A Point Decision For Partially Identified Auction Models

Gaurab Aryal
The Australian National University

Dong-Hyuk Kim
CenSoC, University of Technology, Sydney

ANU Working Papers in Economics and Econometrics
569

January 2012 JEL: C11, C44, D44, E61

ISBN: 086831 569 9

A POINT DECISION FOR PARTIALLY IDENTIFIED AUCTION MODELS*

GAURAB ARYAL† AND DONG-HYUK KIM‡

ABSTRACT. This paper proposes a decision theoretic method to choose a single reserve price for partially identified auction models, such as [Haile and Tamer, 2003], using data on transaction prices from English auctions. The paper employs [Gilboa and Schmeidler, 1989] for inference that is robust with respect to the prior over unidentified parameters. It is optimal to interpret the transaction price as the highest value, and maximize the posterior mean of the seller's revenue. The Monte Carlo study shows substantial gains relative to the average revenues of the Haile and Tamer interval.

Keywords: optimal reserve price, statistical decision theory, partial identification, maxmin expected utility.
JEL classification: C11, C44, D44, E61

1. INTRODUCTION

This paper considers the problem of choosing a reserve price using a sample of transaction prices from English auctions with independent private values (IPV). Under a weak behavioral assumption that the winner always obtains a nonnegative surplus as in [Haile and Tamer, 2003] (HT), the paper proposes a decision rule that selects a single reserve price following the framework of [Gilboa and Schmeidler, 1989] (GS).

[Paarsch, 1997] employs the button auction model, which regards observed bids as latent values, to point identify the valuation distribution and

Date: January 21, 2012.
∗ The authors thank Keisuke Hirano, Amit Gandhi, Laurent Lamy and seminar audience at University of Auckland and University of Queensland. Kim acknowledges the financial support from the Faculty of Business Research Grant, UTS.
†The Australian National University. e-mail:gaurab.aryal@anu.edu.au
‡ CenSoC, University of Technology Sydney. e-mail: dong-hyuk.kim@uts.edu.au .

then chooses a revenue maximizing reserve price (RMRP) implied by the estimated valuation distribution.[1] Many English auctions, however, impose a minimum bid increment, and often bidders raise the price by more than the minimum bid increment (jump bidding), thereby invalidating this identifying assumption.[2] As a result, using this assumption may lead to a severely incorrect inference; see HT.

HT employs an incomplete model instead, only assuming that the bidders neither overbid their values nor let the auction terminate at a price they can profitably overbid. For this partially identified model, HT proposes a set estimator for the RMRP (the HT interval). This approach is robust to any misspecification of bidding behavior. But, the HT interval does not provide a practical guidance for the seller to choose a single reserve price because a significant fraction of the interval can be less profitable than zero reserve price (34–52% ; Table 1).

To solve the seller's problem, this paper, while using the incomplete model, chooses a single reserve price employing the maxmin expected utility framework of GS, thereby complementing HT. GS extends the classic expected utility theory to allow the decision maker to have many equally reasonable distributions over the random vector that affects the payoff. GS shows that if he is uncertainty averse, he behaves as if he maximizes the lower envelop of the equally reasonable expected utilities.[3] Maxmin criteria provides a robust framework to select an optimal policy for partially identified models; see [Song, 2010; Kitagawa, 2010, 2011; Menzel, 2011] for recent applications.

To formulate the seller's problem within this framework, the paper considers the parameter vector that indexes the valuation distribution as the random vector that affects the revenue. The paper then divides the parameters into two groups: identified and unidentified. The former indexes the density of the transaction price and the latter captures any discrepancy between the densities of the transaction price and the highest values.[4] The

[1] Following the convention in the literature, the term 'revenue' refers to the seller's expected revenue where the expectation taken with respect to the values density.

[2] There does not exist an equilibrium with jump bidding under assumptions reasonable in many auction settings [Lellouche and Romm, 2009].

[3] Uncertainty aversion means a decision maker prefers known risks to unknown risks.

[4] Within the setting of [Paarsch, 1997], there would be no need for the second group as the two densities would be the same.

seller is assumed to have a unique prior distribution over the identified parameters, but many reasonable priors over the unidentified ones, where a prior is said to be 'reasonable' if it does not contradict the only behavioral assumption – the transaction price does not exceed the highest value. Although each of these priors conveys different information about bidding behavior, if the seller regards them as equally reasonable, then he should maximize the lower envelop of posterior expectations of revenues.

This paper shows that the lower envelop is achieved by the prior equating the densities of the transaction price and the highest value. This follows from observing that the transaction price distribution is first order stochastically dominated by every highest valuation distribution that is supported by at least one of the reasonable priors. The former then gives the greatest lower bound for the stochastic dominance relation of the marginal distributions of independent value, providing the smallest revenue at every reserve price. Therefore, the method interprets the transaction price as the highest value and implements the classic expected utility framework. Furthermore, the method remains optimal for correlated values under the assumption of [Aradillas-López, Gandhi, and Quint, 2011] that the sample of transaction price identifies the distribution of the second highest values.

The next section describes the seller's problem and section 3 develops an optimal decision rule. Section 4 illustrates typical revenue gains over average revenues of the HT interval via Monte Carlo experiments. Section 5 concludes and an appendix collects all computational details.

2. Auction Models

A single indivisible object is auctioned among $m \geq 2$ risk neutral bidders in an English auction. Each bidder i observes only his values $v_i \geq 0$. Assume $v_1, \ldots, v_m$ are drawn independently from an identical, absolutely continuous distribution P_v with density p_v.[5] The auction starts at zero price and at each time bidders raise the standing price $\tilde{y}$ by at least Δ, the minimum bid increment. Bidding more than $\tilde{y} + \Delta$ is known as jump bidding.

[5]Throughout the paper, the upper (lower) case letter denotes the cdf (pdf).

When no bidder is willing to raise $\tilde{y}$ further, the auction allocates the object to the bidder who offers the final bid at the transaction price $y := \tilde{y}$. Following HT, this paper makes the assumption:

Assumption 1. *(A-HT) The transaction price* $y \leq x := \max\{v_1, \ldots, v_m\}$.

This assumption is weaker than the assumptions in HT, which also requires that the bidders do not let their opponents win at a price they can overbid. The data z_T available to the seller consists of i.i.d transaction prices $(y_1, \ldots, y_T)$ from past T auctions, each with m bidders.

