guilders a year; together this is less than 1% of total turnover. These two extremes basically correspond to big exporting companies and small domestic retail shops, respectively.

2.2 Institutional features of the auctioning process

The AFA uses the Dutch auction to sell lots. The wall in front of the auctioning room contains a large board with a clock and an electronic display of properties of the product to be auctioned (identity of the grower, name of the product, various quality indicators, length of the stem in case of flowers and size of the flower pot in case of plants) as well as properties of the setup of the auction (monetary unit, minimum price, possibly a minimum purchase quantity). 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 flowers some time before the actual auctioning). The auctioneer decides on a starting position for the clock hand which corresponds to an unreasonably high price for the product. He then sets the clock in motion. The value pointed at by the clock hand[2] drops continuously until a buyer stops the clock by pushing the button in front of him. The value pointed at by the clock hand at that moment is the price to be paid by that buyer for a single item. The buyer then announces how many "units" he wants to buy. A "unit" is defined as a fixed amount of single items (e.g., for a particular type of houseplant, the definition of a unit is fixed). The identity of the buyer is shown on the electronic display in front of the room. If the number of units he buys falls short of the supplied number of units then the clock is reset to a very high value, and the process restarts for the left-over units. The auctioneer may decide to stipulate a different minimum purchase quantity than before. This goes on until the whole lot is sold. If the hand of the clock passes the minimum price then the remaining lot is destroyed. Every lot is auctioned in this manner.

The minimum price for a given product is fixed throughout the year (at least, for the time periods from which our data are). For example, for houseplants, the minimum price in 1996 was 25 cents per single plant. The minimum prices are published in the annual codebook which is distributed among buyers and growers (see e.g. Bloemenveiling Aalsmeer, 1996).

Now let us go back one step and consider how the AFA chooses the order of the auctioning of different lots. The AFA uses the term "auction group" to denote a group of products with similar features. The sequence in which auction groups appear at the auction is the same on every day. However, the sequence in which different lots within an auction group appear at the auction is randomized.

The AFA buildings contain four auction rooms. The total number of clocks equals 13. 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 just a couple of seconds. The average number of transactions per day at the Aalsmeer Flower Auction equals about 30,000.

[2]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.

3 First-price auctions

3.1 The theoretical model

In this section we discuss both the economic theory of first-price or so-called Dutch auctions and the Bayesian approach we use to analyze the structural model. We start with a brief overview of the IPV first-price auction model. After that we discuss the Gibbs sampler with data-augmentation, which we use to evaluate the posterior density of the parameter in the model. We end this section with the parameterization of the model.

In a non-experimental setting the Dutch auction is usually the less informative of the four standard auction types. The second-price sealed-bid (or Vickrey) auction is the most informative, i.e. optimal behavior implies bidding the private valuation. Each bidder reveals his private value to the auctioneer. In the English auction the optimal strategy of a bidder is to continue bidding until his private value is reached. In this type of auction the private value of each bidder, except for the winner, is revealed. In both the Dutch and the first-price sealed-bid auction, the equilibrium strategy of a bidder is less easy. To form an optimal bid bidders have to shade their valuations. It is important to stress that the Dutch auction and the first-price sealed-bid auction are strategically equivalent. The amount at which they shade their valuation depends on the behavior of all other bidders. If a bidder shades his valuation with a large amount it decreases the probability to win the auction. However, if he wins the auction, he pays a low price. The equilibrium in these two types of auctions is a symmetric Bayesian-Nash equilibrium. While in a first-price sealed-bid auction the bids of all bidders are revealed to the auctioneer, in a Dutch auction usually only the highest bid is observed.

Consider the case in which a single indivisible object is auctioned. The rules of first-price auctions are simple, the highest bidder obtains the object and pays his bid. To establish the optimal behavior of bidders we use the standard IPV model in which both the seller and the potential buyers are risk neutral. Van den Berg, Van Ours and Pradhan (1999) find no evidence that bidders at the AFA are not risk neutral when studying sequences of auctions of roses.[3] In the remainder, we only limit the discussion to the main results. More extensive discussions of the model can for example be found in surveys by McAfee and McMillan (1987a)

[3]The main argument for risk neutrality is that most bidders do not face strong binding financial constraints, as 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 kind of insurance against losing an auction.

and Wolfstetter (1996).

The IPV model suggests that each bidder exactly knows his valuation of the object to auction and that his valuation is strictly private. Moreover, each bidder knows that the valuations of the other bidders are random draws from a common known distribution function. It seems that the IPV framework is a reasonably accurate description of the valuations of bidders at AFA, as 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 experience, they have an excellent knowledge of the demand functions of the products they sell to the consumers, and these functions 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 highly perishable goods, so there is no scope for extensive re-trading after the auction is held.

The alternative to the IPV model is the common value (CV) model. According to this model, the object to auction has a unique true value similar for all bidders and each bidder gets a private signal, which is used to make an estimate of the true value. More general is the affiliated values (AF) model which nests both the IPV and the CV model (see Milgrom and Weber, 1982). However, as shown by Laffont and Vuong (1996), a structural model based on the AF model is generally unidentified. Milgrom and Weber (1982) suggest that for nondurable consumer goods, like plants, the IPV model suites better than the CV model, which in turn is more convenient if (government) contracts are auctioned.

A seller valuates the object to auction as r (for simplicity we assume $r = 0$). The seller organizes an auctions to maximize the price paid for the object and decides only to sell the object if the price is above some critical value or reservation price v_0. This reservation 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 the AFA, these plants are destroyed.

Suppose there are $n \geq 2$ potential buyers, denoted by $i = 1, \ldots, n$, who are identical ex-ante. Buyer i valuates the object as v_i, which is only privately observed. As the seller does not observe any of the buyers' valuations, the seller considers v_i, $i = 1, \ldots, n$ as independent draws from the same continuous distribution function $F : [\underline{v}, \overline{v}] \longrightarrow [0, 1]$, where $\underline{v}$ can be 0 and $\overline{v}$ can be ∞. The buyers consider the valuations of the other buyers as random realizations from the distribution function $F(\cdot)$.

A bidder only participates in the auction if his valuation of the object is above the reservation price of the seller. This is denoted by the participation function $\xi(v)$, which takes the value 1 if $v \geq v_0$ and 0 otherwise. In case the bidder does not participate in the auction we consider his bid as being equal to 0. Conditional on the number of potential bidders in the auction and the reservation price, the optimal bid of a buyer with valuation v equals

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

Each bidder participating in the auction thus shades his valuation with the amount $\int_{v_0}^{v}\left[\frac{F(x)}{F(v)}\right]^{n-1} dx$. This decreases if the number of bidder n increases or the reservation prices increases. However, increasing the reservation price also increases the probability that none of the bidders has a private valuation above the reservation price, which implies that the object is not sold.

It is easy to show that the expected revenue of the seller is given by

$$\pi(v_0|n) = n\int_{v_0}^{\overline{v}} \left(vf(v) - (1 - F(v))\right) F(v)^{n-1} dv$$

(see Wolfstetter, 1996). Once the seller decides on the type of auction to use, the only remaining 'policy' instrument of the seller is the reservation price v_0.[4] The expected revenue of the seller is optimal if the reservation price satisfies the first-order condition

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

and the second-order condition

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

It should be stressed that the expressions for the optimal reservation price are independent of the number of potential buyers. The reservation 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 reservation price equals the lower bound of the support of the private valuations, $v_0 = \underline{v}$. The uniqueness of a reservation price in the support of v depends on the shape of the distribution function $F(\cdot)$.

