Secret reserve prices in first–price auctions

(This version: May 27, 2014)

Frank Rosar

Department of Economics, University of Bonn, Lennéstr. 37, 53113 Bonn, Germany. Tel.: + 49 228 73 6192. Fax: + 49 228 73 7940. E–mail: email@frankrosar.de.

Abstract

This article offers a theoretical explanation for the use of secret reserve prices in auctions. I study first–price auctions with and without secret reserve price in an independent private values environment with risk–neutral buyers and a seller who cares at least minimally about risk. The seller can fix the auction rules either before or after she learns her reservation value. Fixing the rules early and keeping the right to set a secret reserve price can be strictly optimal. Moreover, I describe the relation of using a secret reserve price to phantom bidding and non–commitment to sell.

JEL classification: D44, D47, D82

Keywords: first–price auction, secret reserve price, phantom bidding, non–connected bid space, risk–averse seller

---

1. Introduction

While reserve prices are often kept secret in practice (e.g., Elyakime et al. (1994), Ashenfelter (1989)), non–standard assumptions are needed to justify their use on a theoretical basis. In the symmetric independent private values auction environment with a regular distribution, risk–neutral buyers and a risk–neutral seller, the optimal mechanism is implemented by any standard auction with an optimally chosen announced reserve price (Myerson (1981)). Secret reserve prices may be used to increase participation in second–price auctions with common values (Vincent (1995)), to credibly signal information in repeated second–price auctions (Horstmann and LaCasse (1997)), to induce more aggressive bidding in first–price auctions with risk–averse bidders (Li and Tan (2000)), and in first–price and second–price auctions with reference–based utility (Rosenkranz and Schmitz (2007)). I offer a further theoretical explanation for why secret reserve prices might be used in first–price auctions. My explanation is based on seller information that improves over time and risk–aversion on the seller's side.

In practice, a seller often fixes and announces the rules of an auction some time before the auction does actually take place. While this is sometimes necessary for exogenous reasons (e.g., because potential buyers need to prepare bids), the seller has normally at least the possibility to announce the rules of the auction some time in advance. During such a time, the seller's information might improve. For example, she might get better informed about her own use value or a new outside option might arise. I explain in this article why there can be a role for using a secret reserve price in a first–price auction when either (1) the seller's information improves for exogenous reasons after she fixes the rules of the auction and before the auction is conducted or (2) the seller is risk–averse and she can endogenously induce a situation in which this happens.

I proceed in two steps. I first analyze in Section 2 the case in which the seller's information improves for exogenous reasons. While I stick to the independent private values model with risk–neutral buyers and a risk–neutral seller, I consider a timing in which the seller has to commit to the rules of a first–price auction before she learns her value. The seller chooses the bid space and whether she keeps the right to set a secret reserve price later on when she is informed. A secret reserve price might be part of the optimal auction rules. The result arises quite naturally in this setting at the cost that it relies on a timing which might seem artificial. Then I show in Section 3 that the (artificial) timing of Section 2 can arise endogenously. If the seller cares about risk, she does under certain conditions prefer to commit to the auction rules early before she is informed to waiting until she is informed and fixing the rules then. By committing to the auction rules early, the seller induces a bidding behavior which does not vary in her own value. In conjunction with a secret reserve price, it can be possible to use this as an instrument to make the induced profit distribution less risky without sacrificing (too much) expected profit.

In the main parts of this article, I employ the simplest model that allows me to demonstrate the effects which drive my results. It relies on a binary seller value and—in the second part of the article—on seller preferences which are lexicographic in expected profit and variance of profit. The assumption of a binary seller value is mainly for technical convenience. The analysis becomes more complicated and the results become less clean when the seller's value is continuously distributed, but the crucial effects extend also to continuous distributions (see Subsection 2.4). The assumption of lexicographic preferences works against my effects. It simplifies however the analysis and, more importantly, it makes analysis and results better comparable to standard auction theory (see Subsection 3.3).

2. Exogenous seller information

2.1. The model

There is a seller of an indivisible object and two potential buyers, buyer 1 and buyer 2. I denote a generic buyer by i and the other buyer by $-i$. The values that the seller and the buyers attribute to the object are realizations of the independently distributed random variables X_0, X_1 and X_2. Let $X := (X_0, X_1, X_2)$. I use lower case letters to denote realizations of these random variables. X_i is distributed according to a cumulative distribution function F with support $[0, 1]$, a continuous and strictly positive density function f and a strictly increasing function $J(x_i) := x_i - (1 - F(x_i))/f(x_i)$. $J(x_i)$ describes the virtual valuation function introduced by Myerson (1981) which is important for many auction theory problems. X_0 is 0 with probability $p \in (0, 1)$ and $z \in (0, 1)$ otherwise. When I denote the indicator variable which describes whether buyer i obtains the object by $q_i \in \{0, 1\}$ and buyer i's payment to the seller by t_i, then buyer i's profit is given by $q_i x_i - t_i$ and the seller's profit is given by $\pi = (1 - q_1 - q_2)x_0 + t_1 + t_2$.

I am interested in first–price sealed–bid auctions with and without a secret reserve price.[1] The timing is as follows: First, the seller chooses a closed bid space $B \subset \mathbb{R}_+$ and whether she will set a secret reserve price in stage 3 ($S = y$) or not ($S = n$). I will denote the lowest admissible bid by r_a and refer to it as announced or open reserve price. The auction rules B and S are observable to the buyers. Second, each player privately learns his value. Third, if $S = y$, the seller chooses a secret reserve price $r_s \in \mathbb{R}_+$. Moreover, each buyer i either submits a sealed bid $b_i \in B$ or does not participate in the auction. I denote non–participation with $b_i = \emptyset$ such that a strategy of a buyer is described by a function $b : [0, 1] \to B \cup \{\emptyset\}$. If $S = n$ (resp. $S = y$),

[1] Why the seller uses a first–price payment rule lies outside of my model. One reason might be that first–price auctions perform well when the seller is risk–averse which is the case in which I will be finally interested in (Waehrer et al. (1998)).

the seller keeps the object when no bid (resp. no bid $b_i \geq r_s$) is submitted. Otherwise, the buyer with the highest bid obtains the object and pays his bid to the seller. Ties between the buyers are broken according to a fair lottery.

Each buyer is risk–neutral such that he strives for maximizing his expected profit. I consider at first the case in which the seller is risk–neutral as well. Later I will discuss the case in which she has lexicographic preferences in expected profit and variance of profit and the case in which she is risk–averse with non–lexicographic preferences. I am interested in undominated Perfect Bayesian Equilibria where participation and bidding behavior is symmetric across buyers (usPBE).[2]

To simplify the exposition of my results, I assume that a buyer participates whenever he is indifferent between participation and non–participation and that he chooses the higher bid whenever he is indifferent between a higher and a lower bid. Moreover, for any $b_i \in \mathbb{R}_+$ I write $b_i > b_{-i}$ to describe the case in which either $b_{-i} = \emptyset$ or $b_{-i} \in [0, b_i)$.

2.2. Strategic bidding behavior and the effect of holes in the bid space

My model generalizes a standard independent private values first–price auction model in three respects: First, the seller's information improves over time. She is better informed at the time the auction is conducted than at the time she designs and announces the auction rules. Second, I explicitly allow the seller to restrict the set of admissible bids further than by setting only an open reserve price. Third, I allow for the possibility that the seller sets a secret reserve price before the auction starts. I explain in this subsection the equilibrium behavior in the subgame which is played after the auction rules S and B are fixed and I describe how the three generalizations affect the analysis.

Consider first how the seller is affected by the secret reserve price decision S for a given behavior of the buyers. If $S = n$, it does not depend on the realization of the seller's value x_0 whether the object is sold or not. The object is sold at the highest bid whenever at least one bid is submitted. If the object is not sold, the seller keeps the object and realizes an expected profit of $\mathbf{E}_X[X_0] =: \overline{x}_0$. By contrast, if $S = y$, the selling decision can depend through the secret reserve price on the seller's private information x_0. She sells the object if the highest bid exceeds the secret reserve–price and keeps it otherwise. Intuitively, the highest bid can be interpreted as a take–it–or–leave–it offer to buy the object and the secret reserve price describes the threshold above which these offers are accepted by the seller. As the value of the secret–reserve price does not feed back on the buyers' bidding behavior, setting $r_s = x_0$ is clearly optimal for the seller. The seller's profit is thus the maximum of the highest bid and her value x_0 if at least one bid is submitted and she realizes an expected profit of $\overline{x}_0$ otherwise. Hence, $S = y$ implies that the seller faces a commitment problem regarding under which conditions she will sell the object, whereas $S = n$ implies that the selling decision is not affected by the seller's value x_0.[3]

The difference in the selling behavior for $S = n$ and $S = y$ may induce for the same bid space B a different behavior by the buyers. Consider thus the problem a buyer i faces when he has value x_i and believes that the other buyer behaves according to a strategy $b : [0, 1] \to B \cup \{\emptyset\}$. If $S = n$, buyer i faces

[2] If the seller sets a secret reserve price, it has a second–price character for her. To exclude equilibria in which the seller sets a secret reserve price which is prohibitively high and in which the buyers submit no bids, I restrict attention to PBE which do not rely on weakly dominated strategies.

