Auctions as Appraisal Mechanisms: Seller Behavior in eBay Auctions*

Jafar Olimov†

January 1, 2013

Abstract

In this paper we study seller behavior using data from eBay auctions of used tractors. We relax the standard assumption that sellers know the distribution functions of items' valuations and find that uninformed and patient sellers use secret reserve prices to run unsuccessful eBay auctions to learn parameters of these unknown distribution functions. We also study determinants of public reserve prices, and find that public reserve prices are simultaneously used to screen out low valuation bidders and to signal the quality of listed items. We find that secret reserve prices and public reserve prices have a strong positive effect on sale prices. We provide a novel theoretical justification for the use of secret reserve prices and show that eBay serves not only as a selling platform but also as an affordable value-appraising mechanism for items whose valuation is not easily available or is costly to obtain. Keywords: auctions, eBay, learning, reserve prices

*I would like to thank Brian Roe, John Kagel, and AEDE seminar participants at OSU. Salaries and research support were provided by the National Institute of Food & Agriculture (grant# 2008-25400-18701), the Ohio Agricultural Research and Development Center, The Ohio State University and the McCormick Program of Agricultural Marketing and Policy, Ohio State University.

†Department of AED Economics, The Ohio State University, 312 Ag. Admin. Building, 2120 Fyffe Road, Columbus, OH 43210, Tel: 614-292-9346; Fax: 614-292-7710; Email: olimov.1@osu.edu

1 Introduction

In this paper we study seller behavior in eBay auctions for high-valued goods where external appraisal values are difficult or costly to obtain. We find that many sellers use eBay auctions not only to sell such items but also to gather information about items' valuations from observed bids in eBay auctions. Sellers use the gathered information to update their beliefs about distribution functions of items' valuations. To gather information about items' valuations from realized bids without necessarily selling the items, sellers in eBay auctions employ high secret reserve prices, because high secret reserve prices prevent auction sales and at the same time do not restrict entry of bidders into auctions. To signal the quality of their items many sellers use public reserve prices together with secret reserve prices. The proposed behavior of sellers on eBay explains several previously unanswered stylized facts in the auctions literature: the widespread use of secret reserve prices by sellers, the presence of repeated auctions, and the low rate of success of eBay auctions for some item categories.

Like other market transaction mechanisms, an auction transfers an item from a seller to a buyer at some specified price. However, unlike other market transaction mechanisms, an auction, and in particular an ascending-price auction, also acts as an information gathering mechanism by inducing competition among buyers and by forcing these buyers to reveal their valuations of the auctioned item. If a seller has no clear idea about the value of her item, is sufficiently patient, and wants to sell at the highest possible price, she clearly gains by running at least one ascending-price auction without selling the item and learning the distribution of realized bids before engaging in an actual sale. The ability to observe realized bids without selling an item can help this uninformed and patient seller to not only discover the demand for her item but also to choose an optimal selling price. In particular, the uninformed and patient seller can run an unsuccessful ascending-price auction, learn the highest bid, and then offer her item in a posted price sale at the highest bid from the unsuccessful auction stage. This way the seller's expected revenue should equal the expected highest valuation minus the cost of an unsuccessful auction. If the seller's cost of an unsuccessful auction is sufficiently small, then the expected payoff from implementing this two-stage sale clearly exceeds the expected second highest valuation - an expected payoff from running a single-stage ascending-price auction.

The question is whether an uninformed and patient seller who wants to learn the distribution of an item's valuation can convince buyers to enter and bid truthfully in the auction stage(s). Electronic selling platforms such as eBay offer an option suitable for this purpose. The rules of an eBay auction allow a seller to set a secret reserve price hidden from bidders. Bidders do not observe a secret reserve price itself, but they can see if a secret reserve price is set until some bid exceeds it. By setting a high secret reserve price, a seller can run one or multiple unsuccessful ascending-price auctions, observe the truthfully revealed valuations, and use this knowledge either in a different selling format on eBay or delist her item from eBay and use the acquired knowledge in a sale elsewhere[1]. In this paper we provide empirical evidence that secret reserve prices could be used exactly for this purpose: as an effective instrument, which allows uninformed and patient sellers to use eBay auctions not only as a selling mechanism but also as an affordable alternative to an otherwise expensive or unavailable appraisal mechanism.

This paper fits into the literature on the empirical estimation of auctions data. However, unlike most studies, we concentrate on the behavior of sellers rather than on the behavior of bidders[2]. In particular, we study how and why sellers set public and secret reserve prices in eBay auctions.

Many studies test public and secret reserve prices in the context of eBay auc-

[1] Specifically, an uninformed and patient seller can use a posted price sale format available on eBay after running an unsuccessful eBay auction and learning parameters of the distribution function of bidders' valuations.

[2] In a recent working paper Einav et al. (2012) also address seller behavior in online markets. They argue that sellers use online platforms such as eBay to experiment with auction parameters and with sale formats.

tions. However, most of them test secret and public reserve prices only in static auctions and ignore the fact that many unsold items are relisted. In general, existing empirical studies find that public reserve prices raise sale prices, while secret reserve prices have either a negative or an insignificantly positive effect on sale prices. For example, Bajari and Hortacsu (2003) study public and secret reserve prices in static eBay auctions for collectible US coins and find a small positive effect of secret reserve prices on sale prices. Katkar and Reiley (2006) employ field experiments and use static eBay auctions to sell Pokemon cards with public and secret reserve prices. They find that the presence of a secret reserve price lowers revenue by 10% and the likelihood of sale by 34%. Lucking-Reiley et. al. (2007) use data from static eBay auctions of collectible pennies and find that, conditional on sale, public reserve prices have a small positive (around 1%) but statistically insignificant effect on final sale prices while secret reserve prices have a sizeable positive (about 15%) and statistically significant effect on final sale prices[3]. Carare (2012) studies revenue effects of only public reserve prices in dynamic auctions for computer processors and finds that optimal public reserve prices in dynamic auctions have a significant positive effect on seller revenue[4].

The theoretical discussion of the functions and determinants of public reserve prices depends on whether bidders' valuations are independently distributed or have a commonly distributed component. If bidders' valuations are independently distributed, then, according to Myerson (1981) and Riley and Samuelson (1981), sellers use a public reserve price to screen out bidders with valuations below some threshold level. If bidders' valuations have a commonly distributed component, then as Milgrom and Weber (1982) and Cai, Riley and Ye (2007) show, a public reserve price serves as a credible signal of the quality of an item on sale.

While the functions and determinants of public reserve prices have been thor-

[3] Lucking-Reiley et. al (2007) study sale effects of public and secret reserve prices only in successful auctions without relistings.

[4] Carare's derivation of an optimal public reserve price is different from the standard Myerson's optimal public reserve price, because Carare's derivation accounts for the possibility of a future relisting.

oughly researched and understood in auction theory, theoretical studies of auctions are silent on the function and determinants of secret reserve prices. We are aware of only two theoretical studies of secret reserve prices, and both these studies look at secret reserve prices as possible alternatives to public reserve prices. Vincent (1995) shows that if bidders' valuations are commonly distributed, then secret reserve prices can generate more revenue than public reserve prices by encouraging entry. Rosenkranz and Schmitz (2007) use prospect theory and find that if bidders' valuations are independently distributed, a public reserve price enters bidders' utility functions, and if bidders' outside options exceed the public reserve price, then a secret reserve price can outperform a public reserve price. However, no empirical analysis of either of these theories is present in the literature.

We can draw several conclusions from the existing theoretical and empirical studies of seller behavior in auctions. First of all, there is a gap in the theoretical literature related to the use and the function of a secret reserve price. The available theoretical studies attribute functions of a public reserve price to a secret reserve price and analyze the effectiveness of a secret reserve price as an alternative to a public reserve price. We are able to address this gap in the literature by relaxing the theoretical assumption that sellers have a perfect knowledge of the distribution function of bidders' valuations and by showing that an uninformed patient seller can use a secret reserve price to obtain information about the distribution function of valuations by running an unsuccessful ascending-price auction. Hence, we argue that a secret reserve price has a completely different purpose than a public reserve price. While a public reserve price is used to screen out low-valuation bidders and to signal the quality of an item on sale, a secret reserve price is primarily used by an uninformed and patient seller to collect information about the distribution function of valuations prior to engaging in an actual sale.