Now, consider a seller with data z_T who wishes to choose a reserve price to maximize his revenue in a future auction. The future auction can take any of the 'standard' auctions where a bidder with zero value expects to pay zero.[6] The valuation distribution remains unchanged and the seller's values for the object is zero. [Myerson, 1981; Riley and Samuelson, 1981] show that the reserve price ρ that maximizes the revenue

$$u(P_v, \rho) = \int_0^\infty \max\{\rho, \xi\} dP_v^{m-1}(\xi) - \rho P_v^m(\rho), \tag{1}$$

solves the following first order necessary condition

$$\rho = \frac{1 - P_v(\rho)}{p_v(\rho)}. \tag{2}$$

Since the seller does not know P_v, he cannot use (2) to determine ρ. For this problem, [Paarsch, 1997] proposes to estimate P_v by treating the observed bids as losers values in the button auction model, and use the point estimates in (2) instead of P_v (a.k.a. 'plug-in' method).[7] When $\Delta > 0$ or there is jump bidding, however, the assumption that bids equal values can be unreasonable. HT shows that this assumption can cause a significant bias in the estimation of P_v even for a correctly specified parametric model.

For this reason, HT only assumes that bidders do not overbid their value and do not let the auction terminate at a price that they can profitably overbid. Then, HT partially identifies P_v and constructs a set estimator for the

[6] A standard auction is an auction where the highest bidder gets the object.

[7] When only transaction prices are observed, the button auction model identifies the distribution of the second highest value, which is sufficient for identifying the valuation distribution when values are i.i.d.

RMRP. These set estimates are robust to any misspecification on bidding behavior with jump bidding and $\Delta > 0$.

The HT interval itself, however, does not completely solve the seller's problem. Moreover, the HT interval includes many reserve prices with revenue lower than with zero reserve price. Table 1 shows that such reserve prices are about 34% to 52% of the HT intervals, each of which is obtained from bid samples of 200 auctions generated from a fixed data generating processes (DGP), i.e., a combination of a valuation distribution in Figures 1 and 2 and $m \in \{3, 5\}$ bidders.[8] In most cases, zero reserve price produces substantially higher revenues than average revenues of the HT intervals, see Table 2. This stems from the asymmetric shape of the revenue function. The revenue gradually increases up to the RMRP, marked in the figures, but it drops sharply thereafter, while the upper limit of the HT interval is much higher than the RMRP.

What should then be the criteria to choose a single reserve price? The next section proposes a solution.

3. Γ MAXMIN SOLUTION

This section develops an optimal point decision rule for the seller. The optimality is associated with the seller's preference ordering. The paper assumes that

Assumption 2. *(A-GS) The seller satisfies the axioms (A.1–A.6) in GS.*

(A-GS) coincides with assumptions in the classic expected utility theory, except it allows the seller to weakly prefers any convex combination of indifferent lotteries to each individual one instead of restricting the combination to be indifferent– *uncertainty aversion*.[9]

This section develops the decision rule for auctions with the IPV and further argues that the rule remains optimal even when values are correlated if the distribution of the second highest value is identified.

[8] Section 4 explains each DGP's in detail.

[9] One interpretation of this is that the decision maker prefers to secure himself against a potential loss from a particular risky asset by spreading the risk over the indifferent assets (as in a portfolio management).

3.1. **Independent Private Values:** Let $\theta \in \Theta$ index the distribution of the transaction price y, and $h \in \mathcal{H}$ capture any discrepancy between the distributions of x and y. Since $P_x(x|\theta, h) = P_v^m(x|\theta, h)$, the paper specifies the cdf of y as

$$P_y(y|\theta) := P_v^m(y|\theta, 0), \tag{3}$$

to make them comparable.

The seller has *only* three sources of information: the data z_T, (A-HT) and his subjective beliefs about θ represented by the prior p_θ over Θ. Consider, however, a hypothetical situation in which he also has a conditional prior $p_h(\cdot|\theta)$ over $\mathcal{H}$ for each $\theta \in \Theta$. In such a situation, [Kim, 2012] posits choosing a reserve price as the seller's decision problem under parameter uncertainty.[10] The paper shows that if the seller behaves rationally in the sense of [Savage, 1954; Anscombe and Aumann, 1963], he would maximize the expected revenue given by

$$E\left[u(\theta, h, \rho)|z_T; p_h\right] := \int_\Theta \int_\mathcal{H} u(\theta, h, \rho) p_{\theta,h}(\theta, h|z_T; p_h) dh d\theta, \tag{4}$$

where the posterior density can be obtained via Bayes theorem

$$p_{\theta,h}(\theta, h|z_T; p_h) := \frac{p_\theta(\theta) p_h(h|\theta) \prod_{t=1}^T p_y(y_t|\theta)}{\iint p_\theta(\theta) p_h(h|\theta) \prod_{t=1}^T p_y(y_t|\theta) dh d\theta}. \tag{5}$$

[Kim, 2012] discusses the optimality of this Bayesian approach from the frequentist perspective, and shows that it can produce substantially higher revenues than the plug-in rule.

When h is not identified, however, this approach can be sensitive to the choice of the conditional prior $p_h(\cdot|\theta)$– since (5) can be written as

$$\begin{aligned} p_{\theta,h}(\theta, h|z_T; p_h) &= p_h(h|\theta) \left[\frac{p_\theta(\theta) \prod_{t=1}^T p_y(y_t|\theta)}{\int p_\theta(\theta) \prod_{t=1}^T p_y(y_t|\theta) d\theta}\right] \\ &= p_h(h|\theta) p_\theta(\theta|z_T), \end{aligned}$$

[10] The plug-in approach does not consider the parameter uncertainty because it regards the point estimate of the valuation distribution as the true distribution. Then, under this hypothesis, the plug-in approach 'certainly' maximizes the seller's revenue.

the impact of $p_h(\cdot|\theta)$ on the solution to maximize (4) does not disappear even when T is large. Prior to HT, the literature had employed the button auction model, which can be viewed as a strong prior that equates the distributions of x and y. As HT shows, such an assumption can lead to a misleading inference, when there is jump bidding or $\Delta > 0$. Even though an econometrician can employ a less informative prior in a hope of using weaker assumptions, it would still bear some information about unverifiable bidding behavior.