[3] Setting a secret reserve price is closely related to placing a phantom bid and to not committing to sell after observing the bids. See Section 4 for a discussion.

the problem either not to participate or to choose a bid $b_i \in B$ to maximize

$$\left[\text{Prob}_X\{b_i > b(X_{-i})\} + \frac{1}{2}\text{Prob}_X\{b_i = b(X_{-i})\}\right] \times (x_i - b_i). \tag{1}$$

He faces a trade–off between a higher probability of winning against the other buyer and a higher rent conditional on winning. If $S = y$, there is an additional effect. A higher bid may then also induce a higher probability with which the seller does actually sell the object besides increasing the probability of having the highest bid. Buyer i's expected profit from submitting a bid $b_i \in B$ is then

$$\left[\text{Prob}_X\{b_i > b(X_{-i})\} + \frac{1}{2}\text{Prob}_X\{b_i = b(X_{-i})\}\right] \times \text{Prob}_X\{b_i \geq X_0\} \times (x_i - b_i). \tag{2}$$

The following lemma summarizes equilibrium properties which follow from standard reasoning and which hold equally for $S = n$ and for $S = y$:

Lemma 1 *Fix any S and any B with $r_a < 1$. If $b : [0,1] \to B \cup \{\emptyset\}$ is part of a symmetric equilibrium of the game which is played after B and S are chosen, then the following is true:*

(a) There is threshold participation behavior and the participation threshold corresponds to the lowest admissible bid r_a.

(b) The buyers' bidding behavior $b(x_i)$ is weakly increasing on $[r_a, 1]$ with $b(r_a) = r_a$.

(c) If $B = [r_a, \infty)$, the buyers' bidding behavior $b(x_i)$ is strictly increasing on $[r_a, 1]$.

(d) If the buyers' bidding behavior $b(x_i)$ jumps at some value $\widehat{x}_i \in [r_a, 1]$ upwards from b_i' to b_i'', then the bids b_i' and b_i'' are associated with different selling probabilities or the equilibrium bidding behavior exhibits pooling on b_i' or on b_i''.

For a detailed proof, see the Appendix. The intuition behind the lemma is as follows: Properties (a) and (b) follow from standard reasoning. Property (c) holds because if $b(x_i)$ prescribed that a buyer chooses a specific bid with a positive probability, the other buyer would have an incentive to deviate from $b(x_i)$ by slightly overbidding this bid for some of his values. This would increase his probability of obtaining the object by a discrete amount while it would reduce his rent conditional on winning only slightly. (d) is the most non–standard property. It states that the bidding behavior can only exhibit a jump when it comes along with a discrete increase in the probability of obtaining the object. On the one hand, $b(x_i)$ may jump at some value $\widehat{x}_i$ from b_i' to $b_i'' > b_i'$ when the equilibrium bidding behavior exhibits pooling on the bid b_i' or on the bid b_i''. By Lemma 1 (c), pooling requires however $B \neq [r_a, \infty)$. That is, bidding must be restricted further than by an open reserve price r_a. On the other hand, $b(x_i)$ may exhibit a jump from b_i' to b_i'' at some value $\widehat{x}_i$ when $S = y$ and $\text{Prob}_X\{b_i'' \geq X_0\} > \text{Prob}_X\{b_i' \geq X_0\}$.[4] For the considered binary distribution of X_0, this means that the bidding behavior may jump once at some value $\widehat{x}_i$ from below z to above z and that it is continuous on the left and on the right of $\widehat{x}_i$. Moreover, if $B = [r_a, \infty)$, a jump can only occur from below z to exactly z. Otherwise, a buyer who bids just above the jump would have a strict incentive to reduce his bid to exactly z as such a bid reduction would by Lemma 1 (c) not reduce his probability of obtaining the object but increase his profit conditional on obtaining it.

When the seller does not set a secret reserve price (i.e., $S = n$) and bidding is only restricted by an open reserve price r_a (i.e., $B = [r_a, \infty)$), deriving the bidding behavior in the symmetric equilibrium

[4] Although it may seem at first glance that the occurrence of a jump for this reason is an artifact of the discrete distribution of X_0, this is not true. See Subsection 2.4 for a discussion.

is standard. Consider the bidding problem of a buyer i with value $x_i \in [r_a, 1]$ when buyer $-i$ behaves according to a strategy $b : [0, 1] \to B \cup \{\emptyset\}$ which is consistent with Lemma 1. Because buyer $-i$'s bidding behavior is continuous (by Lemma 1 (c) and (d)), $b(r_a) = r_a$ (by Lemma 1 (b)) and bidding above $b(1)$ is clearly dominated, choosing a bid $b_i \in B$ directly is equivalent to choosing a value $x'_i \in [r_a, 1]$ which determines the bid $b_i = b(x'_i)$ indirectly. By using this in (1), it follows that a necessary condition for b being part of a symmetric equilibrium is that $x_i \in \arg\max_{x'_i \in [r_a,1]} F(x'_i)(x_i - b(x'_i))$ for any $x_i \in [r_a, 1]$. Weak monotonicity of b (Lemma 1 (b)) implies that b is differentiable almost everywhere. The derivative b_{x_i} of b exists thus almost everywhere and the first–order condition $[f(x'_i)(x_i - b(x'_i)) - F(x'_i)b_{x_i}(x'_i)]_{x'_i = x_i} = 0$ must hold almost everywhere. By rearranging, I obtain $\mathrm{d}/\mathrm{d}x_i(b(x_i)F(x_i)) = x_i f(x_i)$. Continuity of b (which follows from Lemma 1 (d)) implies then that the equilibrium bidding behavior is completely determined by this differential equation and a boundary condition. By integrating both sides of the differential equation from x'_i to x_i, I obtain

$$b(x_i)F(x_i) = b(x'_i)F(x'_i) + \int_{x'_i}^{x_i} \tau f(\tau)\mathrm{d}\tau. \tag{3}$$

Finally, by using the boundary condition $b(r_a) = r_a$ (Lemma 1 (b)) and by applying partial integration to the integral, I obtain that a buyer with value $x_i \geq r_a$ submits the bid

$$\beta_{r_a}(x_i) = x_i - \int_{r_a}^{x_i} F(\tau)/F(x_i)\mathrm{d}\tau. \tag{4}$$

When the seller sets a secret reserve price (i.e., $S = y$), there may be two differences relative to the case in which she does not set a secret reserve price. First, there may exist a value $\widehat{x}_i$ at which the bidding behavior jumps from below z to z. Second, the selling probability is p instead of 1 when a buyer wins with a bid $b_i < z$. Neither of the two differences affects the differential equation which must hold almost everywhere in any symmetric equilibrium, only the relevant boundary condition may be affected. The boundary condition which is relevant for the bidding behavior before a jump occurs is $b(r_a) = r_a$, whereas it is $b(\widehat{x}_i) = z$ for the bidding behavior after a jump. The buyers' bidding behavior follows in both cases from (3) with the respective boundary conditions.

Under which conditions does the bidding behavior exhibit a jump? Before a (potential) jump occurs, the bidding behavior is described by $\beta_{r_a}(x_i)$ like in the case with $S = n$. If the equilibrium bidding behavior exhibits a jump at $\widehat{x}_i$, a buyer with value $\widehat{x}_i$ must be indifferent between bidding $\beta_{r_a}(\widehat{x}_i)$ and bidding z. Because the equilibrium bidding behavior is strictly increasing by Lemma 1 (c), both bids are associated with the same probability of submitting the highest bid. Hence, by bidding z instead of $\beta_{r_a}(\widehat{x}_i)$, a buyer with value $\widehat{x}_i$ increases the probability of selling from p to 1 without affecting the probability of submitting the highest bid. He is indifferent if $p(\widehat{x}_i - \beta_{r_a}(\widehat{x}_i)) = \widehat{x}_i - z$. When I define

$$\sigma_{r_a}(x_i) := p\beta_{r_a}(x_i) + (1 - p)x_i, \tag{5}$$

I can write the indifference condition as $\sigma_{r_a}(\widehat{x}_i) = z$. $\sigma_{r_a}(x_i)$ can be interpreted as the bid which a buyer with value x_i is willing to pay to obtain the object for sure instead of paying the bid $\beta_{r_a}(x_i)$ and getting the object only with probability p. Because $\sigma_{r_a}(x_i)$ is continuous and strictly increasing on $[r_a, 1]$, the indifference condition is satisfied for some value $\widehat{x}_i \in (r_a, 1]$ if and only if $\sigma_{r_a}(r_a) = r_a < z$ and $\sigma_{r_a}(1) \geq z$. The induced bidding behavior is then given by

$$\beta_{r_a,\widehat{x}_i}(x_i) := \begin{cases} \beta_{r_a}(x_i) & \text{if } x_i \in [r_a, \widehat{x}_i) \\ \beta_{r_a}(x_i) + (1 - p)\int_{r_a}^{\widehat{x}_i} F(\tau)/F(x_i)\mathrm{d}\tau & \text{if } x_i \in [\widehat{x}_i, 1] \end{cases} \tag{6}$$

