An Empirical Analysis of Polish Treasury Bill Auctions

Daniel Marszalec[a] *

[a]Department of Economics, University of Oxford,
OX1 3UQ, United Kingdom

Abstract: This paper applies two structural econometric models - the fully non-parametric model of Hortacsu (2002), and the semi-parametric model of Fevrier, Preget and Visser (2002) - to Polish treasury bill auction data. The paper has two main aims: firstly, to compare the performance of discriminatory and uniform-price auctions on the Polish market, and secondly, to evaluate the mutual (in)consistencies of predictions from the two econometric models. I conclude that both of the models considered suggest that the current setup, which uses a discriminatory auction, performs better than a uniform-price auction would, and that both models agree on this conclusion. The number of cases in which predictions from the two auctions are inconsistent is small, though in part this is due to the weakness of the conclusions implied by the Hortacsu model.

As of yet, there are no compelling theoretical models offering an answer as to which auction format is best suited for auctioning shares or treasury-bills. This theoretical ambiguity is mirrored by observing actual auction-mechanism choices across the world: the survey by Brenner, Galai and Sade (2006) covers 48 countries of which 24 use discriminatory price, 9 uniform-price, and further 9 use both. Such variation suggests that the superiority of one auction-format over another is an empirical matter, and results may be country-specific. Though there are a number of structural econometric models analysing the performance of discriminatory and uniform auctions,[2] each of these models has been run on a separate dataset. Hence the degree to which results are model-, rather than data-driven remains an open question.

In this paper I use a dataset on Polish treasury-bill auctions to compare two structural econometric models. The models I compare are Hortacsu (2002), originally run on Turkish auctions, and Fevrier,Preget and Visser (2002), which was initially run on French data. These two models differ both in terms of the economic assumptions that underlie them, as well as the econometric estimation methods that are used: Hortacsu's auction model is based on a private-value assumption, and his estimation procedures are fully non-parametric. Fevrier et al. on the other hand build

*This paper is based on chapter of my D.Phil thesis. I would like to thank my supervisor Paul Klemperer, as well as Steve Bond, Ian Jewitt and Thees Spreckelsen for their comments and criticisms. Financial support from the ESRC and RES is also gratefully acknowledged.

[2]For example Hortacsu (2002) for Turkey, Fevrier & al. (2002) and Armantier & Sbai (2006) for France, Kastl (2006) for Czech Republic and Kim & Ryu (2006) and Kang & Puller (2008) for Korea.

an auction model assuming pure common values, and develop a semi-parametric estimation methodology.

The two main aims of this paper are to, firstly, analyse the Polish dataset using two aforementioned models to determine which auction type is revenue superior, and secondly, to evaluate the mutual consistency of the results that these two models generate. Given the diverging theoretical and methodological assumptions of the two models analysed, if the conclusions from both models are in agreement, that would provide an intuitively compelling reason thinking that the conclusions are data-driven, and not just modelling artefacts. My results suggest that the current auction setup, which uses discriminatory-price auctions, performs no worse than would a uniform-price system. For my application of the Fevrier et al. model, I find that the discriminatory-price auction performs better overall, on both security types. However, in my application of the Hortacsu model, I find significant revenue differences in favour of the discriminatory-price auction in only 10% of my sample, though in no portion of the sample is the uniform-price auction revenue-superior. Thus the results from both model types turn out to be qualitatively consistent.[3]

During the writing of this paper the Polish Ministry of Finance mentioned that they were seriously considering switching from discriminatory to uniform-price auctions, and were interested any academic work that might advise them on this matter. While my work is not explicitly prepared with policy advice in mind, it might nonetheless be a useful starting point. My conclusions, which suggest that the current system using discriminatory auctions is performing at least as well as the alternative would, indicate that if a uniform-price system were to be implemented, other modifications than just the 'auction format' might be necessary to obtain superior revenue results.[4]

In Section 1 I discuss the main differences between the discriminatory and uniform-price auction rules, and in Section 2 I present a generalisation of Wilson's (1979) share auction model, and solve it out under the two differing sets of assumptions from Hortacsu and Fevrier et al. Section 3 describes the institutional setting of the Polish market for treasury bills, and describes the auction participants.[5] In section 4, I discuss how to apply Hortacsu's model to the data, and present the results from estimating his model.[6,7] Section 5, in turn, discusses the estimation of Fevrier et al. model, and presents the results. Comparison of the results from the two models with each other, and past literature, are is carried out in section 6. Section 7 concludes.

1. Description of the Uniform and Discriminatory Pricing Rules

Both the discriminatory and uniform-price auction rules are pricing rules for auctions for multiple items. Frequently the total amount of goods for sale is fixed, but this amount is infinitely divisible, and hence each unit infinitesimally small. In this

[3] That is, both models conclude that discrimnatory-price performs no worse than uniform-price, but the magnitudes in revenue-dominance by the discriminatory auction in those cases when it is 'significantly superior' to uniform price, differ across the two models.

[4] Changing rules for participation, for example, could be one such potential modification.

[5] The feature of 'topup auctions', which is a part of the Polish institutional setup, is discussed in the next Chapter.

[6] Sections 8.2 to 8.1 discuss possible extensions of the Hortacsu methodology, in response to some practical difficulties in applying Hortacsu's method in its basic form.

[7] A discussion of how data from topup auctions can be used to test the validity of estimates from Hortacsu's model has been carried out in Chapter 3 of Marszalec (2011) .

context, a 'bid' submitted by an auction participant is in fact a function, mapping from all possible quantities to prices, or vice versa. In most practical applications of either pricing rule, the auctioneer usually pre-specifies a grid over the set of prices and quantities, and allows the bidders to submit their demands on this grid. For example, the minimum increment on quantity could be 100 bonds, and the degree of precision is specifying a price could be £0.01 to £100 of face-value. Thus rather than continuous demand functions, bidders in these type of auction submit price-quantity pairs, specifying what amount of the good they are willing to buy, at what price.

After receiving all the participant bids, in both auction types, the auctioneer then orders the individual bids in decreasing price order, and picks a cut-off price - usually called the 'stop-out price' - in a way that equates the bidders' demand with the aggregate supply. What differs across the two auction rules, are prices paid by the winning bidders, on the quantities that they do win. In the uniform-price auction, as the name suggests, all winning bidders pay the same per-unit price on all the items they win, and this price equals to the stop-out price. In the discriminatory auction, each accepted bid is paid at full face value, so different bidders will end up paying different prices; indeed, in the case when the same bidder submits multiple winning bids, the price on each of those will differ. A graphical representation of this pricing difference (when demands are approximately continuous) is show in Figure 1.

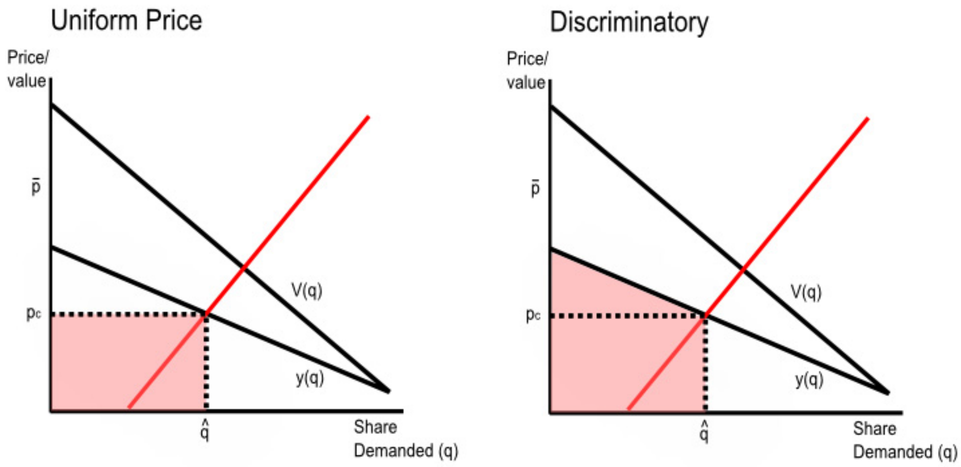

Figure 1: Payments under Uniform-Price and Discriminatory Auction Rules

From the auctioneer's point of view, the total revenue from a uniform price auction is depicted by the shaded rectangle in the left-hand pane of Figure 1: this is simply the total quantity $\hat{q}$ multiplied by the stop-out price p_c. In the discriminatory auction (right-hand pane of the Figure), the revenue equals the $\hat{q} \cdot p_c$ rectangle, plus the shaded triangular area, which depicts the additional revenue (in excess of stop-out price) on all inframarginal accepted bids.

By visually comparing the two areas, it seems obvious that the revenue from discriminatory price auction is higher, and many people who hear about these two rules for the first time question the need for a comparison at all. Yet there is more to this matter than is shown on Figure 1, because it is highly unlikely that the same set of bidders would submit the same individual demand functions under both rules: the incentives for bidding differ significantly across the two rules, hence we should in fact expect different bidding functions to be submitted. In the uniform-price auction, the incentives for shading (i.e. bidding below value) are relatively low, since only the marginal winning bid is paid 'at full face price', and all inframarginal accepted bids are paid at a strictly lower price. In the discriminatory auction, since each accepted bid is paid in full, the bidders need to shade carefully at each portion of their demand curve: any (winning) bid submitted at a price equal to (or very close to) the bidder's true valuation will earn zero (or minimal) profit. Thus incentives to shading are stronger in discriminatory auctions, hence we would expect more aggressive bidding in the uniform price auction. The relevant comparison is then between the stop-out price under the uniform-price rule, and the 'average weighted price' under discriminatory pricing, and the result of this comparison is not intuitively obvious. There is also no current economic theory that would allow for an analytical comparison of the revenue and optimal bidding under these two auction rules, hence data is needed to provide us with a practical answer. In this vein, the next section outlines the economic theory that underpins the econometric models which I use to analyse my data, and provide such an empirical comparison.

2. Economic Theory on Share Auctions

Both econometric models analysed in this paper base on the auction-theoretic model of Wilson (1979). Here I present a generalised version of Wilson's model,[8] and show how each of the two methodologies extends this model to fit their particular set of assumptions.

Assume there are $n \geq 2$ risk-neutral bidders participating in an auction for known aggregate quantity of Q treasury bills. Each bidder i observes a private signal, s_i, and has a (marginal) valuation function of the form $v(q, s_i, s_{-i})$, where q denotes the demanded quantity, and $s_{-i} = \{s_1, .., s_{i-1}, s_{i+1}, ..., s_n\}$ is the vector of signals of $i's$ rivals. While a bidder does not observe the signals of his rivals, we assume that he knows the distribution they are drawn from, as well as n, the total number of bidders. Importantly, at this stage of the model we do not need to assume bidder symmetry, nor any specific parametric dependence between the bidders' signals.

Each bidder's strategy is a function of the price, p, and the bidder's individual signal, s_i; I denote this bidding function by $y_i(p, s_i)$. To solve the model, the strategies are restricted to be differentiable, strictly decreasing functions. The market-clearing price ('stop-out price') p^c is defined as the price where:

$$\begin{aligned}
\sum_{j=1}^{n} y_j(p^c, s_j) &= Q \\
\text{Equivalently} &: \\
q_i(p^c, s_i) &= Q - \sum_{j \neq i} y_j(p^c, s_j)
\end{aligned}$$

The first equation states that at p^c, aggregate demand equals aggregate supply. The second equation re-states this situation from bidder $i's$ perspective: for him,

[8] Notation here mirrors that of Hortacsu (2002).

the stop-out price equates his individual demand function, with his residual supply. Since each bidder only observes his own signal, the stop out price is random from the individual bidders viewpoint. However, since I assumed the bidders know the distribution of their rivals' signals, they can infer the distribution of the stop-out price, conditional on their own information and submitted demand. This distribution is defined as follows:

$$H\left(p|y_i\left(p\right)\right) = \Pr\left(p^c \leq p\right) \tag{1}$$

$$= \Pr\left(y_i\left(p\right) \leq Q - \sum_{j\neq i} y_j\left(p^c\right)\right) \tag{2}$$

Each bidder is then faced with an optimisation problem: to pick a demand function $y_i(p)$ that maximises the expected surplus at every p,conditional on the bidders' own signal and (if applicable) the other bidders' signals that can be inferred at that price. The optimisation problem is thus:

$$\max_{y_i(p)} \int_0^\infty \left(\int_0^{y_i(p)} E_{s_{-i}|p,s_i}\left(v\left(q,s_i,s_{-i}\right) - y_i^{-1}\left(q,s_i\right)\right) dq \right) dH\left(p,y_i\left(p\right)\right) \tag{3}$$

This optimisation programme can then be solved by using calculus of variations to obtain an Euler equation for estimation or other analysis.

From this description, it is evident that there are two key components in Wilson's model which drive the auction outcomes: the specification of the valuation function, and the distribution of the stopout price, conditional on bids. To 'close' the model and apply it to data, we will need to specify a parametrisation for the valuation function, and means of recovering the H-distribution from the data. The latter step is particularly complicated, since the H-distribution itself depends on optimal bidding strategies, and usually will not have a closed-form characterisation.[9] Given this limitation, I cannot construct a likelihood function for H (and consequently the whole model), and so maximum-likelihood modelling will not be feasible. The next two sections present two different methods of meeting these econometric challenges.

2.1. Hortacsu's model - private values

The main model presented by Hortacsu (2002)[10] makes the assumption of symmetric bidders, pure private values and identically and independently distributed signals. These assumptions allow me to re-write the individual valuation function as: $v_i\left(q,s_i,s_{-i}\right) = v\left(q,s_i\right)$: each bidder's valuation now depends on his own signal only, and every bidder's signals are drawn from the same distribution. Consequently, the optimisation programme in (3) becomes:

$$\max_{y_i(p)} \int_0^\infty \left(\int_0^{y_i(p)} \left(v\left(q,s_i\right) - y_i^{-1}\left(q,s_i\right)\right) dq \right) dH\left(p,y_i\left(p\right)\right) \tag{4}$$

[9] At this time, we are not aware of any auction 'share' model in which an analytical expresesion is derived for H, in a non-trival case (i.e. more than 2 bidders).

[10] This paper was subsequently revised, extended and published as Hortacsu and McAdams (2010). For the purpose of my analysis, however, the main econometric methodolgy did not change between these two versions of the paper.

Solving this programme gives an Euler equation (necessary condition) of the form:

$$v\left(y_i\left(p\right), s_i\right) = p + \frac{H\left(p, y_i\left(p, s_i\right)\right)}{\frac{\partial}{\partial p}H\left(p, y_i\left(p, s_i\right)\right)} \tag{5}$$

This equation (implicitly) defines the bidder's bidding function, and the set of all bidders' functions constitute a Bayesian Nash equilibrium of the auction game. The expression suggests at each point where H is strictly positive (i.e. at any price above the lower bound of the support), the bidder will under-report his true valuation, shading it by the $\frac{H(p,y_i(p,s_i))}{\frac{\partial}{\partial p}H(p,y_i(p,s_i))}$. The arguments and proofs showing that this model is non-parametrically identified from the data are supplied in Section 2.2 of Hortacsu (2002), and are thus not re-produced here.

Two further observations need to be mentioned here. Firstly, equation (5) above is a necessary condition for solving the optimisation programme, but it does not guarantee uniqueness. Consequently, even if I recover a valuation function using this equation, it might not be the only one that rationalises my data. Secondly, this equation was derived under the assumption of continuous bidding functions - if there is a finite number of units for sale, the bidder's strategy would also be a (finite) vector of prices rather than a continuous function.

Since actually submitted bid-schedules are finite price-quantity pairs, we need to construct a discrete analogue to equation (5). To construct this analogue, Hortacsu uses the framework of Nautz (1995), which assumes that the price grid is discrete, while the quantity is continuous. Prices on the grid are ordered as $p_0 < p_1 < ... < p_{k+1}$, and each bidders submits a demand schedule $\hat{y}_i : \{q_{i0} \geq y_{i1} \geq ... \geq y_{i(k+1)}\}$, one quantity for each price. The stop-out price in Nautz' framework is then defined as the smallest price at which total demand is just short of total supply. Analogously to the continuous model, in Nautz's setup I define a distribution $H\left(p_k, \hat{y}_i\right)$ which describes the probability that the stop-out price is below p_k, conditional on bidder i submitting the bid-vector $\hat{y}_i$. Given these definitions, the intuitive analogue of (5) becomes:

$$v_i\left(y_{ik}, s_i\right) = p_k + \frac{H\left(p_{k-1}, \hat{y}_i\right)\left(p_k - p_{k-1}\right)}{H\left(p_k, \hat{y}_i\right) - H\left(p_{k-1}, \hat{y}_i\right)} \tag{6}$$

This is the equation I will use in my estimation of individual valuations. While this equation is not precisely correct in the case when not all bidders submit bids at each point in the price grid,[11] in my data this equation provides more sensible estimates.[12]

2.2. Fevrier, Preged and Visser's model - common values

In line with Wilson's (1979) original model, Fevrier et al (2002) assume pure common values. This means that the true value of treasury bills is exactly the

[11] The corrected equation is derived in Appendix 8.2 of Hortacsu.

[12] The 'precisely correct' formulation for recovering v when not all bidders submit bids at each price involves an equation simialr to the discretised Euler-equation, with an extra term, which takes into account the 'missing steps', and also links consecutive steps of the v-function. In this case, the steps of the v-functions are obtained from a system of linear equations.