This paper instead considers a convex set Γ of reasonable conditional priors on h given θ and assumes that the seller regards all the elements in Γ as equally reasonable. A conditional prior $p_h(\cdot|\theta)$ is said to be reasonable if it conforms to (A-HT) for every $\theta \in \Theta$. Formally, Γ is a set of all $p_h(\cdot|\cdot)$ such that, for all $(\theta, h) \in \Theta \times \mathcal{H}$

$$p_\theta(\theta)p_h(h|\theta) > 0 \Leftrightarrow P_x(w|\theta, h) \leq P_y(w|\theta). \tag{6}$$

GS shows that it is optimal for the seller to solve

$$\max_{\rho \in \mathcal{A}} \min_{p_h \in \Gamma} E[u(\theta, h, \rho)|z_T; p_h]. \tag{7}$$

That is, a seller should choose a reserve price that maximizes the revenue in (4) with respect to the most pessimistic prior in Γ.

Definition 1. *A decision rule that solves (7) for every realization of z_T is called the Γ maxmin rule.*

This framework is particularly useful for the partially identified auction model because its policy recommendation is robust to the choice of priors over the unidentified parameter h. Solving (7) is, however, computationally expensive because the 'min' part solves an optimization problem over a space of high dimensional functions for every ρ considered for the maximization problem. [11]

The central result of this paper is that this issue does not arise for the seller's problem. The following proposition establishes that a probability

[11] [Chamberlain, 2000] proposes a computation algorithm for a similar problem, but with a simple utility function.

mass function degenerated at $h = 0$, denoted by δ_0, solves the minimization problem.

Proposition 1. *Under (A-HT) and (A-GS),*

$$\max_{\rho\in\mathcal{A}} \min_{p_h\in\Gamma} E[u(\theta,h,\rho)|z_T;p_h] = \max_{\rho\in\mathcal{A}} E[u(\theta,h,\rho)|z_T;\delta_0] \tag{8}$$

Proof. Consider any $(\theta,h) \in \Theta\times\mathcal{H}$ for which there is some $p_h \in \Gamma$ such that $p_\theta(\theta)p_h(h|\theta) > 0$. Then, (6) implies $P_x(w|\theta,h) \le P_y(w|\theta) \Leftrightarrow P_v^m(w|\theta,h) \le P_v^m(w|\theta,0)$ for all $w \in \Re_+$ and $m \ge 2$. Thus, $\rho P_v^m(\rho|\theta,h) \le \rho P_v^m(\rho|\theta,0)$, for any $\rho \in \mathcal{A} \subset \Re_+$. Moreover, since $P_v^{m-1}(\xi|\theta,h) \le P_v^{m-1}(\xi|\theta,0)$, we have

$$\int_0^\infty \max\{\rho,\xi\}p_v^{m-1}(\xi|\theta,0)d\xi \le \int_0^\infty \max\{\rho,\xi\}p_v^{m-1}(\xi|\theta,h)d\xi.$$

These inequalities then imply $u(\theta,h,\rho) \ge u(\theta,0,\rho)$ (see (1)) and hence

$$E[u(\theta,h,\rho)|z_T,p_h] \ge E[u(\theta,0,\rho)|z_T,p_h] = E[u(\theta,h,\rho)|z_T,\delta_0].$$

Finally, $\delta_0 \in \Gamma$ because it implies $x \le y$ with equality. □

This result implies that choosing the worst prior amounts to treating the transaction price as the highest values.

3.2. **Correlated Private Values:** [Aradillas-López, Gandhi, and Quint, 2011] considers a more general auction model with correlated value for the same type of data set as in this paper but under a stronger assumption that

Assumption 3. *(A-AGQ) The transaction price is equal to the second highest values.*

Let the bidders private values $(v_1,\ldots,v_m)$ be distributed as $P_v(\cdot,\ldots,\cdot|\theta,h)$. Now, the distributions of the highest and second-highest values are not necessarily linked through the identical marginal valuation distribution, in particular $P_x \ne P_v^m$. The revenue function is

$$u(\theta,h,\rho) = \int_0^\infty \max\{\rho,\xi\}dP_y(\xi|\theta) - \rho P_x(\rho|\theta,h), \tag{9}$$

which is more general than (1). Under (A-AGQ), z_T point identifies $P_y(\cdot|\theta)$, but not $P_x(\cdot|\theta,h)$. Hence, (9) can only be partially identified; see [Aradillas-López, Gandhi, and Quint, 2011]. In particular, it is bounded below by

$u(\theta, 0, \rho)$. Now, consider the Γ with property (6). Then, it is straightforward to show that the seller would also solve (8) with the revenue (9) as formalized below.

Proposition 2. *When the value are correlated, under (A-GS) and (A-AGQ), (8) holds true for the revenue defined in (9).*

4. COMPARISONS BETWEEN Γ-MAXMIN & HAILE AND TAMER

This section compares performances of the Γ-maxmin rule with the average revenues associated with the HT interval for four valuation densities in Figures 1 and 2, which also show the associated revenue functions for for bidders $m = 3, 4, 5$. Figure 1 is associated with the density similar to an exponential distribution as well as the long-tailed density.[12] Similarly, Figure 2 is associated with lognormal densities with alternative parameters.[13] Then, for each of these valuation densities, the experiments consider $T \in \{100, 200\}$ sample sizes and $m \in \{3, 5\}$ bidders, leading to a total of 16 experiments. Each experiment, i.e., each triplet of (values density, T, m), conducts 1,000 Monte Carlo replications.

For each experiment, the seller selects a bidder randomly and uses $\Delta_{\tilde{y}} := 0.05 \times \tilde{y}$ as the minimum bid increment rule. The chosen bidder bids exactly $\tilde{y} + \Delta_{\tilde{y}}$, as long as it is less than his value. Each replication uses *only* transaction prices to implement the Γ maxmin approach, but uses *all* bids to implement the HT approach for a comparison with the tightest HT interval. Then the corresponding revenues are computed, where the revenue under HT is defined to be the average revenue across the interval. Note that HT does not propose any particular method to choose the reserve price from the interval.

[12] These densities have the form of (10) with $k = 15$. For the exponential-like density, the parameter values are θ :=(0.3548, 0.2350, 0.1486, 0.0946, 0.0466, 0.0440, 0.0217, 0.0119, 0.0089, 0.0080, 0.0084, 0.0081, 0.0049, 0.0028, 0.0017) and for the long-tailed density, θ :=(0.0748, 0.1403, 0.1871, 0.5145, 0.0009, 0.0009, 0.0009, 0.0009, 0.0009, 0.0009, 0.0009, 0.0750, 0.0009, 0.0009, 0.0002).

[13] The lognormal distributions with $(\mu, \sigma) = (3, 1)$ and (4,1/2) are truncated at the 99-th percentile and rescaled so that their supports are the unit interval. HT employs the lognormal densities that appear in Figure 2 for Monte Carlo studies.

