International Journal of Industrial Organization 31 (2013) 587–602

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Optimal choice of a reserve price under uncertainty☆

Dong-Hyuk Kim

Vanderbilt University, United States

ARTICLE INFO

Article history:
Received 22 February 2012
Received in revised form 9 October 2013
Accepted 14 October 2013
Available online 22 October 2013

JEL classification:
C11
C44
D44

Keywords:
Auction design
Bayesian decision theory
Revenue maximizing reserve price

ABSTRACT

The revenue function for a standard auction is typically asymmetric around the revenue maximizing reserve price. Thus, choosing a reserve price that is smaller than the revenue maximizing reserve price can result in a substantially different loss than choosing one that is larger by the same amount. Therefore, when the revenue function is unknown, it is important to consider uncertainty around the revenue function and its asymmetric structure. For this purpose, I propose a Bayesian decision rule and illustrate its typical revenue gains. I then apply the rule to the bid data from the U.S. timber sales.

1. Introduction

Auctions are widely used to allocate important economic resources such as timber harvesting rights, radio frequency spectra, and offshore oil and gas rights. Since allocation outcomes in these markets may affect a number of economic sectors, policymakers must carefully devise the trading rules. In the auction literature, a particular pair of policy objective and parameter has been studied intensively: revenue maximization and choice of reserve price.[1] Even for this popular topic, however, optimal decision-making under uncertainty has not been formally discussed.

Myerson (1981) as well as Riley and Samuelson (1981) developed the economic theory on optimal auction design, assuming that the policymaker knows the density function of bidders' values (willingness to pay). When the density is unknown, Paarsch (1997) proposed to choose the revenue maximizing reserve price (RMRP) following Riley and Samuelson (1981), but he derived the maximum likelihood estimator (MLE) of the valuation distribution and used the point estimate in place of the true density ('plug-in'). The literature has unanimously used this approach (but not necessarily using the MLE). Examples include Li and Perrigne (2003), Li et al. (2003), Kim and Lee (2009), Krasnokutskaya (2011), Menzel and Morganti (2010), and Roberts (2009).[2]

When a policymaker follows the plug-in rule, however, he acts as if the valuation density equals its point estimate. That is, he ignores uncertainty concerning the valuation density. The concept of uncertainty here differs from what confidence intervals quantify. A confidence interval captures some variation of the plug-in rule (as an estimator of the RMRP) assuming that the valuation density equals its point estimate. It also ignores the uncertainty regarding the valuation density that matters in decision making.

In this paper, I argue that the policymaker can obtain higher revenues by formally considering uncertainty about the revenue and its structure. In particular, I consider a standard auction for a single good with no entry fee and a fixed number of risk neutral bidders whose private values are independently and identically distributed. For this type of auction, the RMRP does not depend on the number of bidders (see Riley and Samuelson, 1981) and, hence, neither does the plug-in rule. Yet, the revenue structure does.

The revenue function is typically a unimodal function of reserve price, which is strictly positive except at the upper boundary of the support of the valuation density. Whenever a bidder is added, the entire

☆ This work is based on chapter 1 of my doctoral dissertation at the University of Arizona. Many thanks to Keisuke Hirano, Gautam Gowrisankaran, Jonah Gelbach, John Wooders, and Mo Xiao. I appreciate useful comments from Gaurab Aryal, Yuanyuan Gu, Mohamad Khaled, and participants at the seminars at University of Arizona and University of Auckland. I also thank the co-Editor, Harry J. Paarsch, and two anonymous referees for useful and insightful suggestions, which substantially improved the quality and the scope of the paper.

E-mail address: dong-hyuk.kim@vanderbilt.edu.

[1] Following the convention in the auction literature, I use the term 'revenue' to refer to the policymaker's expected revenue where the expectation is taken with respect to the valuation distribution.

[2] Haile and Tamer (2003), Quint (2008), and Aradillas-Lopez et al. (2013) have considered incomplete auction models where the valuation distribution is partially identified. They have provided interval estimates of the RMRP using the set estimator of the valuation distribution. In Section 7, I discuss some extensions including such partially identified models.

0167-7187/$ – see front matter
http://dx.doi.org/10.1016/j.ijindorg.2013.10.003

revenue curve except at the boundary shifts up, but the revenue difference between the RMRP and any reserve price that is lower than the RMRP becomes smaller, while it is always zero at the upper boundary. Thus, for a large number of bidders, the revenue function becomes fairly constant below the RMRP, but drops above it. Such an asymmetric revenue structure suggests that a reserve price larger than the RMRP causes a greater revenue loss than does a smaller one. Thus, when a policymaker is uncertain about the revenue function and, therefore, cannot correctly choose the RMRP, he should prefer small reserve prices to large ones. The question is, how small? The answer depends on the amount of uncertainty and the structure of the revenue function. Therefore, by formally considering such decision relevant elements, the policymaker can increase revenues.

To solve the policymaker's problem, I employ a Bayesian decision framework. The method begins with the prior, a probability distribution that reflects the policymaker's beliefs concerning the structural parameter of the valuation density.[3] Employing the Bayes theorem, I update the prior using the bid data. The updated prior, now referred to as the posterior, represents all the perceived uncertainty concerning the parameter after formally considering the initial beliefs and the sample information.

Since the revenue function depends on the valuation density, it is also indexed by the structural parameter. Thus, the posterior quantifies the uncertainty around the revenue function, and can be used for computing the expectation of the revenue at every reserve price. The decision procedure, finally, maximizes this posterior mean of revenue. Savage (1954) as well as Anscombe and Aumann (1963) have argued that this is a coherent behavior under a set of behavioral axioms.