However, this addtional term features a derivative, $\frac{\partial H}{\partial y_{ki}}$, which given the limited variation in my data, cannot be precisely estimated. Consequently, this theoretically 'more correct' method generates erratic predicitons in this dataset. Hortacsu (2002) finds that in cases where $\frac{\partial H}{\partial y_{ki}}$ can be estimated well, the differences from the two methodologies are 'negligibly small', and hence the use of equation 5 should not be problematic in practice.

same to all bidders, so that $v_i(q, s_i, s_{-i}) = V$, for all bidders i. The common value has a distribution function $F_V(w) = \Pr(V \leq w)$. All bidders are assumed to be symmetric, and each of them receives an independent signal, the distribution of which depends on the realisation of V; this conditional distribution is denoted by $F_{S|V}(s|w) = \Pr(S_i \leq s | V \leq w)$.Thus signals are IID, conditional on the true realisation of V. The realisation of each bidder's signal is private information, and not observed by his rivals or the seller, but the two distributions, F_V and $F_{S|V}$ are common knowledge.

Fevrier's derivation of the Euler condition proceeds differently than that of Hortacsu's. To avoid using y_i^{-1}, the inverse bidding function, in the formulation of the optimisation problem, Fevrier et al re-write the 'payment' from a discriminatory auction, as follows. Observe that since $y_i(p, s_i)$ is monotonic decreasing, there is a unique $\bar{p}$ which is the largest price at which y_i is non-negative. Then for any $p \leq \bar{p}$, the following equality holds:

$$\int_0^{y_i(p,s_i)} y_i(q, s_i)\, dq = y_i(p, s_i) * p + \int_p^{\bar{p}} y_i(u, s_i)\, du$$

Substituting this, into equation (3), and recalling that the common value is the same for all bidders, the optimisation programme becomes:

$$\max_{y_i(p,s_i)} \int_0^{\infty} \left(\int_0^{y_i(p)} E_{s_{-i}|p,s_i,V} \left((V-p)\, y_i(p, s_i) - \int_p^{\bar{p}} y_i(u, s_i)\, du \right) dq \right) dH(p, y_i(p), V) \tag{7}$$

Note that the expectations operator and the H-density now both condition on V; this is relevant because the individual signals, s_i, have distribution which depends on V. Using this to find the Euler condition, Fevrier et al obtain:

$$E_{s_{-i}|p,s_i,V} \left((V-p) \frac{\partial H(p, y_i(p, s), V)}{\partial p} - H(p, y_i(p), V) \right) = 0 \tag{8}$$

The H-density and its derivative are evaluated at $y = y_i(p, s_i)$. Since H is non-negative, and $\frac{\partial H}{\partial p} > 0$ this equation, analogously to equation (5), suggests that the bidder participating in a discriminatory auction will again shade his bid. A discussion of the identification properties of this model are presented in Appendix D of Fevrier et al. (2004), where the authors show that the model is parametrically identified, but concede that it is (most probably) not non-parametrically identified.

To facilitate estimation, Fevrier et al. re-write their Euler-equation as

$$E_{s_i,s_j}\{(n-1)[E_{V|S}(V|S_i = s_i, S_j = s_j) - p] * 1_{p^c \leq p} - E_{p^c}[(p - p^c)\, 1_{p^c \leq p}]\} = 0 \tag{9}$$

In the above equation $1_{p^c \leq p}$ is the indicator function, taking the value of 1 when $p^c \leq p$, and zero otherwise. This condition must be satisfied by any optimal bidding function, on the entire support of the stop-out price density. The advantage of formulating the necessary condition as (9) rather than (8) is that it doesn't depend on the H-density, which in itself can be tricky to estimate due to its dependence on the equilibrium bidding strategies themselves.

3. Description of the Polish Treasury Bill Auction System

My dataset contains information on Polish 52-week treasury bill and 2-year bond auctions, covering the period May 2004 to June 2007 for 52-week bills, and September 2003 to June 2007 for 2-year bonds. The raw auction dataset, obtained from the Ministry of Finance (MoF henceforth), has been augmented by adding secondary-market yields for both securities.[13] During the time period covered by my data, the Polish economy was experiencing a period of steady growth, with relatively little financial instability.

Both securities analysed are zero-coupon "pure discount securities". The 52-week bill is the MoF's most important short-term security, primarily used for liquidity management. However, an additional function of the 52-week bill is to provide the basis for calculating the annual coupons on longer-term (10- and 15-year) variable-coupon bonds. The Ministry of Finance auctions off, on average, 1 bn PLN (roughly £175 mln) of these bills weekly. According to a Ministry representative, the two-year zero-coupon bonds were in the past primarily used to balance the 'average duration' of government debt, but have declined in significance over time.[14] Two-year bonds are auctioned monthly, with roughly 2.3 bn PLN (£400 mln) on average, per auction[15].

Bonds and treasury bills have been sold via discriminatory auction in Poland since 1991, with the institutional design drawing heavily on the French system. Initially participation was open to members of the public, but participation rules were modified in 2002 and a 'primary dealer' system was implemented. To obtain a 'primary dealer' status, a bidding institution (e.g. bank or brokerage house) must apply for a licence from the Ministry of Finance, and satisfy the MoF that they are indeed a viable primary bidder. This procedure includes a contest among all the applicants, and the evaluation criteria are announced by the MoF in advance. Once the licence has been granted, the dealers must satisfy certain 'activity rules' to be eligible for a licence renewal.[16] In my dataset, I observe 19 different primary dealers - not all present throughout the whole sample.

Bidders express their demands within the auction by submitting price-quantity pairs, with prices being quoted to the precision of 0.01 PLN per unit of face value, and quantity expressed as the total amount of the securities' face value. Each 52-week bill has face-value of 10 000 PLN, while the two-year bonds have a value of 1000 PLN each; to maintain consistency, my dataset has re-scaled all of the 'face prices' to 100. The minimum quantity for each bid is 100.000 PLN, and if a bidder wishes to participate in a given auction, his total demand must exceed 1 million, which is usually less than 1% of aggregate supply.

The supply of securities at each auction is announced in advance, though the MoF is permitted to withdraw some of the securities if demand is lacking - in practice

[13] This data was also obtained from the MoF.

[14] Since the beginning of the 2008 credit-crunch, the Ministry has introduced new shorter-term securities (5 and 13 weeks), while interest in thet two-year bond has waned further.

[15] This figure has steadily declined over our sample period - the average from the first five auctions is 2.7 bn PLN, while the last five average 1.7 bn PLN. This reflects the declining significance of the two-year bond.

[16] The licencing period is one year at a time. One of the activity rules requires that each licence holder must win at least 5% of the aggregate supplied amount of bonds, during each quarter. This figure applies across all bond maturities, and each maturity has a different 'weight' in contributing to the 5% requirement.

Table 1: Auction summary statistics, by security type

	52 Week	2 Year
Cover	2.7(1.0)	2.7 (1.0)
Submitted bids	103.6 (26.4)	144.5 (36.8)
Satisfied bids	46.8 (24.1)	59.1 (26.1)
Bidders	12.3 (1.0)	12.8 (1.8)
Winners	7.6 (2.2)	8.8 (2.0)
Bids per bidder	8.4 (2.2)	11.5 (3.1)
Maturity (days)	364.0 (N/A)	785.5 (50.4)
Supply (mPLN)	828.2 (246.7)	2155.7 (655.8)

Means reported (st.dev in brackets).

this occurs very rarely.[17] An additional feature of the Polish treasury-bill auctions is the possibility for the Ministry to organise a top-up auction, shortly after the main auction has finished. The legal framework permits the MoF to offer extra 20% of base-auction supply in the top-up auction, at the weighted average price obtained in the main auction.[18,19] Though top-up auctions are permitted in both 52-week and 2-year auctions, the MoF has never used this opportunity on the 52-week bills, while roughly 40% of the 2-year bond auctions are followed by a top-up phase. The econometric models considered in this paper do not take into account the existence of top-up auctions. As of yet, there are no theoretical models I am aware of which would explicitly model the effect of top-up auctions, though Chapter 3 of Marszalec (2011) takes a step towards taking the top-up phase 'seriously', by showing that it may influence the econometric results found in this paper.

The periods spanned by my data contain information on 103 of 52-week bill, and 43 2-year bond auctions;[20] 18 of the 2-year auctions are followed by a top-up stage. The dataset records all the price-quantity pairs actually submitted by the bidders at the auction, as well as the stop-out price (and hence the final allocation of supply to bidders). Some summary statistics on auctions for the two types of securities are provided in Table 1.

While the cover ratio is quite low,[21] this doesn't necessarily follow from 'a low degree of competition', which is how Krawczyk (2006) interprets similar figures. It is more likely that bidders are realistic about the kind of quantities they can expect to win, there is little uncertainty about supply, and the economic environment in Poland was highly stable, so little is gained by submitting excessively optimistic

[17] In my data, all 52-week auctions sell exactly the announced supply, and over 95 percent of 2-year bond auctions sell precisely the announced amount. The ratio of sold-to-announced bonds varied much more in the mid- and late 1990s, when the financial system in Poland was overall more volatile.

[18] The allocation of supply in the to-up auction is based on the proportions allocated to each bidder in the main auction.

[19] The Polish top-up auctions are analogous to ONC2-type 'non-competitive' bids, discussed in Fevrier & al. (2002).

[20] I have data on earlier 52-week bills auctions as well, but the secondary market prices were not available for dates before May 2004.

[21] The 'cover ratio' is measured as a ratio of submitted demand divided by announced supply.

Table 2: Top-up Auction Statistics

	Base Auctions	Top-Up Auctions
Cover	2.7(0.7)	4.2(3.1)
Cover	2.7 (0.7)	4.2 (3.1)
Participants	12.7 (1.9)	7.2 (3.1)
Winners	10.4 (3)	7.2 (3.1)
Supply (mPLN)	2516.7 (633.6)	423.7 (196.8)

Means reported (st.dev in brackets).

demand schedules. We see from the table that roughly 45% of submitted bids are met (fully or in part) in the 52-week auctions, while around 41% are accepted in 2-year auctions. The number of participants in both kinds of auctions looks very similar, though the 2-year auctions produce one more winner on average. This finding is not surprising, given that many of the same bidders participate in both types of auctions. There are, however, bidders on each maturity (52-week and 2y) who only concentrate in bidding on one type of security, though overall these effects do not appear to be large. The number of bids per bidder is above 8 for 52-week bills, and over 11 for 2-year bonds, so it appears sensible to use economic models in which bidders submit 'demand functions' for mutliple items (or shares), rather than single-step models (with unit demands).

Some summary data on the topup auctions is presented in Table 2. In this table, the standard deviation figures are highly misleading, since the data in the topup auctions is heavily skewed: the cover ratio is never less than 1, but is as high as 10 in one auction; thus whenever a topup auction occurs, it is fully subscribed (and usually highly over-subscribed). Roughly 70% of base-auction winners participate in the topup auction on average - those bidders that do not participate are usually the smaller bidders. On average, the topup auction supply was 17% of the base-auction supply, though this average is (again) not very representative: in 12 out of the 18 topup auctions, the MoF offered the full allowed 20% additional supply.

3.1. Ex-ante arguments for common v.s. private values

Looking at the institutional description and statistical summary of the Polish treasury- bill market, it is not clear whether 'private values' or 'common values' is the correct theoretical paradigm. Fevrier et al (2002) justify the 'common value' assumptions by arguing that under the 'dealer system', bidders in the treasury auctions are re-selling to the same secondary market. Given that Poland's institutions closely resemble those of France, Fevrier's argument could be carried over. Conversely, Hortacsu suggests that 'private values' are applicable in Turkey since many bidders participating in the auctions face institutionally regulated 'liquidity requirements', and the exact state of individual liquidity reserves is private information. There is no corresponding legislation in Poland - liquidity reserves of primary bidders are not regulated.[22] However, there are other sources of 'private information' which

[22] While pension funds are required by law to hold a certain percentage of their assets in government-backed securities, these funds do not participate in the auction directly; they can, however, submit customer-bids via a primary dealer.

could justify using Hortacsu's model on Polish data. For example, primary dealers usually meet with their customers prior to the auction to discuss, in advance, what kind of bids to submit on the customer's behalf. Since for these bids the transaction price between the dealer and customer are fixed in advance, at least a portion of each dealer's demand is based on information particular him. Hence a private-values approach could also be justified.

Discussions with the dealers themselves do did not yield a consensus as to what the correct paradigm should be - if anything, the picture appears even more complex that my dichotomy suggests.[23] The dealers admit that usually 50-80% of bids that are submitted are in fact forwarded customer bids - but sometimes the dealers change these bids at their own risk.[24] Furthermore, it is possible that customers perform 'order splitting' in order to hide some of their private information from the dealers - so the same client may submit different portions of his demand curve through different dealers. In this context, it is (even theoretically) hard to say whether private, or common, values apply: even if each bank serves the same customer pool, each dealer may be observing a different part of each customer's demand. It is important to note that these customer bids do not fit in exactly with the notion of the 'secondary market' as used by Fevrier - these customers submit bids *before* the auction, rather than requesting the securities ex-post. Furthermore, in addition to the larger customers, each bank has a pool of small 'private individual' clients, for whom order splitting is not necessary - this also points towards a private-value component.

As of yet, there doesn't exist an econometric methodology that would allow to formally test for common values against private values in the kind of data I have.[25] Similarly, I am not aware of any complete 'auction for shares' model which would allow both private and common value components to be present at the same time. In this paper I will remain agnostic as to which paradigm is 'true' - but I will evaluate as to whether the predictions of the two models are consistent with each other. If they do, the theoretical disagreement will have limited significance to the practitioner.

3.2. Description of bidders

The generalised shares-auction model presented in Section 2 does not need assumptions on bidder symmetry. Both of the models I analyse assume independence, and while Hortacsu's model, prima facie, does not invoke symmetry,[26] Fevrier's derivation hinges crucially on such an assumption. Since signals are not observed in my data, I cannot directly test these assumptions, but looking at reduced-form statistics of the bidders could give us an idea. A per-bidder summary of selected

[23]For this purpose, interviews with 8 of the primary dealers were carried out - this included the 4 largest ones, and 4 smaller dealers.

[24]For example, if the customer requests a given quantitiy at a price that the dealer thinks is too high, the dealer may reduce the price he actually submits on this bid, and keep some proportion of the 'savings' for himself. The risk that the dealer is taking is that the modified bid may not be accepted at all in a case when the original bid would have been accepted. In this case, the dealer is still contract-bound to provide his client with the appropriate amount of securities.

[25]Hortacsu & Kastl (2008) develop a test for common values in a 'shares' type auction, but that test is specific to their data format. Unfortunately, few datasets contain the kind of information that is necessary to apply the Hortacsu & Kastl test.

[26]But there is an implicit reliance on symmetry in Hortacsu's generation of the H-function. See Section 8.1 for a modification that considers asymmetry.

Table 3: Bidder statistics on 52 Week Bills

Name	Active in	Total Bids	% Demand Satisfied	% won if active	% won overall
2	103	1487	25.97	11.6	11.4
9	103	1308	27.7	11	11.3
4	101	1012	31.95	12.2	11.2
11	101	464	27.43	9.6	9.6
10	102	1149	22.57	8.5	9
5	102	1090	26.16	8.3	8.8
13	93	490	13.7	7.9	7.6
6	103	1016	25.48	7.9	7.5
12	100	612	26.61	8.1	7.3
7	101	266	12.09	6.4	6
3	102	1049	16.27	5	4.7
1	103	492	13.66	4.4	4.7
15	22	108	7.98	2.7	0.7
14	34	121	6.22	1	0.4
16	1	3	1.67	1.7	0

Total amount of 52-week bills sold was 85.3 bnPLN.

statistics is provided in Tables 3 and 4.

There are a few caveats to interpreting this table. Since primary dealer status is not permanent, some bidders participate if fewer auctions than others due to late entry.[27] Secondly, not all bidders participate equally (some not at all) in both types of auction. Keeping these points in mind, there still appears to be much heterogeneity among bidders. For 52-week auctions, even if I look at all the bidders who participate in 100 auctions or more, the average percentages won vary a lot, for example, bidder 4 wins on average 12.2% of the aggregate supply if he participates in an auction, while bidder 1 wins only 4.4%. While it would not be surprising to find a discrepancy of this magnitude *within* a single auction (and maintain assumptions of symmetry), the fact that the discrepancy persists (on average) across 100 auctions suggests that asymmetries may be significant.

A similar pattern can be seen in the columns for 2-year bonds - e.g. bidder 9 wins 16.1% of bills in auctions where he participates, while bidder 11 only wins 3.0%. While these reduced-form statistics do not necessarily tell us that the bidders' *signals* are from different distributions, they do seem imply a degree of asymmetry, hence I will check, where possible, whether including asymmetries significantly changes results from the symmetric model.[28]

[27]For example, bidder 20 enters late, and only into 2y bond auctions, while bidder 19 is present in early 2y auctions, but later exits.

[28]This will be possible when using the Hortacsu method, but not in the FPV model.

Table 4: Bidder Statistics on 2 Year Bonds

Name	Active in	Total Bids	% Demand Satisfied	% won if active	% won overall
9	42	853	42.2	16.1	17.9
12	42	468	35.4	15.2	14.3
4	41	459	27.5	11.4	12.8
5	42	715	22.3	9.5	11
6	42	546	23.9	9.3	9.6
13	42	605	34.4	6.5	7
2	42	402	13.9	6.4	7
10	40	496	14.3	7	6.9
3	41	498	17.7	7.6	6.3
1	42	571	16.1	5.9	5.8
11	42	211	7.9	3	3.2
14	15	57	18.2	2.4	0.8
16	14	32	4	1.7	0.5
7	19	45	2	0.8	0.4
15	12	66	6.3	0.9	0.3
8	8	25	1.8	1.2	0.3
18	7	11	16.4	2.8	0.4
20	1	2	15	12.7	0.1
19	3	5	0	0	0

Total amount of 2-year bonds sold was 90.5 bnPLN.