The existing empirical studies of seller behavior in auctions are inconclusive about the impact of secret reserve prices on seller revenue and sale prices. We test revenue effects of secret reserve prices under the assumption that secret reserve prices are used not as screening or signaling devices but as a mechanism through which uninformed and patient sellers are able to gather information about items' valuations. We find that under our specification, secret reserve prices have a strong positive effect on sale prices. We run a simulation exercise using the actual data from sales of used tractors in eBay auctions, where we test the price effect of imposing a secret reserve price in repeated auctions under different belief-updating rules. We find that as long as the cost of relisting is not too high and an uninformed seller updates a secret reserve price given the information from previous unsuccessful auctions, the sale price in an auction with an appropriately set secret reserve price exceeds the sale price in an auction without a secret reserve price[5].

Our final contribution is related to determinants of public reserve prices, secret reserve prices and buy-it-now prices on eBay. We are not aware of any empirical studies with structural tests of determinants of secret and public reserve prices using eBay data. We find that sellers of used tractors on eBay use public reserve prices both to screen out low valuation bidders and to signal the quality of listed tractors. Hence, many empirical studies, which a priori assume that public reserve prices are used exclusively for screening, overlook an important component defining seller behavior. We also test for various determinants of secret reserve prices and find that the size of a secret reserve price is determined by the highest bid from previously run unsuccessful second-price auctions and by average highest bids observed in auctions for similar items. In addition, we test for determinants of buy-it-now (BIN) prices and find that buy-it-now prices are almost exclusively determined by highest bids from previous unsuccessful auctions. To test for determinants of public reserve prices, secret reserve prices, and BINs, we employ a completely non-parametric approach to minimize the number of identification assumptions.

In the next section we discuss the data. In section 3 we define the model of

[5]In the simulation exercise we account for the possibility that an item may go unsold if the secret reserve price is not met after multiple relistings. We assume that the sale price is zero if the item does not sell after the fifth relisting.

a seller and a bidder behavior. In section 4 we discuss the estimation strategy. In section 5 we discuss determinants of public reserve prices, secret reserve prices, and BINs. In section 6 we test determinants of sale prices.

2 Data

To study the behavior of sellers on eBay we use the data on used tractors sold in eBay auctions between 11/17/04 and 5/30/07. Our data consists of two components: (1) a complete account of all auction sales on eBay between 11/17/04 and 5/30/07 without bidder characteristics, and (2) bidder characteristics and actual bids for a portion of auctions held on eBay between 11/17/04 and 5/30/07. Since we don't have realized bids and bidder characteristics for a majority of auctions held on eBay, for the analysis of a seller behavior we construct a smaller sample of auctions with complete information on bids and bidder characteristics[6]. To give an idea about the market for tractors on eBay, we first present information about all auction sales of tractors held on eBay between 11/17/04 and 5/30/07. However, when we discuss tractor characteristics, we present information only about those tractors, for which we have complete information on bids and bidder characteristics[7].

A cursory look at outcomes of sales of tractors in eBay auctions suggests that many auctions are not successful and a vast number of tractors are not sold on eBay. This is a mystery, especially, if we recall that auctions have an advantage over other sale formats (such as posted price sales) in terms of a higher probability of success. There is a reasonable suspicion that many of these unsold tractors are

[6] There is a sizeable share of auctions with missing information about tractor characteristics such as horse power and year of production. We use tractor model numbers to recover these parameters from outside sources. The main source of the outside data is www.tractor-data.com.

[7] All the auctions with missing bidder characteristics are auctions taking place in 2006 and 2007. We find a bias in the sub-sample of auctions with known bidder characteristics: for example, the average tractor horse power in the sub-sample of auctions with known bidder chracteristics is less than the average horse power in the full sample by 1.43 HP. This difference is statistically significant at 1% level with t-statistic of 15.421.

relisted over and over again. We are able to match tractors from multiple listings and construct a dataset of tractors with multiple relistings[8]. We present the data on all repeated and single auction sales of tractors held on eBay between 11/17/04 and 5/30/07 in Table 1.

Table 1. Auction sales of tractors on eBay between 11/17/04 and 5/30/07.

Listing patterns	Number listed	Number sold	Percent sold
Single listing	23253	13251	56.99
2 listings	4031	1448	35.91
3 listings	1069	344	32.19
4 listings	404	117	28.96
5 listings	197	60	30.46
6 listings	102	30	29.41
7 listings	61	20	32.79
8 or more listings	111	34	30.63
Total number of unique tractors	29228	15304	52.36

Notes: Single listing tractors might have been listed multiple times prior to 11/17/04, when the earliest observations in our data first appear.

We can notice two regularities by looking at Table 1. First, many tractors are delisted from eBay auctions after being unsuccessfully auctioned[9]. In fact, only about a half of all tractors are successfully sold even if we account for relistings.

[8] When deciding whether it is the same tractor with multiple listings or different tractors, we use information on seller id's, whether the prior sale was successful, tractor characteristics, and engine hours.

[9] By "delisted" we mean unsuccessfully auctioned tractors that are not sold in auctions again. It is possible that some unsuccessfully auctioned tractors are relisted on eBay under a different sale format (such as, for example, an eBay posted price sale). In this dataset we do not have information on eBay posted price sales.

The highest share of successfully sold tractors is found among tractors which were listed only once[10]. The second interesting regularity is that the share of sold tractors across multiple relistings is stable and on average does not diverge from a 30% success rate.

Next, we present the data on times between listings for tractors listed twice. Table 2 shows that a half of unsuccessfully sold tractors are relisted within 3 days after the end of an unsuccessful sale and almost 30% of unsuccessfully sold tractors are relisted within the same day. The maximum time between listings for tractors listed twice is 540 days, or about a year and a half. The average time between listings is 17.36 days.

Table 2. Times between listings for tractors listed twice.

Duration between relistings	Number of 2-stage auctions	Percent
less than 1 day	1201	29.79
1-3 days	811	20.12
4-7 days	559	13.87
8 days to 2 weeks	439	10.89
2 weeks to 1 month	525	13.02
from 1 month to 2 months	262	6.50
more than 2 months	234	5.81
Total	4031	100.00

Sellers on eBay can choose the duration of their auctions. In our data the maximum duration of auctions is 10 days and the minimum duration of auctions is less than 1 day. More precisely, 30% of auctions last less than 7 days, 53% of all auctions last exactly 7 days and about 18% of auctions last 10 days.

We present the data on various auction parameters used by sellers of tractors on eBay in Table 3. Sellers on eBay have a choice to set three parameters before running an auction: a public reserve price (PRP), a secret reserve price (SRP),

[10] Note that tractors with 1 listing might have been listed multiple times prior to 11/17/04, when the earliest observations in our data first appear.

and a buy-it-now price (BIN)[11]. A public reserve price defines the minimum bid from which participating bidders must start. A secret reserve price defines the unobservable reserve price, which the highest bid must exceed to successfully end an auction. If the highest bid is below a secret reserve price by the end of an auction, the item goes unsold and the auction ends unsuccessfully. Although a secret reserve price is not observable, bidders can see an indicator "Reserve Not Met" if there is an unmet secret reserve price in an eBay auction listing. The secret reserve price indicator disappears, once a bid exceeds it. A seller can adjust her secret reserve price at any time before the last 12 hours of an active auction. In the last 12 hours of an auction an active secret reserve price cannot be changed or deactivated.

Table 3. Parameters in eBay auctions.

Auction parameter	Observability	eBay fee for tractors
Public reserve price	Fully observable	\$5
Secret reserve price	Unobservable price, observable indicator	\$5
Buy-it-now price	Observable prior to auction, unobservable later	if listing $\leq$ 50 items/month, free, otherwise \$0.25/listing

A BIN is a publicly observable temporary sale price at which any bidder can purchase an item before the start of an auction. Any indicator of a BIN or its presence disappears in an auction without a secret reserve price, once any bidder makes a bid. If a BIN is used together with a secret reserve price, then the BIN indicator is observable until a secret reserve price is met. We present a description

[11] On eBay a public reserve price is called a starting price, while a secret reserve price is called a reserve price. To avoid confusion, in the rest of the paper we adopt the standard terminology from auctions literature rather than the terms used on eBay.

of these three parameters and fees for using them in auctions for tractors[12].

In the next two tables we present the data on the use of different auction parameters across sales with one and two listings. We can see that at least a third of sellers use a public and a secret reserve price at the same time. Further, more than a half of sellers of tractors with two listings and 46% of sellers of tractors with a single listing employ secret reserve prices alone or in combination with a public reserve price and a BIN. We can also see that the use of secret reserve prices is more than 6% higher in the first listing than in the second listing for tractors listed twice and that the use of a public reserve price and a secret reserve price in the first listing (all unsuccessful auctions) among tractors listed twice is much higher than the use of a public and a secret reserve price for tractors listed only once.

Table 4. The use of different auction features for tractors listed only once.