[4]If all bidders are risk neutral and behave optimally, the revenue equivalence theorem proves that the four most often used auction mechanisms (Dutch, English, first-price sealed-bid and second-price sealed-bid) generate the same expected revenue (see for surveys McAfee and McMillan, 1987a, and Wolfstetter, 1996).

3.2 The empirical model

We start this subsection with a brief outline of the type of data we have. Our database contains information about auctions of houseplants at the AFA. For a single auction we observe the usual information revealed in Dutch auctions, the winning bid, the reservation price and the actual number of bidders. In addition, we observe all losing bids slightly below the winning bid. If the winning bid equals $b_{(1)}$, we observe this bid and all bids made in the interval $(b_{(1)} - \delta, b_{(1)})$, for some known value $\delta > 0$. All remaining (unobserved) bids are smaller than $b_{(1)} - \delta$. We deal with the length of the interval δ in Subsection 4.2. The reservation prices are very low at the AFA. In general, we never observe that complete lots are destroyed. Therefore, we assume that the lower bound of the support of the distribution function of private values $\underline{v}$ equals the reservation price v_0. Hence, the potential number of buyers is similar to the observed number of participants in an auction.

It is important to stress that the structural analysis is straightforward in case we would observe all private values. However, the private values are latent as we observe bids. But, as shown in the previous subsection there is a one-to-one relation between the buyer's private value and his actual bid. The relation is given in equation (1), $b = \beta(v)$, implying that given the bid and the shape of the distribution function of private values we can compute the corresponding private value. For many specifications of the distribution function of private values, the support of the distribution function of bids is bounded, and the upper bound depends on the structural parameters (see e.g. Laffont, Ossard and Vuong, 1995). Because bounds are very informative, from an efficiency point of view it is attractive to use inference that uses the information captured in the upper bound. In general, 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).

We use a sampling Bayes approach. 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. Since we compute the exact posterior distributions of the parameters, we do not have to rely on asymptotics. It is common knowledge that the results (i.e. the shape of the posterior distribution) are not very robust as they are extremely sensitive to outliers in the data. Outliers can be caused by measurement errors or extreme behavior. In our setting, observing extreme outliers is not very likely. Our data are from an administrative database. Furthermore, all bidders are very experienced in the auctioning process, which most likely excludes unusual high bids.

We assume that all auctions are independent. This implies that before each new auction each bidder draws a new private value independent of private values drawn at earlier auctions. Not every auction of houseplants in the data is identical ex-ante. Typically the distribution of private values $F(\cdot)$ is unknown and differs 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 $F(\cdot)$ can be uniquely characterized by a vector of unknown parameters θ and the set of known covariates x, $F(\cdot) = F(\cdot|x, \theta)$. To achieve nonparametric identification it is necessary to observe the bids of all bidders in every auction (see Guerre, Perrigne and Vuong, 1999). Our data do not meet this requirement, which requires us to make some parametric assumptions, like in Donald and Paarsch (1996) and Laffont, Ossard and Vuong (1995). Under this assumption and some regularity conditions concerning continuity of the distribution of valuations, Donald and Paarsch (1996) establish (parametric) identification in case only the winning bid is observed. We also observe some losing bid and thus the data reveal a part of the distribution function of private values. Consequently, we observe more information than necessary to achieve parametric identification, but our data are not sufficient informative for nonparametric identification. We return to this issue in Section 5, when we investigate the identifying power of observing the losing bids within a simulation study.

We only observe bids in an interval close to the winning bid. For the unobserved bids we know that these are below the lower bound of the interval, $b_{(n)} - \delta$. We use data augmentation methods to sample the latent private values from the distribution function $F(v|v \leq \beta^{-1}(b_{(n)} - \delta))$ (see Tanner and Wong, 1987). Having sampled a set of private values for all bidders in each, we can simply evaluate the distribution function $F(v|\theta, x)$. Using the Gibbs sampler we sample a new set of values for the parameter vector θ. To do so we need to specify a prior distribution of the set of parameters. We use a noninformative prior, which we return to in the next subsection. When sampling the new vector of θ, we have to ensure that this lies within the feasible parameter space bounded by the observed bids. In particular, after sampling a new set of values for θ, the observed bids must be within the support of the bid distribution function. Using the method of Gibbs sampling with data augmentation we can construct a Gibbs sequence of values for the vector of parameters θ, which we use to evaluate the (marginal) posterior distribution (see Casella and George, 1992, and Gelfand and Smith, 1990).

Now consider the support problem that arises because the support of bid distribution function is bounded (even if the support of the distribution function of the private valuations is not bounded). The lower bound of the support is v_0, which is known. The upper bound equals $\beta(v)$ in the limit v to ∞,

$$\begin{aligned}
\lim_{v\to\infty}\beta(v) &= \lim_{v\to\infty} v - \int_{v_0}^{v}\left[\frac{F(x)}{F(v)}\right]^{n-1} dx = \lim_{v\to\infty} v_0\frac{F(v_0)^{n-1}}{F(v)^{n-1}} + \frac{\int_{v_0}^{v} xf(x)F(x)^{n-2}dx}{F(v)^{n-1}} \\
&= v_0F(v_0)^{n-1} + \int_{v_0}^{\infty} xf(x)F(x)^{n-2}dx \\
&= \mathrm{E}\left[\max(V_{n-1}, v_0)\right]
\end{aligned}$$

where V_{n-1} is the largest order statistic of $n-1$ draws from $F(\cdot)$ (see also Laffont, Ossard and Vuong, 1995). This expectation depends on all structural model parameters. The constraint is slightly simplified because the lowerbound of the support of the distribution function of private values is assumed to equal v_0. Thus by definition V_{n-1} exceeds v_0. Because the bids are ordered, it is sufficient to impose the restriction

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

When sampling parameters we have to take into account that a bid can never exceed $\mathrm{E}\left[V_{n-1}\right]$.

3.3 Parameterization

In the IPV first-price auction model with risk-neutral bidders as described above, the only unknown component of the model is the distribution function of private values. We assumed that this (parametric) distribution function has the following properties: (*i*) all heterogeneity between auctions is covered by a vector of observed characteristics x, (*ii*) the density function is continuous and (*iii*) the lower bound of the support equals to the reservation price v_0. We take the distribution function of private values to be a transformed beta distribution, with density function

$$f(y|a,b,c) = \frac{(y/c)^{(\alpha-1)}(1-y/c)^{(\beta-1)}}{cB(\alpha,\beta)} \qquad 0 \leq y \leq c \tag{3}$$

where α, β and c are unknown parameters and $B(\alpha,\beta)$ is the beta function which ensures that the density integrates to 1. The density function is transformed such that it has support from v_0 up till an unknown finite upper bound (the length of the support equals c). We rewrite y as $y = v - v_0$. We allow for heterogeneity between auctions of different plants by allowing the upper bound of the support to depend on the observed characteristics of the auction. We specify c as $c = \exp(x'\gamma) - v_0$. Under these assumptions the valuations v have support on $(v_0, \exp(x'\gamma))$.[5]

The beta distribution is relatively flexible. The density is symmetric if $\alpha = \beta$. The uniform distribution is a special case ($\alpha = \beta = 1$). The shape at the lowerbound of the support of the density function is determined by α and at the upperbound by β. Close to the lowerbound the density increases (decreases) if $\alpha > 1$ ($\alpha < 1$). Similar $\beta < 1$ ($\beta > 1$) implies that the density increases (decreases) close to the upperbound. Subsequently, if both α and β are smaller than 1 the density is U-shaped, and hump-shaped if these both parameters exceed 1. The density is strictly increasing (decreasing) if $\alpha > 1$ and $\beta < 1$ ($\alpha < 1$ and $\beta > 1$). The expectation equals $c\alpha/(\alpha+\beta)$ and the variance $c^2\alpha\beta/((\alpha+\beta)^2(\alpha+\beta+1))$.