4. An Application of the Hortacsu Model

Hortacsu's approach to estimating discriminatory auctions is fully non-parametric - his analysis does not make any assumptions about the shape of the distributions which bidder's signals, or values, follow. The advantage of this approach is that any conclusions the model produces cannot be criticized as being driven by such parametric assumptions. A significant disadvantage of this method, however, is that it will not allow me to explicitly solve for optimal bidding strategies in a uniform-price auction. To generate any kind of comparability between the two auction types, I will thus have to use strong assumptions as to bidders' behaviour in the uniform price auction, given their valuations. In particular, I will assume that the bidders bid 'truthfully', in the sense that they submit the valuations which I recover via equation (5) as their actual bids in a uniform-price auction. As Ausubel and Cramton (2002) have shown, in most circumstances this assumption is not valid, and even in uniform-price auctions some shading will be present. However, truthful bidding in a uniform-price auction provides us with an upper bound on the revenue we can expect from such an auction; with optimally calculated (shaded) bidding functions, the equilibrium price could be lower than the truthful-bidding price, but not higher.[29] Hence I will have a (one-sided) bound on the uniform-price revenue,

[29] Kastl(2006) shows that this statement is not, strictly speaking, correct - and that it is possible that bidders in an uniform-price auction would submit bids above their true valuation. The Kastl argument, for a bidder i, is the following: "Suppose my valuation at q_k is $v_k(q_k)$, but it is unlikely

against which to compare the discriminatory auction: if the discriminatory weighted average price is above the truthful-bidding price, then we should conclude that the discriminatory auction is revenue-dominant.[30] In the converse case, we cannot tell: even if the truthful bidding equilibrium is at a higher price than the discriminatory price equilibrium, it might still be the case that effects of shading (demand reduction) would depress uniform-price auction revenue.

4.1. Simulating the Distribution of the Stopout Price

The intuitive idea behind Hortacsu's resampling approach to constructing H follows naturally from the use of Bayesian equilibrium in solving the economic model and the underlying assumptions of statistical independence. Since in a Bayesian equilibrium each bidder 'knows' the strategy of all of his opponents (conditional on the signal they receive), he can calculate the aggregate demand of all his opponents for every possible combination of their signals - thus also his own residual supply, and the stop-out price, conditional on that particular set of signals and his own demand curve. Furthermore, since the signal distributions are assumed to be common knowledge, each possible combination of signals can be weighted appropriately, whereby the distribution of the stop-out price can be obtained - this is H, as defined in equation(2). In the setting of the problem, this approach is feasible for the bidders, but not feasible for the econometrician: we don't know which distribution the signals are drawn from, nor can we calculate (in this case) the relationship between the signals and the bidding functions.

Hortacsu proposes a resampling (bootstrap) estimator which allows us to construct H, without having to make any parametric assumptions. Since I have assumed private values, each bidder's submitted demand curve is a function of his signal only, so given independence of signals, the independence of submitted demand schedules follows immediately. Furthermore, since these demand schedules were actually submitted in the 'true' auction, they should constitute equilibrium strategies (under the stated assumptions). From the point of view of calculating a stop-out price, the only information I need is every bidder's (submitted) demand function - the signals themselves are only instrumental. I can then simulate the distribution of the stop-out price by repeatedly re-drawing an appropriate number of demand curves and running a discriminatory auction and recording its stop-out price. If I repeat this procedure sufficiently many times then, given the IID assumptions, the distribution

---

that the stop-out price occurs at this step. Then I might as well submit a $p_k > v(q_k)$, and make it more likely that I do win q_k shares in case the stop-out price is in fact $v(q_k)$". In short, this way the bidder feels that he is less likely to be rationed if he bids more aggressively.

Intuitively, we would expect that such agressive bidding would be most prevalent at price-quantity combinations which are unlikely to be marginal. Kastl's data-based examples show that over-bidding does indeed occur, but only at the first step. In Kastl's data this may be significant, since bidders submit step-functions consisting of only 2.3 steps on average.

In my data, bidders submit much more 'steps' in their bid functions, and the quantities submitted by bidders at their first few steps are miniscule. The likelihood that these bids would be marginal is very small. Thus it is unlikely that 'over-bidding' would occur at the stopout price in the uniform-price auction. Finally, the MoF has commented that on 52w auctions, it often does not ration the marginal bidder - and rather accepts his bid in full (even if this increases aggregate supply by a little, ex-post). For all these reasons, it is we would not expect Kastl's argument to apply in my data to a significant extent.

[30] The two relevant prices that are compared here are the uniform-price auction 'stop-out price', and the discriminatory auction 'weighted average price' (which is different, and higher than, the discriminatory price stop-out price).

of the simulated stop-out prices will tend to the true distribution H.[31] Since the H-distribution is conditional on each individual bidder's submitted demand schedule, and the equilibrium strategies are specific to each auction, we need to generate a separate H-function for each bidder, in each auction. The full procedure for obtaining one such H-distribution is as follows:

- 1. Fix an auction (say l), and a bidder (say i). Find n_l - the number of bidders in auction l.
- 2. For auction l, draw, with replacement, $(n_l - 1)$ actual submitted bid schedules.
- 3. Aggregate these bids together, and use them to construct a residual supply function for bidder i.
- 4. Intersect the simulated residual-supply function with bidder i's actual submitted demand schedule, and record the equilibrium stop-out price.
- 5. Repeat steps 2 to 4 B times to generate a sample of B simulated stop-out prices.
- 6. Use the simulated stop-out prices to generate a histogram and a 'cumulative histogram' - this generates H.

In the above algorithm, B denotes the number of 'bootstrap repetitions' - and I would expect to obtain a more accurate shape for H as B increases. Figure 2 below graphically shows the representation of this re-sampling procedure, with B =10. An example of a simulated H-distribution is given below, in Figure 3, with B set to 10 000.

The left-hand panel depicts the bidder's submitted demand curve, and the simulated density at each price at which he submitted a bid. These densities have been aggregated in the right-hand panel, and show the cumulative density - H(p|y).

It is apparent from the sketch of the re-sampling procedure that it may cause complications when I proceed to use H to back-out individual valuations. From equation 5, we see that for the FOC to be defined everywhere, $h(.) = \frac{dH(.)}{dp}$ must be non-zero everywhere. This implies that we should observe some 'simulated density' at every point in the support of H - so each price should (at least once) appear as the stop-out price of a simulated auction. However, we might expect some very low, or very high, prices never to feature as stop-out prices: at very low prices there may always be excess demand, while at high prices there may always be excess supply. If so, there may be a range of prices over which H is 0, and a range of prices all for which H is 1 - on both of these intervals, h=0, whereby the FOC in equation 5 is not defined. A similar situation may occur if some bidder(s) submit bid-schedules with a group of prices and quantities very close to each other. Despite repeated sampling, it is possible that the simulated residual supply never intersects with the relevant section of the bidder's submitted bid-schedule, and so some prices are never 'hit' in

[31] The relevant theorems for showing the consistency of this bootstrap estimator can be found in Hortacsu (2002), Appendix 8.4

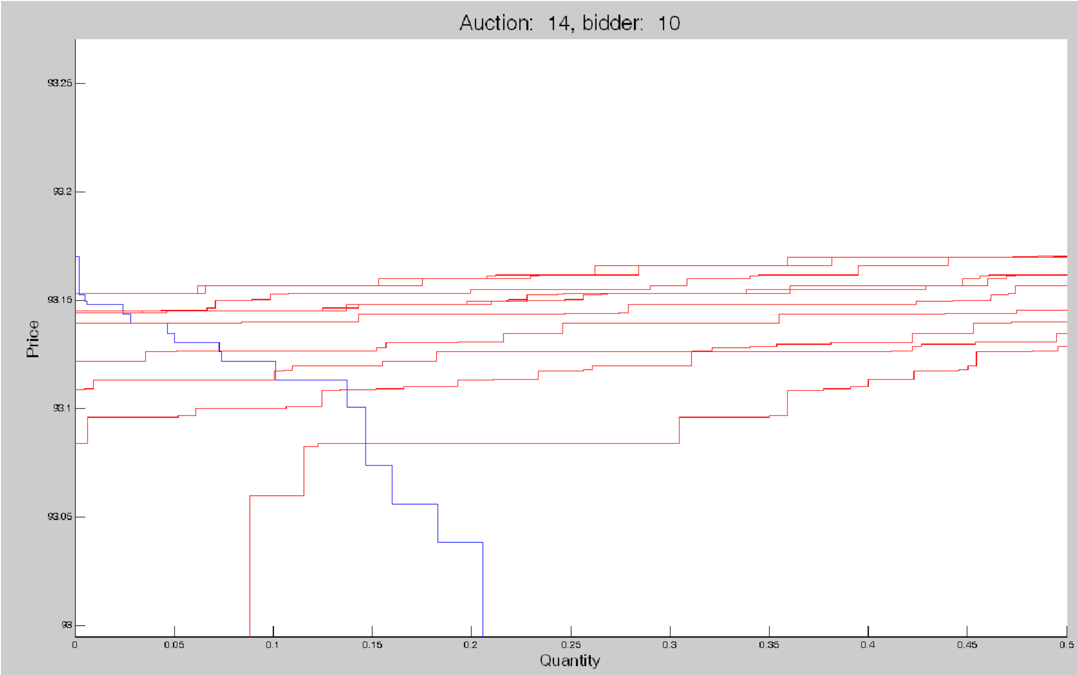

Figure 2: Illustration of the Hortacsu procedure for simulating H(p|y)

equilibrium . Both of these problems are illustrated in Figure 4. Here we see that for very low prices, H is flat (so h = 0), and there is also a flat portion in H on the interval p$\in$ $[95.44, 95.64]$.

This zero-density problem is not explicitly discussed in Hortacsu (2002), whence there is no indication as to how the author would deal with this issue. One alternative is to attempt to modify the density estimator in a way that makes the zero-density problem less likely - this is discussed in Appendix A. Another viable approach is to use interpolation to calculate a valuation for those quantities where I find zero density, based on correctly inferred valuations near by. The method of interpolation that I select will, however, influence the conclusions of my later analysis, so I need to pick a method which does not make (too) strong assumptions, while not weakening the conclusions of my analysis. The method I use is the following: if at a given price-quantity pair (p, q) the density $h(p)$ is zero, I look for the largest $q' < q$, and I set $v(q) = v(q')$. In other words, I assume that the marginal valuation remains the same between q' and q.[32] The two main reasons for selecting this procedure is that firstly, it preserves monotonicity. Given that the solution to my economic model assumes monotonic equilibrium bid-functions, assuming that valuations are (weakly) monotonic seems appropriate.[33] Secondly, this assumption is likely to make my conclusions stronger, rather than weaker, in the following sense: we know that the Hortacsu method will only give me an upper-bound on uniform-price revenue, and

[32] We have made the additional assumption that if the p at which zero-density occurs is so high that there is no $q' < q$ for which a value estimate exists, we assume that $v(q) = 100$, that is, value equals the maximum possible value of the bills.

[33] Indeed, if we were to pick an imputation method which *induces* non-monotonicity, we would need a convincing argument as to why that model should hold, and how it could generate a monotonic equilibrium bid function.

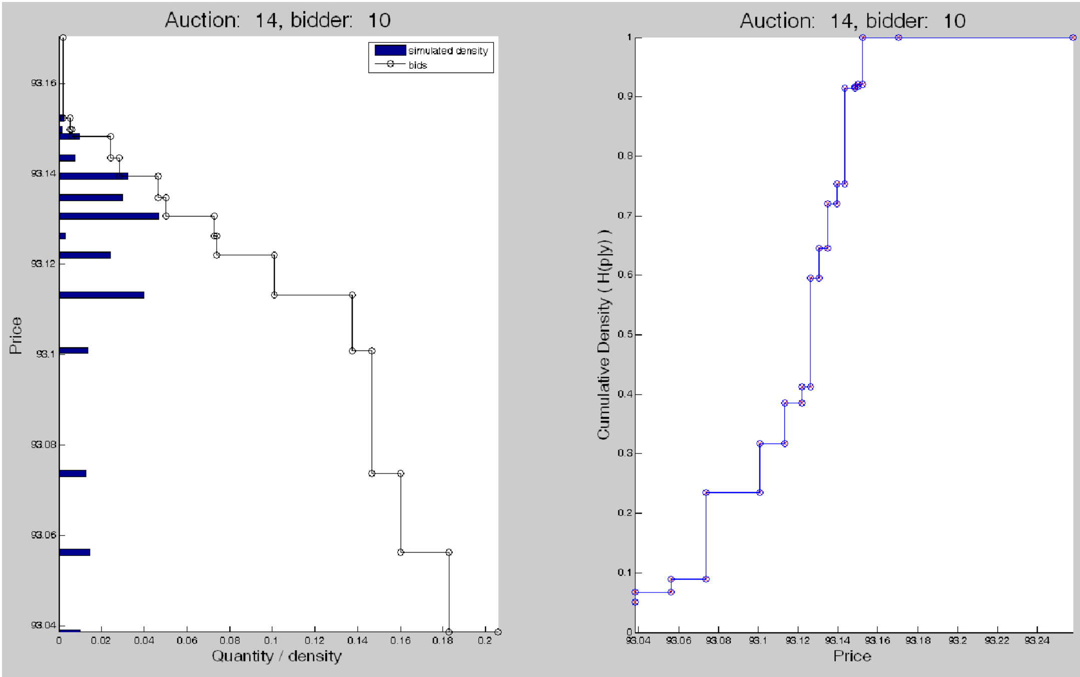

Figure 3: Example of a simulated H(p|y) distribution

this method of imputation is likely to bias the bidders' value function upwards. The upward bias in individual valuations will carry over to an upward bias in the uniform-price best-case revenue. If even given this bias, I find that the discriminatory auction is revenue-dominant, this conclusion is unlikely to be driven by my assumptions on interpolation. Thus I am less likely to find that the discriminatory auction revenue dominates - but if I do reach this finding, I should be relatively confident that it is not an artifact of my interpolation method.[34]

A second potential worry from Hortacsu's re-sampling method of generating H is that it does not guarantee that H/h is monotonic. If this monotonicity fails, then it is possible that the recovered valuations will not be monotonic either. This would imply complications for calculating the counterfactual revenue from truthful bidding in the uniform price auction: observe that the FOC gives a valuation conditional on 'cumulative', not 'marginal', demands. Thus at a certain per-unit price, the bidder may want to buy a both a quantity q_1 and a quantity $q_2 > q_1$, but not any amount in between. When I aggregate non-monotonic valuation functions across bidders to obtain the aggregate demand function, I may find that it intersects the (perfectly inelastic) supply curve more than once. While this feature is unappealing, the assumptions made in defining the stop-out price in Section 2.1, following Nautz (1995), gives a natural way of selecting as to which intersection to pick as the stop-out price: this will be the lowest price at which aggregate demand just falls short of aggregate supply.

[34]If we used the a converse method of interpolation (i.e. biasing v down whenever there is zero denisty), then we would be more likely to 'find' a significant difference between discriminatory and unifrom-price outcomes. However, these may be an artifact of the interpolation, since the uniform-price upper-bound is now biased downwards.

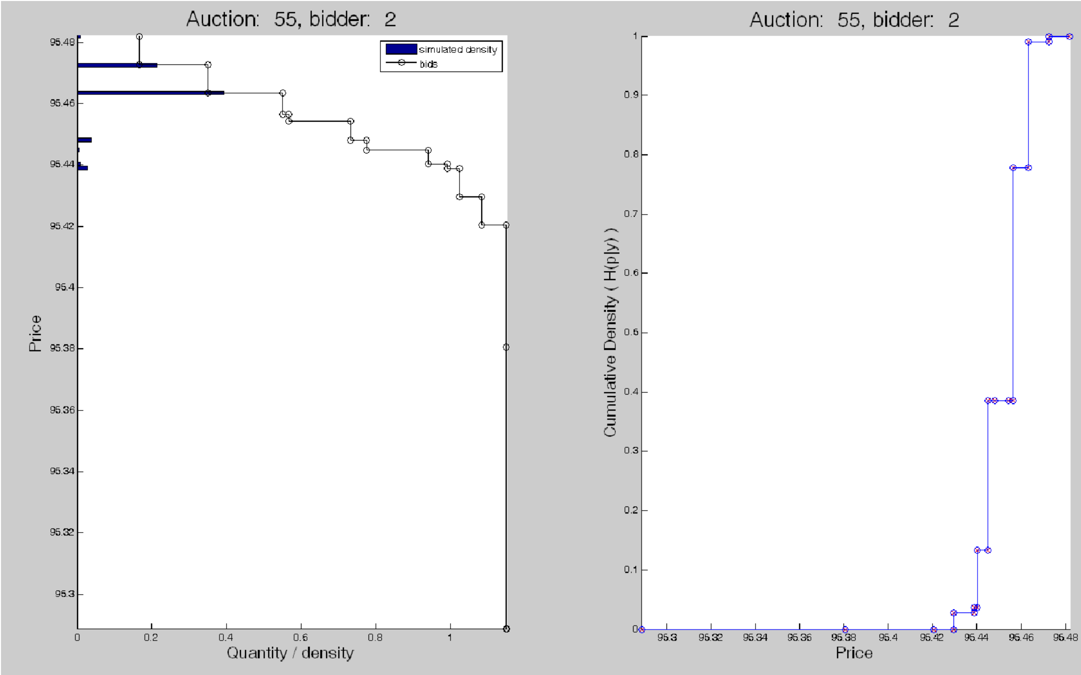

Figure 4: Example of the 'zero-density' problem

4.2. Inference in the Hortacsu Model

Since in the Hortacsu method no 'parameters' are estimated per-se, inference is done through creating 'error bounds' around the inferred quantities of interest. The two which are of most interest are the individual valuations (at pre-specified quantities), and the upper bound on the stop-out price in a uniform auction. Error bounds on both are calculated through simulation, though the methods used are different.