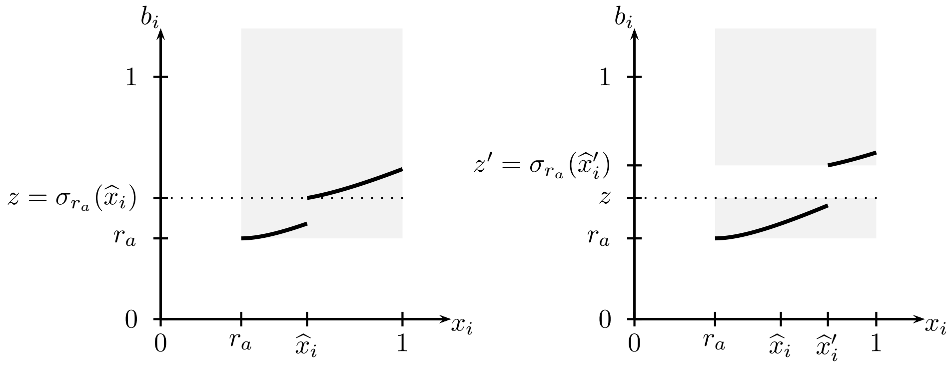

(a) No hole: $B = [r_a, \infty)$ (b) Hole: $B = [r_a, z - \epsilon] \cup [z', \infty)$

Figure 1: Bidding behavior for $S = y$ [$X_i \sim U[0,1]$, $p = 1/2$, $z = 1/2$, $r_a = 1/3$, $z' = \sigma_{r_a}(4/5)$]

with $\widehat{x}_i = \sigma_{r_a}^{-1}(z)$.[5] Because the values r_a and $\widehat{x}_i$ at which the buyers start bidding and at which the bidding behavior jumps from below to above z, respectively, completely determine the equilibrium bidding behavior, I use them as subscripts of the function which describes the bidding behavior. The discussion in the text implies the following result:[6]

Lemma 2 *Suppose $r_a < z$ and $\sigma_{r_a}(1) \geq z$ and consider the subgame that is played after S and $B = [r_a, \infty)$ are chosen. Then a buyer with value $x_i \in [r_a, 1]$ submits in any symmetric equilibrium the bid $\beta_{r_a}(x_i)$ if $S = n$ and the bid $\beta_{r_a,\widehat{x}_i}(x_i)$ with $\widehat{x}_i = \sigma_{r_a}^{-1}(z)$ if $S = y$.*

Figure 1(a) illustrates the bidding behavior that arises when the supposition of the lemma holds. The set of admissible bids is indicated by the grey area. There are two ways in which the supposition of the lemma can be violated. If $r_a \geq z$, any admissible bid always exceeds the seller's value implying that the buyers' problem coincides for $S = n$ and $S = y$. If $\sigma_{r_a}(1) < z$, z is so high that it is not attractive to any buyer type to bid above z in order to increase the selling probability. In both cases, the same bidding behavior is induced for $S = n$ and for $S = y$.

An optimal first–price auction is often interpreted as a first–price auction with an optimal open reserve price. However, in practice, there is normally no good reason why a seller should not be able to restrict bidding further than by setting only an open reserve price. For example, she might choose $B = \{0.3, 0.35, 0.4, 0.5, 0.6, 0.8, 1\}$ instead of $B = [0.3, \infty)$.[7] While the restriction to bid spaces of the form $B = [r_a, \infty)$ is known to be without loss of generality for certain important environments like that considered in Myerson (1981), it is a priori not clear whether this is also the case for my environment. It is therefore important for the analysis in this article to understand how the bidding behavior is affected by holes in the bid space.

Consider first $S = n$. Any hole in the bid space which affects the bidding behavior must in this case introduce some kind of pooling. There are only two possibilities: First, some buyer types choose bids above

[5]See the Appendix for a derivation of (6) from the differential equation and the boundary conditions.

[6]I derive merely necessary conditions for the symmetric equilibrium behavior in the text. Proving sufficiency does not provide any additional insights and can follow along the lines of Proposition 2.2 in Krishna (2009).

[7]Such a restriction of the set of admissible bids resembles the restriction of bidding that arises in (English) open outcry auctions. The seller (or the auctioneer as her agent) announces in such auctions the possible bids successively. The possible bids differ by discrete increments.

the hole. This means that there is a jump in the bidding behavior which can by Lemma 1 (d) only come along with some kind of pooling. Second, all buyer types bid below the hole. Bidding is then effectively restricted by a binding bid cap which can also only be part of an equilibrium if there is pooling on the bid cap. Holes introduce thus noise in the allocation of the object without affecting when the object is sold.

Consider now $S = y$. Holes in the bid space can then also have a different effect. By forbidding intermediate bids, the seller can force some buyer types to bid lower and others to bid higher (relative to their behavior without such a restriction). As this can affect under which conditions the secret reserve price is binding, holes in the bid space can be used to mitigate the commitment problem which is associated with the secret reserve price. To see this more clearly, suppose the bidding behavior induced by $B = [r_a, \infty)$ exhibits a jump. By forbidding bids just above z, say by choosing $B = [r_a, z - \epsilon] \cup [z', \infty)$ with $z' > z$ and a small positive ϵ, the jump occurs at $\widehat{x}_i' = \sigma_{r_a}^{-1}(z') > \widehat{x}_i$. A buyer bids then more often below z, but if he bids above z, he bids higher. See Figure 1(b) for an illustration. That is, the seller can use a hole in the bid space as an instrument to affect where the jump occurs. The trade–off associated with the design of such a hole is comparable to the trade–off associated with an announced reserve price. Intuitively, such a hole in the bid space allows the implementation of multiple reserve prices in a single auction. r_a (resp. z') is the reserve price for a buyer who is satisfied with getting (resp. eager to get) the object with a low (resp. high) probability conditional on submitting the highest bid.

I know from the analysis so far how any strictly increasing bidding behavior which exhibits a jump at $\widehat{x}_i'$ must look like. It remains to investigate for which thresholds $\widehat{x}_i'$ it is possible to construct a bid space such that the induced bidding behavior is indeed strictly increasing and jumps at $\widehat{x}_i'$. There are two conditions: First, where the jump occurs can be affected by forbidding bids just above z. This allows it however only to induce jumps that occur later relative to the case where bidding is only restricted by the open reserve price r_a. This imposes a lower bound on the set of inducible jump points $\widehat{x}_i'$: $\widehat{x}_i' \geq \sigma_{r_a}^{-1}(z)$. Second, the bids $\beta_{r_a}(x_i)$ that buyer types on the left of $\widehat{x}_i'$ are willing to submit must be below z. Otherwise, the intended bidding behavior of these buyer types would be conflicting with the restriction of the bid space just above z. Even though it might be possible to design the bid space such that a jump occurs in this case at $\widehat{x}_i'$, the induced bidding behavior would exhibit some pooling on the left of $\widehat{x}_i'$ and would thus not be the strictly increasing bidding behavior $\beta_{r_a,\widehat{x}_i'}(\widehat{x}_i')$. I obtain thus also an upper bound on the set of jump points $\widehat{x}_i'$ that are inducible with a strictly increasing bidding behavior: $\widehat{x}_i' < \beta_{r_a}^{-1}(z)$.

The following lemma summarizes which strictly increasing bidding behavior can be induced for $S = y$ by some bid space B when the supposition of Lemma 2 holds.

Lemma 3 *Suppose $r_a < z$ and $\sigma_{r_a}(1) \geq z$. Any bidding behavior $\beta_{r_a,\widehat{x}_i}(x_i)$ with $\widehat{x}_i \in [\sigma_{r_a}^{-1}(z), \beta_{r_a}^{-1}(z))$ can be induced by choosing $S = y$ and by appropriately designing the bid space. If $\widehat{x}_i \in (\sigma_{r_a}^{-1}(z), \beta_{r_a}^{-1}(z))$, inducing $\beta_{r_a,\widehat{x}_i}(x_i)$ requires a non–connected set of admissible bids.*

2.3. Optimal mechanism design and first–price auctions

In this subsection, I explain why there exist parameters $(z, p) \in (0, 1)^2$ such that $S = y$ is strictly optimal for the seller. I proceed in two steps. First, I use standard mechanism design results to derive an upper bound $\mathcal{U}_b(z, p)$ on the expected profit that the seller can obtain from any first–price auction with and without a secret reserve price. Then I use the insights about the equilibrium bidding behavior in first–price auctions from the preceding subsection to explain why the upper bound is under certain conditions attainable by a first–price auction and why a secret reserve price (possibly in conjunction with a non–connected set of admissible bids) is necessary to attain it.