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 $c = c^*(\alpha+\beta)$) and the exponential distribution ($\beta \to \infty$, $\alpha = 1$ and $c = c^*(\alpha+b)$) (see McDonald, 1984). McDonald (1984) also specifies a generalized beta distribution by adding an additional parameter. Obvious this allows for more flexibility. As an alternative one could also consider flexible densities based on polynomials (e.g. Gallant and Nychka, 1987). An application of the beta distribution is given by Heckman and Willis (1977) who use this distribution to estimate female labor force participation probabilities

The noninformative prior of the vector of parameters is

$$p(\alpha, \beta, \gamma) = \frac{1}{\alpha\beta}$$

Specifying the private values distribution function to be a beta distribution function and choosing a noninformative prior, we use the Gibbs sampler to sample a sequence of values from α, β and γ. So given some values β_t and γ_t and a set of simulated private values v_t, we can sample a new value α_{t+1} from the conditional density function

$$f(\alpha_{t+1}|\beta_t, \gamma_t) = C\left(\prod_{i=1}^{I}\prod_{j=1}^{b_i} \frac{\Gamma(\alpha_{t+1}+\beta_t)(v_{jit}-v_0)^{\alpha_{t+1}-1}}{\Gamma(\alpha_{t+1})(\exp(x_i'\gamma_t)-v_0)^{\alpha_{t+1}-1}}\right)\frac{1}{\alpha_{t+1}}$$

where C is some constant ensuring that the density integrates to one, and where

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

$\Gamma(\cdot)$ is the gamma function. We do the same for β and γ

$$f(\beta_{t+1}|\alpha_{t+1},\gamma_t) = C\left(\prod_{i=1}^{I}\prod_{j=1}^{b_i}\frac{\Gamma(\alpha_{t+1}+\beta_{t+1})(\exp(x_i'\gamma^t)-v_{jit})^{\beta_{t+1}-1}}{\Gamma(\beta_{t+1})(\exp(x_i'\gamma_t)-v_0)^{\beta_{t+1}-1}}\right)\frac{1}{\beta_{t+1}}$$

and

$$f(\gamma_{t+1}|\alpha_{t+1},\beta_{t+1}) = C\left(\prod_{i=1}^{I}\prod_{j=1}^{b_i}\frac{(\exp(x_i'\gamma_{t+1})-v_{jit})^{\beta_{t+1}-1}}{(\exp(x_i'\gamma_{t+1})-v_0)^{\alpha_{t+1}+\beta_{t+1}-1}}\right)$$

By sampling from these densities we can sample Gibbs sequences $\alpha_0, \beta_0, \gamma_0, \ldots, \alpha_T, \beta_T, \gamma_T$, which we can use to evaluate the marginal posterior densities (see Gelfand and Smith, 1990).

Having determined the full specification of the model, we can focus on the importance of imposing reservation prices. We investigate this by comparing the expected revenue in case the reservation price is optimal and the expected revenue if there is no reservation price, i.e. we focus on the percentage increase in expected revenue when moving from an auction without reservation price to an auction with an optimal reservation price. In fact is suffices to examine a simple numerical example Because the expected revenue depends on the (potential) number of bidders participating in the auction in a nonlinear way, the percentage increase in revenue is a function of the number of bidders as well.

The setup we choose is as follows. We consider a beta distribution function with support on 25 cent to 100 cents. Given the values of α and β we can compute for a given number of bidders the percentage increase in expected revenue when increasing the reservation price from 25 cents to the optimal reservation price. Figures 1 to 3 show for α equals to 0.5, 1 and 2, respectively, what the maximum percentage increase in expected revenue is if β is also 0.5, 1 and 2. The effect of imposing a reservation price depends very much on the value of the parameters (shape of the density). Note that if $\alpha = 0.5$ and β equals 1 or 2, the optimal reservation price is equal to 25 cents and thus increasing the reservation price has a negative effect on the expected revenue. The percentage increase in expected revenue never exceeds 10%. In general, the effect of imposing a reservation price is higher if β is small, which implies that the density function is increasing close to the upper bound of the support. Reservation prices are particularly effective if the density is U-shaped (both α and β are smaller than 1). Reservation prices are only important if the number of bidders is low. Only in case of extreme values of the parameters of α and β imposing a reservation price when the number of bidders is above 5 generates (small) extra expected revenue. In general, when the number of bidders is higher than 5 reservation prices are not a very efficient instrument for generating additional expected revenue. In general, the impact of the reservation price on the expected revenue of the seller converges very fast to 0 as the number of bidders goes to infinity.

4 Description of the data

4.1 The data set

In this section we give an overview of the data we use in the empirical analysis. We start this section with an outline of the database. In Subsection 4.2 we discuss the observed losing bids. Finally, the results of some reduced-form analyses on the winning prices are presented in Subsection 4.3.

The results on optimal bidding behavior and the optimal reservation price presented in Subsection 3.1 largely depend on the assumption that the object to auction is unique and indivisible. However, for the estimation of the model a sequence of observations in necessary. More precise, for statistical inference repeated realizations from the same data generating process are required. Therefore, we modify the data such that we obtain a series of auctions of indivisible goods, which are as homogenous as possible. In the ideal case, we observe a particular type of houseplant which always has the same quality and is only supplied by a single grower.

Our database describes auctions taking place in one of the auctioning halls of the AFA. In this particular hall, there are four auctioning clocks on which the auctions of houseplants take place. Just before the auctioning starts, which is half past 6 a.m., bidders enter the 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. 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 participants. 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 participants may cause biases in the empirical analyses. To have a reliable measure of the number of participants in the auction, we restrict the database to auctions taking place during the first hour after the auctioning starts.

In the reduced database around 2000 lots are auctioned consisting of 332 different types of houseplants. We observe 826 lots which have the minimum purchase quantity equal to the size of the lot. This implies that the lots are thus indivisible and that the first winner has to buy all units. By considering only these lots, we avoid that bidders can not anticipate on a sequence of auctions to sell the complete lot and there is no buyer's option. In this restricted database only 25 types of houseplants appear more than 10 times.

From this restricted database we create three different subsamples. These consist of the types of houseplants which are auctioned the most often, which 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'. Furthermore, we only focus on the supply of the 'large' growers of these types. We consider a grower to be a large grower of a given type of houseplant if he supplied this type more than 5 times during the observation period.