To calculate standard errors on individual valuations, Hortacsu uses a 'jack-knife-then - bootstrap' of Efron (1992). The variation in individual values is obtained by calculating a series of 'perturbed' H-distributions, and then re-calculating the individual values (v_i's) based on these perturbed distributions. The 'perturbation' is obtained by excluding, in turn, each of the participating bidders from the stage of simulating H. That is, we re-calculate H based on drawing the re-sampled demand functions from all bidders except (say) bidder 1; then we do the same, but exclude bidder 2, and repeat this re-calculation until each bidder has been excluded once. This will give me a set of n_l different H-functions, and consequently n_l re-calculated valuations at for each (actually submitted) quantity. Using the notation of Hortacsu, let $\hat{v}_B^l(y_{ik})$ be the estimated marginal valuation for bidder i at quantity y_{ik} in the l-th auction, based on an H-distribution obtained through B replications of the bootstrap procedure. Then the variance estimate will be:

$$Var_{jackknife}\left(\hat{v}_B\left(y_{ik}\right)\right) = \frac{n_l}{n_l - 1}\sum_{j=1}^{n_l}\left(\hat{v}_{B(j)}^l\left(y_{ik}\right) - \frac{1}{n_l}\sum_{i=1}^{n_l}\hat{v}_{B(j)}^l\left(y_{ik}\right)\right)^2$$

with $v_{B(j)}^l\left(y_{ik}\right)$ denoting the estimate of i's marginal valuation at y_{ik} based on an H-distribution which was simulated from re-samples which did not contain the submitted bidding-functions of bidder j.[35] This gives me a standard deviation estimate for each step in the valuation function, for each bidder and each auction. An example of recovered valuation-function and its associated error bounds (± 2 standard deviations) is shown in Figure 5, for bidder 3 in auction 21 .

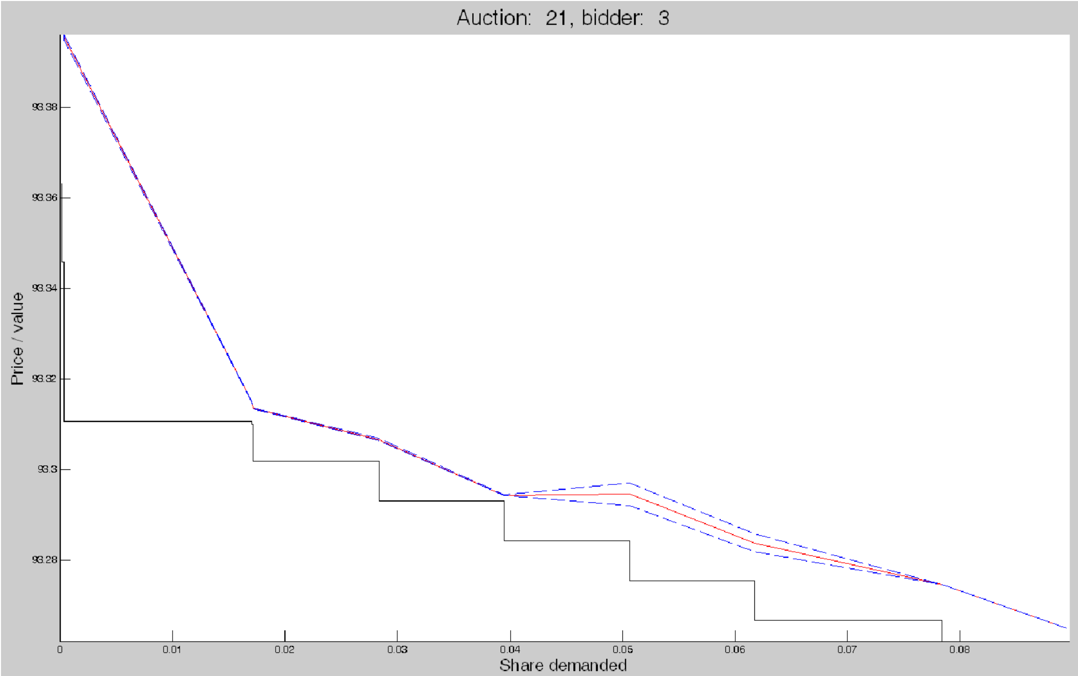

Figure 5: Bid-function, with corresponding $v(p)$ function, and confidence-intervals

To obtain error bounds on the stop-out price from a uniform auction, I simulate a series of uniform-price auctions by re-drawing the estimated valuation functions within each auction - according to an algorithm exactly analogous to that used in Section 4.1 to estimate H. Thus for each auction l, I draw n_l recovered valuation functions (for that auction),[36] and use those valuations to obtain a stop-out price. I repeat this procedure 10 000 times, and use the data to find the simulated standard deviation, and 'simulated/empirical' 95% confidence intervals. I can perform this same calculation for the discriminatory price revenue, and the revenue-difference (for each auction), which will allow us to see directly which auction is revenue-dominant.

4.3. Results from the basic Hortacsu model

The complete per-auction results from estimating Hortacsu's model are presented in Appendix B in Chapter 2 of Marszalec (2011)[37]. A subset of these results, covering 5 of the 103 auctions on 52-week bills are presented in Table 5. In this table

[35] The essence of this procedure is to simulate the impact of the presence of bidder j on the H-distribution of bidder i, and hence on the inferred values of bidder i. Repeating this procedure gives us a set of value estimates at each step - and we use these to calculate the poinhtise varaince.

[36] This is equivalent to randomly drawing the appropriate number of participants, within each auction, and having them play one round of the auction, assuming that their bidding and valuation functions remain unchanged.

[37] The full results are also available from the author upon request.

columns 2 to 4 show revenues obtained based on the 'actually submitted' demand functions, and the inferred valuations therefrom - these results have 'no variance' as they are obtained based on a single iteration of the auction only. The columns on simulated revenues have been obtained using 10 000 re-samples of bidding and valuation functions (for each auction), and the standard errors are calculated based on these re-samples. Thus columns 5 to 7 depict the 'simulated' counterparts of columns 2 to 4; here the standard deviation is presented in parentheses next to the estimate itself, and the simulated 95% confidence interval is shown in square-braces, below the estimate.

Table 5: Per-auction results for Hortacsu Model

No.	DP	UP	RevDiff	DP Sim	UP Sim	RevDiff Sim
1	93.58	93.48	0.10	93.58(0.02)	93.50(0.10)	0.08(0.09)
				[93.55;93.61]	[93.40;93.67]	[−0.06;0.18]
2	93.43	93.42	0.00	93.42(0.01)	93.41(0.02)	0.02(0.02)
				[93.40;93.44]	[93.35;93.44]	[−0.00;0.06]
3	93.46	93.46	-0.01	93.45(0.01)	93.47(0.06)	−0.01(0.05)
				[93.42;93.47]	[93.39;93.52]	[−0.05;0.03]
4	93.51	93.51	0.00	93.50(0.01)	93.50(0.02)	−0.00(0.01)
				[93.48;93.52]	[93.47;93.55]	[−0.03;0.02]
5	93.58	93.55	0.03	93.57(0.01)	93.57(0.20)	0.00(0.20)
				[93.55;93.60]	[93.53;93.63]	[−0.03;0.04]

Means reported. St.dev in parentheses, simulated 95 % CI in square brackets below.
Numbers in columns 2 to 4 are realised results, columns 5 to 7 are simulated.

Particular attention needs to be paid to the 'standard-deviation' figures that I have estimated, since they may be misleading. In Hortacsu (2002), Table 3, standard deviations are provided as a 'guide' to the significance of the revenue difference - the 'rule of thumb' (RoT) being that any difference larger than two standard deviations is 'significant at 95% level'. Though Hortacsu finds deviations of 14% on average, none of these are 'significant', using his rule of thumb. In the two Figures (6 and 7) I compare the behaviour of this 'rule of thumb' relative to the 'simulated' 95% confidence intervals for the sample of 52-week bills. The revenue difference is calculated by subtracting the uniform price revenue from the discriminatory price revenue. If, at a given auction number, the 'zero' intercept is within the confidence bound, I cannot reject revenue equivalence; if the whole confidence interval is above zero, discriminatory price is revenue-dominant (and uniform price is dominant in the converse case).

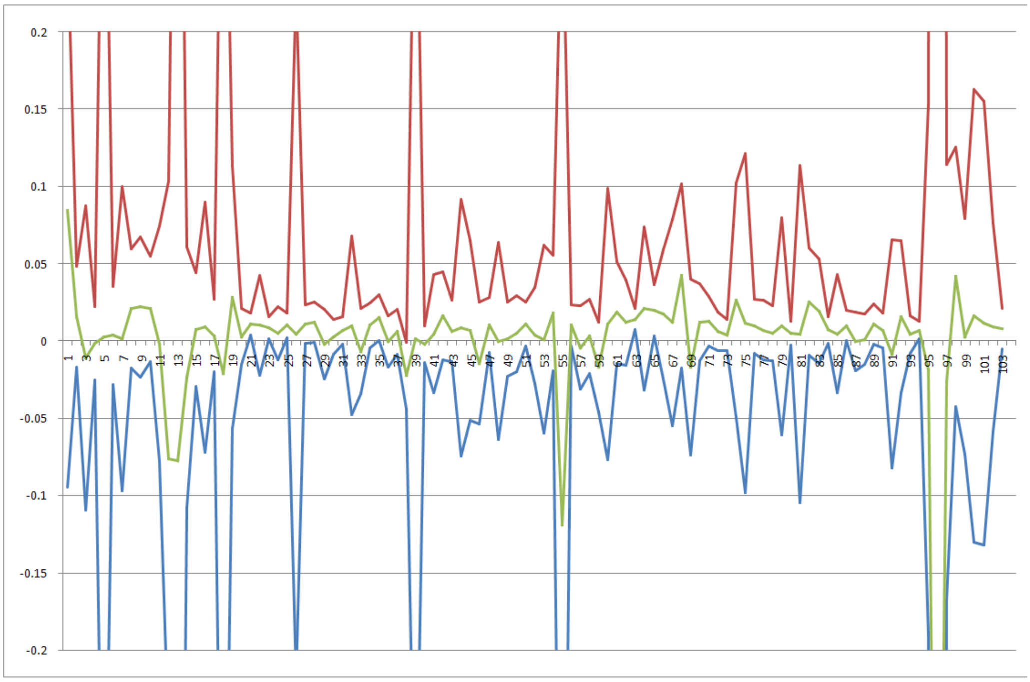

Figure 6: Per-auction revenue differences for 52w bills in the Hortacsu model, with ROT confidence intervals

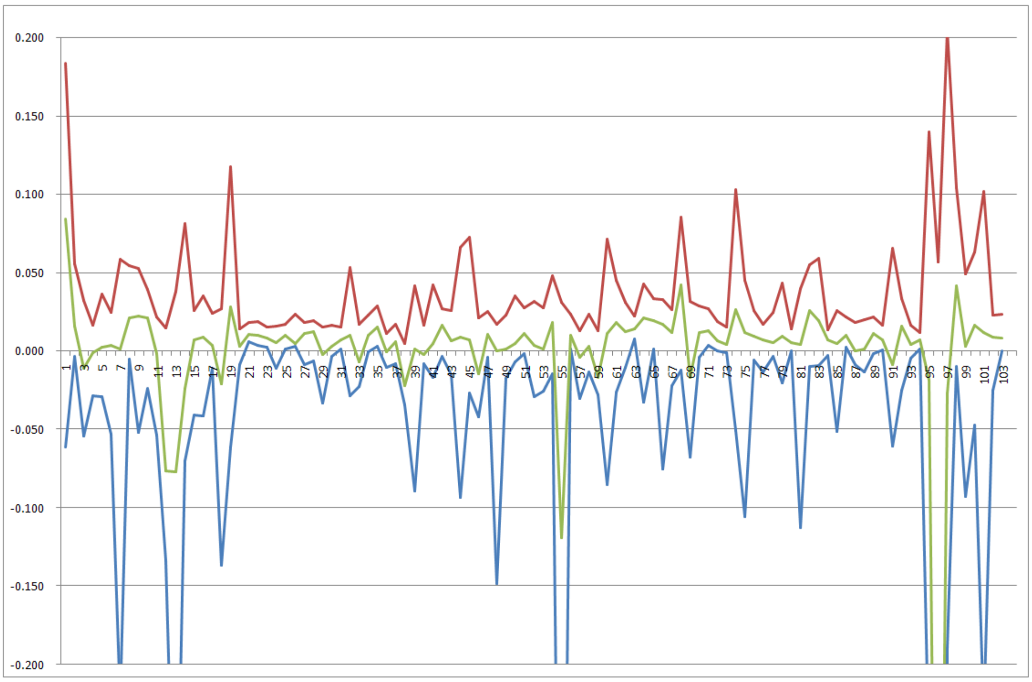

Figure 7: Per-auction revenue differences for 52w bills in the Hortacsu model, with simulated confidence intervals

Both figures have the same scaling, and it is clear that Figure 6 which depicts 'rule of thumb' confidence intervals shows much wider confidence intervals than the simulated intervals in Figure 7 - the large standard deviations in Figure 6 are driven by a small number of large (simulated) revenue differences, which in turn makes it hard to reject 'revenue equivalence'. Another feature of these confidence intervals - more evident from Table 6, than from the figures - is that the simulated confidence intervals are asymmetric, while those assumed by the 'rule of thumb' are not. Overall, however, the revenue differences appear to be very small, and visually it may be difficult to tell which auction format dominates and how many times. Thus a summary of the per-auction revenues is given in Table 6.

The figure reports average prices of two kinds: realised and simulated. The 'realised' DP-price for an auction is calculated by using exactly those bidding functions that bidders submitted when the auction took place - the 'realised' UP-price is then calculated by using the valuation functions of these same bidders, obtained using the Hortacsu method. Since these prices are calculated from a single iteration of the auction, they have no variance. The 'simulated' prices are obtained by randomly re-sampling the appropriate number of bidding and valuation functions (within each auction), and then calculating the relevant prices. The standard deviations are from these re-sampled data.

In terms of realised prices it appears that the results from 52-week auctions strongly favour the discriminatory auction, in that it revenue-dominates in 83 cases, as opposed to the uniform-price's dominance in 20 auctions only. There is less difference between the two auctions in the sample of 2-year bond auctions, with discriminatory-price dominating 24 times, compared to the uniform-price's 18. Note that on both maturities, the average price differences across the whole sample are very small: in the case of the 52-week bills, the DP average price is higher by 0.009% and for 2-year bonds the UP upper bound is roughly 0.02% above the discriminatory-price revenue. In the case of 2-year bonds, the results on the UP upper bound being above the DP-price are driven by a handful of auctions, in which uniform-price does significantly better than discriminatory-price. This explains the discrepancy between the discriminatory-price being dominant in a larger number of auctions, but the uniform-price upper-bound being higher than the discriminatory-price on average.

Table 6: Per-auction summary for Hortacsu model

	52 Weeks	2 years
Number of Auctions	103	42
DP Realised Price	94.79	89.08
UP Realised Price	94.78	89.10
DP Simulated Price	94.79 (1.16)	89.08 (2.21)
UP Simulated Price	94.79 (1.17)	89.46 (2.52)
Realised Rev. Dominance	DP: 83 / UP: 20	DP: 24 / UP: 18
Rev. Dominance - RoT	DP: 8 / UP: 1	DP: 0 / UP: 0
Rev. Dominance - Empirical Bounds	DP: 15 / UP: 0	DP: 1 / UP: 0

Means reported (st.dev in brackets).

Table 7: Revenue Sequence Comparison: Hortacsu Model

	52 Weeks	2 years	Pooled
Number of Auctions	103	42	145
Face Value total	85.30	90.54	175.84
DP Revenue total	80.85 (0.00)	80.65 (0.00)	161.50 (0.00)
UP Revenue total	80.85 (0.02)	80.99 (0.43)	161.85 (0.43)
Revenue Difference total	-0.00 (0.02)	-0.35 (0.43)	-0.35 (0.43)
Simulated CI for Rev. Diff.	[-0.08 : 0.01]	[-1.16 : 0.00]	[-1.16 : 0.01]

Amounts in bnPLN.
Means reported (st.dev in brackets).

Looking at simulated results, when following the 'rule of thumb' methodology, as Hortacsu does, I find that the discriminatory price is significantly revenue-dominant in 8 of the 52-week bill auctions, while the uniform-price revenue-dominates only in one; for the remaining 94 auctions I cannot reject revenue equivalence. In the sample of 2-year bonds, the RoT method never rejects revenue-euqivalence. When using simulated 95% confidence intervals, however, I find that the results become (slightly) more favourable for the discriminatory auction: for 52-week bills it is now dominant in 15 auctions (while uniform-price never is), and I also reject revenue-eqivalence in favour of the discriminatory auction in one of the 2-year auctions.

To test for overall revenue dominance across the sample, I also simulated 10 000 repetitions of the 'sequences' of auctions for each maturity. To obtain such a sequence, I simulate revenue once for each auction and aggregate across all auctions - this constitutes one re-sample, and is then repeated to obtain a sample of sequences. The results are given in 7.