This procedure is shown to be optimal under the average risk principle, which is a widely used frequentist criterion. More generally, Berger (1985) argued that any rational decision method must correspond to some type of Bayesian method. Any debate regarding Bayesians versus frequentists is, therefore, irrelevant to what I claim here: use an optimal decision rule to solve a decision problem. From a frequentist perspective, using repeated sampling, I demonstrate the revenue gains under the Bayes rule relative to the *Bayesian* plug-in rule for a range of data generating processes (DGPs). Since two approaches become similar as the sample size grows, the Bayes rule is more valuable for small sample problems.

The Bayes rule exploits some prior beliefs about the valuation density. Since it selects a reserve price without completely determining the valuation density, the method shares the spirit of the Wilson doctrine, which argues that a mechanism should not depend on implausibly detailed information. For example, the valuation density at an auction; see Wilson (1987).

In the next section, I describe the auction model under consideration and motivate a formal decision framework. In Section 3, I propose the Bayesian decision rule, while in Section 4, I discuss the concepts of uncertainty and optimality. I provide evidence from a Monte Carlo study and analyze a sample from the U.S. timber sales in Sections 5 and 6, respectively. I conclude in Section 7, and collect computational details in an appendix.

2. Policymaker's problem: revenue maximization

An indivisible object is to be allocated to one of N risk-neutral, expected-utility maximizing bidders. Bidders' values, $x_1,\ldots, x_N$, are independently drawn from an absolutely continuous distribution characterized by a probability density function (pdf) f having bounded support $[0,\overline{x}]$. The values are private information.

An auction is said to be *standard* if its rule allocates the object to the bidder with the highest bid, provided his bid is not lower than reserve price ρ. I consider a class of standard auctions at which a bidder with zero value expects to pay zero. For every auction in this class, Myerson, 1981 as well as Riley and Samuelson, 1981 have argued that a Bayes–Nash equilibrium with a symmetric strictly increasing bidding strategy yields the same revenue (revenue equivalence principle), and if the policymaker's value is zero, the revenue at ρ is given as

$$u(f,\rho;N) := N\rho[1-F(\rho)]F(\rho)^{N-1} + N(N-1)\int_{\rho}^{\overline{v}} y[1-F(y)]F(y)^{N-2}f(y)dy \tag{1}$$

where $F(x) := \int_0^x f(t)dt$ is the cumulative distribution function (cdf) of the valuation distribution.[4]

Assume that every bidder follows the equilibrium bidding strategy associated with Eq. (1). For example, it is optimal for a bidder with $x \geq \rho$ to follow

$$\beta_I(x|N,f) := x-\int_{\rho}^{\overline{v}}\left[\frac{F(y)}{F(x)}\right]^{N-1}dy \tag{2}$$

at a first-price, sealed-bid auction, and it is

$$\beta_{II}(x|N,f) := x \tag{3}$$

at a second-price, sealed-bid auction.[5] In light of the revenue equivalence principle, I focus on the problem of choosing ρ to maximize the revenue, abstracting from the details of a specific auction rule.

If f were known, then the policymaker would maximize Eq. (1). Let $\lambda_f(x):=f(x)\ /\ [1-F(x)]$ denote the hazard rate function. When the seller's valuation is zero, the necessary condition for maximizing expression (1) is then

$$\rho\lambda_f(\rho) = 1. \tag{4}$$

If $\lambda_f(\cdot)$ is increasing, Eq. (4) is also sufficient; see Krishna (2002). Note that the RMRP only depends on f, but not on N. Thus, let

$$\rho_R(f) := \arg\max_{\rho} u(f,\rho;N) \tag{5}$$

denote the RMRP under f. In this paper, I consider a policymaker who does not know f, but has a sample of bids that is informative about f.

For choosing a reserve price, Paarsch (1997) developed the plug-in rule. He estimated the valuation density using the MLE and proposed to choose $\rho_R\left(\hat{f}\right)$ where $\hat{f}$ denotes the point estimate of f. While the article did not spell out, presumably, the argument is that the MLE of a function of the parameters is just the function of the MLE of the parameters.

Since Paarsch (1997), the empirical auction literature has employed the plug-in rule only, but often using estimates derived by procedures other than the MLE. Notice that such an approach does not fully exploit the shape of the function in Eq. (1), which is useful when the policymaker cannot correctly choose $\rho_R(f)$. In Figs. 1 and 2, I depict a set of valuation densities (panels (a) and (b)) and their revenue functions for the number of bidders $N = 3$, 4, and 5 (panels (c) and (d)). In Fig. 1, panels (a) and (c) are associated with a density that is similar to an exponential distribution, and panels (b) and (d) with a long-tailed density.[6] In Fig. 2, I depict densities that are similar to a lognormal with alternative parameters.[7]

[3] The prior beliefs may come from similar sales, the policymaker's experience, and/or a widely held common sense. For example, the policymaker for timber sales in Montana may form a prior using some results from timber auctions in California.

[4] The revenue function of the policymaker with value $x_0 > 0$ is $u(f,\rho;N) + x_0F(\rho)^N$.

[5] It is optimal to bid any price less than ρ, if $x < \rho$.

[6] These densities have the form of Eq. (13) with $k = 15$. For the exponential-like density, the parameter values are $\boldsymbol{\theta}:$ =(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, $\boldsymbol{\theta}:$ =(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).

[7] The lognormal distributions with $(\mu,\sigma) = (3,1)$ and $(4,1/2)$ are truncated at the 99th percentile and rescaled, so their supports are the unit interval.