Suppose that the seller can choose general mechanisms to allocate the object and that she is not subject to a commitment problem regarding how the outcome of the mechanism depends on her own value x_0. By the revelation principle, it is then without loss of generality to restrict attention to direct mechanisms and to equilibria where each buyer reveals his value truthfully. A direct mechanism (q,t) consists of two components: an allocation rule $q : [0,1] \times [0,1] \times \{0,z\} \rightarrow \{(q_1,q_2) \in [0,1]^2 | q_1 + q_2 \leq 1\}$ which describes the probabilities with which each of the two buyers obtains the object; and a payment rule $t : [0,1] \times [0,1] \times \{0,z\} \rightarrow \mathbb{R}^2$ which describes the buyers' payments to the seller. The seller's profit is given by $(1 - q_1(x) - q_2(x))x_0 + t_1(x) + t_2(x)$ and buyer i's profit is given by $q_i(x)x_i - t_i(x)$. The derivation of the mechanism which maximizes the seller's expected profit is standard. I give in the subsequent paragraphs a brief intuition for how it is derived, a detailed derivation can for example be found in Chapter 5 of Krishna (2009).

The seller strives for maximizing her expected profit by choosing a direct mechanism (q,t) for which each buyer is willing to reveal his private value voluntarily. This means that each buyer must be willing to participate rather than to realize an outside option of zero (individual rationality) and that each buyer must prefer to reveal his true value x_i over pretending to have any other value x_i' (incentive compatibility). When I define $\overline{q}_i(x_i) := \mathbf{E}_X[q_i(X)|X_i = x_i]$ and $\overline{t}_i(x_i) := \mathbf{E}_X[t_i(X)|X_i = x_i]$, I can write buyer i's interim expected profit when she has value x_i and announces to have value x_i' as $\overline{q}_i(x_i')x_i - \overline{t}_i(x_i')$. Incentive compatibility is satisfied with a binding individual rationality constraint if and only if $\overline{q}_i(x_i)$ is weakly increasing and $\overline{t}_i(x_i) = \overline{q}_i(x_i)x_i - \int_0^{x_i} \overline{q}_i(\tau)\mathrm{d}\tau$. Intuitively, the seller can extract the entire expected value that the buyer obtains from the object, $\overline{q}_i(x_i)x_i$, but she has to leave him an information rent $\int_0^{x_i} \overline{q}_i(\tau)\mathrm{d}\tau$ in order to incentivize him to voluntarily reveal his private value x_i.

By using the structure of the buyers' interim expected payments to rewrite the seller's expected profit, I obtain

$$\mathbf{E}_X[X_0 + \sum_{i=1}^{2} q_i(X)(J(X_i) - X_0)]. \tag{7}$$

The derivation of this representation of the seller's expected profit is standard (see for example Section 5.2.1 in Krishna (2009)). The optimal auction literature refers to $J(x_i)$ commonly as virtual valuation. $J(x_i)$ is strictly increasing for any regular distribution and can be interpreted as the value that is extractable from a buyer with real value x_i. Intuitively, by selling to a buyer with value x_i with a higher probability, it becomes more attractive for this buyer to pretend to have the value x_i when his real value is actually higher. In order to incentivize nevertheless truth–telling by higher buyer types, the seller must leave those types a higher informational rent. The virtual value corresponds to buyer i's actual value net of the effect that selling to buyer i at value x_i has on this buyer's ex ante expected information rent.

The most notable property of the representation of the seller's expected profit that I have given in (7) is that it depends only through the allocation rule q on the mechanism. By ignoring the monotonicity constraint (which is necessary for incentive compatibility) and by maximizing (7) through the choice of an allocation rule, I obtain an upper bound $\mathcal{U}_b(z,p)$ on the expected profit that the seller can obtain from any mechanism. The allocation rule which attains the upper bound follows from comparing the seller's actual value with the buyers' virtual values. More specifically, it is attained if the object is allocated to the buyer with the highest value if $\max\{x_1, x_2\}$ exceeds $J^{-1}(z)$ and if it is kept by the seller otherwise. It follows that the upper bound is given by $\mathcal{U}_b(z,p) = \mathbf{E}_X[\max\{X_0, J(X_1), J(X_2)\}]$ and that any first–price auction which induces this allocation (and for which individual rationality constraints are binding) attains this upper bound. As individual rationality constraints are in my setting binding for any first–price auction, this leaves me with the question whether some first–price auction induces the desired allocation.

For the considered setting in which the seller's information improves after she designs the bid space, the allocation induced by a first–price auction without a secret reserve price cannot depend on the realization of the seller's value. It is thus not possible to attain the upper bound on the seller's expected profit with $S = n$. By contrast, the allocation induced by a first–price auction with a secret reserve price depends on the realization of the seller's value. Necessary for implementing the desired allocation is that the induced bidding behavior is strictly increasing. Otherwise, the object would not always be sold to the buyer with the highest value when it is sold. By introducing a hole in the bid space, the seller can affect which strictly increasing bidding behavior is induced. Lemma 3 characterizes which strictly increasing bidding behavior can be induced. The desired allocation is implemented under the following two conditions: First, the participation threshold must be at $J^{-1}(0)$. This can always be achieved by choosing an open reserve price $r_a = J^{-1}(0)$. Second, the bidding behavior must jump at $\widehat{x}_i = J^{-1}(z)$ from below z to above z. By Lemma 3, this can be achieved if $J^{-1}(z) \in [\sigma_{r_a}^{-1}(z), \beta_{r_a}^{-1}(z))$. If $J^{-1}(z) = \sigma_{r_a}^{-1}(z)$, a bidding behavior with the desired properties is induced by the connected bid space $B = [J^{-1}(0), \infty)$. If $J^{-1}(z) \in (\sigma_{J^{-1}(0)}^{-1}(z), \beta_{J^{-1}(0)}^{-1}(z))$, a hole in the bid space is necessary for inducing a bidding behavior with the desired properties. A specific optimal bid space is then given by $B = [J^{-1}(0), \beta_{J^{-1}(0)}(J^{-1}(z))] \cup [\sigma_{J^{-1}(0)}(J^{-1}(z)), \infty)$. The discussion in the text implies the following result:

Proposition 1 *If $\beta_{J^{-1}(0)}(J^{-1}(z)) < z$ and $\sigma_{J^{-1}(0)}(J^{-1}(z)) \geq z$, the seller strictly prefers the optimal first–price auction with a secret reserve price ($S = y$) over the optimal first–price auction without a secret reserve price ($S = n$). Moreover, the optimal first–price auction with a secret reserve price implements under this supposition the generally optimal mechanism.*

An intuition for the result is the following: From an ex ante perspective, the seller wants that the object is less often sold when her value turns out to be low than when it turns out to be high. While such a dependence is not achievable if the seller uses only an open reserve price to which she commits herself in advance before she is informed, it could, in principle, be achievable if she keeps the right to set a secret reserve price after she is informed. The problem is that she is subject to a commitment problem concerning which secret reserve price she will set after she is informed. If the supposition of Proposition 1 holds, it is however possible for her to manipulate the bid space (and therewith the induced bidding behavior) in such a way that her commitment problem does not become binding. In that respect, the restriction of the set of admissible bids serves in my framework as a commitment device for the seller.

Four remarks are in place.

Remark 1: The supposition of Proposition 1 is satisfied for a non–empty set of parameters. For any F and any p there exist z–values for which the supposition of the proposition holds: I have $\sigma_{J^{-1}(0)}(J^{-1}(0)) = J^{-1}(0) > 0$ and $\sigma_{J^{-1}(0)}(J^{-1}(1)) = p\beta_{J^{-1}(0)}(1) + (1-p)1 < 1$. Continuity of $\sigma_{J^{-1}(0)}$ implies that there exists some $z \in (0,1)$ such that $\sigma_{J^{-1}(0)}(J^{-1}(z)) = z$. As $\beta_{J^{-1}(0)}(J^{-1}(z)) < \sigma_{J^{-1}(0)}(J^{-1}(z))$ for any $z \in (0,1)$, both conditions hold simultaneously when $\sigma_{J^{-1}(0)}(J^{-1}(z)) = z$. For uniformly distributed buyer values the two conditions become $z > (\sqrt{7} - 1)/3$ and $z \leq (\sqrt{p^2 + 2p + 4} - 4)/(p + 2)$. The set of all parameters (z, p) for which the supposition of Proposition 1 holds is illustrated by the grey area in Figure 2.[8]

[8]For any $p \in (0,1)$ there exists an interval of z–values for which the supposition holds. Further, there exists $z' \in (0,1)$ such that for any $z \in (z', 1)$ there exists an interval of p–values for which the supposition holds. The strategy of proof in Proposition 1 relies on the possibility to induce a strictly increasing bidding behavior which exhibits a jump. As it is by Lemma 1 (d) necessary for this that z is strictly larger than the open reserve price $r_a = 1/2$, the strategy of proof in Proposition 1 can only work for $z > 1/2$.