The section reports the average percentage gain of the Γ-maxmin rule over the HT interval. The experiments specify the distribution of the transaction price using the Bernstein cdf:

$$P_y(y|\theta) = \left[\sum_{j=1}^{k} \theta_j \text{Beta}(y|j, k-j+1)\right]^m, \tag{10}$$

where $\theta \in \Delta_{k-1}$, the $k-1$ dimensional unit simplex, i.e., $\theta_j \geq 0$ for all $j = 1, \ldots, k-1$ and $\sum_{j=1}^{k-1} \theta_j \leq 1$, and $\text{Beta}(\cdot|a, b)$ denotes the beta cdf with parameters a and b; see [Petrone, 1999a,b] for a nonparametric Bayesian method that uses (10). This paper employs the model with $k = 15$ with the uniform prior over Δ_{14}.

Table 3 summarizes the main results. Each column stands for the valuation densities and each row for the number of bidders m with different sample sizes T, so that each cell of the table shows the percentage revenue gain. For example, the first row and the first column is associated with $(P_v, T, m) = (\text{Exponential-like}, 100, 3)$ and that the revenue gain of the Γ-maxmin approach over the HT interval is around 20.16%. With five bidders, this gain is around 3.71% (second row). The third and fourth rows collects the results when $T = 200$.

Figure 3 explains these revenue gains. Each panel depicts the distributions of the reserve prices chosen by Γ minimax approach (heavy solid) and the lower and upper bounds for the HT interval (light solid) along with the revenue function. Upper (lower) panels are with $T = 100$ ($T = 200$), and the left (the right) are with $m = 3$ ($m = 5$). The left-upper panel shows that the upper bound of the HT interval is distributed around 0.5, while the revenue function indicates that all the reserve prices larger than approximately 0.3 produces lower revenues than zero reserve price. This implies that a significant portion (33.95%; Table 1) of the HT interval is less profitable than zero reserve price. On the other hand, the Γ maxmin rule is distributed over the area in which the revenue is increasing, selecting higher reserve prices than the lower bound of the HT interval. As a result, the Γ maxmin rule produces larger revenues than the average revenue of the HT interval.

This pattern is commonly observed from all the experiments; see Figures 3 - 6. The rest of the table shows significant revenue gains of the Γ maxmin rule, suggesting that Γ maxmin approach can provide a practical policy recommendation.

5. Conclusion

The literature, since [Paarsch, 1997], has proposed various procedures to determine a reserve price for the revenue maximizing seller of an auction. In particular, for the English auction, the exact procedure depends on the type of data (all bids vs only transaction prices) and on the behavioral assumptions (button auction model vs. incomplete models). All of these procedures except Kim [2012], however, view the seller's decision problem as essentially an estimation problem: estimate the valuation distribution, use these estimates to obtain the RMRP, and study asymptotic properties of the estimates (either a point or a set). Extending the formal Bayesian decision method of Kim [2012], and using GS this paper proposes a solution to choose a single reserve price for an English auction with partially identified valuation distribution. This paper shows that it is optimal to employ the Bayesian method interpreting the transaction prices as the highest values.

Appendix A. Computation

Each Monte Carlo experiment obtains a sample $\theta^1, \ldots, \theta^S$ from posterior distribution using the Metropolis-Hastings algorithm. For this initial step, a flat prior over Δ_{k-1} and a sample z_T^1 is employed for constructing the posterior. Each replication conducts all the inference applying the importance sampling method to $(\theta^1, \ldots, \theta^S)$, the prior for the experiment, and a new sample z_n^l from the given DGP. This section illustrates computational details.

A.1. **Sampling from the Posterior with a flat prior.**

For each pair of (P_v, m) and for the first Monte Carlo replication, The Metropolis Hastings algorithm draws random parameters from the posterior with the sample z_T^1 and the prior given by $p(\theta) = \prod_{j=1}^k p_j(\theta_j)$ with $p_j(\theta) \propto 1$ for $j = 1, \ldots, k$. Let θ^s denote the s-th sample from the Metropolis-Hastings algorithm.

For the experiments with the exponential-like and the long-tailed densities, The true parameter value is used as the initial value for the algorithm, and for the log-normal like densities, a vector of $1/k$ is used for the initial value. Then, at the s-th Metropolis Hastings step, the algorithm updates θ^s component by component from $j = 1$ to k. Let $\theta^{j,s} := (\theta_1^s, \dots, \theta_{j-1}^s, \theta_j^{s-1}, \dots, \theta_k^{s-1})$. Then, the algorithm draws a candidtate $\tilde{\theta}_j \sim q_j(\cdot|\theta^{j,s})$, the proposal density for the j-th component of θ. Let $\tilde{\theta}^{j,s} := (\theta_1^s, \dots, \theta_{j-1}^s, \tilde{\theta}_j, \theta_{j+1}^{s-1}, \dots, \theta_k^{s-1})$, and set $\theta_j^s = \tilde{\theta}_j$ with probability

$$\min\left\{\left[\frac{p_j(\tilde{\theta}_j)}{p_j(\theta_j^{s-1})}\right]\left[\prod_{t=1}^{T}\frac{p_v(y_t^1|\tilde{\theta}^{j,s})}{p_v(y_t^1|\theta^{j,s})}\right]\left[\frac{q_j(\theta_j^{s-1}|\tilde{\theta}^{j,s})}{q_j(\tilde{\theta}_j|\theta^{j,s})}\right], 1\right\} \tag{11}$$

and $\theta_j^s = \theta_j^{s-1}$, otherwise. For each k, we iterate the algorithm $200,000$ times recording every 200-th iteration. Among these 2,000 draws, we employ the last 1,000 for the implementation of the decision rules ($S = 1,000$).