Features used	Number	Percent
Auctions with a public reserve price (PRP) (>$100)	15573	66.97
Auctions with a secret reserve price (SRP)	10705	46.04
Auctions with a buy-it-now price (BIN)	5208	22.39
Auctions with PRP and SRP	7544	32.44
Auctions with PRP, SRP and BIN	1925	8.28
Auctions with SRP and BIN	2861	12.30
Auctions without any features	4135	17.78
Total number of tractors listed once	23253	100.00

[12] eBay fees for using public/secret reserve prices and BINs depend on an item category.

Table 5. The use of different auction features for tractors listed twice.

Features used	First listing		Second listing	
	Number	Percent	Number	Percent
Auctions with a public reserve price (PRP)(>$100)	3763	93.35	3301	81.89
Auctions with a secret reserve price (SRP)	2314	57.41	2064	51.20
Auctions with a buy-it-now price (BIN)	1250	31.01	1375	34.11
Auctions with PRP and SRP	2094	51.95	1506	37.36
Auctions with PRP, SRP and BIN	643	15.95	533	13.22
Auctions with SRP and BIN	724	17.96	751	18.63
Auctions without any features	45	1.12	136	3.37
Total number of tractors listed twice	4031	100.00	4031	100.00

In the next table and two figures we present the data on tractor characteristics, seller ratings, and the number of bidders in those auctions, for which we have complete information about all realized bids and all bidders. Table 6 shows that the average number of bidders in auctions with complete bidder information is about 8.5 with a minimum of 1 and a maximum of 32. Since we don't have complete information on bids for all auctions, we cannot identify auctions with zero entry from auctions with missing bidder characteristics. This is why we do not have auctions with zero entry in our sample and the minimum number of bidders in our sample is 1.

Table 6. Tractor characteristics in auctions with complete bidder information.

	Number of observations	Mean	Standard Deviation	Min	Max
Number of bidders	12429	8.469	5.282	1	32
Seller feedback score	12429	201.652	937.296	0	80282
Tractor age	12429	22.212	16.749	0	80
Tractor horse power	12429	42.076	29.446	10	150

Table 6 shows that the average age of tractors in our data is more than 22 years and 68% of tractors range in age from 5.5 years to 38.9 years. The distribution of tractors by age is given in Figure 1. By looking at the figure we can see that there are two clearly identified modes. The first mode is located at the age of 1 and the second mode is located at the age of 25. We think that the bimodal distribution of tractor ages can be explained by practices of manufacturers of tractors. Manufacturers of tractors usually update their production lines from 4 to 10 years. As a result, sellers of used tractors on eBay have a stronger incentive to sell their tractors once new generations of tractors become available. The mode at the age of 1 clearly captures sales of the latest generations of tractors, while the mode at the age of 25 captures sales of tractors of previous 2 or 3 generations, when the difference between tractor characteristics from the current generation and the past generations becomes more critical.

Figure 1. Distribution of tractors by age.

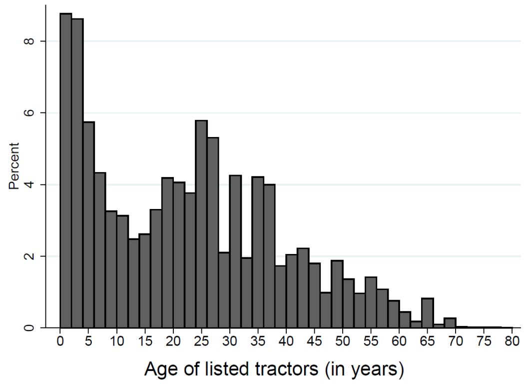

Figure 2. Distribution of tractors by horse power.

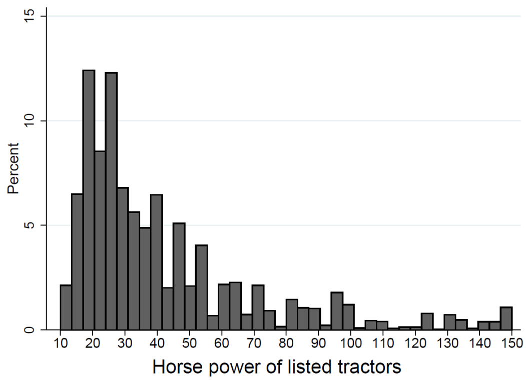

We present the distribution of tractors by horse power in Figure 2. In general, tractors can be separated into three categories by horse power: lawn tractors, small or compact utility tractors and utility tractors. Lawn tractors rarely exceed 20 HP in power and are used for mowing lawns or carrying relatively light loads. Small or compact utility tractors range in power from 20 HP to 40 HP and are the most popularly used tractors for everyday farming needs. Utility tractors range in power from 40 HP to 300 HP and are used for cropping, construction, and other heavy duty tasks. The distribution of tractors by horse power in Figure 2 shows that the most frequently listed tractors are either lawn tractors with HP below 20 or compact utility tractors with HP between 20 and 40.

3 Model and Predictions

In this section we discuss the model of bidder and seller behavior on eBay. The bidder behavior in eBay auctions has been extensively discussed (see Hasker and Sickles (2010) for the latest survey), and it can be described as "proxy bidding." To participate in proxy bidding in an eBay auction, a bidder has to specify the maximum willingness to pay. A computer then raises bids in increments from the minimum starting price set by the seller and up to the value indicating the maximum willingness to pay specified by the bidder. A winning bidder is notified by E-mail about the winning price after the end of an auction. Bidders can raise values of their willingness to pay or to make bids personally at any time during an auction.

Given the rules of proxy bidding, the key feature of an eBay auction is that bidders incur minimum participation costs. Since bidding on eBay does not require bidders' physical participation and is free, a typical bidder does not have to face any participation costs aside from the search costs.

Since bids increase and a winning bidder pays an increment above the second highest bid in an eBay auction, we use a theoretical model of a static ascending second-price auction to model bidder behavior[13]. We further assume that bidders' valuations have a finite common support $[0, v]$ and are identically and independently distributed. The independence assumption is used to capture the fact that we study a market of used goods with an unlikely possibility of a post-auction resale[14]. We further allow that bidders engage in "sniping," or that bidders place

[13]By the static ascending second-price auction we mean a standard second-price auction without resale or relisting and satisfying the assumptions of Myerson (1981).

[14]Unfortunately, it is impossible to identify if tractors are bought for resale, because we do

their bids in the last minutes of an auction closing time[15]. According to Ockenfels and Roth (2006), the main consequence of sniping is that some bidders are not able to place their highest bids before an auction closing time. As a result, the realized bids of non-winning bidders may not reflect the bidders' valuations of an item on sale.

According to bidding strategies in an ascending second-price auction with sniping, the winning bidder in an eBay auction for used tractors bids the second-highest valuation or some value below the second-highest valuation[16]. All other bidders bid their true valuations or some values below them. As a result, if an eBay auction is unsuccessful and there is no winner, the seller and the bidders observe true valuations of all participating bidders or some values below them, and if an auction is successful and the item is sold, the seller and the bidders observe true valuations or some values below them of all participating bidders but the winner. Hence, we state our first two assumptions about bidder behavior in eBay auctions.

Assumption 1. Bidders' cost of participation is zero.

Assumption 2. Bidders' valuations are identically and independently distributed on a finite support $[0, v]$.

It is harder to formalize seller behavior on eBay, because in the theoretical auction literature seller behavior is analyzed under the same informational assumptions as bidder behavior. In particular, it is usually assumed that seller and bidder valuations are realizations of the same distribution function, and that the seller and bidders know this distribution function but do not know each others' realized valuations. Under the assumption that the seller knows the distribution function of bidders' valuations, the sole role of a profit-maximizing seller is to

not observe tractors' serial numbers and we do not have data on the fate of purchased or unsold tractors outside eBay.

[15] We find that 36.23% of auctions in our sample have at least one bid placed in the last minute of an auction closing time.

[16] Due to sniping, there is a positive probability that the second-highest bidder is not able to bid her true valuation and the winning bid is below the second-highest valuation.

calculate and set an optimal public reserve price before the start of an auction.

As we can see the actual eBay auction format gives a much richer set of actions to a seller. In addition to choosing a sale format, which is beyond the scope of discussion of this paper, in a standard eBay auctions a seller can set a public reserve price, a secret reserve price and a buy-it-now (BIN) price.

The function of a public reserve price depends on whether bidders' valuations are independently distributed or have a commonly distributed component. If bidders' valuations are independently distributed, then according to Myerson (1981) and Riley and Samuelson (1981), a public reserve price screens out low valuation bidders. If bidders' valuations have a commonly distributed component, then Cai, Riley, and Ye (2007) show that a public reserve price is a credible signal of the value of the item on sale.

The second parameter available to a seller is a secret reserve price. An important difference between a secret reserve price and a public reserve price is that a secret reserve price does not restrict bidders' entry, and bidders with valuations below a secret reserve price can freely participate in an auction.