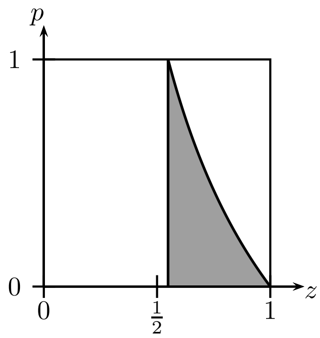

Figure 2: Parameters (z, p) for which the supposition of Proposition 1 holds when $X_i \sim U[0,1]$

Remark 2: Why is there no role for holes in the bid space in first–price auctions without a secret reserve price? The allocation induced by first–price auctions without a secret reserve price cannot condition on the realization of the seller's value. An upper bound on the seller's expected profit from any first–price auction without a secret reserve price can thus be obtained by maximizing (7) through the choice of an allocation rule $q(X)$ which does not condition on X_0. As this implies $\mathbf{E}_X[q(X)|X_1, X_2] = q(X)$, (7) becomes $\mathbf{E}_X[\overline{x}_0 + \sum_{i=1}^{2} q_i(X)(J(X_i) - \overline{x}_0)]$ by applying the Law of Iterated Expectations. This expression is maximized when the object is allocated to the buyer with the highest value if $\max\{x_1, x_2\}$ exceeds $J^{-1}(\overline{x}_0)$ and if the object is kept by the seller otherwise. The upper bound is thus given by $\mathcal{U}_b(\overline{x}_0) := \mathbf{E}_X[\max\{\overline{x}_0, J(X_1), J(X_2)\}]$ and it is attained by a first–price auction without a secret reserve price and with a bid space $B = [J^{-1}(\overline{x}_0), \infty)$. Necessary for attaining the upper bound is that the induced bidding behavior is strictly increasing. Because holes in the bid space would by the reasoning in the preceding subsection either induce a bidding behavior which exhibits pooling or they would have no effect on the induced bidding behavior (see Lemma 1 (d)), there is no role for holes in the bid space when there is no secret reserve price.

Remark 3: Comparison of $S = n$ with $S = y$ when there is only an open reserve price. The optimal first–price auction with a secret reserve price often exhibits a non–connected set of admissible bids. However, even if the seller does not want to restrict bidding further than by setting an open reserve price, $S = y$ is at least weakly optimal for any F, any p and any z: If $r_a < z$, this is implied by $\beta_{r_a}(x_i) \leq \beta_{r_a, \widehat{x}_i}(x_i)$ for any $x_i \geq r_a$ (see (4) and (6)). If $r_a \geq z$, the optimal secret reserve price has no bite. The bidding behavior coincides then in the cases with and without a secret reserve price.

Remark 4: Lexicographic risk–aversion on the seller's side. If the seller had preferences which are lexicographic in expected profit and the variance of profit, Proposition 1 would still hold. This is because the seller's profit has the same variance for any combination of S and B that attains the upper bound on the seller's expected profit $\mathcal{U}_b(z, p)$. The same auction rules would be optimal and the same bidding behavior would be induced.

2.4. Continuous seller valuations

Two effects were crucial in the analysis of my model with a binary distribution of the seller's value: First, for a given bid space B, $S = y$ may induce a more aggressive bidding behavior than $S = n$. Second, the bidding behavior induced by $S = y$ and $B = [r_a, \infty)$ may exhibit a jump. In this subsection, I explain why both effects are not specific to the case in which the seller's value is discretely distributed.

First effect. When the seller's value is distributed according to the binary distribution which I considered until now, the bidding behavior is only for certain parameter constellations and for certain bid spaces more aggressive under $S = y$ than under $S = n$. That is, the bidding behavior is *not always* more aggressive. When the seller's value is distributed according to a cumulative distribution function G with support $[0, 1]$, the additional effect in (2) relative to (1) makes the bidding behavior under $S = y$ *always* more aggressive.[9] In that respect, the first effect arises more generally when the seller's value is continuously distributed.

Second effect. At first glance, it might seem to be an artifact of the discrete nature of the seller's value that the bidding behavior induced by $S = y$ and $B = [r_a, \infty)$ may exhibit a jump. Interestingly, this is not the case. I demonstrate this in the remainder of this subsection with a specific example which will be useful for some arguments later on.

Fix any $(z, p) \in (0, 1)^2$ for which the bidding behavior induced by $S = y$ and $B = [J^{-1}(0), \infty)$ exhibits a jump in the binary case and consider the following distribution: X_0 is uniformly distributed on $[-\delta, 0]$ with probability p and it is uniformly distributed on $[z - \delta, z]$ with probability $1 - p$. The so constructed cumulative distribution function is for any small positive δ continuous and converges pointwise towards the cumulative distribution function in the binary case with parameters (z, p) as $\delta \to 0$.

Suppose $\delta < z - J^{-1}(0)$ and consider $S = y$ in combination with $B = [J^{-1}(0), \infty)$. The condition on δ ensures that the lowest participating buyer type $J^{-1}(0)$ has a value which is strictly smaller than $z - \delta$. Note first that the bidding behavior which is induced in the continuous case may only differ from that in the binary case if some buyer type x_i' is willing to submit a bid $b_i' \in (z - \delta, z)$. However, if such a buyer type chose a marginally higher bid, his probability of getting the object would increase at least at rate $(1 - p)/\delta$ while his profit conditional on obtaining the object would decrease at rate -1. If δ is sufficiently small, he could increase his expected profit by increasing his bid. Hence, a bid $b_i' \in (z - \delta, z)$ cannot be consistent with optimal behavior when δ is sufficiently small.[10] It follows that the same bidding behavior as in the binary case is induced. As the bidding behavior exhibits a jump in the binary case, it exhibits a jump in the continuous case.

Setting a secret reserve price can also in the continuous case be strictly optimal for the seller. As in the binary case, the seller can use a hole in the bid space to affect where a jump occurs. Such a hole can also serve in the continuous case as a strategic tool to increase her expected profit. I construct now a bid space with a hole and I identify conditions under which this bid space in combination with $S = y$ is strictly better for the seller than any bid space in combination with $S = n$.

I continue to consider the case in which X_0 is distributed according to the continuous distribution with parameters (z, p, δ). Suppose that (z, p) satisfies the supposition in Proposition 1 and suppose that $\delta \in (0, \min\{z - \beta_{J^{-1}(0)}(J^{-1}(z)), \mathcal{U}_b(z, p) - \mathcal{U}_b(\overline{x}_0)\})$ (with $\mathcal{U}_b(z, p)$ and $\mathcal{U}_b(\overline{x}_0)$ denoting the bounds which I have derived in the preceding subsection for the binary distribution with parameters (z, p)).[11] As the

[9] This follows from standard mechanism design reasoning. For example, if $B = [0, \infty)$, standard mechanism design results can be used to show that the bidding behavior induced by $S = y$ is $b^y(x_i) := x_i - \int_0^{x_i} F(\tau)/F(x_i) \cdot G(\tau)/G(x_i) \mathrm{d}\tau$ whereas the bidding behavior induced by $S = n$ is $b^n(x_i) := x_i - \int_0^{x_i} F(\tau)/F(x_i) \mathrm{d}\tau$. $G(\tau)/G(x_i) < 1$ for any $\tau < x_i$ directly implies that $b^y(x_i) > b^n(x_i)$ for any $x_i > 0$.

[10] What does "sufficiently small" mean in this case? Let $\widehat{x}_i$ denote the value at which the bidding behavior jumps in the binary case. When a buyer with this value decreases his bid from z to $z - \delta'$ with $\delta' \in (0, \delta]$ in the binary case, the selling probability drops from 1 to p. When a buyer does this in the continuous case, the selling probability drops only from 1 to $p + (1 - p)(\delta - \delta')/\delta'$. Bidding z is nevertheless optimal for a buyer with value $\widehat{x}_i$ if $0 \in \arg\max_{\delta' \in [0,\delta]} [p + (1 - p)(\delta - \delta')/\delta](\widehat{x}_i - z + \delta')$. As this problem is strictly concave, $\delta = 0$ constitutes a maximum if the derivative of the objective function, $-(1 - p)/\delta \cdot (\widehat{x}_i - z + \delta') + [p + (1 - p)(\delta - \delta')/\delta] \cdot 1$, is non–positive at $\delta' = 0$. This is the case if $\delta \leq (1 - p)(\widehat{x}_i - z)$.

[11] Note that δ is from a non–empty set. $z - \beta_{J^{-1}(0)}(J^{-1}(z)) > 0$ follows from the supposition of Proposition 1. $\mathcal{U}_b(z, p) -$

supposition in Proposition 1 holds, $S = y$ in combination with some bid space B_{bin} is optimal in the binary case and the induced bidding behavior exhibits a jump from below z to above z. The supposition $\delta < z - \beta_{J^{-1}(0)}(J^{-1}(z))$ ensures that this bidding behavior does not rely on bids from $(z - \delta, z)$. It follows that by choosing $S = y$ and $B = B_{bin} \cap (z - \delta, z)$ the bidding behavior that is induced in the binary case can also be induced in the continuous case. By construction of the continuous distribution, this bidding behavior implies an expected profit for the seller which is in the continuous case by at most δ smaller than in the binary case. Hence, the seller's expected profit from the optimal first–price auction with a secret reserve price exceeds $\mathcal{U}_b(z, p) - \delta$. The construction of the continuous distribution implies further that the expected profit which the seller obtains from any first–price auction without a secret reserve price is smaller than $\mathcal{U}_b(\overline{x}_0)$. As $\delta < \mathcal{U}_b(z, p) - \mathcal{U}_b(\overline{x}_0)$ by my supposition, it follows that the constructed first–price auction with a secret reserve price is strictly better for the seller than any first–price auction without a secret reserve price.