Since we have seen from Table 6 that revenue-differences are very small on average, it is unsurprising that Table 7 cannot reject revenue equivalence across the whole sample, for both bond maturities. When interpreting these results, it is essential to remember that the uniform-price estimates are 'best-case' estimates, which do not take account of any possible shading that bidders may perform under those auction rules. Since the confidence intervals on both auction-type revenues (and the revenue difference) are small, it would only take a small degree of shading in equilibrium to render the uniform-price auction revenue-inferior much more widely.

4.3.1. The zero-density problem and Non-monotonicity

The setup of Wilson's share auction model in section 2 doesn't assume monotonicity of the marginal valuation function - only the equilibrium bidding schedule is assumed monotonic.[38] It appears sensible to require that marginal valuations should be monotonic and (at least weakly) decreasing - this appears to be a standard assumption of multi-object auctions. For Hortacsu's model, there is an additional reason why monotonic marginal valuations are attractive: they ensure that there is only one 'equilibrium' in the uniform-price auction with truthful bidding. Unfortunately my inferred valuation functions are rarely monotonic over the whole support. Most commonly, however, the monotonicity violations occur at low demanded quantities/ high submitted bids - these are the points at which H (and h) are particularly

[38]In Wilson's original setup there is an implicit assumption of constant marginal valuation.

badly estimated.

Even if individual valuation schedules are not monotonic, it possible that the aggregated valuation function would be, which would still allow me to find a unique stop-out price in the uniform auction. However, out of 145 auctions, I find 3 cases in which the aggregation does not help, and the aggregate valuation function crosses aggregate supply more than once. A larger number of auctions manifest aggregate valuation functions which are not monotonic, but in those cases the non-monotonicity occurs at lower (aggregate) values of q, hence 'far away' from the stop-out price/quantity. In particular, for aggregate demands close to Q, (50% to 150% of Q) the valuation function is usually monotonic[39] - and hence usually the truthful-bidding stop-out price is unique. These types of results appear to be close in spirit to the general idea of Hortacsu's model. Since H is generated by simulation at the submitted sets of prices, irrespective of the quantity bid (at this price), the number of points in the support of H will be quite large. But many of these prices (especially very high and very low ones) are far off from what could (reasonably) be expected to be the stop-out price - and in fact when I simulate H, I find that indeed at these prices are rarely 'hit'. So h is very sensitive at these price points, and can induce large variation in the inferred v. But at prices which are closer to the stop-out price, H is estimated more precisely, and variation in H/h is smoother. Thus even if non-monotonicity manifests at the aggregate level, it doesn't appear to influence the stop-out price.

5. An Application of the Fevrier, Preget and Visser Model

In this section I illustrate how FPV combine the theoretical solution to their auction model from Section 2.2, with additional statistical assumptions, to complete a model that can be applied to data. Let l be the subscript for variables specific to the l-th auction out of a total of L auctions, so $l \in \{1, ..., L\}$. In general, it does not seem sensible to assume that each auction is exactly identical - the number of bidders, N_l, may vary across auction, as may other auction-specific characteristics - which I denote by Z_l. FPV assume that the pairs of random variables (N_l, Z_l) , for $l = 1, ..., L$ are independently and identically distributed. Next, assume that the true value of the bonds in the l-th, V_l, is independent of N_l, but dependent on Z_l. Conversely, an individual i's signal in auction l, S_{il} is independent of N_l, though it does depend on both V_l and Z_l. Conditional on Z_l, the realisations of the t true values, $V_1, ..., V_l$, are assumed independent. Furthermore, FPV assume that the signals $S_{il}, ..., S_{iL}$ are independent conditional on (V_l, Z_l), and also that S_{il} is independent of $S_{il'}$ conditional on Z_l and $Z_{l'}$, for all $l \neq l'$.

To model the distributions of V and S which satisfy the above assumptions, FPV assume a parametric framework. This allows them to restrict their attention (in estimation) to a finite-dimensional vector of parameters, which characterises the distributions of interest, rather than modelling these distributions non-parametrically. Thus denote the distribution of V_l, conditional on $Z_l = z$ as $F_{V|Z}(.|z, \theta_1)$, where θ_1 is the set of parameters characterising the distribution. Analogously, denote the distribution of S_{il}, conditional on $V_l = v$ and $Z_l = z$ as $F_{s...}(...\theta_2)$, with θ_2 denoting parameters specific to this distribution. Given these two distributions, I can recursively derive the distribution of the signal S_{il} itself, conditional only on

[39] There is one exception, in 145 auctions.

$Z = z$ - denote this $F_{s|Z}(..\theta)$, with $\theta = (\theta_1, \theta_2)'$. I assume that the true value of the parameters describing my data is θ^0, and it is these parameters I wish to estimate.

Notice that thus far, we are still unable to estimate the model - the values of interest (the signals) are not directly observed. The next step in solving the FPV model is to link the S_{il} to quantities we can observe (or measure) from the data. To this end, I assume that each bidder has an optimal bidding strategy of $q(p, s, n, z|\theta)$ - this denotes the quantity which the bidder demands at price p, in an auction where his signal is s, auction-specific covariates take the value z, and the true parameters equal θ^0. Recall that in Section 2.2, we have assumed that the equilibrium bidding functions are strictly monotonic - this means they admit an inverse, $x^{-1}(p, s, n, z|\theta)$. Given this notation, define $G(q|n, z, p)$ as the distribution function of the equilibrium demand, conditional on Z=z, V=v. Using this definition, and the assumptions on the dependence structures of V, S and Z:

$$\begin{aligned} G(q|n, z, p) &= \Pr\left(q\left(p, S_{il}, N_l, Z_l, \theta^0\right) \leq q|N_l = n_l, Z_l = z_l\right) \\ &= \Pr\left(q\left(p, S_{il}, n_l, z_l, \theta^0\right) \leq q|N_l = n_l, Z_l = z_l\right) \\ &= \Pr\left(S_{il} \geq q^{-1}\left(p, S_{il}, n_l, z_l, \theta^0\right) \leq q|Z_l = z_l\right) \\ &= 1 - F_{S|Z}\left(q^{-1}\left(p, S_{il}, n_l, z_l, \theta^0\right)|z, \theta^0\right) \end{aligned}$$

The argument from the first line to the second follows from conditioning on $N_l = n_l$, step from second to third line follows from strict monotonicity of the optimal strategy q (the final step follows by definition of $F_{S|Z}$).

We thus have a connection between the distribution of (equilibrium) demands and the distribution of signals. Since I observe individual bidder's bids, and these are assumed to be part of their equilibrium demands, I can use this data to construct a non-parametric estimate of the distribution function G. I follow FPV in doing this non-parametrically using kernel-density estimation:

$$\hat{G}(q|n, z, p) = \frac{\sum_{l=1}^{L} K\left(\frac{n-n_l}{h_N}, \frac{z-z_l}{h_Z}\right)\frac{1}{n_l}\sum_{i=1}^{n_l} 1_{\left(x_{ilp} \leq x\right)}}{\sum_{l=1}^{L} K\left(\frac{n-n_l}{h_N}, \frac{z-z_l}{h_Z}\right)} \tag{10}$$

where $K(.,.)$ is a multivariate kernel, h_N is the bandwidth parameter associated with the number of bidders, and h_Z is a vector of bandwidths (one for each per-auction covariates to be modelled).

Now that I have an empirical connection between the F and G distributions, I need to re-write equation 9 in terms of G, which will make the Euler equation a function of only the data and the θ, the parameters I wish to estimate:

$$\begin{aligned} &0 = E\{w(N_l, Z_l) \cdot (N_l - 1) \cdot 1_{[P_l^0 \leq p]} \cdot \\ &\cdot [E[V_l|S_{1l} = F_{S|Z}^{-1}\left(1 - \hat{G}\left(q_{1lp}|n_l, z_l, p\right)|z_l, \theta, ...\right), ..., \\ &S_{n_l l} = F_{S|Z}^{-1}\left(1 - \hat{G}\left(q_{1n_l p}|n_l, z_l, p\right)|z_l, \theta, ...\right), N_l = n_l, Z_l = z_l] - p)\} - \\ &\qquad - E[w(N_l, Z_l)\left(p - P_l^0\right) \cdot 1_{[P_l^0 \leq p]}]\} \end{aligned}$$

The corresponding empirical moment is described by:

$$m_t\left(q_{1lp_t},...,q_{n_L Lp_t},n_1,...,n_L,p_1^0,...,p_L^0,z_1,...,z_L,p_t,\theta\right)$$
$$=\frac{1}{L}\sum_{l=1}^{L} w\left(n_l,z_l\right)\cdot\left(n_l-1\right)\cdot 1_{[P_l^0\leq p]}\cdot$$
$$\cdot\{E[V_l|S_{1l}=F_{S|Z}^{-1}\left(1-\hat{G}\left(q_{1lp}|n_l,z_l,p\right)|z_l,\theta,....\right),...,$$
$$S_{n_l l}=F_{S|Z}^{-1}\left(1-\hat{G}\left(q_{1n_l p}|n_l,z_l,p\right)|z_l,\theta,....\right),N_l=n_l,Z_l=z_l]-p\}$$
$$-\frac{1}{L}\sum_{l=1}^{L} w\left(n_l,z_l\right)\cdot\left(p_t-p_l^0\right)\cdot 1_{[P_l^0\leq p]}$$

This is a moment condition that I can use for GMM estimation, and in theory should be satisfied for any p$\in[0,\infty)$, giving a continuity (i.e. infinity) of moment conditions. In practice, we perhaps should not expect the submitted bids to satisfy this condition for any p in this interval - in fact, bidders may have a better 'guess' as to the subset of prices to which they are responding, and the Euler equation may behave erratically if I pick a price that is wildly different from the observed stop-out prices.[40] I follow FPV to evaluate these moment condition at the actual stop-out prices, estimate θ by minimising the sum of weighted squared moments:

$$\hat{\theta}=Arg\min_{\theta}\sum_{t=1}^{T} m_t^2\left(q_{1lp_t},...,q_{n_L Lp_t},n_1,...,n_L,p_1^0,...,p_L^0,z_1,...,z_L,p_t,\theta\right)$$

where $w\left(n,z\right)$ is the weighting function. This gives me T=145 moment conditions.

5.1. Parametric Specification of the FPV model

Having shown how the H-distribution can be replaced by the G-distribution, I now need to explicitly select a parametric specification for $\mathrm{F}_{V|Z}$ and $\mathrm{F}_{S|VZ}$, and relate these to my G-distribution. FPV have picked these distributions in a way that allows them to analytically calculate the equilibrium in a uniform-price auction, thus enabling a direct comparison between the performance of the discriminatory and uniform-price auctions, and I follow their design. The distribution of the true value, conditional on auction-specific covariates $Z_l=z_l$, and the parameter vector θ_l is assumed to follow a distribution that is a combination of the gamma and Weibull distributions:

$$F_{v|z}\left(v|z_l,\theta\right)=\int_0^v \gamma u^{\gamma-1}\frac{\beta^\alpha}{\Gamma\left(\alpha\right)}u^{\gamma(\alpha-1)}\exp\left(-\beta u^\gamma\right)du$$

where $\Gamma\left(.\right)$ is the gamma function, and $\alpha_l=\left(1,z_l\right)*\alpha$ and $\beta_l=\left(1,z_l\right)*\beta$. Here α and β are x-by-1 vectors, so that α_l and β_l can vary from auction to auction. When

[40] Note that even this consideration does not get rid of the continuity of moment conditions that exist.

$\gamma = 1$, the above distribution collapses into gamma distribution, with parameters α_l and β_l. On the other hand, if $\alpha = 1$, the distribution collapses into a Weibull distribution with the parameters γ and β_l.

To close the model, FPV assume the signal distribution, conditional on $V_l = v$, to follow the exponential distribution, parametrised by γ :

$$F_{S|VZ}(s|v, z_l, \gamma) = 1 - \exp(-sv^{\gamma})$$

Given these assumptions, FPV can calculate (analytically) the expected value, $E(V_l|s_1)$ as:

$$\begin{aligned} E[V_l|S_{1l} &= F_{S|Z}^{-1}\left(1 - \hat{G}(q_{1lp}|n_l, z_l, p)|z_l, \theta, ...\right), ..., \\ S_{n_l l} &= F_{S|Z}^{-1}\left(1 - \hat{G}(q_{1n_l p}|n_l, z_l, p)|z_l, \theta, ...\right), N_l = n_l, Z_l = z_l] = \\ &= \frac{\Gamma\left(n_l + \alpha_l + \frac{1}{\gamma}\right)}{\Gamma(n_l + \alpha_l)}\left(\beta_l + \sum_{i=1}^{n_l}\beta_l\left(\frac{1}{\hat{G}^{\frac{1}{\alpha_l}}(x_{ilp}|n_l, z_l, p)} - 1\right)\right)^{-\frac{1}{\gamma}} \end{aligned} \quad (11)$$

we now have a complete econometric model, which can be applied to the data.[41]

The key contribution of FPV's mode is that given the parametric assumptions, we can compute the optimal bidding strategies in a uniform-price auction, and hence also obtain an analytical expression for the stop-out price. Supposing that the true values of the gamma-exponential configuration are $\theta^0 = \left(\alpha^0, \beta^0, \gamma^0\right)$, we have:

$$\begin{aligned} p_l^{0-uniform} &= \frac{1}{1 + \frac{1}{\gamma^0}} E\{V_l|S_{1l} = s_{1l}, ..., S_{nl} = s_{nl}, Z_l = z_l\} \quad (12) \\ &= \frac{1}{1 + \frac{1}{\gamma^0}} \frac{\Gamma\left(n_l + \alpha_l^0 + \frac{1}{\gamma^0}\right)}{\Gamma\left(n_l + \alpha_l^0\right)}\left(\beta_l^0 + \sum_{i=1}^{n_l} s_{il}\right)^{-\frac{1}{\gamma}} \end{aligned}$$

Replacing the true parameter values by their estimated counterparts, I obtain:

$$\widehat{p_l^{uniform}} = \frac{1}{1 + \frac{1}{\hat{\gamma}}} \frac{\Gamma\left(n_l + \hat{\alpha}_l + \frac{1}{\hat{\gamma}}\right)}{\Gamma(n_l + \hat{\alpha}_l)}\left(\hat{\beta}_l + \sum_{i=1}^{n_l}\hat{\beta}_l\left(\frac{1}{\hat{G}^{\frac{1}{\alpha_l}}(x_{ilp}|n_l, z_l, p)} - 1\right)\right)^{-\frac{1}{\hat{\gamma}}}$$

From equation (12) we see that the uniform stop-out price is a constant fraction of the (true) common value. Furthermore, since the FPV specifies that γ is fixed across all auctions, the stop-out price will always be the same fraction of the valuation. The FPV paper provides no justification as to why γ should be a constant, while α and β vary across auctions - though most likely this was done primarily for convenience and a lower computational burden. In my estimation, covariates will also be included in the $\gamma-$vector.

5.2. Inference and Practical Issues with Estimation of the FPV Model

The FPV estimator presented here belongs to a class of two-step estimators, as discussed by Newey and McFadden (1989): the first stage consists of the non-parametric estimation of the G-distribution, while the second stage covers the actual GMM estimation of the θ parameters, given the G-distribution that has been previously

[41] The derivation is presented in FPV (2002), appendix C.

calculated. The calculation of standard errors on the estimated parameter $\hat{\theta}$ is highly involved, and the formulas are not discussed here - the relevant derivations are in the appendix of Fevrier et al (2002).

In addition to testing for the significance of the parameters in the gamma-exponential configuration, I will simulate confidence bounds around stop-out price in the uniform price auction, analogously to what I did in Section 4.2 for Hortacsu's model. For each auction l, I will draw n_l of the recovered individual signals, calculate the V_l from these signals, and consequently the stop-out price. I will repeat the procedure 10 000 times for each auction, and hence obtain the appropriate standard deviations and 'simulated/empirical' error bounds. Finally, I repeat the procedure to evaluate the uniform-price auction over the whole sample: I simulate each auction l once, for $l \in [1, 145]$, and aggregate the uniform-price revenues. I repeat this sequence 10 000 times, and compute the standard deviations and simulated confidence intervals.

Before proceeding to present the results of my estimation, the issue of scaling needs to be mentioned. From Section 5.1 we know that the gamma-exponential configuration gives $E(V^\gamma) = \frac{\alpha}{\beta}$, but if we are expecting V to be near 100, and γ to be (much) larger than 1, the difference in scaling of α and β will be very large. Formally, the FPV model is flexible enough to accommodate scaling changes,[42] but in practice, the magnitudes that (certain) scaling imposes on the estimated parameters is beyond the numerical precision of the estimation software. For example, in FPV's results the β parameters are of the order 10^{-24}, which is beyond the numerical precision of Matlab. My estimation will thus proceed by first re-scaling the data to unity (i.e. face value of bonds is 1, rather than 100), carrying out the estimation and inference, and finally re-scaling the simulation results back by 100, for easier comparability with the results from Hortacsu's model.[43]

The per-auction covariates that I include in my z_l vector include the secondary market price (for both security types) and the maturity of the security, calculated in days from day of issue until maturity (for 2 year bonds). Since the 52-week bills all have the same maturity, the maturity regressor is not included in the model for the 52-week security, and a measure of 'inflation expectations 1 year ahead' is included instead. I do not include nominal yield among my regressors, since both of the analysed securities pay no coupon. Thus my z_l vector is of dimensions 1×2 for both security types. For the kernel function, I follow FPV in using a multiplicative Epanechnikov kernel, so that $K\left(\frac{n-n_l}{h_N}, \frac{z-z_l}{h_Z}\right) = K_1\left(\frac{n-n_l}{h_N}\right) * K_1\left(\frac{z-z_{1l}}{h_{1Z}}\right) * K_1\left(\frac{z-z_{2l}}{h_{2Z}}\right)$, with $K(u) = 0.75\left(1-u^2\right) 1_{|u|\leq 1}$. Bandwidths are selected using the same rule-of-thumb as in FPV, such that:

$$h_{z_i} = 2.214 \frac{\sigma(z_i)}{L^{\frac{1}{7}}}$$

5.3. Results from the FPV Model

For the estimation of the parameters for the FPV model, I ran a separate estimation for each security type.[44] The reason for this is twofold. Firstly, the available

[42] That is, there is a one-to-one mapping from parameters in a model estimated with un-scaled data, and one with scaled data.