The last parameter available to a seller on eBay is a buy-it-now (BIN) price. A seller can set a price at which the item on sale can be sold immediately before the beginning of an auction. Once the first bid greater than the secret reserve price (when applicable) is made, the BIN and any indicator of its presence disappear and no bidder can observe the BIN later.

Before we continue discussing seller choice of reserve prices and a BIN, we need to specify how bidder behavior is affected by these parameters. Before we proceed, we assume that bidders and the seller are risk-neutral.

Assumption 3. Bidders and the seller are risk-neutral.

The impact of a public reserve price on bidder behavior has been extensively discussed in the literature, and under assumption 2 it amounts to limiting the participation of bidders with valuations below some threshold level. If assumption

2 is violated and bidders' valuations have a commonly distributed component, then a public reserve price signals the quality of an item on sale and raises bids of participating bidders. Nevertheless, whether assumption 2 is violated or not, when a public reserve price is present, no bidder with a valuation below the public reserve price should enter an auction.

Mathews (2003) has discussed the optimality of temporary BINs used on eBay and the effect of such BINs on bidder behavior[17]. He shows that in a particular case when bidders and a seller are risk-neutral, there is an equilibrium in which a profit-maximizing seller sets a BIN equal to the upper bound of the support of the distribution function of bidders' valuations, while bidders bid without taking the BIN option. When the seller is risk-averse, in equilibrium the seller sets the BIN below the upper bound of the support of valuations, while bidders take the BIN option with a positive probability. Using Mathews' result, we conclude that under risk-neutrality, the presence of a BIN in an eBay auction should not affect bidders' participation rates and bidding strategies.

However, whenever a BIN is used together with a very high secret reserve price, a temporary BIN effectively turns into a permanent BIN, since by eBay rules a BIN is active as long as the secret reserve price is not met. In this case, Hidvegi et. al. (2006) show that a bidder's equilibrium strategy depends on whether her valuation is above or below an active BIN. If the bidder's valuation is below an active BIN, then the bidder should truthfully bid up to her valuation. If the bidder's valuation is above an active BIN, then the bidder should bid truthfully until winning an auction or until reaching some threshold value after which the bidder should take the BIN.

The theoretical discussion of a secret reserve price and its effect on bidder behavior has received only limited attention in the theoretical literature. The

[17] The difference between a temporary BIN used in eBay auctions and a permanent BIN used in other selling formats is that a temporary BIN disappears once an auction starts, while a permanent BIN stays available during the course of the whole auction. For discussion of a permanent BIN, see, for example, Hidvegi et. al. (2006).

available studies view a secret reserve price as an alternative to a public reserve price, and assume that a secret reserve price plays the same screening role as a public reserve price, however, without limiting the entry. For example, Vincent (1995) argues that if bidders are risk-averse, then in a common-value environment a secret reserve price may have an advantage over a public reserve price by encouraging entry.

Based on assumption 1 that bidders' participation costs in a standard eBay auction are zero, we also argue that the presence of a secret reserve price should not affect bidders' entry as long as there is a non-zero probability of winning an auction. The next question is whether participating bidders have a different bidding strategy in the presence of a secret reserve price. The answer to this question depends on whether participating bidders acquire any additional information by observing the presence of an active secret reserve price and whether bidders obtain any additional payoff in an auction with a secret reserve price.

We argue that the presence of a secret reserve price in a single-stage ascending second-price auction does not reveal any additional information to bidders nor does it give any additional payoff above the standard expected bidder payoff[18]. Hence, if bidders participate in a single-stage ascending second-price auction with a secret reserve price, we should not observe any changes in the bidder behavior from an equilibrium bidder strategy in a single-stage ascending second-price auction without a secret reserve price. However, if the same bidders participate in several auctions for the same item and with secret reserve prices in one or more listings, the informational structure of bidders and bidders' expected payoffs should change. In particular, if bidders participate in a repeated ascending second-price auction with a secret reserve price in each stage, (a) the bidders are able to update their beliefs about the secret reserve price at the end of each unsuccessful stage, and (b) the bidders should take into account an expected discounted payoff from participating

[18]By a single-stage ascending second-price auction we mean a standard static second-price auction without resale or relisting and satisfying the assumptions of Myerson (1981).

in all auction stages under an evolving information structure. To simplify analysis, we assume that bidders do not participate in more than one stage of an eBay auction with multiple relistings.

Assumption 4. Bidders participate only in one stage of auctions with multiple relistings.

Given assumptions 1, 2, 3, and 4 we are ready to define bidder strategies in an eBay auction with one or more relistings in the presence of a public reserve price, a secret reserve price, and a BIN.

Proposition 1. *Let $[0, v]$ be a finite support of the distribution function of bidders' valuations, $r \geq 0$ be a public reserve price, $s \geq 0$ be a secret reserve price, $b \geq 0$ be a buy-it-now price, and N be a number of bidders. Then under assumptions 1, 2, 3, and 4, the following holds:*

a) Conditional on a bidder's valuation exceeding the public reserve price r, it is a weakly dominant strategy for the bidder to enter an auction if the lower bound of the support of a secret reserve price is the public reserve price r.

b) Equilibrium bidding strategies in an eBay auction with r,s,b, N bidders and sniping coincide with equilibrium bidding strategies in an eBay auction with r, without s, N+1 bidders, and sniping.

In the proof we concentrate on bidders' decision to participate only in the presence of a secret reserve price. For the discussion of bidders' strategies in the presence of a public reserve price, see Myerson (1981). For the discussion of bidders' strategies in the presence of a temporary BIN and a permanent BIN, see respectively Mathews (2003) and Hidwegi et. al. (2006).

Proof:

Under symmetry of distribution functions of bidders' valuations, let N be the number of bidders in an auction, $v_i \in [0, v]$ be a valuation of a representative bidder i, $b_i \in [0, v_i]$ be an equilibrium bid of bidder i, and assume that $v_i \geq r$. Since bidder i does not observe a secret reserve price, bidder i treats the secret reserve price s as a random variable with a support $[\underline{s}, \overline{s}]$. Under the assumption that the lower bound of the support is $\underline{s} = r$, costless bidding, and if bidder i follows an equilibrium bidding strategy in a second-price auction as in Myerson (1981), bidder i's expected payoff from entry is $(v_i - b_i)\text{Prob}(b_i \geq b_{j \neq i}, b_i \geq s, b_i \geq r) \geq 0$. If bidder i does not enter, then bidder i's payoff is zero. Hence, it is a weakly dominant strategy for bidder i to enter an auction if the lower bound of the support of the secret reserve price s is $\underline{s} = r$.

To see that part (b) is true, note that the formulation of the problem with a secret reserve price, a public reserve price and N bidders in Proposition 1 is identical to the formulation of a problem with a public reserve price and $N+1$ bidders, where the $(N+1)$st bidder's valuation is $s \in [\underline{s}, \overline{s}]$. Hence, bidder i's equilibrium bidding strategy in an ascending second-price auction with sniping, N bidders, a public reserve, and a secret reserve coincides with an equilibrium bidding strategy in an ascending second-price auction with sniping, $N+1$ bidders, a public reserve, and no secret reserve. ■

Given the equilibrium strategies of bidders in the presence of a secret reserve price, we can discuss the seller equilibrium behavior. We relax the standard assumption that the seller knows the distribution function of bidders' valuations and assume that the seller neither perfectly knows the distribution function of bidders' valuation nor knows its support.

Assumption 5. A seller does not have perfect knowledge of the distribution function of bidders' valuations and its support.

If the seller does not know the distribution function of bidders' valuations, then she cannot set an optimal public reserve price and an optimal BIN. Hence, it is natural to expect that the seller may want to gather some information about the distribution function of bidders' valuations, if the cost of such information is sufficiently low and the future discount factor is one. The natural place to gather such information is an ascending auction itself. An ascending-price auction is the best source of information about the distribution function of bidders' valuations as opposed to a descending-price auction or a posted price sale, because in an ascending-price auction unsuccessful bidders reveal the most accurate information about their valuations. The main drawback of an ascending auction is that the winning bidder may not reveal her valuation[19]. However, if the item does not sell and there is no winner, even in the presence of sniping a seller is able to gather the most accurate information about valuations of all participating bidders. Hence, if the cost of running an unsuccessful ascending-price auction is small and an uninformed seller cares only about learning the distribution function of bidders' valuations, then the seller has an incentive to run an unsuccessful ascending-price auction to learn the distribution function of bidders' valuations to form a belief about the distribution function of bidders' valuations.