3. Endogenous seller information

3.1. The augmented model

I introduce now an augmented model in which the time at which the seller commits to the auction rules B and S is endogenized by modifying two aspects of the model I introduced in Subsection 2.1. First, the seller can either commit to the auction rules early before she learns her value ($R = e$) or late after she learns it ($R = l$). The buyers observe when the rules are fixed and can thus infer whether the seller was already informed then. Second, the seller cares minimally about risk.[12] She has lexicographic preferences in her expected profit and the variance of her profit.

I will argue in this section that committing to the auction rules early can be an instrument for the seller to reduce the risk she faces. By considering preferences where the seller cares only minimally about risk, I reduce the usefulness of such a risk reduction. In that respect, the assumption of lexicographic preferences will work against my effects. However, lexicographic preferences improve the tractability of my model and have the advantage of making my model better comparable with standard auction models. In particular, the optimal auction without a secret reserve price corresponds for lexicographic risk–aversion to the optimal auction in Myerson (1981). This makes it transparent when and why the optimal auction with a secret reserve price differs structurally from this reference case. By contrast, for non–lexicographic risk–aversion, it is much less well understood how the optimal auction without a secret reserve price looks like. This makes it much harder to explain when and why structural differences arise through a secret reserve price.

3.2. Risk reduction through early fixation

I consider in this subsection the case in which the supposition of Proposition 1 holds. When the seller decides to commit to the auction rules early ($R = e$), the auction design problem corresponds to the problem I have analyzed in Section 2 except for that the seller cares now minimally about risk. As this difference does neither affect which auction rules are optimal nor which bidding behavior is induced (see Remark 4 after Proposition 1), it follows from the analysis in Section 2 that the seller chooses a first–price auction with a secret reserve price and that the bidding behavior $\beta_{J^{-1}(0),J^{-1}(z)}(x_i)$ is induced. The induced bidding behavior does not depend on the seller's value x_0. See Figure 3(a) for an illustration.

$\mathcal{U}_b(\overline{x}_0) > 0$ follows from the analysis in the preceding subsection.

[12]See Waehrer et al. (1998) for a discussion of auctions with risk–neutral buyers and a risk–averse seller.

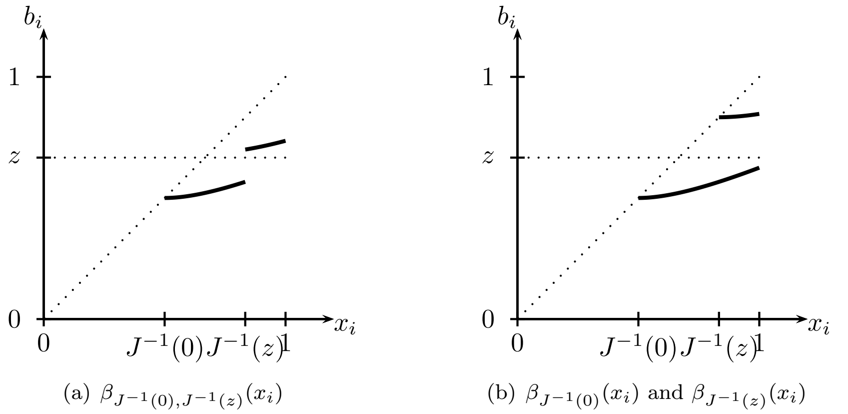

Figure 3: Bidding behavior under $R = e$ and $R = l$ [$X_i \sim U[0,1]$, $p = 1/2$, $z = 2/3$]

The auction design problem which arises when the seller waits with the auction design until she is informed ($R = l$) is an informed seller problem. The expected profit that a certain seller type obtains from a certain auction may depend on whether there is separation (i.e., the auction is chosen by her alone) or whether there is pooling on this auction (i.e., the auction is chosen by both seller types). The low value seller may benefit from pooling with the high value seller, but the high value seller never benefits from pooling with the low value seller.[13] As an unraveling argument applies, only a full separation equilibrium exists. The high value seller chooses $B = [J^{-1}(z), \infty)$ and the low value seller chooses $B = [J^{-1}(0), \infty)$. The choice of S does not matter as a secret reserve price would in neither case be binding. The induced bidding behavior depends on the seller's value. It is $\beta_{J^{-1}(0)}(x_i)$ if $x_0 = 0$ and $\beta_{J^{-1}(z)}(x_i)$ if $x_0 = z$ as illustrated by the lower and the upper curve in Figure 3(b), respectively.

As $R = e$ and $R = l$ induce the same allocation, they imply the same ex ante expected profit for the seller. Moreover, if a sale occurs to a buyer with a value $x_i \in [J^{-1}(0), J^{-1}(z))$, it is for the same price. However, if a sale occurs to a buyer with a value $x_i \in [J^{-1}(z), 1]$, the price under $R = l$ is a mean–preserving spread of the price under $R = e$:

$$\begin{aligned} & p\beta_{J^{-1}(0)}(x_i) + (1-p)\beta_{J^{-1}(z)}(x_i) \\ = \; & p(x_i - \textstyle\int_{J^{-1}(0)}^{x_i} F(t)/F(x_i)\mathrm{d}t) + (1-p)(x_i - \int_{J^{-1}(z)}^{x_i} F(t)/F(x_i)\mathrm{d}t) \\ = \; & x_i - \textstyle\int_{J^{-1}(0)}^{x_i} F(t)/F(x_i)\mathrm{d}t + (1-p)\int_{J^{-1}(0)}^{J^{-1}(z)} F(t)/F(x_i)\mathrm{d}t \\ = \; & \beta_{J^{-1}(0),J^{-1}(z)}(x_i) \end{aligned} \tag{8}$$

The first equality follows from (4) and the third equality follows from (6). The price distribution induced by $R = e$ has thus a lower variance than the price distribution induced by $R = l$. As the allocation is the

[13]The low value seller may benefit from pooling on auction rules with a secret reserve price for which the high value seller's secret reserve price is sometimes binding. E.g., the low value seller would be strictly better off if the auction rules $S = y$ and $B = [J^{-1}(0), \infty)$ were chosen also by the high value seller than when they are chosen by her alone (this follows from (4) and (6)). However, for any auction rules for which the high value seller's secret reserve price is sometimes binding, the high value seller is strictly worse off than under some auction rules for which the reserve price is higher than her value, i.e., for which a secret reserve price is never binding for neither seller type. For such auctions, any belief about the seller's value implies the same bidding behavior. It follows that it is optimal for the high value seller to choose $B = [J^{-1}(z), \infty)$ with $S \in \{n, y\}$ for any system of beliefs. Given that the high value seller behaves in this way, the only behavior of the low value seller that can be consistent with equilibrium behavior is choosing the auction rules that are optimal for her under separation, $B = [J^{-1}(0), \infty)$ with $S \in \{n, y\}$.

same in both cases, this implies that the profit distribution induced by $R = e$ has a lower variance than the profit distribution induced by $R = l$. This implies the following result:

Proposition 2 *If $\beta_{J^{-1}(0)}(J^{-1}(z)) < z$ and $\sigma_{J^{-1}(0)}(J^{-1}(z)) \geq z$, it is strictly optimal for the seller to commit to the auction rules early ($R = e$) but to set a secret reserve price later on ($S = y$).*

A rough intuition is the following: When the seller decides on the auction rules late after she is informed about her value, she chooses a higher open reserve price when her value turns out to be higher. As a consequence, the induced bidding behavior varies in her value. By contrast, when she commits to the auction rules early, the induced bidding behavior does not vary in her value. Committing to the auction rules early can thus serve as an instrument to induce a less variable bid distribution. On the other hand, by committing to the auction rules early, the seller foregoes to learn information which might be useful for the maximization of her profit. A risk–averse seller faces thus a trade–off between a less risky profit distribution ($R = e$) and an at least weakly larger expected profit ($R = l$). As the optimal first–price auction with a secret reserve price does by the reasoning in Section 2 not imply any sacrifice in expected profit when the supposition of Proposition 2 holds, committing to the auction rules early and setting a secret reserve price later on is clearly optimal for her in this case.