[43] This same approach was used by Castellanos & Oviedo (2006), who apply the FPV model to Mexican data.

[44] The estimation was done in Matlab, and the optimization module SolvOpt of Kuntsevich & Kappel (1997) was used as a component of the overall estimation prodcedure.

Table 8: FPV Estimated Parameters - 52 Week Bills

	α	β	γ
Constant	622.3* (0.4097)	547.4e3* (186.1e3)	0.8866 (0.9905)
Secondary Market Price	-235.9* (0.5886)	-488.7e3* (179.9e3)	159.4* (1.104)
Inflation Expectations	1,131.9* (0.3829)	-28.7e3* (6.230e3)	-47.3* (0.6099)

Estimates reported (std.error in brackets).
Denoted with * if difference from 0 rejected at 95% confidence level.

Table 9: FPV Estimated Parameters - 2 Year Bonds

	α	β	γ
Constant	9,9967.7* (0.0004)	7,808.1e5* (3,024.5e5)	0.1627 (4.417)
Secondary Market Price	-159.4* (0.001)	-9,026.6e5* (3,458.4e5)	60.5* (4.417)
Maturity	1,155.8* (0.001)	0.7636e5* (0.2968e5)	-0.0305* (0.0008)

Estimates reported (std.error in brackets).
Denoted with * if difference from 0 rejected at 95% confidence level.

covariates for the estimation vary for the two security types: while on 2-year bonds 'maturity' is a relevant variable, this is not the case for 52-week Bills.[45] For 52-week bills, I have included a covariate for 'inflation expectations one year ahead' instead, since this is relevant for the pricing of zero-coupon bonds.[46] So for each kind of security, I estimate a model with an intercept and two covariates.[47] Secondly, given that the two securities differ in the role they play in an investors' portfolios, there is no a-priori reason to expect the same parameters to hold across both types.

While FPV assume γ to be a constant, my formulation allows γ to vary according to the same auction-level covariates at α and β. The assumption of a constant γ implies that bidders 'shade' by exactly the same proportion in each auction, and from an econometric point of view, such an assumption is not necessary (nor is such a restriction suggested by the underlying economic theory).

For the 52-week bill sample, automatic bandwidth selection picked $h_n^{52week} = 3.71$, $h_{\text{seconary-price}}^{52week} = 0.278$, and $h_{\text{inflation}}^{52week} = 0.712$, while for the 2-year bonds the bandwidths were: $h_n^{2year} = 2.32$, $h_{\text{seconary-price}}^{2year} = 0.0249$, and $h_{\text{maturity}}^{2year} = 65.37$.

Parameter estimates from my model are presented in Tables 8 and 9.

All parameters in the $\alpha-$. and $\beta-$ vectors are significantly different from zero, for

[45] The reason why 'maturity' varies across acutions for 2-year bonds is that usually four auctions are held for the same 'line' of bond. So while nominally all these bonds are classified as 2-year securities, the effective maturity in certain auctions is less. In practice, the first set of bonds issued in a given line has a maturity of 2 years, the second issue a maturity of 1 year 11 months, etc.

For 52-week bills, the maturity is the same in all auctions, and hence doesn't provide any explanatory power on this subset of data.

[46] In FPV's paper, the inflation expectations are not included, since all the securities they consider are interest-bearing, and the yield variable is included in the model.

[47] The data on secondary market prcies was obtained from the Polish Ministry of Finance, and the data on inflation expectations was supplied by Reuters.

both 52-week bills and 2-year bonds. For the $\gamma-$vector, only the constant tests as insignificant for both security types - other elements of the γ vector are significant. Similarly to previous papers using the FPV methodology, such as Fevrier et al. (2002) and Castellanos and Oviedo (2006), the average value of γ is significantly larger than 1.[48] Based on these parameters, I can simulate per-auction revenues from the two auction formats, as well as the inferred common value. Examples of this are given in Table 10

Table 10: Per-auction results for Fevrier Model

No.	DP	UP	Val	RD	DPSim	UPSim	ValSim	RDSim
1	93.58	88.6	96.96	4.98	93.58(0.02)	88.6(0)	96.96(0)	4.98(0.02)
					[93.55;93.61]	[88.59;88.6]	[96.95;96.96]	[4.96;5.01]
2	93.43	88.22	96.59	5.21	93.42(0.01)	88.22(0)	96.59(0)	5.21(0.01)
					[93.4;93.44]	[88.21;88.22]	[96.59;96.6]	[5.19;5.22]
3	93.46	87.76	96.16	5.7	93.45(0.01)	87.76(0)	96.16(0)	5.69(0.01)
					[93.42;93.47]	[87.76;87.77]	[96.15;96.17]	[5.66;5.71]
4	93.51	88.47	95.33	5.04	93.5(0.01)	88.47(0)	95.33(0)	5.03(0.01)
					[93.48;93.52]	[88.47;88.47]	[95.33;95.33]	[5.01;5.05]
5	93.57	88.98	95.82	4.59	93.57(0.01)	88.98(0)	95.82(0)	4.59(0.01)
					[93.55;93.6]	[88.98;88.99]	[95.81;95.82]	[4.57;4.61]

Means reported. St.dev in parentheses, simulated 95 % CI in square brackets below.
Columns 2 to 5 show realised results, columns 6 to 9 are simulated.
RD denotes "revenue difference"

A complete summary of the per-auction results is relegated to the Appendix C in Chapter 2 of Marszalec (2011)[49], while (weighted) average performance of both auction types is summarised in Table 11.

Analogously to Table 6, the 'realised' prices and values are calculated based on a single iteration of the auction, using bidding functions exactly as submitted in the actual auction - the bidders' signals from this single iteration are used to obtain the UP-price, and the common value. Simulated prices and values were obtained by randomly re-sampling (within each auction) the bidding functions and signals for the bidders in that auction, and hence finding the relevant prices and the common value.

The discriminatory auction performance is exactly the same as in the Hortacsu model, as in both cases this auction is treated the same way. Estimates of the average price from the uniform-price auction are now lower than the discriminatory price revenue for both 52-week and 2-year securities, with the average (expected) valuation being higher than the expected revenue from both types of auction, as we would expect. On 52-week bills, the discriminatory auction is revenue-dominant in all auctions when looking at both realised and simulated prices. For the 2-year bonds, discriminatory auction revenue-dominates in 41 cases on both realised and

[48] The average γ across all 52-week bill auctions is 38.0 with standard error of 1.8; the average across all 2-year bond auctions is 30.5 with standard error 0.3. In both cases, t-test for $\gamma = 1$ reject with p-values less than 0.0005.

[49] The full results are also available from the author upon request.

Table 11: Per-auction summary for the FPV Model

	52 Weeks	2 years
Number of Auctions	103	42
DP Realised Price	94.84	89.14
UP Realised Price	92.91	87.00
Realised Value	96.04	89.86
DP Simulated Price	94.83 (1.12)	89.13 (2.16)
UP Simulated Price	92.91 (2.53)	87.00 (2.29)
Simulated Value	96.04 (1.27)	89.86 (2.23)
Realised Rev. Dominance	DP: 103 / UP: 0	DP: 41 / UP: 1
Rev. Dominance - RoT	DP: 103 / UP: 0	DP: 41 / UP: 1
Rev. Dominance - Empirical Bounds	DP: 103 / UP: 0	DP: 41 / UP: 1

Means reported (st.dev in brackets).

Table 12: Revenue Sequence Comparison, FPV Model

	52 Weeks	2 years	Pooled
Number of Auctions	103	145	248
Face Value total	85.30	90.54	175.84
DP Revenue total	80.85 (0.00)	80.65 (0.00)	161.50 (0.00)
UP Revenue total	79.10 (0.00)	78.75 (0.00)	157.86 (0.00)
Value Total	81.85 (0.00)	81.33 (0.00)	163.19 (0.00)
Revenue Difference total	1.75 (0.00)	1.89 (0.00)	3.64 (0.00)
Simulated CI for Rev. Diff.	[1.75 : 1.75]	[1.89 : 1.90]	[3.64 : 3.65]

Amounts in bnPLN.
Means reported (st.dev in brackets).

simulated prices, against the uniform-price auction's dominance in 1 auction only. Notably, in the FPV model results the dominance conclusions are not much altered depending on whether I use the 'rule of thumb' or 'simulated confidence intervals'. The reason for this is that the influence of re-sampling in the FPV model is very small (i.e. the variance of the simulated uniform-price revenue is small), while the average revenue difference is (relatively) large - so even if there was slight 'asymmetry' in the precise confidence intervals, this appears insignificant.

Given the estimated parameters, I can also estimate revenue results aggregated over time, just as I did for the Hortacsu model. The results are presented in Table 12.

This table confirms the intuition from Table 11 that on the Polish data, the discriminatory price auction brings (significantly) higher revenue. On 52-week bills, the revenue difference is 1.75 bnPLN (i.e. 2.2%), while on 2-year bonds it is 1.89 bnPLN (approximately 2.3%). The discriminatory auction also appears to be very effective at extracting value from the bidders: on 52-week bills, the bidders' surplus is only 1.22% of their valuation, and even less on 2-year bonds, reaching 0.83%. Even though this surplus is small, it is still positive, suggesting that bidders manage to account for winners' curse, and don't persistently overbid. Given these estimates, the FPV model suggests that for the Polish setting, the discriminatory price auction out-performs the uniform-price auction.

6. Interpretation and Comparison

Since the two models that I am analysing are not directly comparable (i.e. they are not nested as sub-cases of some more general model and/or one another), all the comparisons I can perform are inherently 'reduced-form': I can compare the output from the simulations in both models. For this comparison, I use results from the basic version of Hortacsu's model - as can be seen in the Appendix A, the basic results are robust to sensible perturbations.[50]

Firstly, while my revenue differences obtained from the Hortacsu model are much smaller than Hortacsu (2002) himself obtains, I can still reject revenue-equivalence in a few auctions, while Hortacsu (2002) can never reject this hypothesis. Given that my 'interpolation' procedure for backing-out $v(q)$ is biased to over-estimate v, and thus favour the uniform-price auction, the fact that the discriminatory auction is revenue dominant (in almost all the cases when I can find dominance) would strongly favour the discriminatory auction. In many auctions where the uniform-price revenue is above discriminatory price revenue, the variance in revenue difference is too large to register the difference as 'significant'.

My results in the FPV model suggest realised revenue-differences which are two orders of magnitude larger than those in the Hortacsu model, and in these also favour the discriminatory price auctions. While in the Hortacsu model the 'realised price' results suggest higher average revenues in the uniform-price auction, this is not the case for the FPV model. We should bear in mind, however, that the Hortacsu uniform-price results are calculated on the basis of 'truthful bidding', and hence an upper bound - the fact that it is only slightly above the DP-price (and not statistically significantly so) suggests that the discriminatory-price auction is in fact doing well.

I also find that the distinction between using the 'rule of thumb' or simulated confidence intervals matters in the Hortacsu model, but not in FPV. Part of the reason for this is the asymmetry of the Hortacsu confidence-intervals - the simulated confidence intervals in the FPV model are more symmetric. Additionally, the revenue difference is much larger in the first place in the FPV model (relative to the standard errors), so even if there is (slight) asymmetry, it will not matter as much.

In terms of revenue, the degree of agreement between the two methodologies depends on which indicator I use. Looking at realised revenues, the two methods disagree in 35 out of 145 auctions, which is just under 25%. When looking at simulation results, I need to distinguish between when the two models 'disagree' and when they are mutually inconsistent. To fix terminology, say that the results

[50] The results change slightly if I model bidder asymmetries - and in this case they become more favourable to the discriminatory-price auction.

When smoothing of the H-functions is performed, the changes are also slight, but favour the uniform-price auction. The results appear to change a lot if I introduce moving–average type expectations. Both of these modifications, however, generate some problems for the functioning of the model itself. For details, refer to Sections 8.1 and 8.3.

Table 13: Summary of mutual consistency of Hortacsu and FPV models

	Realised Revenues	Rule of Thumb	Simulated CI
Conclusions Differ	35	137	129
Conclusions Inconsistent	35	1	0

from the two models 'differ' if they don't suggest the same dominance outcome, and say that the results are 'inconsistent' if their conclusions are opposite. The difference here is how I treat the cases when either model can't produce a definite dominance conclusion.[51] A summary of the disagreement counts is presented in Table 13. While the two models disagree frequently, the main reasons for this is the large number of auctions for which the Hortacsu model cannot decide on revenue dominance. The number of auctions where the models are mutually inconsistent is tiny: this occurs only once,[52] when looking at dominance results generated using the 'rule of thumb'.

Thus conclusions from these two models are not in conflict. Firstly, in those cases where both FPV and Hortacsu predict superiority of the discriminatory auction, there's is nothing to conflict. In the remaining cases, the truthful-bidding equilibrium in the Hortacsu model may sometimes give more revenue than discriminatory auction, but in the vast majority of cases the revenue difference is not significantly different from zero. For the one case that the Hortacsu model shows the UP-price upper bound as 'significantly above' the DP-price, how close to 'truthful bidding' bidders would actually bid in the uniform-price auction is an open question.

A complete summary of average prices and revenues is provided in Table 14. Curiously, for 52-week bills, the average discriminatory auction price is above the secondary-market equivalent, though by only a miniscule amount. This result is driven by a handful of small bids at very high prices - these constitute roughly 10% of all submitted bids. Potentially, this may also be a data problem: firstly, the MoF did not report wether the secondary market prices were bid or ask quotes, and since the margin is so small, this could be enough to explain the difference. Secondly, from discussions with the banks bidding in these auctions, it emerged that the secondary-market prices for 52-week bills are particularly inaccurate, since the liquidity on these bonds is very low, and most of them are kept to term.[53] However, the 'common value' estimated using the FPV model is above both the discriminatory price, as well as all the uniform-price upper-bound in the Hortacsu model - so overbidding due to winners' curse is unlikely.

In the Hortacsu case, the uniform-price upper-bound is below the discriminatory-price revue on the 52-week bills, and slightly above on the 2-year bonds and the pooled dataset. For the cases where the uniform-price bound out-performs the discriminatory price, the difference is very small, which would suggest that the discriminatory auction is still very effective at revenue-extraction. These findings are

[51] So if one model cannot decide revenue dominance, while the other can, I say that they "differ" on this particular auction. On the other hand, if one model suggests that UP dominates DP, while the other says that DP dominates UP, I say that these conclusions are "inconsistent".

[52] For auction 38, which took place on 2nd of July, 2005.

[53] Despite the fact that the secondary market is illiquid, all dealers are required (by contract) to quote prices for all maturities of bonds that they hold - even if there are (actually) no bonds for sale at this 'market price'. The illiquidity is particularly important on 52-week bills, but less so for the 2-year bonds.

Table 14: Summary of mutual consistency of Hortacsu and FPV models

	52 weeks	2 years	Pooled
Number of Auctions	103	42	145
Realised DP Price	94.79	89.08	91.85
Realised Market Price	94.78	89.56	92.09
Hortacsu Realised UP	94.78	89.10	91.86
FPV Realised UP	92.91	87.00	89.87
FPV Realised Value	96.04	89.86	92.86
Simulated DP Price	94.79 (1.16)	89.08 (2.21)	91.85 (3.36)
Hortacsu Simulated UP Price	94.79 (1.17)	89.46 (2.52)	92.04 (3.36)
FPV Simulated UP Price	92.91 (2.53)	87.00 (2.29)	89.87 (3.04)
FPV Simulated Value	96.04 (1.27)	89.86 (2.23)	92.86 (3.36)
Simulated DP Total Revenue	80.85 (0.00)	80.65(0.00)	161.50 (0.00)
Hortacsu Simulated Total UP Revenue	80.85 (0.02)	80.99 (0.43)	161.85 (0.43)
FPV Simulated Total UP Revenue	79.1 (0.00)	78.75 (00)	157.86 (0.00)
FPV Simulated Total Value	81.85 (0.00)	81.33 (0.00)	163.19 (0.00)
Secondary Market Total Value	80.84	81.09	161.93

Means reported (st.dev in brackets).
Total revenues & values in bnPLN.

Table 15: Summary of Past Results on T-bill models

Author(s)	Method	Country	Result
Hortacsu (2002)	H	Turkey	DP better by 14%, not significant
Kastl (2006)	H-type	Cz. Rep	UP price may exceed "truthful bidding"
Kang & al. (2008)	H	Korea	DP better by 0.1%, significant
FPV (2002)	FPV	France	DP Better by 2%, significant
Castellanos & al. (2006)	FPV	Mexico	UP better by 2%, significant
Armantier & al. (2006)	other ⋆	France	UP better by 4.8%, significant
This Paper	H	Poland	DP significantly better in 11% of data
This Paper	FPV	Poland	DP better by 2.2 %, significant

⋆ The model used by Armantier & Sbai is parametric and includes risk aversion.

supported by the FPV model, where the margin between the estimated common value, and the discriminatory (average) price is small.

The conclusions of my paper, as compared with the past literature on the subject is described in Table 15 - the table indicates the direction of revenue superiority, and whether the results were statistically significant.