The use of an ascending-price auction in this fashion relies on three key elements: the low cost of an unsuccessful auction in terms of a physical cost and foregone time, a sufficient entry in an auction, and the ability to run an auction without selling an item. All three elements are available to a seller on eBay. The cost of running an eBay auction varies across item categories. In our particular case, the cost of listing a tractor is \$20, the cost of using a public or a secret reserve price is \$5, and the cost of using a BIN never exceeds \$0.25. Hence, the total cost of unsuccessfully auctioning a tractor on eBay with a secret reserve price at most amounts to \$25.25[20].

The assumption that sellers who use auctions to learn distributions of valuations do not discount the future is based on the observation that we should observe a self-selection of patient sellers into such auctions, since rational sellers should expect a low probability of success in auctions with secret reserve prices. This claim

[19]It is a weakly dominant strategy to reveal valuation for the winner only in an ascending sealed-bid second-price (Vickrey) auction.

[20]The cost of an unsuccessful auction with a secret reserve price does not include a final value fee, which a seller has to pay if the item sells. For tractors sold in eBay auctions the final value fee is %1 of the sale price with a maximum of \$250.

is further reinforced by the empirical finding that about 63% of relists occur within 7 days after an unsuccessful listing. Given these considerations we are ready to state our next assumption.

Assumption 6. A seller has the future discount factor of 1 and faces negligible costs of running an unsuccessful ascending-price auction on eBay.

The second key element is that bidders enter an eBay auction with a secret reserve price and always bid according to equilibrium strategies of a single-stage ascending second-price auction with sniping. By part (a) of Proposition 1, it is a weakly dominant strategy to enter an auction with a secret reserve price, and by part (b) of Proposition 1, participating bidders use equilibrium bidding strategies of a single-stage ascending second-price auction with sniping.

Lastly, the ability of an uninformed seller to run an ascending-price auction without selling an item relies on the availability of a secret reserve price in eBay auctions. In particular, the ability to set an unmet secret reserve price allows an uninformed seller to run an ascending second-price auction without selling an item and without limiting entry. In the next proposition we state how a seller should set an equilibrium secret reserve price.

Proposition 2. *Let $[0, v]$ be an unknown support of the distribution function of bidders' valuations, $s \geq 0$ be a secret reserve price, and $b \geq 0$ be a BIN price. Then under assumptions 1, 2, 3, 4, 5, and 6, in equilibrium an uninformed seller, who does not discount the future, sets a secret reserve price equal to the maximum of the highest order statistic of the distribution function of bidders' valuations.*

Proof: First, assume that the seller knows the support of the distribution function $[0, v]$ but does not know the distribution function itself. Then, according to Mathews (2003), a risk-neutral seller should set a temporary BIN equal to v, while a risk-averse seller should set a temporary BIN to some value below v. Hidvegi et al. (2006) show that whether a seller and bidders are risk-neutral or not, a seller should set a permanent BIN above or equal to the expected highest valuation of a symmetric bidder.

Next, by Proposition 1, the most complete revelation of bidders' valuations in an ascending-price auction is achieved if the seller sets a secret reserve price s above or equal to v, $s \geq v$. Under the assumption that the distribution function of bidders' valuations is discrete, the seller obtains a strictly positive expected payoff by setting the secret reserve s equal to v. If the distribution function of bidders' valuations is continuous, then it is a weakly dominant strategy for an uninformed and a patient seller to set a secret reserve price s equal to v. There is an infinite number of mixed-strategy equilibria in which an uninformed and patient seller sets a secret reserve price s to v with a non-zero probability. There is no pure strategy equilibrium where the seller sets the secret reserve price s above v with probability 1, because in this case no bidder enters in equilibrium.

Since the seller does not know v, she should set a secret reserve price equal the sample analogue of v, which is the maximum of the highest order statistic. ■

We need to emphasize several important premises and consequences of Proposition 2. First of all, an uninformed seller should set a secret reserve price equal to the maximum of the highest order statistic only if the cost of relisting in terms of time and resources is zero. If the seller faces a sizeable cost of relisting, then there is the usual trade-off between acquisition of an additional information and foregoing profits from selling earlier. If this is the case, a seller is better off lowering a secret reserve price to capture expected profits from selling an item earlier.

Second, under assumptions of Proposition 2, an uninformed seller should set a secret reserve price to the maximum unsuccessful bid across all unsuccessful prior auctions. This result holds independently of whether bidders engage in sniping or not. Even if some bidders are not able to place their highest bids with a positive probability due to sniping, in the absence of any other information about the distribution function of bidders' valuations the highest unsuccessful bid across all unsuccessful auctions is the most accurate statistic of the upper bound of the support of valuations. Further, note that the maximum unsuccessful bid across all unsuccessful prior auctions is a random variable with a support $[r, v]$, where r is a public reserve price and v is the upper bound of the support of the distribution function of bidders' valuations. Hence, by part (a) Proposition 1 in equilibrium a seller should observe an non-zero entry in an eBay auction with a secret reserve price.

Third, the seller's choice of an optimal BIN depends on her attitude toward risk and whether a BIN is used together with a secret reserve price or not. If the BIN is used without a secret reserve price, then by rules of eBay the BIN is essentially a temporary posted price, and according to Mathews (2003), a risk-neutral seller should set the BIN equal to v, while a risk-averse seller should set the BIN equal to some value below v[21]. Under the assumption that the seller has no knowlededge of the upper bound of the support v, we assume that the seller should use the maximum of the highest order statistic as a proxy for v.

If the BIN is used together with a secret reserve price, then the BIN essentially turns into a permanent posted price, since by rules of eBay, whenever a BIN and a secret reserve price are used together, the BIN is active as long as the secret reserve price is not met. According to Hidvegi et. al. (2006), in this case if either the seller or buyers are risk-averse, the seller should set the BIN at least to the bidders' expected highest valuation.

Last, when a BIN and a secret reserve price are used together, there is an equilibrium in which a BIN and a secret reserve price are simultaneously set to the upper bound of the support of the distribution function of bidders' valuations or to its sample analogue. In this case bidders can infer the size of a secret reserve prices by observing the BIN. As a result, a secret reserve price is no longer a random variable from the standpoint of bidders and entry strategies of bidders should change. However, by rules of an eBay auction, a seller can change a secret

[21] Note that if the BIN is used without a secret reserve price, it disappears once the first bid is made.

reserve price at any time prior to the 12 hours of an auction closing time. This gives bidders an incentive to participate in an auction with a secret reserve price and a BIN even if their valuations are below the BIN. We test whether bidders infer sizes of secret reserve prices from BINs by estimating the effect of a simultaneous presence of a BIN and a secret reserve price on entry. We find that the simultaneous use of BINs and secret reserve prices has a small though statistically significant negative effect on entry[22].

4 Estimation Strategy and Identification Problems

To analyze seller behavior in eBay auctions, we construct determinants of public and secret reserve prices from distributions of bids, which sellers observe in eBay auctions. We assume that sellers form their beliefs about distribution functions of bidders' valuations by observing bids in eBay auctions of similar items or from unsuccessful eBay auctions for their own items[23]. To avoid making assumptions about sellers' beliefs of distribution functions of bidders' valuations, we use a non-parametric density estimator with a normal kernel to estimate distribution functions of bidders' valuations.

To resolve the issue with the lack of data given the non-parametric approach we make the following assumptions. First, we reduce the number of dimensions according to which bidders form their valuations of a tractor down to just three. These three dimensions are an age of a tractor, a horse power of a tractor, and a brand of a tractor. Second, we discretize each dimension to allow for a sufficient number of observations for a non-parametric estimation. In particular, we allow

[22]For the results of this test see Table A2 in the appendix.

[23]Einav et. al. (2012) make a similar argument - they argue that many sellers use online markets not only to sell their items per se, but also for experimenting with an optimal selling strategy by relisting their items multiple times with varying auction parameters or sale formats.

only three categories of age: 0-10 years old, 11-30 years old, and above 30 years old. Similarly, we allow only three categories of horse power: up to 20 HP, 21-40 HP, above 40 HP, and we allow only two categories of brand: whether a tractor is manufactured by John Deere or not. Thus, totally we have 18 categories of tractors.

The discretization of tractors in this way is driven by the types of tractors in each category. The division by age reflects whether a tractor on sale is of the current generation, the past two or three generations, or is an antique tractor.

The division of tractors by horse power captures the intended purpose of tractors. Tractors up to 20 HP are smaller lawn tractors such as mowers, while tractors with 21-40 HP are compact utility tractors used for everyday farming needs. Tractors with horse power above 40 HP are utility tractors used for heavier tasks such as cropping or construction.

The division of tractors by John Deere brand and non-John Deere brands is driven by the popularity of John Deere tractors among farmers and the large number of John Deere tractors in the data.