Some characteristics of the three subsamples are presented in Table 1. The prices are measured in price per unit. All lots of houseplants in all three subsamples are of the highest quality code. The subsample of Begonias is the largest with 64 observations. The Begonias are also auctioned on the most days, 11 days of the 14 days on which we have data. The Begonias and Compactas are both supplied by 3 growers. Opposite to the subsample Begonias, the name 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. This implies that the grower is very important for the price even is the quality codes are similar. Because both suppliers of Compactas are also suppliers of Camillas, we can see whether growers have a general reputation or a reputation with respect to a particular type of houseplant. The average number of bidders lies around 50 for all subsamples and within each subsample we only observe auctions taking place on Monday until Thursday.

4.2 The losing bids

The database does not only contain information on the winner of the auction, but also some of the losing bids are registered. This is unusual for Dutch auctions. The AFA registers bids made up till one second after the winner stopped the auctioning clock. In this short time period all bids made by the losers of the auction are registered. 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 reservation 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 full database, 94% of the auctions contain information on losing bids. The average number of losing bids per auction equals 4.3, while the maximum number of observed losing bids in a single auction is 26.

In Figure 4 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 5 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.[6] Note that both in the histogram and the kernel densities we observe a relatively large drop around 0.2 seconds.

There exists some psychological literature on the speed at which individuals are capable to stop performing some (planned) actions (see for example Logan and Cowan, 1984). Although none of the experiments we found has the exact design of Dutch auctions, most experiments indicate that the speed of reaction (the time necessary to stop a planned task) lies between 0.25 and 0.3 seconds. Because, the bidders at the AFA are very well trained we choose the cut-off point at 0.2 seconds, which is a arbitrary choice. This means that we consider bids made within 0.2 seconds after the winner as actual losing bids. We ignore all other losing bids.

From now on we only consider losing bids within the 0.2 seconds time interval after the winning bid has been made. In the subsamples of Begonias, Camillas and Compactas 100%, 93% and 97% of the auctions contain information about losing bids. For the Begonias in total we observe 11% of all losing bids and the maximum number of losing bids observed in a single auction equals 13. For Camillas and Compactas these number are 8% and 11, and 7% and 9, respectively.

[6]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).

4.3 Some preliminary analyses

In this subsection we perform some reduced-form analyses to get preliminary insight in the covariates that determine the winning bid at an auction. We regress the number of bidders participating, the grower and the day of the week at which the auction takes place on the logarithm of the price. We use a linear specification and ordinary least squares as the estimation method.

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 and negative for the Compactas, but in all cases insignificant. The economic theory predicts that a higher number of bidders increases bids due to increasing competition. The grower seems to be the most important covariate on the price paid at the auction. 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 has a better reputation than grower 5. Even though the type of houseplants are very strict 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 unambiguous. Whereas, prices for Begonias are lowest on Wednesday, they are highest on this day for Compactas.

5 Simulation study

In Subsection 3.2 we addressed the issue of identification of first-price auctions. On the one hand, Donald and Paarsch (1996) show that parametric identification can be achieved if only the winning bid is observed. On the other hand, Guerre, Perrigne and Vuong (1999) prove that first-price auctions are only nonparametrically identified if the bids of all bidders in the auction are observed. Both Donald and Paarsch (1996) and Guerre, Perrigne and Vuong (1999) assume that also the number of bidders and the reservation price are known.

As described in the previous section, our data are not only informative on the winning bid. Also bids close to the winning bid are observed. This implies that our data are more informative than necessary to achieve parametric identification. However, for nonparametric identification it is not sufficient to observe only part of the losing bids. To investigate the identifying power of these losing bids we perform a simulation study. In the simulation study we do not consider the distribution of private values, instead we focus on the distribution of bids. It may be clear that if the distribution of bids is known, it is possible to derive from this the distribution of private values (this has formally been proven by Guerre, Perrigne and Vuong, 1999).

In the simulation study we set the number of auctions in a sample equal to 50. Furthermore, we take the distribution of bidders in the simulation study approximately the same as in our data sets. We draw the bids from a given distribution function. Thus, we neglect the economic model. After drawing the bids of all bidders we construct three different samples. The first sample only contains the winning bid, the second sample is also informative on the losing bids close to the winning bid and the third sample includes the bids of all bidders. 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 the bids of all losing bidders. Using these samples we estimate the distribution function of bids. We repeat this 100 times. To be as close as possible to the economic model, we take the distribution function of bids to be truncated on the right-tail. We consider three distribution functions from which we draw the bids, (*i*) a uniform distribution function, (*ii*) a truncated exponential distribution function and (*iii*) a truncated lognormal distribution function.

To estimate from the observed data the distribution of bids, we take the beta distribution function (see equation (3)). In Figure 6, 7 and 8 we show the average estimated densities of the simulation studies. Each figures 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 sample containing all bids. The estimated density based on the winning bid and some of the losing bids lies close to the estimated density based on the sample with all bids. It is important to note that the fit improves enormously if we observe only a small part of the losing bids. According to these results it seems hard to approximate an underlying density if one only observes a sequence with highest observations.

6 Structural analysis

In this section we discuss the results of our structural empirical analysis. We start by focusing on the marginal posterior distributions of the parameters in the distribution function of private values. We have computed the posterior distribution both using the information on the losing bids and ignoring this information. After that we consider the reservation prices and the impact of the reservation prices on the expected revenue.

We have sampled a Gibbs sequence consisting of 2600 iterations, i.e. values for the parameter α, β and γ. Two Gibbs sequences are sampled, one using information on the losing bids and one ignoring this information. The first 100 iteration were used as an initial period to reach the equilibrium of the Gibbs sequence. After that we used only every 25$^{\text{th}}$ iteration.[7] To evaluate the marginal posterior distribution functions of the parameters, we have used the marginal densities given in Subsection 3.3. As is shown in Casella and George (1992), the marginal posterior densities can be approximated by

$$\hat{f}(\alpha) = \sum_{t=1}^{T} f(\alpha|\beta_t, \gamma_t)$$

and similarly for β and γ. The marginal posterior densities of the parameters for the auctions of Cammilas are given in Figures 9–15. Each figure shows a graph obtained when using the information on the losing bid and not using this information.

Before discussing the covariate effects, we first focus on the shape of the density function of private values, $f(v)$. The shape is determined by the values of the parameters α and β. In general, in the posterior distribution function the support of α is between 0 and 1 and of β above 2. This implies that the The density function of bids $f(v)$ is downward sloping on the support of v. Many bidders have a private value close to the reservation price. This suggests that only a small proportion of the bidders has actual interest in buying the houseplants.

By comparing Figure 11 with Figure 12, it can easily be seen that the reputation of Grower 4 is better than the reputation of Grower 5. Even though both growers supply plants with the same quality codes, the prices paid for plants supplied by Grower 4 are higher. Obviously, reputation of the grower is important at the flower auction.

On Tuesday higher prices are paid for the Camillas than on the other days of the week. The modes of the marginal posterior distributions for the covaraite effects of auctioning on Wednesday and Thursday are both almost 0. On these days the prices are on average almost the same as on Monday.