A.2. **Proposal density.** For the proposal density q_j, we employ a Gaussian density with mean θ_j^{s-1} and variance σ_j^2 that is truncated so that $\tilde{\theta}^{j,s}$ belongs to Δ_{k-1}. Then, $\tilde{\theta}_j \in [0, \bar{\theta}_j]$ where the upper for $\tilde{\theta}_j$ is given by

$$\bar{\theta}_j := \sum_{a=1}^{j-1}\theta_a^s + \sum_{a=j+1}^{k-1}\theta_a^{s-1}$$

and also let $\overline{\Phi}(\theta_j^s) := \Phi\left((\bar{\theta}_j - \theta_j^{s-1})/\sigma_j\right)$, and $\underline{\Phi}(\theta_j^s) := \Phi\left(-\theta_j^{s-1}/\sigma_j\right)$. Then, q_j has the form of

$$q_j(\tilde{\theta}_j|\theta^{j,s}) = \left\{\frac{\phi\left((\tilde{\theta}_j - \theta_j^{s-1})/\sigma_j\right)}{\overline{\Phi}(\theta_j^s) - \underline{\Phi}(\theta_j^s)}\right\} 1\left(\tilde{\theta}_j \in [0, \bar{\theta}_j]\right) \tag{12}$$

The inverse CDF method draws $\tilde{\theta}_j$ from (12). Moreover, since $p_j(\cdot) \propto 1$ for all $j = 1, \dots, k-1$ and $\phi(\cdot)$ is symmetric about zero, (11) simplifies

$$\min\left\{\left[\prod_{t=1}^{T}\frac{p_v(y_t^1|\tilde{\theta}^{j,s})}{p_v(y_t^1|\theta^{j,s})}\right]\left[\frac{\overline{\Phi}(\theta_j^s) - \underline{\Phi}(\theta_j^s)}{\overline{\Phi}(\tilde{\theta}_j) - \underline{\Phi}(\tilde{\theta}_j)}\right], 1\right\}$$

The algorithm uses $\sigma_j = 1.54$ for $s \leq 20$ and $\sigma_j = 1.54 \times \text{stdev}\left(\theta_j^1, \theta_j^1, \dots, \theta_j^{s^*}\right)$ for $s > 20$ with $s^* := \min(s, 0.4 \times S)$. This adaptive method is similar to Haario, Saksman, and Tamminen [2005].

A.3. **Replication 2 to 1,000: Importance Sampling.** This paper approximates the posterior mean of the revenue (15) or (17) using the importance sampling. Each l-th Monte Carlo replication constructs the importance weight

$$\omega_B^l(\theta) := \frac{l_n(z_T^l|\theta)}{l_n(z_T^1|\theta)} \tag{13}$$

approximates the Bayes action by

$$\hat{\rho}_B^l(z_T^j; m) := \arg\max_{\rho \in \mathcal{A}} \sum_{s=1}^{S} \left\{ \frac{\omega_B^l(\theta^s)}{\sum_{t=1}^{S} \omega_B^l(\theta^t)} \right\} u(\theta^s, \rho; m) \tag{14}$$

where $\theta^1, \ldots, \theta^S \sim p(\theta|z_n^T) \propto l_T(z_T^1|\theta)$.

A.4. **Revenue Approximation.** The seller's revenue in (1) can be written as

$$\begin{aligned} u(P_v, \rho; m) \quad := \quad & m\rho(1 - P_v(\rho))P_v(\rho)^{m-1} \\ & + m(m-1)\int_{\rho}^{\bar{v}} y(1 - P_v(y))P_v(y)^{m-2} p_v(y) dy \end{aligned} \tag{15}$$

The trapezoid rule approximates all the integrals of (15) using the $J = 1,001$ equidistant reference points on the unit interval, $x_0, x_1, \ldots, x_{1001}$ with $x_j = j/1000$. Let $\mathcal{A} := \{x_j\}_{j=0}^{1,000}$, the set of all feasible reserve price. With a slight abuse of notation, f and F denotes J dimensional vector of the pdf and cdf of the values density evaluated at each $x_j \in \mathcal{A}$, respectively. For this purpose, many statistical softwares (e.g., Matlab) evaluate the pdf and the cdf for the beta distribution and the (log)normal distribution. Let $\tilde{A} := (\tilde{a}_1, \ldots, \tilde{a}_J)'$ with $\tilde{a}_j := 0$ if $j = 1$, and otherwise

$$\begin{aligned} \tilde{a}_j \quad := \quad & \tilde{a}_{j-1} + \frac{F_j - F_{j-1}}{2(x_j - x_{j-1})} \\ \approx \quad & \tilde{a}_{j-1} + \int_{x_{j-1}}^{x_j} F(\alpha|\theta) d\alpha \approx \int_0^{x_j} F(\alpha|\theta) d\alpha \end{aligned}$$

Let $A := (a_1, \ldots, a_J)' := \tilde{a}_J \iota_J - \tilde{A}$. Then, for each j

$$a_j \approx \int_{x_j}^{1} F(\alpha|\theta) d\alpha \tag{16}$$

Define element-wise operators, $C \times D := (c_1 d_1, \ldots, c_J d_J)'$ and $C^g := (c_1^g, \ldots, c_J^g)'$ for any dimension conformable vectors $C = (c_1, \ldots, c_J)'$ and $D = (d_1 \ldots, d_J)'$. Then, the j-the element of

$$u := (m \cdot \iota_J) \times \left(x \times (1 - F) \times F^{(m-1)} + A \right) \tag{17}$$

approximates the revenue (15) at each $x_j \in \mathcal{A}$ under the values density f.

REFERENCES

ANSCOMBE, R., AND R. AUMANN (1963): "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 199–205. 6

ARADILLAS-LÓPEZ, A., A. GANDHI, AND D. QUINT (2011): "Identification and Inference in Ascending Auctions with Correlated Privated Values," *Working Paper, University of Wisconsin-Madison*. 3, 8

CHAMBERLAIN, G. (2000): "Econometric Applications of Maxmin Expected Utility," *Journal of Applied Econometrics*, 15, 625–644. 7

GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153. 1

HAARIO, H., E. SAKSMAN, AND J. TAMMINEN (2005): "Componentwise Adaptation for High Dimensional MCMC," *Computation Statistics*, 20, 265–273. 12

HAILE, P. A., AND E. TAMER (2003): "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy*, 111, 1–51. 1

KIM, D.-H. (2012): "Optimal Choice of a Reserve Price under Uncertainty," *Mimeo*. 6, 11

KITAGAWA, T. (2010): "Estimation and Inference for Set-Identified Parameters Using Posterior Lower Probability," *Working Paper, CeMMAP, UCL*. 2

——— (2011): "Intersection Bounds, Robust Bayes, and Updating Ambiguous Beliefs," *Working Paper, CeMMAP, UCL*. 2

LELLOUCHE, D., AND A. ROMM (2009): "Information Effects of Jump Bidding in English Auctions," *Working Paper*. 2

MENZEL, K. (2011): "MCMC Methods for Incomplete Structural Models of Social Interactions," *Working Paper, NYU*. 2

MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73. 4

PAARSCH, H. (1997): "Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales," *Journal of Econometrics*, 78, 333–357. 1, 2, 4, 11