3.3. Continuous seller valuations and non–lexicographic preferences

My technique of proof in the preceding subsection relies on the binary distribution of X_0 and on the lexicographic preferences on the seller's side. For lexicographic risk–aversion the result does not extend to continuous distributions. To get an idea why, consider again the continuous distribution which I introduced in Subsection 2.4. By a reasoning analogous to that in Subsection 2.3, an upper bound on the seller's expected profit from any mechanism is given by $\mathcal{U}_c(z, p, \delta) := \mathbf{E}_X[\max\{X_0, J(X_1), J(X_2)\}]$. Although $\mathcal{U}_c(z, p, \delta)$ can under the supposition of Proposition 2 for $R = e$ be approximated by first–price auctions with a secret reserve price as $\delta \to 0$, the attained expected profit is for any small positive δ strictly smaller than $\mathcal{U}_c(z, p, \delta)$.[14] By contrast, the upper bound $\mathcal{U}_c(z, p, \delta)$ is under $R = l$ attained for any δ. Hence, the seller strictly prefers $R = l$ over $R = e$ although this implies only a slightly higher expected profit when δ is small.

The result extends however to cases with a continuous distribution of X_0 where the seller's preferences exhibit non–lexicographic risk–aversion: Observe first that lexicographic risk–aversion works against my effect. When I continue to consider a binary distribution of the seller's value but allow for non–lexicographic risk–aversion, the seller still prefers $R = e$ over $R = l$ whenever both induce the same expected profit, but she may then also prefer $R = e$ over $R = l$ when $R = e$ implies a strictly lower expected profit. Second, observe that while it is not possible to construct continuous distributions such that $R = e$ implies the same expected profit as $R = l$, it is possible to construct continuous distributions such that $R = e$ implies only a slightly lower expected profit but for which the induced variance of profit is much lower. Hence, for continuous distributions and non–lexicographic preferences $R = e$ can be optimal.

[14] The reason why $\mathcal{U}_c(z, p, \delta)$ can be approximated but can not be attained for $R = e$ is the following: By the reasoning in Subsection 2.4, $\mathcal{U}_c(z, p, \delta) \in [\mathcal{U}_b(z, p) - \delta, \mathcal{U}_b(z, p)]$ for any δ below some threshold. Since $\mathcal{U}_b(z, p)$ can be approximated as $\delta \to 0$ by the reasoning in the same subsection, it follows that $\mathcal{U}_c(z, p, \delta)$ can be approximated as $\delta \to 0$. On the other hand, $\mathcal{U}_c(z, p, \delta)$ cannot be attained for any fix δ as for any $x_0 \in [z - \delta, \delta]$ a different set of admissible bids is necessary to avoid the seller's commitment problem. It is thus impossible to attain $\mathcal{U}_c(z, p, \delta)$ when the seller commits to the set of admissible bids before she is informed about her value.

4. Discussion

The relationship between secret reserve price, phantom bidding and non–commitment to sell. I have interpreted a secret reserve price in this article as a secret minimum bid $r_s \in \mathbb{R}_+$ which is chosen *before* the auction is conducted. A secret reserve price is strongly related to two other instruments which a seller may use: A phantom bid can be interpreted as a bid $b_0 \in B$ which the seller submits *within* the auction, possibly through a third party or a fake identity. Moreover, when the seller cannot commit to sell or decides that she does not want to commit to sell, she chooses whether to accept the highest bid or to keep the object *after* she learns the bids submitted in the auction. The three instruments imply in my modeling framework the same commitment problem for the seller.[15] In that respect, my model can also be interpreted as one about phantom bidding or about non–commitment to sell.

Reinterpretation of my results. This relationship between the different instruments allows me to interpret my main result in the following way: Committing to the auction rules early but not committing to sell after observing the bids can be part of the optimal first–price auction for a seller who is risk–averse and whose information improves over time (even when waiting until she is better informed and committing to sell is possible for her). Likewise, committing to the auction rules early but submitting possibly a phantom bid later on can be optimal.

The role of the three instruments in applications. Secret reserve prices, phantom bidding and non–commitment to sell are frequently observed in practice. However, while secret reserve prices and phantom bidding seem to be important in online auctions, non–commitment to sell seems to be more important in procurement problems.

A seller who offers an object for sale via an online auction platform has often the opportunity to set a secret reserve price besides setting an open reserve price. Further, it is often possible for her to affect the auction outcome through a phantom bid. While the possibility of phantom bids is normally not communicated by the seller, it is clear to all parties that the placement of phantom bids via a fake identity or a third party can normally not be prevented. My analysis applies to situations in which the auction rules can be fixed a significant amount of time before the auction closes and in which the seller can either place a phantom bid or can set a secret reserve price which is updatable until the auction closes. My model provides a rationale for why it may in such situations be optimal for the seller to use an open reserve price in conjunction with either a phantom bid or a secret reserve price. Phantom bidding and secret reserve price can be used interchangeably. An open reserve price has however to be used in conjunction with either of the two other instruments.

A further possible application is a procurer's make–or–buy decision:[16] Bid preparation normally requires time. Bidders may need to build prototypes, design blueprints, negotiate with second tier suppliers, and so on. Procurement auctions are thus announced some time before the auction actually takes place. Information typically improves over time. E.g., the seller might find out whether she is able to produce herself or whether

[15] There are two small technical differences associated with phantom bidding. First, if the seller's phantom bid coincides with the highest bid submitted by the buyers, the allocation decision is taken according to a lottery. By contrast, I assumed that the object is sold if the highest submitted bid coincides with the seller's secret reserve price. Second, if the highest admissible bid in B is strictly smaller that z, the object is sold with a positive probability when the seller's value is $x_0 = z$, she uses a phantom bid and the highest submitted bid by the buyers is the highest admissible bid. By contrast, when she sets a secret reserve price or does not commit to sell, the object is not sold in this case. Both differences are inconsequential for my conclusions.

[16] Note that the effects in a reverse auction/procurement set–up are analogous to the effects in the forward auction set–up which I have analyzed in this article.

she requires an outside source. She may either commit to buy from a supplier or keep the possibility to produce herself. Fixing the auction rules early but keeping the possibility to produce herself can be optimal for a risk–averse seller.

The reason why secret reserve prices and phantom bidding are often used in online auction problems whereas non–commitment to sell is rather used in procurement problems may be due to reasons which lie outside of my model. For example, an auction platform may be able (and willing) to enforce that the seller does actually sell whenever a bid is submitted which exceeds her reserve price(s). This includes that she has to sell to "herself" when a phantom bid wins. Secret reserve price and phantom bidding may thus be the only instruments a seller can use in online auctions to induce the effects described in this article. On the other hand, there is typically very little commitment in procurement situations. A procurer can virtually always find a reason to finally produce in–house instead of buying from a supplier. It is thus not surprising that non–commitment is frequently observed in procurement. A procurer may even have strong reasons not to use a phantom bid even if this is in principle possible. If he uses a phantom bid, this might eventually be learnt by the suppliers as there is typically much industry information available. This in turn might have a negative impact on the trust within the procurer–supplier relationships and it might have negative consequences for the collaboration in the procurement of other parts or when the same part has to be procured for the next time.

5. Conclusion

For a setting in which a risk–averse seller's information about her value improves over time, I find that committing to the rules of a first–price auction early but keeping the right to set a secret reserve price later on can be optimal for the seller. Moreover, it can be optimal for her to force the buyers to choose extreme bids by forbidding intermediate ones. Similar conclusions can be drawn for situations in which the seller can use a phantom bid or in which she may decide to keep the object after observing the submitted bids. I discuss my results in the light of procurement and online auction problems.

Acknowledgements

This article arose from Chapter 3 of my doctoral dissertation "First–price Auctions, Seller Information and Commitment to Sell" at the University of Mannheim. I would like to thank Hans Peter Grüner, Benny Moldovanu, Konrad Stahl, Nora Szech, the editor, and anonymous referees for helpful comments. Financial support by the Deutsche Forschungsgemeinschaft through Graduiertenkolleg Risk and Liquidity in Financial, Commodity, and Factor Markets and through SFB–TR 15 is gratefully acknowledged.

Appendix

Proof of Lemma 1.

Suppose b specifies a symmetric equilibrium of the game induced by B and S. Define

$$Q(b_i) := \begin{cases} \left[\text{Prob}_X\{b_i > b(X_{-i})\} + \frac{1}{2}\text{Prob}_X\{b_i = b(X_{-i})\}\right] & \text{if } S = n \\ \left[\text{Prob}_X\{b_i > b(X_{-i})\} + \frac{1}{2}\text{Prob}_X\{b_i = b(X_{-i})\}\right] \times \text{Prob}_X\{b_i \geq X_0\} & \text{if } S = y \end{cases}$$

to unify the proofs for $S = n$ and $S = y$. $Q(b_i)$ describes the probability with which a buyer obtains the object when he submits a bid b_i and the other buyer behaves according to the strategy b. Moreover, define

$$U(x_i) := \begin{cases} Q(b(x_i))(x_i - b(x_i)) & \text{if } b(x_i) \in B \\ 0 & \text{if } b(x_i) = \emptyset \end{cases} .$$

$U(x_i)$ is the expected equilibrium profit of a buyer with value x_i. I divide the proof in two parts. I first show that that the four properties must hold when $Q(r_a) > 0$ and I show then that they must also hold when $Q(r_a) = 0$.