6.1 Optimal reservation prices

In this subsection we focus on the reservation prices used at the flower auction. At the moment the data were collected, there was almost no differentiation of reservation prices between different types of plants. The AFA beliefs that the

[7]This is based on analyzing the correlations within the sequence. The correlation between α_t and α_{t-s} is close to 0 for $s \geq 20$. This is also the case for β and γ.

reservation prices are low. This is confirmed by the data, which do not show any winning bids close to the reservation price. Since, we assumed that the distribution function of private values does not have support below the current reservation price, we can only study the consequences of increasing in the reservation prices.[8] But as argued above, the reservation prices are most likely too low to be optimal.

For each of the parameter values obtained in the Gibbs sequence, we can compute the optimal reservation prices. We compare the current situation in which there is only a single reservation prices with a situation where there is full differentiation of reservation prices. This means that we allow for different reservation prices for each grower on each day of the week. As is also stressed in Subsection 3.3 reservation prices are not very effective if $\alpha < 1$ and $\beta > 1$.

Although the reservation price is too low to be optimal, the expected revenue increases with less than 0.1% if the reservation prices are set optimally as compared to the current situation. The main reason for this small increase is that there is a relatively large number of buyers active at AFA. The expected revenue is computed under the assumption that the same number of bidders participate in the auction after an increase in reservation prices.

7 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 the 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 the AFA. Growers, may get different prices for their plants, even if these plants have the same quality code. Although the current reservation prices at the AFA are too low to optimize the expected revenue, increasing the reservation prices hardly

[8]Because the data are collected using auctions with reservation prices, it is impossible to identify the shape of the distribution function of valuation below the reservation price. Without making arbitrary assumptions it is not possible to investigate the consequences of a decrease in reservation prices.

generates any additional revenue. This is mainly caused by the large number of auction participants.

References

Bloemenveiling Aalsmeer (1994), *Annual report* (In Dutch).

Bontemps, C., J-M. Robin and G.J. van den Berg (2000), Equilibrium search with continuous productivity dispersion: theory and non-parametric estimation, *International Economic Review*, forthcoming.

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 (1993a), Piecewise pseudo-maximum likelihood estimation in empirical models of auctions, *International Economic Review* 34, 121–148.

Donald, S.G. and H.J. Paarsch (1993b), Maximum likelihood estimation when the support of the distribution depends upon some or all of the unknown parameters, Working Paper, University of Western Ontario, London, Canada.

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.

Elyakime, B., J-J. Laffont, P. Loisel, and Q. Vuong (1994), First-price sealed-bid auctions with secret reservation prices, *Annales d'Economie et de Statistique* 34, 115–141.

Gallant, A.R. and D.W. Nychka (1987), Semi-nonparametric maximum likelihood estimation, *Econometrica* 55, 363–390.

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.

Guerre, E., I. Perrigne and Q. Vuong (1999), Optimal nonparametric estimation of first-price auctions, *Econometrica*, forthcoming.

Heckman, J.J. and R.J. Willis (1977), A beta-logistic model for the analysis of sequential labor force participation by married women, *Journal of Political Economy* 85, 27–58.

Hong, H. (1998), Maximum likelihood estimation for job search, auction and frontier production function models, Working Paper, Princeton University, Princeton.

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.

Lancaster, T. (1997), Exact structural inference in optimal job-search models, *Journal of Business and Economic Statistics* 15, 165–179.

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

Levin, D and J.L. Smith (1996), Optimal reservation prices in auctions, *Economic Journal* 106, 1271–1283.

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.

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.

Paarsch, H.J. (1992), Deciding between the common and private value paradigms in empirical models of auctions, *Journal of Econometrics* 51, 191–215.

Paarsch, H.J. (1997), Deriving an estimate of the optimal reserve price: An application to British Columbian timber sales, *Journal of Econometrics* 78, 333–357.

Reiss, P.C. and F.A. Wolak (1999), Structural econometric modeling in Industrial Organization, in J.J. Heckman and E.E. Leamer (eds.), *Handbook of Econometrics, Volume 5*, forthcoming North-Holland, Amsterdam.

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

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., J.C. van Ours and M.P. Pradhan (1998), Declining prices in the sequential Dutch flower auction of roses, Working Paper, Tinbergen Institute, Amsterdam.

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

	Begonia	Camilla	Compacta
Number of observation	64	29	30
Number of days	11	8	9
Number of growers	3	2	3
Average price (in cents)	206	172	157
	(21)	(94)	(92)
Minimum price	150	55	53
Maximum price	250	320	320
Average number of bidders	53	53	49
	(13)	(8.2)	(12)

Explanation: Standard deviations are given in parentheses.

Table 1: Some characteristics of the datasets.

	Begonia		Camilla		Compacta	
Number of bidders	0.0017	(0.0012)	0.0068	(0.0047)	−0.0064	(0.0038)
Grower 1	5.25	(0.065)				
Grower 2	5.33	(0.082)				
Grower 3	5.28	(0.067)				
Grower 4			5.04	(0.27)	5.92	(0.22)
Grower 5			3.84	(0.24)	4.36	(0.16)
Grower 6					5.48	(0.21)
Monday	0		0		0	
Tuesday	−0.048	(0.029)	0.14	(0.088)	−0.039	(0.11)
Wednesday	−0.15	(0.061)	0.0086	(0.086)	0.095	(0.092)
Thursday	−0.12	(0.036)	−0.12	(0.099)	−0.022	(0.088)

Explanation: Estimated standard errors are given in parentheses.

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

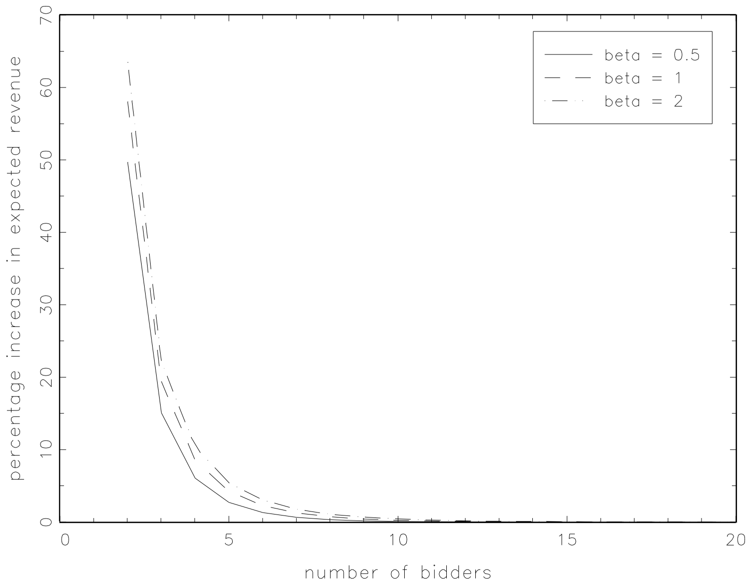

Figure 1: The percentage increase in expected revenue if the auction sets an optimal reservation price instead of no reservation price as a function of the number of bidders ($\alpha = 1/2$).