PETRONE, S. (1999a): "Bayesian Density Estimation Using Bernstein Polynomials," *Canadian Journal of Statistics*, 27, 105–126. 10

——— (1999b): "Random Bernstein Polynomials," *Scandinavian Journal of Statistics*, 26, 373–393. 10

RILEY, J., AND W. SAMUELSON (1981): "Optimal Auctions," *American Economic Review*, 71, 381–392. 4

SAVAGE, L. (1954): *Foundation of Statistics*. Wiley, reissued in 1972 by Dover, New York. 6

SONG, K. (2010): "Point Decisions for Interval-Identified Parameters," *Working Paper*. 2

TABLE 1. Proportion (%) of ρ in HT with $u(0,\cdot) > u(\rho,\cdot)$

Bidders m	Exponential -like density	Longtail density	LogNormal (3,1)	LogNormal (4,1/2)
3	33.9450	41.0042	44.2857	46.3710
5	50.5227	46.7213	51.7094	45.8763

TABLE 2. Percentage Gain in Revenue $\rho = 0$

Bidders m	Exponential -like density	Longtail density	LogNormal (3,1)	LogNormal (4,1/2)
3	-0.8729	2.4371	2.2681	2.0731
5	3.4608	1.9613	2.5633	0.3006

TABLE 3. Percentage Revenue Gain of the Γ maxmin rule

Bidders m	Exponential -like density	Longtail density	LogNormal (3,1)	LogNormal (4,1/2)
$T = 100$				
3	20.1632	2.5702	8.4945	2.5867
5	3.7097	0.8163	3.7629	0.2331
$T = 200$				
3	5.9121	5.0051	7.3387	2.5821
5	5.0957	2.1929	3.8419	0.3489

FIGURE 1. Valuation Densities and Revenue Functions

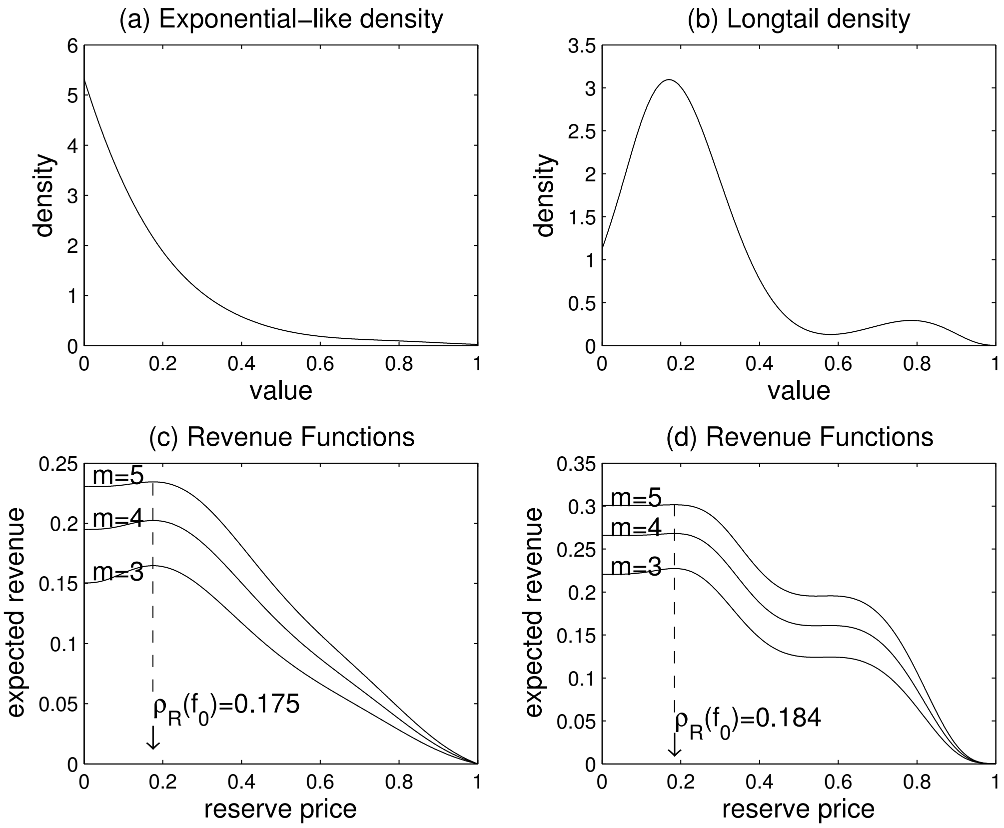

Panels (a) and (c) plot the exponential like density function and associated revenue functions for alternative number of bidders m. Panels (b) and (d) similarly for the longtail values density. On panels (c) and (d), $\rho_R(f_0)$ indicates the revenue maximizing reserve price.