As already remarked, Hortacsu's results show large revenue superiority for discriminatory price, but it is not statistically significant. The Kastl (2006) results are only tangentially related to my results, since his paper is unable to calculate the performance of a discriminatory-price auction. However, the results of Kang and Puller (2006) mirror closely my application of the Hortacsu model: the revenue differences are much smaller than in Hortacsu, but (overall) more significant. One caveat that needs to be considered here is that Kang and Puller benchmark both the uniform-price and discriminatory-price auction against the common benchmark of a 'multi-unit Vickrey auction' - so the comparison is not 'direct' as it is in my case.[54] Similarly to my analysis, Kang and Puller also use simulated confidence-intervals, rather than standard-deviations only, which assists them finding significant differences.[55]

All of the other applications of the FPV model, or other semi-parametric models, yield conclusions of significant differences. The magnitude of my estimated revenue differences are similar to those of Castellanos and Oviedo (2006), who directly apply the FPV model to Mexican data; however, the direction of dominance is reversed.

There is a disagreement between the FPV and Armantier and Sbai (2006) papers, which both use French data. The cause of this disagreement may follow from the different economic models that underlie the econometrics. While the FPV estimation is based on Wilson's (1979) economic model, the economic theory behind the structural model in Armantier and Sbai (2006) builds on that of Wang and Zender (2002). The implementation of the model with risk-aversion and asymmetry is complex since the equilibrium cannot be solved for 'exactly', and hence must be numerically approximated. Given the significant computational and programming burden of this modification, I have not tested the Armantier and Sbai model on my data.

Overall, however, the magnitudes of revenue differences from my application of the two models are not wildly different from other (statistically significant) results obtained in other countries, suggesting that the performance of these models is relatively stable. Hence their mutual comparison can be considered meaningful.

When interpreting the results from both models, it is also essential to emphasise the limitations. In particular, the results tell us that the discriminatory auction performs no worse than the uniform-price auction, and very likely is revenue-dominant, if all other aspects of my reference market are kept fixed. Thus from the MoF's viewpoint, switching to uniform-price auctions would be unattractive, and could lead in up to 2.2% increase in debt-servicing costs (as suggested by the FPV model).

If a switch of auction method were enacted, it is possible that other factors aspects of the auction setup would also be changed. For example, Back and Zender (2001) mention that collusion may be more easily sustained under uniform-price rules - and the MoF might want to change the market rules to reduce likelihood of such outcomes. Within the current legislative framework, one tool readily available is more active management (restriction) of supply.[56] A more radical remedy would be to change the participation rules in the auction - for example, by making a larger number of dealer licences available (at lower administrative cost), or abolishing the dealer system altogether. In this last case, a particularly strong effect may manifest if new bidders (i.e. ones that did not bid as 'clients' before) enter the market, in

[54] The reason for a comparison of this type is that Kang & Puller observe (over different time-periods) data from uniform-price and discriminatory auctions, for Korea. Since the Hortacsu model doesn't admit an 'optimal bidding for either auction type, thir alternative is to use an auction where 'truthful bidding' is optimal - the Vickrey auction satisfies this condition.

[55] The use of simulated confidence-intervals in this paper was implemented before the publication of the Kang & Puller (2006) paper - so this modification appears to have been arrived at independently in these two papers.

[56] In Back & Zender's (2001) model, endogenous supply selection eliminates the 'collusive seeming' equilibria.

addition to the 'disaggregation' of client demands. Any of these variations on the auction setup would reduce the direct applicability of my analysis.

7. Conclusions

This paper has applied two structural econometric models of share auctions to a dataset of Polish treasury-bill auctions. To my knowledge, this is the first consistent application of the Hortacsu and FPV models in the Polish context, as well as the first time that these two models have been directly compared on a single dataset. Both models suggest that the discriminatory-price setup, currently used, is not revenue-dominated by the uniform-price auction. In the FPV model, the discriminatory auction had a revenue advantage of roughly 2.2%, and was statistically significant in all but one of individual auctions, and also statistically significant on the whole revenue sequence. While the Hortacsu model didn't confirm revenue-dominance for the discriminatory auction on aggregate, such dominance was nonetheless found on 11% of individual auctions. Although there was widespread disagreement among the (simulated) predictions of the two models, these were driven largely by the weakness of the conclusions available from he Hortacsu model, which only gives us a one-sided revenue bound for the uniform-price auction. The degree of outright inconsistency was tiny, and in terms of realised-price based comparisons, the discrepancies were moderate, covering 25% of the auctions. It thus appears that both models are useful in analysing the treasury-bill market, despite using differing theoretical and econometric assumptions.

It could also be argued that using two different methodologies on the same dataset has reduced the likelihood that my conclusions are model-driven. Both models suggest that on the Polish data, the discriminatory-price auction performs (at least weakly) better than uniform-price, and there is nothing inherent in the setup of either model that pre-disposes them to favour one auction type over the other. Thus when we 'let the data speak', it is telling us that discriminatory-price auctions are revenue-superior.

In terms of policy advice, my research would recommend keeping the current system, if it is auction revenue that is of primary importance to the MoF. However, given the flexible mandate of the MoF on treasury-bill auctions, combined with the possibility of new participation rules, if the auction setting changes together with a change in auction rules, my results should be interpreted with caution.

Future work in this area will attempt to procure data on 5- and 10-year bonds also, since these constitute a much larger, and more liquid, market than either of the two zero-coupon papers considered in this paper. Those bonds are also coupon-bearing, hence could exhibit different auction behaviour from the data used thus far. Such augmentation would also be particularly useful for checking the influence of the (accuracy of) secondary-market prices on the behavior of the FPV model.

8. Appendix A: Extensions of the Hortacsu Model

This appendix explores three kinds of possible modifications to the Hortacsu model. The modifications concern the modelling of bidder asymmetries, trying to deal with the zero-density problem by using kernel smoothing, and finally, implementing different kind of modelling of bidder expectations.

8.1. Accounting for Asymmetry

The basic resampling algorithm specified in section 4.1 re-draws the bidding functions of each bidder with equal probability. It is thus possible (say) that drawing 11 re-sampled bidding functions might give us six times the bids of bidder 1 and five times those of bidder 5. If all bidders indeed are symmetric, this is not a problem, but as Section 3.2 suggested, there appear to be differences in the size of bidders present. Even if we end up drawing the relative right amount of each bidders' bidding functions across all re-samples, we might still often get 'wrong proportions' of different types of bidders within an individual auction.

To see if correcting for asymmetry is important in my data, I can modify the re-sampling procedure to always draw 'the right proportion' of each kind of bidder. For this purpose, I classify each of the bidders into one of three groups, based on their average 'proportion of bids satisfied'. I then modify the re-sampling algorithm from in the following way:

- Fix an auction l - for this auction, find n_l, the total number of bidders, and n_a, n_b and n_c - the numbers of bidders in each of groups A, B and C, respectively. Fix a bidder index, i.
- Let m be the group to which bidder i belongs; draw $n^m - 1$ bids from the set of bids submitted in group m, and draw exactly n^k for all groups k other than m.

Using the modified procedure will always draw the correct proportions of each bidder type. A similar modification can be made when simulating the average stop-out price in the uniform price auction - so that again I have the correct proportions of each bidder type participating. The Figure 8 below compares the inferred valuation functions for bidder 10, in auction 5; the standard valuation function is in RED, while the BLUE one has been obtained by taking asymmetry into account. The differences in the shape of these two curves are slight, and there is no persistent valuation 'shift' is visible (for most pairs of compared valuation functions, the functions intersect at least once).

I have re-run the whole estimation procedure using these modified re-sampling routines, and results are reported in Table 16.[57]

Comparing this with Table 6, we see that the average realised revenue from discriminatory-price auctions has remained unchanged (up to 2 decimal points). Looking at realised prices, for 52-week bills,the conclusions are unchanged with discriminatory price-being dominant in 80 auctions, but for 2-year bonds discriminatory auction is revenue dominant in two fever auctions than before. Simulated uniform-price revenues on 52-week bills have decreased slightly compared to the 'symmetric'

[57] Full listing of per-auction results is available from the author on request.

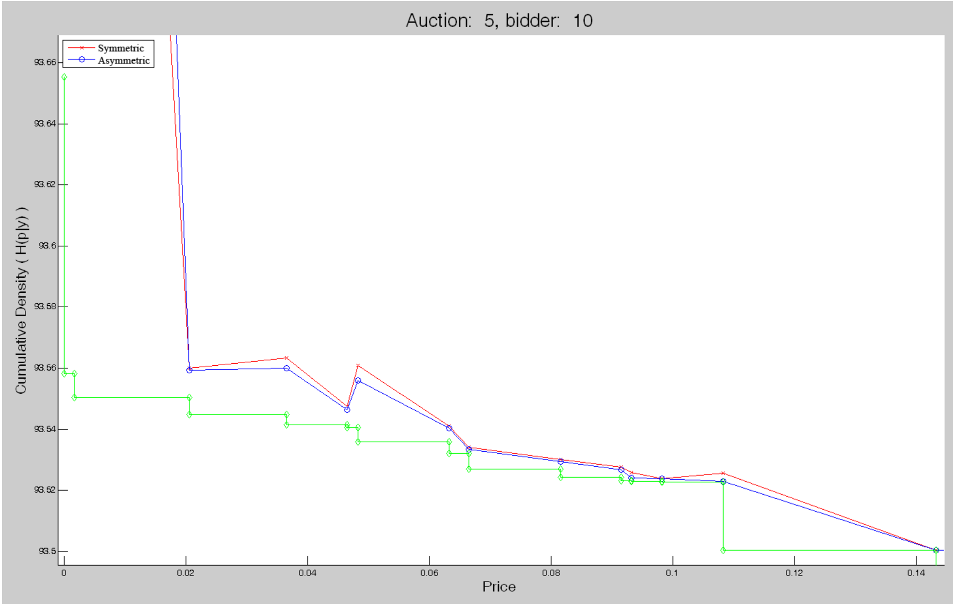

Figure 8: Effects of incorporating 'asymmetries' into resampling

Table 16: Per-auction summary for Hortacsu model, with asymmetries

	52 Weeks	2 years
Number of Auctions	103	42
DP Realised Price	94.79	89.08
UP Realised Price	94.78	89.10
DP Simulated Price	94.79 (1.16)	89.08 (2.21)
UP Simulated Price	94.78 (1.17)	89.24 (2.21)
Realised Rev. Dominance	DP: 83 / UP: 20	DP: 26 / UP: 16
Rev. Dominance - RoT	DP: 15 / UP: 0	DP: 2 / UP: 0
Rev. Dominance - Empirical Bounds	DP: 26 / UP: 0	DP: 3 / UP: 0

Means reported (st.dev in brackets).

benchmark, while for 2-year bonds the uniform-price revenue has decreased proportionately much more. Though on 2-year bonds the uniform-price bound is still above discriminatory-price revenue, the margin is only 0.27%. Overall, on simulated revenues it appears that incorporating asymmetry into the resampling procedure has favoured the discriminatory auction. The number of cases in which this auction is dominant for the 52-week bills has increased both when using the 'rule of thumb' (from 8 to 15), or empirical error bounds (from 15 to 26). Similarly, the 'rule-of-thumb' results for the 2-year bonds are now more favourable for the discriminatory price auction (up from 0 to 2), when I use empirical error bounds, the number of auctions in which discriminatory auction dominates has gone up from 24 to 26. Thus on a per-auction basis, modelling asymmetries reinforces the superior performance of the discriminatory price-auction, relative to uniform-price.

The 'revenue sequence' results in Table 17 below show that on aggregate, asym-

Table 17: Revenue Sequence Comparison: Hortacsu Model with Asymmetries

	52 Weeks	2 years	Pooled
Number of Auctions	103	42	145
Face Value total	85.30	90.54	175.84
DP Revenue total	80.85 (0.00)	80.65 (0.00)	161.50 (0.00)
UP Revenue total	80.85 (0.01)	80.80 (0.20)	161.65 (0.20)
Revenue Difference total	0.00 (0.01)	-0.15 (0.19)	-0.15 (0.19)
Simulated CI for Rev. Diff.	[-0.03 : 0.01]	[-0.50 : 0.01]	[-0.50 : 0.02]

Amounts in bnPLN.
Means reported (st.dev in brackets).

metries have changed little. On aggregate, I sill cannot reject revenue equivalence of the two auction types.

8.2. Smoothing of the H-distribution

The simplest way to introduce some density to those parts of H which are previously 'flat' (i.e. where h is zero) is to introduce kernel smoothing. Smoothing is widely used in non-parametric density estimation, and discussed in for example Pagan and Ullah (1999). In my context, to obtain smoothed estimates of H, I can take a set of re-sampled stop-out prices - call this set $\{p_j^*\}$ - and then go through each price in the support of H, applying the following formula to each step:

$$f(p_i) = \frac{1}{B * \lambda_K} \sum_{j=1}^{B} K\left(\frac{p_i - p_j^*}{\lambda_K}\right) \tag{13}$$

where K is the kernel function, and λ_k is the bandwidth. Hortacsu's original procedure is nested in this approach - the kernel function is just the indicator function $K\left(\frac{p-p_j}{\lambda_k}\right) = 1_{\left(\frac{p-p_j}{\lambda_k}\right)=0}$, with $\lambda_k = 1$. When a non-degenerate kernel is used, a 'hit' on a stop-out price of p^* also generates a bit of density to on prices close to p^*. So if previously there was a set of prices which were clustered closely together, but only one of them would be hit in equilibrium, smoothing would disperse a bit of density to the other points also, which would solve non-existence problem of the valuation function at these points.

There are two reasons to expect kernel smoothing to be of limited use in my context. Firstly, if I use a kernel which has finite support (e.g. Epanechnikov), it is unlikely that I will get much of an improvement in estimating prices which are at the extreme ends of the support: if at these points there was no density to begin with, then even if I 'smooth' the small amounts that are present in this part of the distribution, I should not expect much of a change. The second worry is more fundamental: if I do get some density at points where there previously were none, this addition is 'artificial' in the sense that now I have density in places where, according to the *economic* foundations of my model, there shouldn't be any. What amount of density is added will depend on the parametrisation (shape) of my kernel smoother, and this may drive the shape of the v-functions. Since the density estimate features as the denominator of Euler equation, small variations in this density may have a very large impact to the valuation inferred at this point. There is no reason to expect that smoothing would induce monotonicity - so at the points where I previously couldn't 'construct' an inferred valuation, I may be able to do so now, but the estimate I get can behave badly.

To check the impact of kernel-smoothing on my results, I applied formula 13 to my estimates, and used the normal kernel with automatic optimal bandwidth selection.[58] The left-hand panel in Figure 9 below shows an example of a smoothed compared to unsmoothed estimates of H, while the right-hand panel compares the resultant valuation functions.

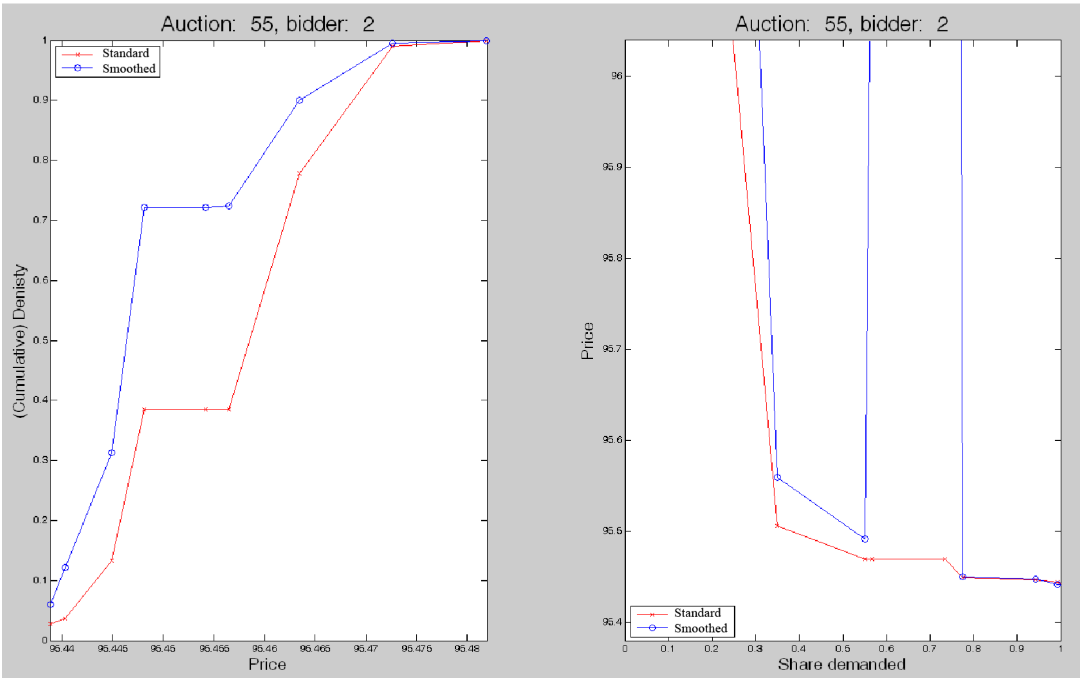

Figure 9: The effect of smoothing H(p|y)

In the LHS of the figure, we see that the smoothed distribution (BLUE) has fixed the zero-density problem of the standard distribution (RED) around the interval [96.45, 96.65] - where the cumulative density was previously flat, it is now (very slightly) sloping upwards. The effects of this can be seen in the RHS of the figure - the values corresponding to the unsmoothed distribution (in red) have had to use my 'interpolation' assumption over this interval, and hence the valuation function is flat - the valuation function corresponding to the smoothed distribution (in blue) has used equation 5 to back-out the valuation function at these points, and is thus not flat. However, the amount of 'smoothed' density over this range is small, and hence the estimates of v are erratic and very large.