To obtain an untruncated distribution of valuations we use bids only from auctions with a public reserve price of at most $100[24]. In addition, we ignore multi-stage auctions where there are two or more identical bidders participating in more than one stage of a multi-stage auction. We do so to remove the strategic component from the analysis, since repeated bidders and a seller may engage in a strategic dynamic behavior[25]. To account for the censoring of the highest valuations in successful auctions we run a two-step procedure[26]. First, we

[24] We do not consider auctions without high public reserve prices, because sellers most likely ignore such auctions as an accurate source of information about distribution functions of bidders' valuations given that such auctions truncate bids and act rather as posted price sales.

[25] We recognize that by ignoring auctions with repeated bidders we may overlook an important component from the analysis. However, we leave the analysis of the repeated bidder and seller interaction for a follow up paper.

[26] Note that by proposition 1, in a successful auction with sniping the winning bidder never reveals her true valuation and bids the second highest valuation or some value below it, while all other bidders reveal their valuations or some values below them.

non-parametrically estimate distribution functions of bidders' valuations without accounting for the censoring problem. Next, we replace highest bids in successful auctions with expected valuations calculated by using non-parametrically estimated distribution functions from the first step[27]. Given the updated highest bids, we non-parametrically re-estimate distribution functions.

We use the resulting estimated distribution functions for each tractor category to calculate determinants of auction parameters set by sellers in eBay auctions. It is important to emphasize that these estimated distribution functions are not unbiased estimates of true distributions of valuations which bidders have, and the question of whether sellers are able to control for the bias at least partially is left open for a future study. Nevertheless, in the appendix we test for major determinants of the bias. We find that the presence of secret reserve prices reduces entry of bidders and particularly of low-valuation bidders. As a consequence, there is an upward bias due to under-representation of valuations in the left tail of distributions of valuations, which sellers should observe in auctions with secret reserve prices. Second, we find that bidders' valuations are not independently distributed and contain a commonly distributed component. As a result, the highest bids in eBay auctions overstate their underlying valuations, which also contributes to the upward bias. Lastly, due to sniping, realized unsuccessful bids do not always reflect underlying valuations, since some bidders may not be able to place their bids before auction closing times. Hence, the presence of sniping results in the downward bias. We leave the question of which form of the bias dominates the distributions of valuations obtained by sellers from eBay auctions for a future work.

[27] We use the standard expression for an expectation of a truncated variable to calculate an expected highest valuation in a successful auction: $E[v] = Prob(v \leq b)*b+Prob(v > b)*E[v|v > b]$, where $E[v]$ is the expected highest valuation in a successful auction, b is the highest bid in a successful auction, $Prob(v \leq b)$ is the non-parametrically estimated distribution function of bidders' valuations, and $E[v|v > b]$ is the expectation of the highest valuation conditional on exceeding the highest bid in a successful auction.

5 Determinants of Auction Parameters

5.1 Determinants of public reserve prices

Myerson (1981) and Riley and Samuelson (1981) show that when bidders' valuations are independently distributed, it is always optimal for a seller with a positive valuation of an item on sale to set a non-zero public reserve price and exclude some low-valuation bidders. The authors further show that the optimal public reserve price should satisfy the following closed form expression: $r_i^* = x_i + \frac{1}{\lambda(r_i^*)}$, where r_i^* is the optimal public reserve price for an item i, x_i is the seller's valuation of the item on sale, and $\lambda(r_i^*)$ is the hazard rate associated with the distribution function of bidders' valuations and evaluated at the seller's optimal public reserve price. When bidders' valuations have a commonly distributed component, Cai, Riley and Ye (2007) show that public reserve prices can be used as a credible signal of the value of the item on sale.

The test of an inter-dependent versus independent distribution of valuations in the appendix suggests that bidders' valuations have a commonly distributed component. Hence, public reserve prices in auctions for used tractors are used both to exclude low-valuation bidders and to credibly signal values of tractors on sale.

If public reserve prices are used together with secret reserve prices and if uninformed sellers use secret reserve prices to learn distribution functions of bidders' valuations, it is hard to justify the use of public reserve prices as a screening device. However, if public reserve prices can serve as a credible signal of quality, then uninformed sellers can obtain a more accurate right tail of distribution of valuations by setting high public reserve prices and attracting higher-valuation bidders. Next, we explicitly state the hypotheses about functions of public reserve prices in eBay auctions for used tractors.

Hypothesis 1 (public): *If bidders' valuations have both an independently and a commonly distributed component and if sellers use eBay auctions to sell, then* *the sellers employ public reserve prices to screen out low valuation bidders and to credibly signal the value of their items on sale.*

Hypothesis 2 (public): *If bidders' valuations have both an independently and a commonly distributed component and if sellers use auctions to gather information rather than sell, then the sellers use public reserve prices exclusively to credibly signal the value of their items on sale.*

In our test of functions of a public reserve price we include determinants accounting for the screening effect and the signalling effect. To account for the screening effect we include one over the hazard rate estimated at the public reserve price, where the hazard rate is calculated from non-parametrically estimated distribution functions of bidders' valuations. To account for the signalling effect we include the average bid among the determinants of a public reserve price. Since there is a reasonable suspicion that the public reserve price may lower the average bid by limiting entry of low valuation bidders, and our estimates may suffer from a simultaneity bias, in our estimates we use the maximum bid as an instrument for the average bid. In addition, we include tractor characteristics such as tractor age, horse power and brand to capture sellers' valuations of tractors on sale.

To test whether public reserve prices are used exclusively for signalling when sellers use public and secret reserve price at the same time, we interact signalling and screening determinants of public reserve prices with a dummy for a secret reserve price. The results of our test are presented in Table 7.

In both models in Table 7, when secret reserve prices are not used, the determinants of a public reserve price accounting for screening and signalling effects are positive and statistically significant. However, when public reserve prices are used together with secret reserve prices, the coefficient for one over the hazard rate becomes negative, violating the screening assumption, while the determinant for a signalling effect in the fixed effects IV model more than doubles in size in comparison to the case without a secret reserve price. We interpret this result as an evidence that when sellers use auctions to gather information about bidders' valuations and employ secret reserve prices, the sellers use public reserve prices exclusively to signal the quality of their tractors on sale and to obtain valuations of their tractors primarily from higher-valuation bidders. However, when sellers use auctions exclusively to sell their items and do not empoy secret reserve prices, the sellers use public reserve prices both to screen out low valuation bidders and to signal the quality of their items.

Table 7. Determinants of public reserve prices.

Variable	IV		Fixed Effects IV	
$\frac{1}{\text{Hazard rate}_j(\text{r}_i)}$*(1-Secret res. dummy$_i$)	0.172**	(0.088)	0.283***	(0.031)
Average bid$_i$*(1-Secret res. dummy$_i$)	0.257**	(0.107)	0.131***	(0.022)
$\frac{1}{\text{Hazard rate}_j(\text{r}_i)}$*Secret res. dummy$_i$	-0.432***	(0.039)	-0.333***	(0.022)
Average bid$_i$*Secret res. dummy$_i$	0.563***	(0.028)	0.501***	(0.014)
Engine HP$_i$	15.834***	(2.446)	17.821***	(1.903)
Age$_i$	-60.628***	(7.915)	-58.539***	(4.436)
John Deere$_i$	918.560***	(134.717)	890.793***	(122.691)
Constant	2516.408***	(337.854)	2480.307***	(126.929)
Overall R^2	0.644		0.609	
Number of groups			4499	
Number of observations	8268		8268	

Notes: The dependent variable is the public reserve price above $100;
Max. bid$_i$*(1-Secr. res. dummy$_i$) and Max. bid$_i$*Secr. res. dummy$_i$ are used as instruments for Aver. bid$_i$*(1-Secr. res. dummy$_i$) and Aver. bid$_i$*Secr. res. dummy$_i$;
Observations are clustered by seller ID in FE model; Robust standard errors in parentheses; ***, **-statistical significance at 1%, 5%;

5.2 Determinants of secret reserve prices

To test the purpose of a secret reserve price, we suggest three possible hypotheses. Under the first hypothesis, a secret reserve price is an alternative to a public reserve price and is set to screen out bidders with valuations below the threshold level indicated by the size of a secret reserve price[28]. Under this hypothesis a secret reserve price is a function of one over the hazard rate estimated at the secret reserve price and a seller's valuation of the item on sale[29]. The main argument against this hypothesis is the prevalence of a simultaneous use of public and secret reserve prices in eBay auctions.