FIGURE 2. Valuation Densities and Revenue Functions

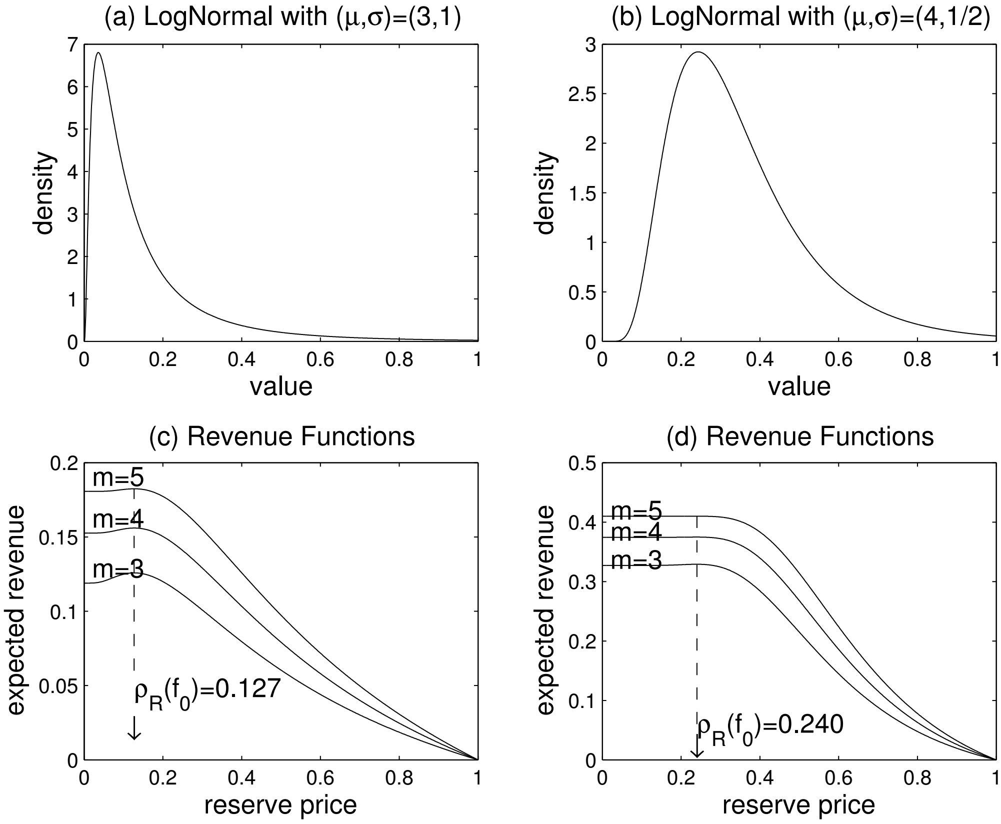

Panels (a) and (c) plot the lognormal density with $(\mu, \sigma) = (3, 1)$ and associated revenue functions for alternative number of bidders m. Panels (b) and (d) similarly for $(\mu, \sigma) = (4, 1/2)$. On panels (c) and (d), $\rho_R(f_0)$ indicates the revenue maximizing reserve price.

FIGURE 3. Exponential-like Distribution

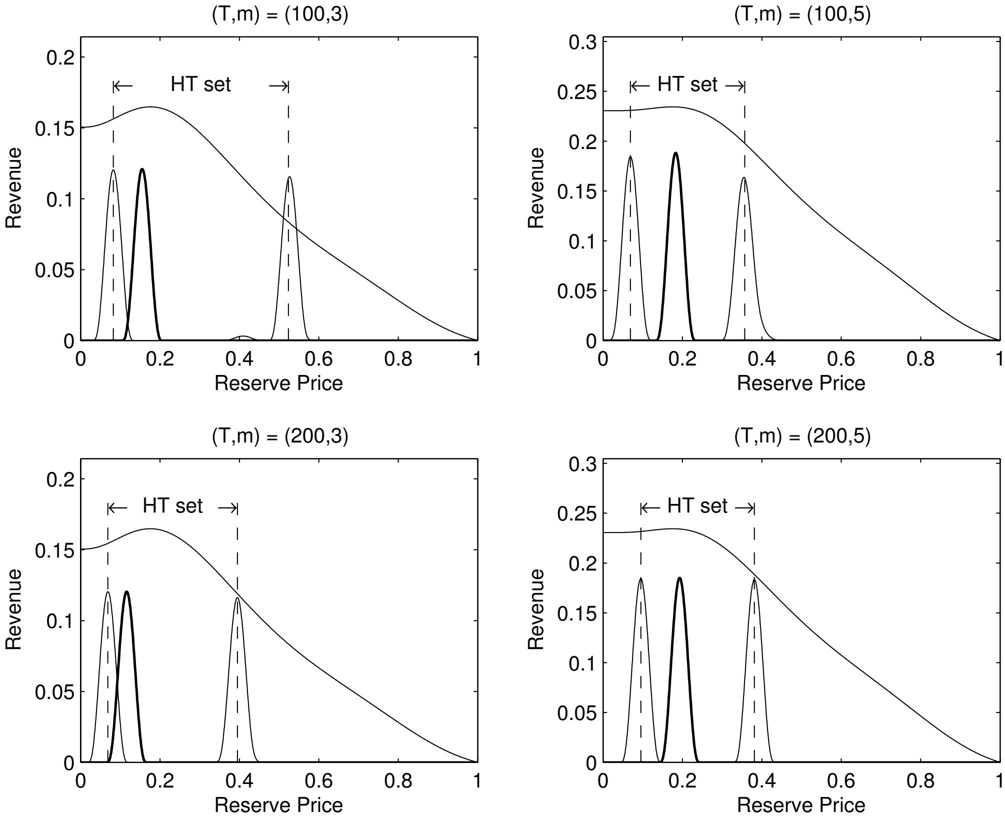

The revenue function is plotted along with the distributions of the lower and the upper bounds for the HT set (light lines) and the reserve prices chosen by the Γ-minmax (heavy).

FIGURE 4. Longtail Distribution

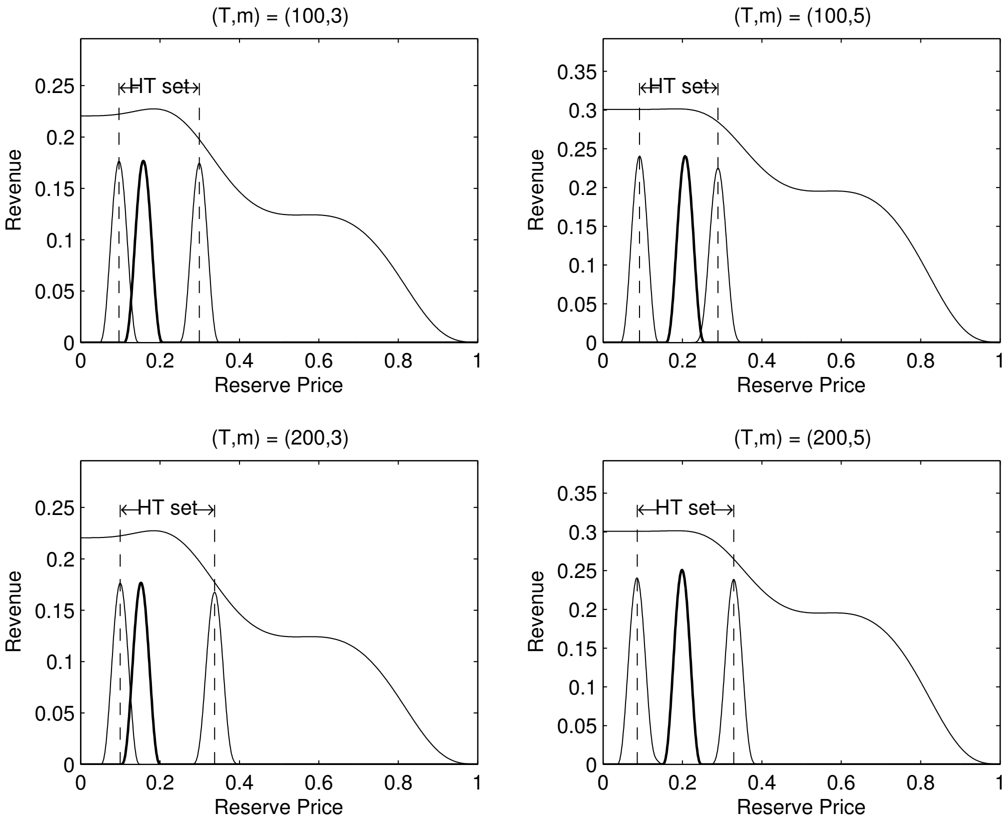

The revenue function is plotted along with the distributions of the lower and the upper bounds for the HT set (light lines) and the reserve prices chosen by the Γ-minmax (heavy).