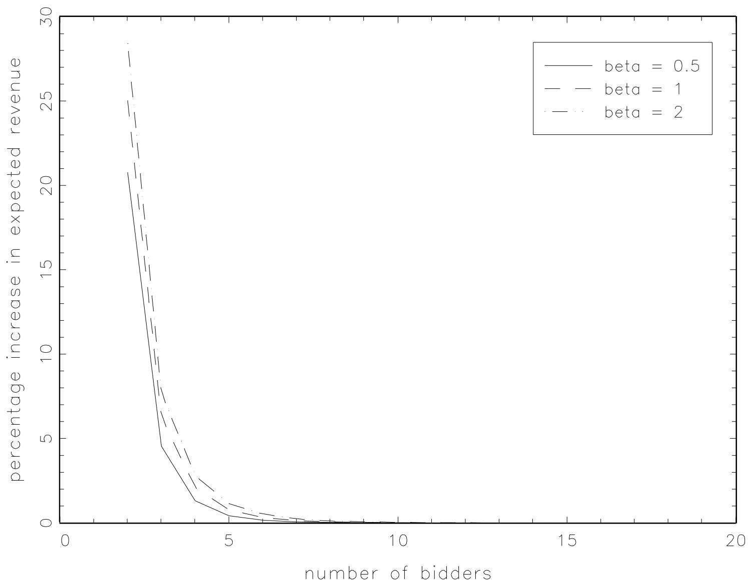

Figure 2: The percentage increase in expected revenue if the auction sets an optimal reservation price instead of no reservation price as a function of the number of bidders ($\alpha = 1$).

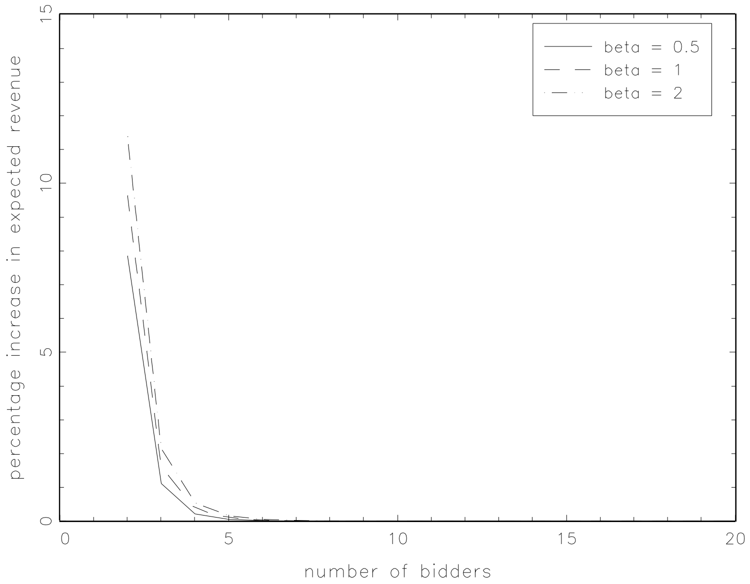

Figure 3: The percentage increase in expected revenue if the auction sets an optimal reservation price instead of no reservation price as a function of the number of bidders ($\alpha = 2$).

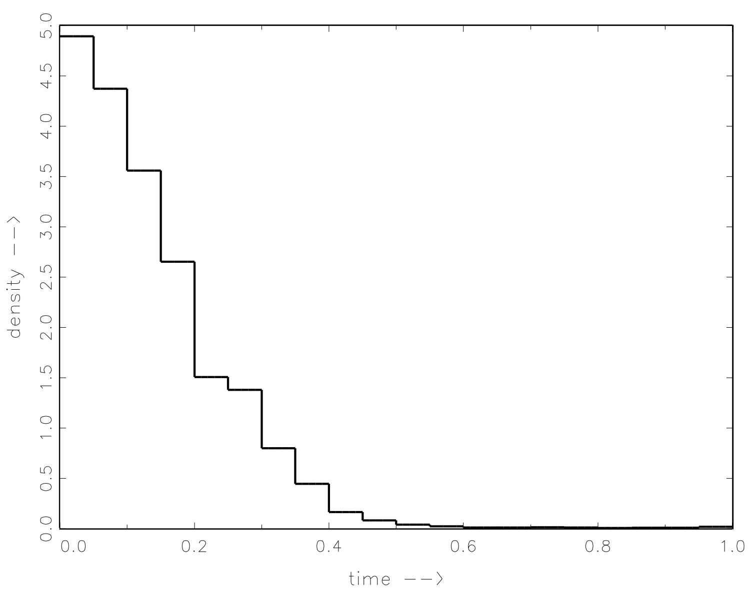

Figure 4: Scaled histogram of the time (in seconds) between the moment that the highest bidder bids and potential other bidders are observed to bid.

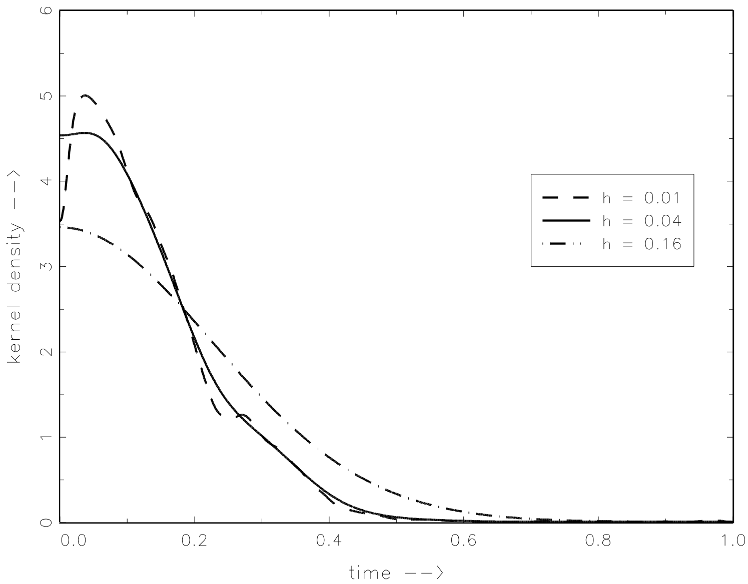

Figure 5: Kernel estimates of the time (in seconds) between the moment that the highest bidder bids and potential other bidders are observed to bid (for different values of the bandwidth (h)).

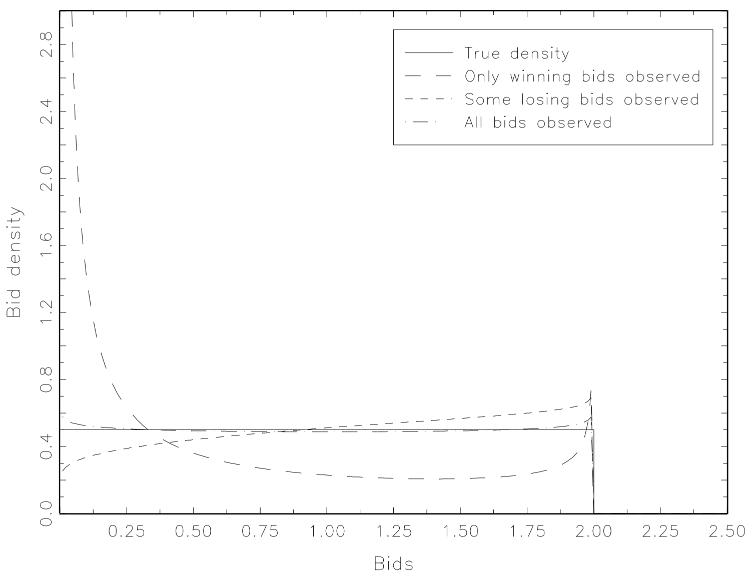

Figure 6: The true density and estimated densities of the simulation study with an uniform distribution.