Hypothesis 1 (secret): *If bidders' valuations have both an independently and a commonly distributed component, sellers use secret reserve prices exclusively to screen out low valuation bidders.*

Under the second hypothesis a very patient informed seller uses a secret reserve price to sell at the highest possible price. In this case an informed seller sets a secret reserve price at the known upper bound of the support of the distribution function hoping that the highest-valuation bidder enters an auction. Under this hypothesis, an informed patient seller uses an auction as a posted price sale where the sale price is revealed through bidding.

There are two main arguments against this hypothesis. The first one is the prevalence and stability of the share of unsold items in eBay auctions. If sellers had a perfect knowledge of the distribution function of bidders' valuations and were sufficiently patient, then we would observe an increase in the share of sold tractors with the number of relistings. However, we see that the share of sold tractors of about 30% is relatively stable across relistings. The second argument against this hypothesis is the availability of a cheaper posted price sale format on

[28] The view of a secret reserve price as an alternative to a public reserve prices has received the most attention in the literature (see, for example, Katkar and Reiley (2006), Lucking-Reiley et. al. (2007) for empirical tests).

[29] Under this hypothesis a secret reserve price cannot serve the signalling function of a public reserve price, because bidders do not observe it.

eBay, which an informed seller should prefer to a more costly auction format with a secret reserve price[30].

Hypothesis 2 (secret): *If bidders' valuations have both an independently and a commonly distributed component, informed and patient sellers use secret reserve prices to sell at the highest possible price.*

Under the third hypothesis, a patient uninformed seller uses a secret reserve price to learn parameters of the distribution functions of bidders' valuations and to sell an item at the highest possible price given the knowledge of these parameters. In this case the equilibrium secret reserve price is set at the seller's belief about the upper bound of the support of the distribution function of bidders' valuations. As the uninformed patient seller observes unsuccessful bids, she updates her belief about the support and other parameters of the distribution function and recalculates the secret reserve price. By Proposition 2 the seller's strategy under this hypothesis is to set the secret reserve price at the maximum of the observed highest order statistic or at the highest bid across all unsuccessful previous auctions.

Hypothesis 3 (secret): *If bidders' valuations have both an independently and a commonly distributed component, uninformed and patient sellers use secret reserve prices to learn parameters of the distribution function bidders' valuations.*

First, we test hypothesis 1 against hypotheses 2 and 3. Under all three hypotheses, an introduction of a secret reserve price should lower the likelihood of a successful sale. However, under hypothesis 1, the negative effect of a secret reserve price on the probability of sale should be of a comparable size as the effect of a public reserve price. The regression results in Table 8 show that a dollar increase in the public reserve price has a small negative and statistically insignificant effect on the probability of sale, while a dollar increase in the secret reserve price has a much stronger negative and statistically significant effect on the probability of

[30] The cost of an unsuccessful posted price sale on eBay is $0.50 as opposed to a US 25.25 dollar cost of running an unsuccessful auction with a secret reserve price and a BIN. Of course, this argument holds under the assumption that the pool of buyers in auctions and posted price sales is the same.

sale.

In the next test we explicitly test possible determinants of secret reserve prices. In our sample we have data on secret reserve prices only in unsuccessful auctions. In successful auctions the data on secret reserve prices is replaced by sale prices. Further, the secret reserve prices in our data are the ones that were active in the last 12 hours of auction closing times. Hence, the secret reserve prices in our data do not necessarily equal the secret reserve prices set initially at auction start times[31]. To account for the missing secret reserve prices in successful auctions, we use average secret reserve prices calculated for each of the 18 categories of tractors.

Table 8. Determinants of a successful sale (pooled data from all auctions).

Variable	Probit Marginal Effect	Standard Error
Secret reserve$_i$	-0.00005***	0.000002
Public reserve$_i$	-0.000001	0.000005
BIN price$_i$	-0.000001	0.000001
John Deere$_i$	0.019	0.015
Age$_i$	0.0007	0.0005
Engine HP$_i$	-0.0004*	0.0002
Seller feedback score$_i$	0.00005*	0.00002
Number of bidders$_i$	0.021***	0.002
Average bid$_i$	0.00003***	0.00001
Number of observations		13051

Notes: The dependent variable is Sale (Y/N); Standard errors are clustered by Seller ID; ***,*-statistical significance at 1%, 10%;

To construct a sample for the test, we pool together tractors from all relistings. We ignore single-stage auctions with secret reserve prices, because some single-

[31] Note that by eBay auction rules a seller can update a secret reserve price at any time before the last 12 hours of an auction closing time.

stage auctions are relistings of tractors listed before the data was collected and whose determinants we do not observe.

Lastly, since our test of determinants of secret reserve prices is based on the assumption that sellers do not have knowledge about distribution functions of bidders' valuations, it is necessary to specify initial beliefs of sellers about these distribution functions. We assume the following process of formation of beliefs. Initially, uninformed patient sellers set their secret reserve prices by observing highest bids in auctions of similar items. Once they are able to run their own unsuccessful auctions and gather information about valuations of their own tractors, they update their secret reserve prices by incorporating the data from unsuccessful auctions. To account for the initial beliefs of sellers we average highest bids from auctions in each of the 18 categories of tractors and include these values among determinants of secret reserve prices.

In our test of determinants of secret reserve prices we include the following regressors: $\{\widehat{b}_{i,t-k}, \overline{\overline{b}}_j, \frac{1}{\lambda_j(s_{i,t})}, x_i, m_i\}$, where $\widehat{b}_{i,t-k}$ is the highest bid from all previous unsuccessful auctions for tractor i with the highest bid made during an unsuccessful auction listing $k \in \{1,2,3,4\}$, $\overline{\overline{b}}_j$ is the average highest bid for a tractor in category j, $\frac{1}{\lambda_j(s_{i,t})}$ is one over the estimated hazard rate for a tractor in category j and calculated at the secret reserve price $s_{i,t}$, x_i is the set of parameters accounting for tractor heterogeneity such as tractor's age, horse power and brand, and $m_i \in \{1,2,3,4\}$ is the number of times tractor i was previously listed.

The results of both models in Table 9 allow us to reject hypothesis 1 that sellers use secret reserve prices to screen out low valuation bidders. In the OLS model, the screening effect represented by one over the hazard rate is negative and statistically insignificant, while in the fixed effects model the screening effect is statistically significant but negative, which violates the screening assumption.

Further, in both models by far the most significant determinants of secret reserve prices are the highest bids from previous unsuccessful auctions, which capture sellers' acquisition of information from running unsuccessful eBay auctions for their individual items, and average highest bids for tractors in a similar category, which capture sellers' acquisition of information from observing eBay auctions for similar items.

To see that secret reserve prices are used by uninformed or partly informed sellers for the purposes of learning, note that in Table 9 the coefficient for the number of relistings is statistically significant and negative. In other words, there is a downward adjustment in secret reserve prices with an additional relisting. This contradicts hypothesis 2 that informed sellers use secret reserve prices solely to sell at the highest possible prices, since such informed sellers should not adjust secret reserve prices based on information acquired from running additional unsuccessful relistings.

Table 9. Determinants of secret reserve prices.

Variable	OLS		Fixed Effects	
Highest $\text{Bid}_{i,t-k}$	0.589***	(0.058)	0.521***	(0.133)
Average Highest Bid_j	0.553***	(0.119)	0.576***	(0.197)
$\frac{1}{\text{Hazard rate}_j(s_{i,t})}$	-0.145	(0.136)	-0.463**	(0.204)
Age_i	-16.007**	(7.358)	-78.649**	(32.886)
Engine HP_i	17.961***	(6.754)	42.246*	(22.774)
John Deere_i	1300.824***	(342.253)	2297.449**	(904.939)
Number of relistings_i	-380.236**	(154.887)	-349.718**	(158.671)
Constant	-986.682**	(477.178)	374.094	(750.065)
Overall R^2	0.759		0.738	
Number of groups			489	
Number of observations	782		782	

Notes: The dependent variable is secret reserve price $s_{i,t}$; Robust standard errors in parentheses; Observations and errors are clustered by seller ID in FE model; ***,**,* - statistical significance at 1%, 5%, 10%;

To explicitly test hypothesis 2 that informed sellers use secret reserve prices exclusively to sell at the highest possible prices against hypothesis 3 that uninformed sellers use secret reserve prices to learn, we regress current secret reserve prices on current and past highest bids. If sellers are informed, then the variation in the current secret reserve price should be better explained by the current highest bid then by the past highest bid. However, if sellers are uninformed, then the variation in the current secret reserve price should be better explained by the past highest bid.

Table 10. Informed versus uninformed sellers.