FIGURE 5. Lognormal (3,1)

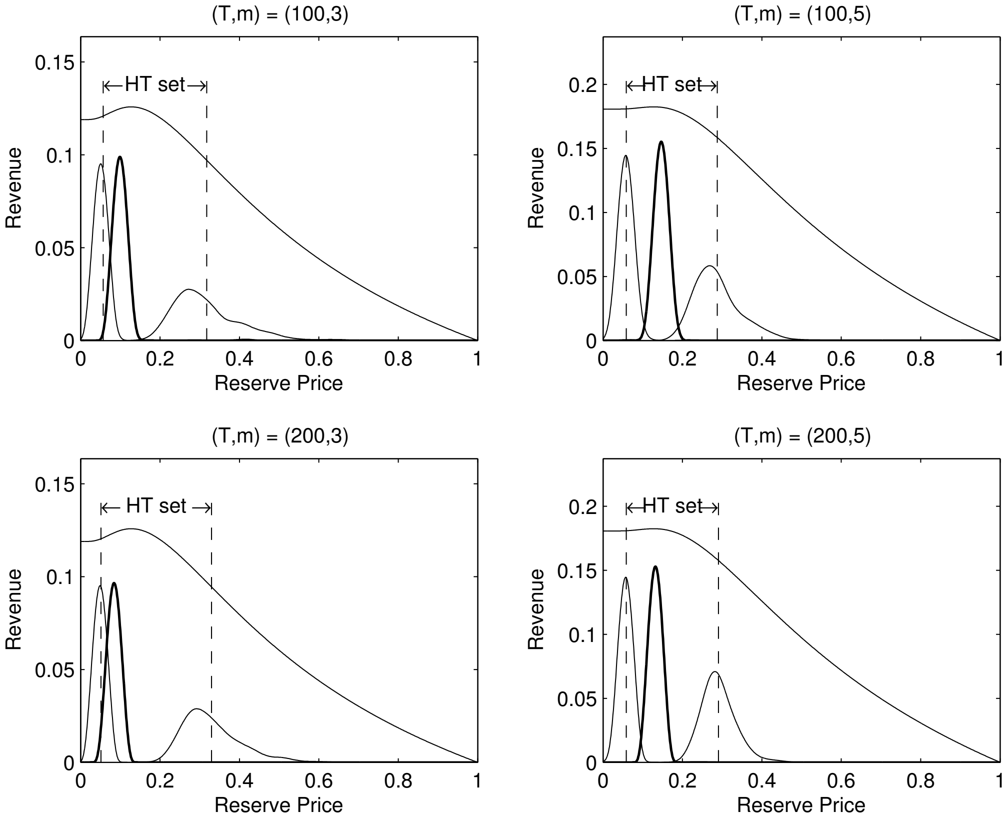

The revenue function is plotted along with the distributions of the lower and the upper bounds for the HT set (light lines) and the reserve prices chosen by the Γ-minmax (heavy).

FIGURE 6. Lognormal (4,1/2)

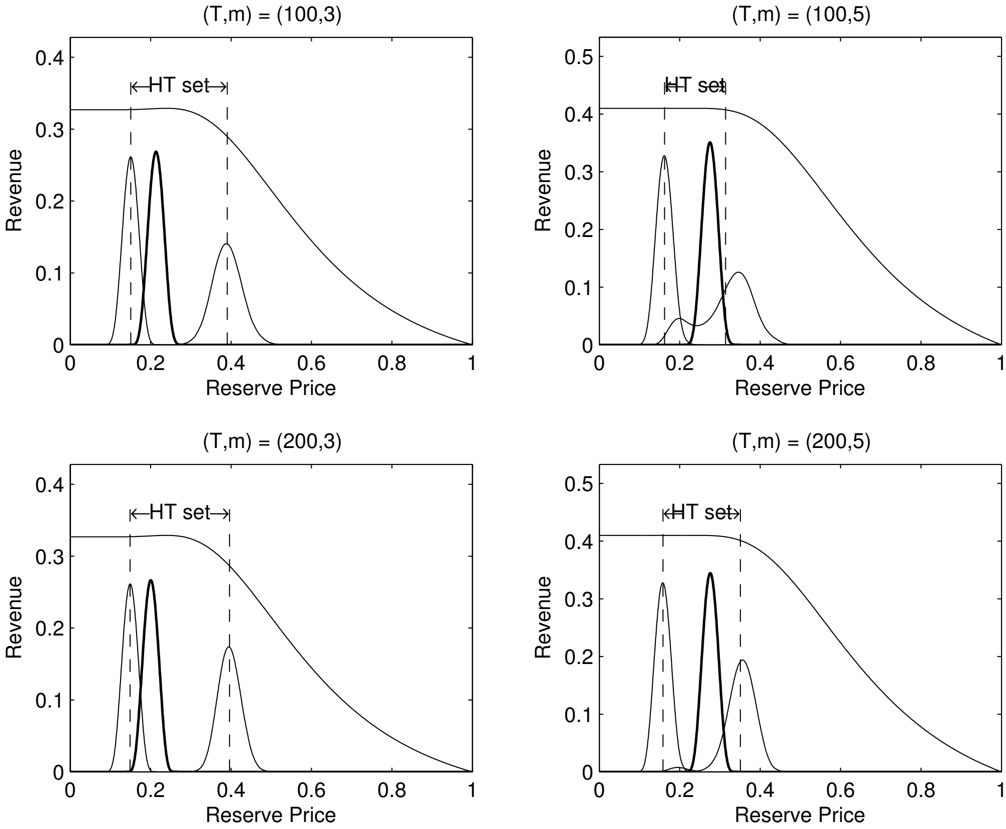

The revenue function is plotted along with the distributions of the lower and the upper bounds for the HT set (light lines) and the reserve prices chosen by the Γ-minmax (heavy).