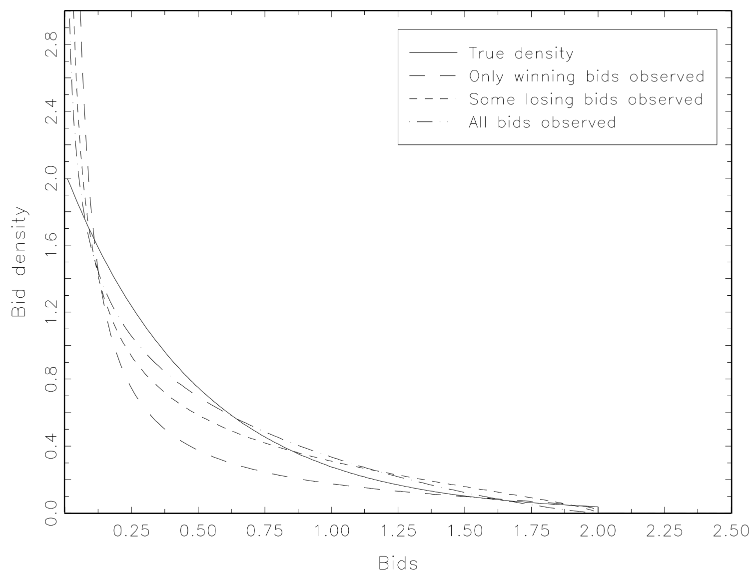

Figure 7: The true density and estimated densities of the simulation study with a truncated exponential distribution.

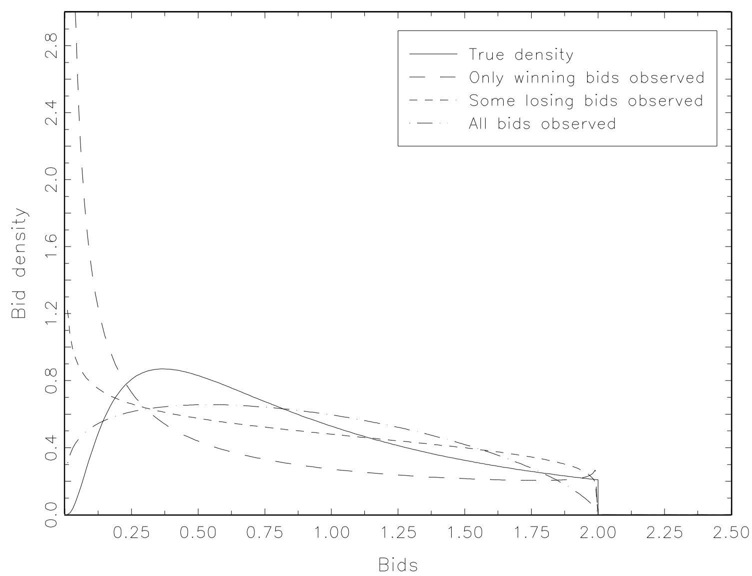

Figure 8: The true density and estimated densities of the simulation study with a truncated lognormal distribution.

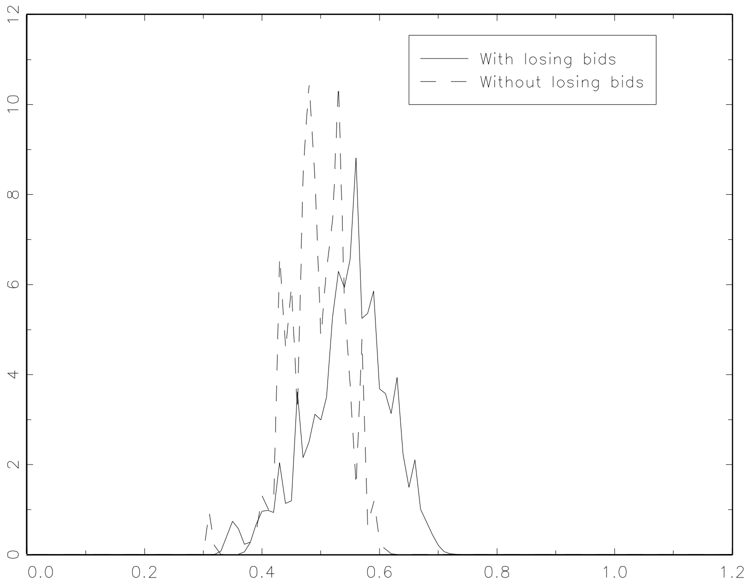

Figure 9: Marginal posterior density for the parameter α in the model of auctions of Camillas.

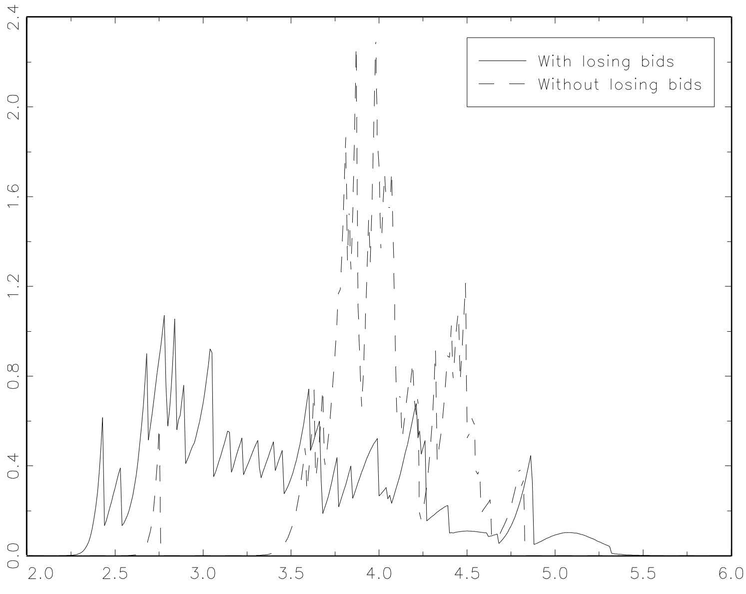

Figure 10: Marginal posterior density for the parameter β in the model of auctions of Camillas.