Part I: Suppose $Q(r_a) > 0$.

(a) Any buyer type $x_i > r_a$ obtains at least an expected profit of $Q(r_a)(x_i - r_a) > 0$ from participation and has thus a strict incentive to participate. On the other hand, any buyer type $x_i < r_a$ has a strict incentive not to participate as participation would lead to a strictly positive probability of obtaining the object and a strictly negative profit conditional on obtaining it. This is (a).

(b.i) Consider $x_i', x_i'' \in [r_a, 1]$ with $x_i' < x_i''$. Incentive compatibility implies that $U(x_i') \geq Q(b(x_i''))(x_i' - b(x_i'')) = U(x_i'') + Q(b(x_i''))(x_i' - x_i'')$ and that $U(x_i'') \geq Q(b(x_i'))(x_i'' - b(x_i')) = U(x_i') + Q(b(x_i'))(x_i'' - x_i')$. By combining both inequalities, I obtain $Q(b(x_i'')) \geq (U(x_i'') - U(x_i'))/(x_i'' - x_i') \geq Q(b(x_i'))$. This implies that $Q \circ b$ is weakly increasing on $[r_a, 1]$. Since Q is strictly increasing on $b([r_a, 1])$, I obtain $b(x_i'') \geq b(x_i')$. This is the first part of (b).

(b.ii) A buyer with value r_a has to be indifferent between non–participation and participation with bid $b(r_a)$. Since $Q(r_a) > 0$ implies that $Q(b_i) > 0$ for any $b_i > r_a$ and since $(r_a - b_i) < 0$ for any $b_i > r_a$, $b(r_a) = r_a$. This is the second part of (b).

(c) Assume that b is only weakly increasing. By (b), there exist then $x_i', x_i'' \in [r_a, 1]$ with $x_i' < x_i''$ and $b_i' \in B$ such that $b(x_i) = b_i'$ for any $x_i \in (x_i', x_i'')$. Individual rationality implies $(x_i - b_i') > 0$ for any $x_i \in (x_i', x_i'')$. This in turn implies that any buyer type $x_i \in (x_i', x_i'')$ has a strict incentive to slightly overbid b_i'. This would increase his probability of obtaining the object by a discrete amount while it would decrease his profit conditional on obtaining it only marginally. Since $B = [r_a, \infty)$ implies that it is also feasible to overbid any bid $b' \in B$ slightly, I obtain (c).

(d) Suppose there exists $\widehat{x}_i \in [r_a, 1]$ such that b jumps upwards from b' to b''. Standard reasoning implies that a buyer with value $\widehat{x}_i$ must be indifferent between b' and b''. Indifference requires $Q(b')(\widehat{x}_i - b') = Q(b'')(\widehat{x}_i - b'')$. Observe now that since Q is weakly increasing with $Q(r_a) > 0$, $Q(b') > 0$ and $Q(b'') \geq Q(b')$. This observation together with $\widehat{x}_i - b'' < \widehat{x}_i - b'$ implies that the indifference condition can only hold if $Q(b'') > Q(b')$. This directly implies (d).

Part II: Suppose $Q(r_a) = 0$.

(a) Since $\text{Prob}_X\{b_i \geq X_0\} > 0$ for any b_i, $\text{Prob}_X\{b(X_{-i}) = \emptyset\} = 0$ is necessary for $Q(r_a) = 0$. That is, each buyer participates with probability one. Suppose there exist $x', x'' \in [0, 1]$ with $x' < x''$ such that $b(x') = b(x'') = \emptyset$. Participation with probability one implies then that there exists $x''' \in (x', x'')$ such that $b(x''') \in B$ and $Q(b(x''')) > 0$. Individual rationality implies $Q(b(x'''))(x''' - b(x''')) \geq 0$. This implies in turn that a buyer with value x'' would obtain a strictly positive expected profit from submitting the bid $b(x''')$ contradicting that a buyer with value x'' does not participate. It follows that there can only be threshold participation behavior. Since $Q(r_a) = 0$, the participation threshold can only be zero. Since a participation threshold of zero can only arise in equilibrium if $r_a = 0$ (otherwise at least some bidders with values in $(0, r_a)$ had a strict incentive not to participate), the participation threshold is given by r_a. This is (a).

(b) Since Q is even when $Q(r_a) = 0$ strictly increasing on $b([r_a, 1])$, the proof that $b(x_i)$ is weakly increasing is as in Part I. It remains to argue that $b(r_a) = r_a$. Assume to the contrary that $b(r_a) > r_a$. b being weakly increasing implies that a buyer with value $x_i \in (r_a, b(r_a))$ obtains the object with a strictly positive probability and has a strictly negative profit conditional on obtaining it. This contradicts that $b(r_a) > r_a$ can be true in equilibrium. Hence, $b(r_a) = r_a$.

(c) The proof of this property is like in Part I.

(d) For the same reason that $b(r_a) > r_a$ cannot be true (see the proof of (b)), b cannot jump at r_a. For any jump at $\widehat{x}_i > r_a$, (b) implies that $Q(b') > 0$. This allows me to apply the same reasoning as in Part I again.

Derivation of (6).

$\beta_{r_a,\widehat{x}_i}(x_i) = \beta_{r_a}(x_i)$ for any $x_i \in [r_a, \widehat{x}_i)$ is obvious since both functions follow from the same differential equation and the same boundary condition. It remains to argue how $\beta_{r_a,\widehat{x}_i}(x_i)$ looks like for $x_i \in [\widehat{x}_i, 1]$. $\beta_{r_a,\widehat{x}_i}(x_i)$ is in this case determined by (3) with $x_i' = \widehat{x}_i$ and $b(x_i') = z$. I have thus the following:

$$\begin{aligned}
\beta_{r_a,\widehat{x}_i}(x_i) &= z\frac{F(\widehat{x}_i)}{F(x_i)} + \int_{\widehat{x}_i}^{x_i} \tau \frac{f(\tau)}{F(x_i)} \mathrm{d}\tau \\
&= (p\beta_{r_a}(\widehat{x}_i) + (1-p)\widehat{x}_i)\frac{F(\widehat{x}_i)}{F(x_i)} + \left[\tau \frac{F(\tau)}{F(x_i)}\right]_{\tau=\widehat{x}_i}^{\tau=x_i} - \int_{\widehat{x}_i}^{x_i} \frac{F(\tau)}{F(x_i)} \mathrm{d}\tau \\
&= p(\beta_{r_a}(\widehat{x}_i) - \widehat{x}_i)\frac{F(\widehat{x}_i)}{F(x_i)} + x_i - \int_{\widehat{x}_i}^{x_i} \frac{F(\tau)}{F(x_i)} \mathrm{d}\tau \\
&= -p\int_{r_a}^{\widehat{x}_i} \frac{F(\tau)}{F(\widehat{x}_i)} \mathrm{d}\tau \frac{F(\widehat{x}_i)}{F(x_i)} + x_i - \int_{\widehat{x}_i}^{x_i} \frac{F(\tau)}{F(x_i)} \mathrm{d}\tau \\
&= \beta_{r_a}(x_i) + (1-p)\int_{\widehat{x}_i}^{x_i} \frac{F(\tau)}{F(x_i)} \mathrm{d}\tau
\end{aligned}$$

The second equality follows from applying partial integration to the integral, from using that $z = \sigma_{r_a}(\widehat{x}_i)$ and from using the definition of $\sigma_{r_a}(\widehat{x}_i)$ (see (5)). The fourth and the fifth equality follow from using the definition of $\beta_{r_a}(\widehat{x}_i)$ and of $\beta_{r_a}(x_i)$, respectively. See (4). The fifth equality establishes that $\beta_{r_a,\widehat{x}_i}(x_i)$ is also for $x_i \in [\widehat{x}_i, 1]$ as stated in (6).

References

Ashenfelter, O. (1989): "How Auctions Work for Wine and Art," *Journal of Economic Perspectives*, 3, 23–36.

Elyakime, B., J. J. Laffont, P. Loisel, and Q. Vuong (1994): "First-Price Sealed-Bid Auctions With Secret Reservation Prices," *Annals of Economics and Statistics*, 34, 115–142.

Horstmann, I. J. and C. LaCasse (1997): "Secret Reserve Prices in a Bidding Model with a Resale Option," *American Economic Review*, 87, 663–684.

Krishna, V. (2009): *Auction Theory*, Academic Press, second edition ed.

Li, H. and G. Tan (2000): "Hidden Reserve Prices with Risk–Averse Bidders," Working Paper.

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

Rosenkranz, S. and P. W. Schmitz (2007): "Reserve Prices in Auctions as Reference Points," *The Economic Journal*, 117, 637–653.

Vincent, D. R. (1995): "Bidding Off the Wall: Why Reserve Prices May Be Kept Secret," *Journal of Economic Theory*, 65, 575–584.

Waehrer, K., R. M. Harstad, and M. H. Rothkopf (1998): "Auction Form Preferences of Risk-Averse Bid Takers," *The RAND Journal of Economics*, 29, 179–192.