The smoothing doesn't always help: there are some points in the middle of the support of H where I previously had zero density, and smoothing has not fixed this.

[58] Matlab was used to perform these calculations.

Table 18: Per-auction summary for Hortacsu model with Smoothed H

	52 Weeks	2 years
Number of Auctions	103	42
DP Realised Price	94.79	89.08
UP Realised Price	94.81	89.50
DP Simulated Price	94.79 (1.16)	89.08 (2.21)
UP Simulated Price	94.85 (1.21)	89.58 (2.73)
Realised Rev. Dominance	DP: 78 / UP: 25	DP: 24 / UP: 18
Rev. Dominance - RoT	DP: 6 / UP: 0	DP: 1 / UP: 0
Rev. Dominance - Empirical Bounds	DP: 10 / UP: 0	DP: 0 / UP: 0

Means reported (st.dev in brackets).

Also, smoothing does little to influence H at the extremes of the price distribution.[59] When I compare the zero-density point ratios with and without smoothing (overall) I find that the improvement is significant. Using non-smoothed H's, I get 'zero density' in roughly 30.9% of points in the support of H, while using smoothing I find zero-density in 15.3% of points. Thus alleviates the zero-density problem, but does not completely eliminate it. As expected, smoothing is of little help at the edges of the support of H - where there is no density to 'smooth over' in the first place.

The conclusions with respect to revenue-dominance have changed slightly by using smoothed densities, as can be seen from Table 18. Looking at realised prices, the discriminatory-price auction does worse on the 52-week bills: number of cases when it is dominant has gone down from 83 to 78. There is no change on the 2-year bond sample. Looking at simulated prices, the uniform-price auction now does better, and this slight increase in the average uniform-price stop-out price is due to the effect shown on Figure 9. Some previously 'monotonised' parts of the v-function that were 'marginal' have, due to smoothing, shifted upwards - hence a higher stop-out price. However, the influence is minor. The improved performance of the uniform-price auction is reflected in the per-auction dominance numbers on the 52-week sample: discriminatory price is revenue-dominant only 6 times using the RoT criterion (down from 8 cases previously), and only 10 times using simulated confidence intervals (down from 15 cases). Curiously, there is no corresponding shift on the 2-year bond sample: for the 2-year bonds, I can reject revenue-equivalence only in one case using the RoT and not at all using simulated confidence intervals.

Looking at revenue sequences, as shown in Table 19, we still cannot reject revenue equivalence on either individual security type using the 'rule of thumb'. However, when looking at simulated confidence intervals, both on the 2-year and the pooled sample, the unfirom-price revenue is above discriminatory-price, though the margin is still very small, at 0.32%. It thus appears that smoothing the H-distribution favours the uniform-price auction, however, this improvement can be attributed to an artifact of the model: the smoothing induces numerous significant (upwards) non-monotonicities in the individual valuation functions in those places where I had previously used a monotonic interpolation which was already chosen in such a way

[59] Figures showing these two results are not shown here, but available upon request.

Table 19: Revenue Sequence Comparison: Hortacsu Model with Smoothed H

	52 Weeks	2 years	Pooled
Number of Auctions	103	42	145
Face Value	85.30	90.54	175.84
DP Revenue	80.85 (0.00)	80.65 (0.00)	161.50 (0.00)
UP Revenue	80.91 (0.04)	81.11 (0.26)	162.02 (0.26)
Revenue Difference	-0.06 (0.04)	-0.46 (0.26)	-0.52 (0.26)
Simulated CI for Rev. Diff.	[-0.15 : 0.00]	[-1.00 : -0.01]	[-1.06 : -0.06]

Amounts in bnPLN.
Means reported (st.dev in brackets).

as to bias the results in favour of the uniform-price auction. Sometimes the re-sampling algorithm picks valuation functions in a way that makes such a portion of the valuation curve 'marginal', and this pushes up the stop-up price compared to before.

8.3. Moving-Aaverage - Type Expectations

One response to solving the zero-density problem follows from questioning the 'precise' validity of Bayesian equilibrium as applied to the auction itself. While Bayesian equilibrium assumes that each bidder responds optimal to the 'true' distribution of opponents' valuations in the current auction, in practice the bidders may not possess such accurate information - in fact, they might be responding to some 'compound' distribution of current and past rival valuations. A way of trying to implement this line of argument is to introduce some pooling of auctions into the resampling procedure. I can do this by expanding the set of 'eligible bidders' from whose demand functions I draw at the re-sampling stage; I can re-write step 2 of my resampling procedure from section 4.1 as:

- 2a. Draw, with replacement, $(n_l - 1)$ actual demand bid schedules from auctions $l, (l-1), \ldots, (l-k)$

In this context, the variable 'k' specifies how many steps back I want to 'pool'; in a very loose sense, this parameter can be seen 'moving-average' order of the bidder's expectations. Pooling of this sort has been used by Kastl (2006), since his data was originally too sparse per-auction to allow a good simulation of H based on bids submitted in a single auction only. In this section, I will use k=3, since this gives me (approximately) the same number of observed bid-functions to re-sample as in Hortacsu's Turkish context.

The reason I might expect that this modification alleviates the zero-density problem follows from the fact that the prices at which bidders submit their bids vary across auctions. Thus if I pool the bid functions across a number of auctions,we would expect a 'denser' collection of bid across most parts of the support of the price distribution. Due to this, I should expect the re-sampled residual-supply curves to intersect with the current bidder's demand function at a wider range of prices than before, so some prices which were never hit before may now appear as stop-out prices at least occasionally.

Overall, this method is almost as successful as kernel smoothing in reducing the number of zero-density points in the support of H: using the standard procedure, 30.9% of relevant points have zero density - using the MA-resampling procedure, this number falls to approximately 17 %. When I compare the values recovered using H as simulated here, and the original one,we find that the differences in inferred valuations are large - much larger than obtained via kernel smoothing. The reason for this is twofold: firstly, for those points where I previously had zero density, I used 'monotonicity' arguments to 'interpolate' the value function. Now that some of these points have a non-zero density, I can use equation 5 to back-out values at these points - but this sometimes leads to (additional) violations of monotonicity and hence deviates a lot from the standard inferred marginal valuation, at this point. The effect here is exactly analogous to that caused by the smoothing of H.

Secondly, and more worryingly, some of this discrepancy may be due to bias: if the auctions that we are pooling across are 'not very similar', then the MA-based re-sampling procedure will generate a biased estimate of H. In particular, if the price trend is upwards over time, then the MA-procedure will bias the new H-distribution to be first-order-stochastically-dominanted (FOSD'd) by the single-auction based one (conversely if the price trend is downwards). An example of this is provided in figure 10, where in each figure the left-hand pane shows the H-distribution obtained from single-auction resampling in RED, and the one obtained via MA-type resampling in BLUE, while the right-hand pane shows the MA-window plotted on top of the weighted stop-out price over time.

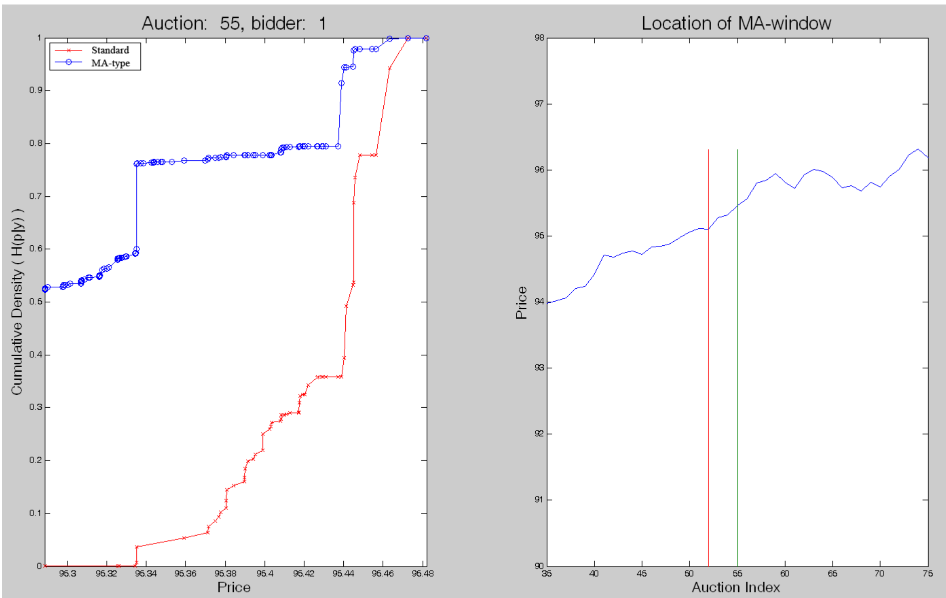

Figure 10: Example of FOSD in the case of MA-type resampling

From the two above figures we indeed see that if the price trend is upwards, the MA-procedure generates an H-distribution that FOSD'd by the single-auction equiv-

Table 20: Results of Kastl's regression-based test

	52 weeks	2 years	Pooled
	q(t-1)	q(t-1)	q(t-1)
Constant	114.1 (1.5)	106.7 (1.2)	112.0 (1.1)
Regression Coefficient	11.1 (10.5)	20.6 (8.1)	14.2 (7.8)
R^2	0.0009	0.013	0.0019
N	1237	512	1749

Estimates reported (st.error in brackets).

alent. A quick check of this intuition could be carried out, for example, by running a series of one-sided Wilcoxon tests, with the 'alternative hypothesis' implying FOSD according to whether the price trend (over the 4 periods) has been increasing or decreasing. Running this testing procedure on the simulated distributions generates rejection (in the expected direction) on roughly 50% of the distributions at the 95% confidence level.[60]

If there is intertemporal dependence between the distribution of bidders' signals across auctions, that will undermine my estimates for two reasons. Firstly, the assumption of independent and identically distributed signals will not hold, whereby the arguments for bootstrap consistency (when estimating H) would no longer apply and the whole estimation theory would be invalid. Secondly, if such intertemporal dependence induces anything like a 'common value' element, or affiliation, to my bidders valuations, then the assumptions underlying my revenue estimates from the uniform auction are no longer appropriate, and would need to be modified (as in, for example, Hortacsu (2002), Section 6.2).

Kastl (2006), who uses pooling in the context of a model very similar to Hortacsu (2002), applies a test for intertemporal dependence. Since the tests I am interested in are applied to the bidders' individually observed signals, I must first estimate them using additional assumptions on top of Hortacsu's methodology. The assumption Kastl makes is that the value is linear and separable in both signal and quantity, in particular:

$$v_i\left(q_i|s_i\right) = s_i + \beta q_i$$

From this assumption, β (common for all bidders) can be estimated by using the first two recovered valuations for each bidder, and once β has been estimated the s_i can be calculated. Subsequently, Kastl regresses $\hat{q}_i^{l-1}$ on s_i^l, and looks at the significance of the regression coefficient as an indicator if independence; if the coefficient is insignificant, we cannot reject independence. In Kastl's sample, the test fails to reject.[61]

I performed the same tests as Kastl, and the results are summarised in Table 20, below.

We see that the test based on lagged won quantities does not reject independence on the 52-week sample, or on the pooled sample of 52-week and 2-year data; it

[60] Detailed results available upon request.

[61] Kastl also applies a second test to his data, to evaluate whether affliation is present or not. However, his test relies on the uncertainty of supply, which is not a feature of my data - hence I omit the application of this test.

Table 21: Per-auction summary for Hortacsu model, with MA3 Expectations

	52 Weeks	2 years
Number of Auctions	100	39
DP Realised Price	94.84	88.94
UP Realised Price	95.53	91.25
DP Simulated Price	94.83 (1.16)	88.93 (2.23)
UP Simulated Price	95.70 (3.05)	92.41 (7.86)
Realised Rev. Dominance	DP: 29 / UP: 71	DP: 8 / UP: 31
Rev. Dominance - RoT	DP: 5 / UP: 8	DP: 0 / UP: 5
Rev. Dominance - Empirical Bounds	DP: 9 / UP: 43	DP: 1 / UP: 22

Means reported (st.dev in brackets).

Table 22: Revenue Sequence Comparison: Hortacsu Model with MA Expectations

	52 Weeks	2 years	Pooled
Number of Auctions	100	39	139
Face Value	85.30	90.54	175.84
DP Revenue	78.05 (0.00)	74.38 (0.00)	152.43 (0.00)
UP Revenue	78.76 (0.31)	77.30 (1.94)	156.05 (1.96)
Revenue Difference	-0.71 (0.31)	-2.91 (1.94)	-3.62 (1.96)
Simulated CI for Rev. Diff.	[-1.34 : -0.30]	[-10.15 : -1.21]	[-10.84 : -1.78]

Amounts in bnPLN.

Means reported (st.dev in brackets).

does, however, reject independence on the sample of 2-year data only. These results suggest that I should be particularly cautious in interpreting the 2-year results. However, there are good reasons to have strong reservations about the importance of this test in the first place: the assumption of linear valuations is arbitrary and restrictive, and may in fact be driving the test results. It also does not sit very well with the (otherwise) non-parametric methodology of the Hortacsu model. Finally, the assumption to use the first two recovered valuations may be cause severe bias in the Polish context. Bearing in mind these caveats, I present the results based on MA-type expectations in table21:[62]

This table suggests a very large departure from the results of section 4.3. In terms of realised prices, the uniform-price auction is now dominant in 71 of the 52-week bill auctions (up from 20 previously), and in 31 of the 2-year bond auctions (up from 18). Whereas previously simulated per-auction dominance statistics only favoured uniform-price auctions once, now this auction type is more often dominant using both the RoT criterion as well as simulated confidence intervals. The results from simulated revenue-sequences are presented in Table 22.

In terms of simulated revenue sequences, the uniform-price upper bound now is revenue-dominant on all samples, with the revenue difference being 0.09% on 52-week bills and 3.9% on 2-year bonds (and 2.4% on the pooled sample). The fact that

[62] Per-auction comparison results can be obtained from the author on request.

revenue differences are bigger in 2-year bonds than on 52-week bills is consistent with other incarnations of the Hortacsu model. However, the usefulness of the precise results themselves is dubious. Given the results from Kastl's test, we should have most faith in results from the 52-week sub-sample, but even here the conclusions are contaminated by the H-functions being mis-estimated by pooling bidders across auctions (due to the moving-average formulation). These problems are sufficiently severe to significantly undermine the validity of the results from this section.

REFERENCES

Armantier, Olivier & Sbai, Erwann, 'Estimation and comparison of treasury auction formats when bidders are asymmetric', *Journal of Applied Econometrics*, vol. 21, no. 6, 745–779 (2006).

Ausubel, Lawrence M. & Cramton, Peter, 'Demand Reduction and Inefficiency in Multi-Unit Auctions', *working paper, University of Maryland* (2002).

Back, K. & Zender, J.F., 'Auctions of divisible goods with endogenous supply', *Economics Letters*, vol. 73, no. 1, 29–34 (2001).

Efron, B., 'Jackknife-after-bootstrap standard errors and influence functions', *Journal of the Royal Statistical Society. Series B. Methodological*, vol. 54, no. 1, 83–127 (1992).

Efron, B. & Tibshirani, RJ, *An Introduction to the Bootstrap* (Chapman and Hall/CRC, Boca Raton, 1998).

Fevrier, P., Preget, R. & Visser, M., 'Econometrics of Share Auctions', *mimeograph, University of Chicago and INRA* (2002).

Fevrier, P., Preget, R. & Visser, M., 'Econometrics of Share Auctions', *mimeograph, University of Chicago and INRA* (2004).

Hortacsu, A., 'Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market', *Typescript, Department of Economics, University of Chicago* (2002).

Hortacsu, A. & Kastl, J., 'Testing for Common Valuation in the Presence of Bidders with Informational Advantage', *working paper presented at the Penn State auctions conference, April 2008* (2008).

Hortaçsu, A. & McAdams, D., 'Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market', *The Journal of Political Economy*, vol. 118, no. 5, 833–865 (2010).

Kang, Boo-Sung & Puller, S.L., 'The Effect of Auction Format on Efficiency and Revenue In Divisible Goods Auctions: A Test Using Korean Treasury Auctions', *The Journal of Industrial Economics*, vol. 56, no. 2, 0022–1821 (2008).

Kastl, J., 'Discrete Bids and Empirical Inference in Divisible Good Auctions', *working paper, Stanford University* (2006).

Krawczyk, M., 'Estimating Demand in Treasury Auctions: A Normal Cumulative Distribution Function Approach', *working paper, CREED, University of Amsterdam* (2006).

Kuntsevich, A. & Kappel, F., 'SolvOpt: The solver for local nonlinear optimization problems', *Institute for Mathematics, Karl-Franzens University of Graz* (1997).

Marszalec, D., *Essays on Auctions* (D.Phil Thesis, submitted to the University of Oxford, 2011).

Newey, W. K. & McFadden, D., 'Large sample estimation and hypothesis testing', in Engle, R. F. & McFadden, D. eds., *Handbook of Econometrics* (1986).

Pagan, A. & Ullah, A., *Nonparametric Econometrics* (Cambridge University Press, 1999).

Ryu, K., Kim, G.R. & Oh, S., 'Discriminatory vs Uniform Price Auction: Auction Revenue Comparison in the Case of the Korean Treasury Auction Market', *Econometric Society 2004 Far Eastern Meetings*, no. 539 (2004).

Wilson, R., 'Auctions of Shares', *The Quarterly Journal of Economics*, vol. 93, no. 4, 675–89 (1979).

Zender, Jaime F. & Wang, James J.D., 'research articles : Auctioning divisible goods', *Economic Theory*, vol. 19, no. 4, 673–705 (2002).