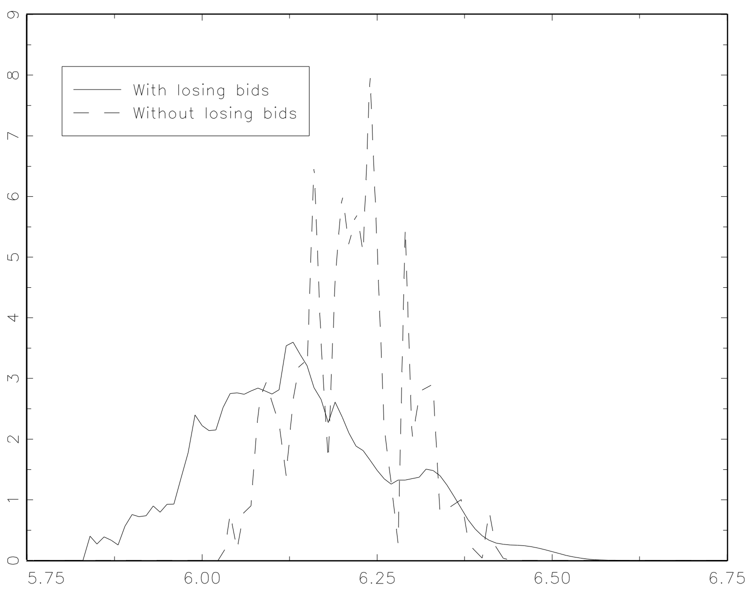

Figure 11: Marginal posterior density for the covariate effect of supply by Grower 4 in the model of auctions of Camillas.

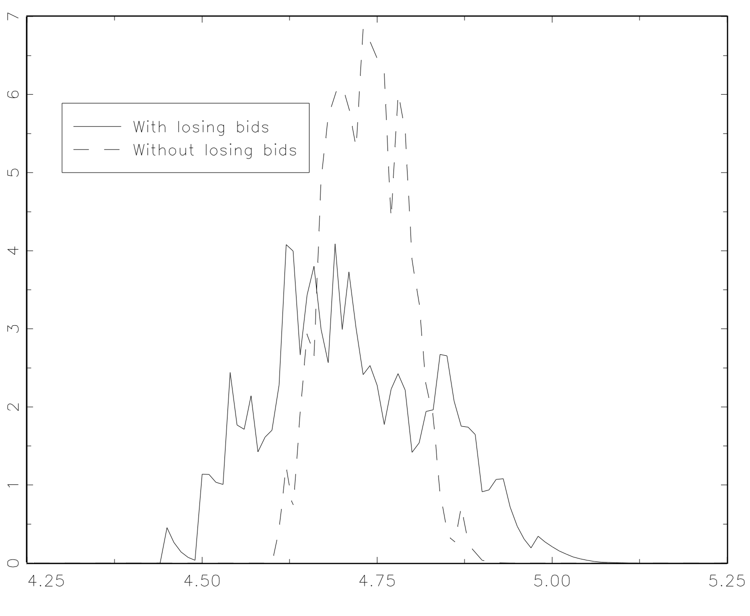

Figure 12: Marginal posterior density for the covariate effect of supply by Grower 5 in the model of auctions of Camillas.

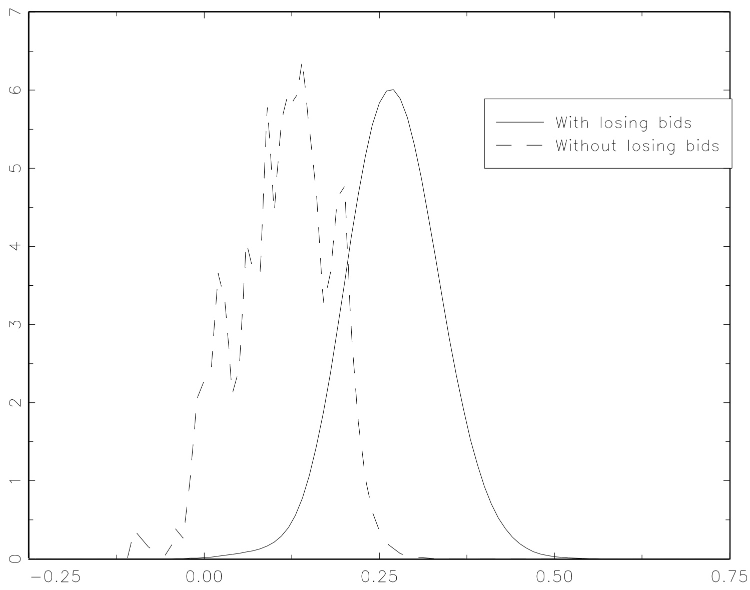

Figure 13: Marginal posterior density for the covariate effect of supply on Tuesday in the model of auctions of Camillas.

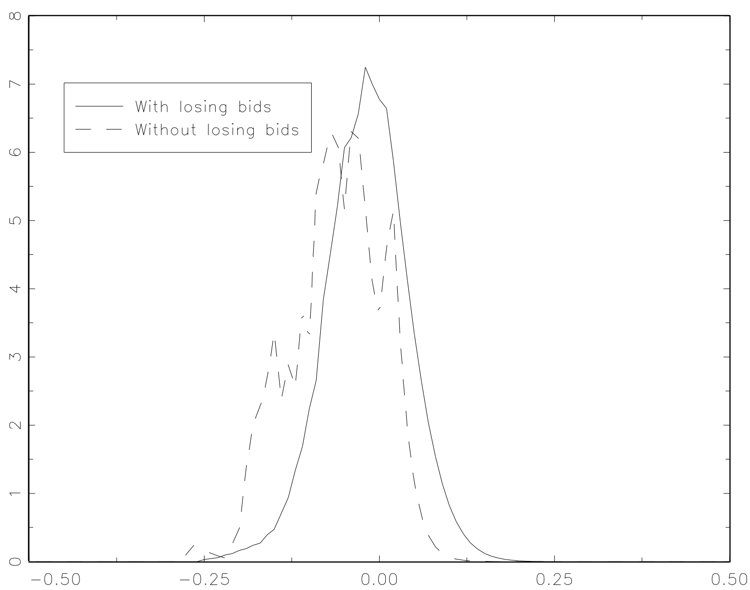

Figure 14: Marginal posterior density for the covariate effect of supply on Wednesday in the model of auctions of Camillas.

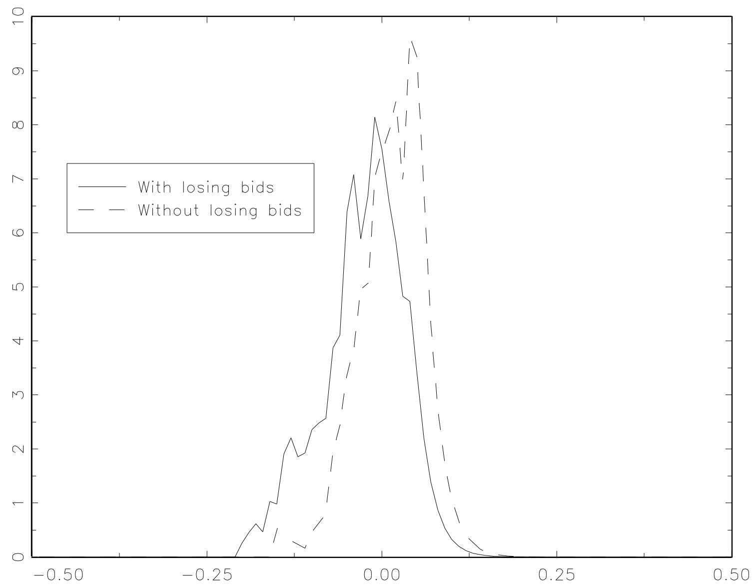

Figure 15: Marginal posterior density for the covariate effect of supply on Thursday in the model of auctions of Camillas.