Variable	OLS Coefficient		OLS Coefficient	
Highest Bid$_{i,t}$	0.013	(0.008)	0.017	(0.019)
Highest Bid$_{i,t-k}$	0.587***	(0.057)		
Average Highest Bid$_j$	0.445***	(0.092)	1.001***	(0.080)
Number of relistings$_i$	-369.971**	(153.222)	37.384	(222.001)
Age$_i$	-12.466*	(7.015)	-68.589***	(8.793)
Engine HP$_i$	15.409**	(6.733)	51.711***	(6.862)
John Deere$_i$	1164.441***	(352.464)	2534.878***	(377.694)
Constant	-814.972*	(477.040)	-1419.283**	(598.874)
R^2	0.761		0.591	
Number of observations	782		782	

Notes: The dependent variable is secret reserve price$_{i,t}$; Robust standard errors in parentheses; ***,**,*-statistical significance at 1%, 5%, 10%;

The variation in the current secret reserve price in Table 10 is better explained by the past highest bid than by the current highest bid. In fact, in both models in Table 10, the current highest bid is statistically insignificant. Hence, there is more evidence to accept hypothesis 3 than hypothesis 2. However, since there is a strong correlation between highest unsuccessful bids across different periods, it is hard to make a definite claim about exact determinants of a secret reserve price. In addition, since initially uninformed sellers become more informed once they run unsuccessful auctions or observe bidding behavior in auctions for similar items, it is virtually impossible to separate the behavior of a completely informed seller from the behavior of a partially informed seller by observing only the variation in secret reserve prices. Guided by these considerations, we do not make a definite claim that secret reserve prices are used exclusively to learn and that sellers who use secret reserve prices are completely uninformed.

5.3 Determinants of BINs

To test determinants of BINs, we invoke theoretical results of Mathews (2003) and Hidvegi et. al. (2006). The size of a BIN in an eBay auction depends on whether the BIN is permanent or temporary. According to Mathews (2003), a risk-neutral seller should set a temporary BIN equal to the upper bound of the support of the distribution function of bidders' valuations, while a risk-averse seller should set a temporary BIN to some value below the upper bound. According to Hidvegi et. al. (2006), a risk-neutral and a risk-averse seller should set a permanent BIN at least to the expected highest valuation. Since the upper bound of the support of valuations is strictly above the expected highest valuation, we conclude that a permanent or a temporary BIN in an eBay auction should be bounded above by the upper bound of the support of valuations and bounded below by the expected highest valuation.

Under the assumption that a seller does not have perfect information about the distribution functions of bidders' valuations, and therefore, cannot set the temporary BIN to the upper bound of the support of bidders' valuations, we assume that such an uninformed seller uses the highest bid from previous unsuccessful auctions to determine the upper bound. To account for the expected highest valuation, we use an average highest bid for each of the 18 tractor categories. Further, we interact these determinants with secret reserve prices to see whether sellers set BINs differently in auctions with secret reserve prices.

To test the determinants of BINs, we pool together auctions with 2 or more listings. We ignore single-listing auctions, because for these auctions we do not observe the highest unsuccessful bids from previous listings. We presents the results of our test in Table 11.

Table 11. Determinants of BINs.

Variable	OLS	Fixed Effects
Highest $\text{Bid}_{i,t-k}$	1.013*** (0.032)	0.715** (0.358)
Average Highest Bid_j	0.151*** (0.058)	0.256 (0.389)
Highest $\text{Bid}_{i,t-k}$*Secret reserve_i	0.123*** (0.046)	0.368 (0.345)
Average Highest Bid_j*Secret reserve_i	-0.042 (0.059)	-0.232 (0.346)
Constant	185.674 (266.492)	1183.553 (822.096)
Overall R^2	0.936	0.933
Number of groups		219
Number of observations	354	354

Notes: The dependent variable is $\text{BIN}_{i,t}$; Robust standard errors in parentheses; Observations and errors are clustered by seller ID in FE model; ***,** - statistical significance at 1%, 5%;

The results in Table 11 show that independently of whether a secret reserve price is used or not, by far the most important determinant of a BIN in terms of size and significance is the highest bid from previous unsuccessful auctions. Hence, we conclude that in eBay auctions sellers tend to set BINs closer to their beliefs about upper bounds of bidders' valuations rather than to their beliefs about expected highest valuations.

It is important to emphasize that in our tests of determinants of auction parameters the determinants of BINs and secret reserve prices are identical. We show that the variation in both BINs and secret reserve prices is largely explained by the variation in previous highest bids and expected highest valuations. The main difference is in the weights of these two determinants. While the variation in secret reserve prices is explained almost equally by the variation in previous highest bids and expected highest valuations, the variation in BINs is almost exclusively determined by the variation in the previous highest bids only. Given these findings, we conclude that the data on BINs can serve as a relatively accurate approximation of the data on secret reserve price, whenever the data on secret reserve prices is not available or impossible to obtain.

6 Determinants of Sale Prices

To test the effect of secret reserve prices on sale prices, we pool together tractors which were sold after up to 5 relistings on eBay. We test sale effects of reserve prices only for those tractors which were sold on eBay and discard those tractors which were listed but never sold on eBay. We do so, because we do not observe final selling prices of unsold tractors as sellers of such tractors might be using other selling formats or other selling platforms after delisting from eBay. We also omit observations with more than 5 relistings, because sellers of such tractors probably use the eBay platform for purposes other than selling or learning bidders' valuations [32].

To test sale effects of a secret reserve price, we decompose the total sale effect of a secret reserve price into an information-acquisition effect and any other possible effect. To capture the information-acquisition effect of a secret reserve price, we include among the regressors the total number of times a secret reserve price has been used. If secret reserve prices are used to acquire information about bidders' valuations, then the frequency of the use of a secret reserve price should serve as an indication of the degree of seller knowledge about the distribution function of bidders' valuations. To capture any other possible sale price effect of a secret reserve price, we include a secret reserve price at the time of a successful auction listing. In addition, among the determinants of sale prices we include the total number of relistings, a public reserve price at sale, a BIN at sale, and tractors

[32] We find that several online stores use eBay as an advertising outlet by constantly relisting an item and providing links to their websites in item description areas.

characteristics such as age, horse power and brand categories.

Table 12 shows that the non-information-acquisition effect of a secret reserve prices, indicated by the coefficient of a secret reserve price at sale, is positive, statistically significant, and comparable in scale to positive sale effects of a public reserve price and a BIN. The information-acquisition effect of a secret reserve price, indicated by the number of times a secret reserve price has been used, also has a sizeable positive and a statistically significant sale effect. Hence, contrary to previous empirical studies of secret reserve prices, we provide statistical evidence that the use of a secret reserve price has a strong positive effect on a sale price[33].

We believe that previous studies have different results, because they test the impact of secret reserve prices under the assumption that sellers do not relist items after an unsuccessful sale. Since the presence of a secret reserve price is highly correlated with the probability of relisting, in the absence of an explicit control for relistings the positive impact of a secret reserve price on a sale price is neutralized by the negative effect of relistings. In fact, if we sum up coefficients of the number of times a secret reserve price has been used and the number of relistings in Table 12, the resulting small positive difference becomes statistically insignificant[34].

To make sure that there is no endogeneity problem and the use of secret reserve prices is uncorrelated with an unobserved heterogeneity in the quality of tractors, we regress OLS residuals from Table 12 on the set of regressors in Table 12. We find that residuals are not correlated with any of the regressors[35].

[33] See, for example, Bajari and Hortacsu (2003) and Katkar and Reiley (2006).

[34] The difference between these two coefficients in OLS model has an F-statistic of 0.55 and a p-value of 0.459. In FE model, the difference has an F-statistic of 0.39 and a p-value of 0.534.

[35] We run an additional endogeneity test by regressing a secret reserve dummy on tractor characteristics such as tractor age, horse power, engine hours and eleven brand categories. We find that tractor age and horse power are the only statistically significant determinants of the use of a secret reserve price. The rest of tractor characteristics are not statistically significant suggesting that there is no unobserved tractor heterogeneity correlated with the use of secret reserve prices.

Table 12. Determinants of sale prices.

Variable	OLS	Fixed Effects
[Number of times a secret reserve price is used]$_i$	0.120*** (0.022)	0.103*** (0.034)
[Number of relistings]$_i$	-0.106*** (0.029)	-0.087** (0.041)
[Secret reserve at sale]$_i$	0.00002*** (0.000003)	0.00001*** (0.000005)
[Public reserve at sale]$_i$	0.00005*** (0.000002)	0.00005*** (0.000007)
[BIN at sale]$_i$	0.00002*** (0.000001)	0.00002*** (0.000004)
Age$_i$	-0.027*** (0.0007)	-0.025*** (0.002)
Engine HP$_i$	0.008*** (0.0003)	0.009*** (0.001)
Constant	8.744*** (0.040)	8.629*** (0.076)
Overall R^2	0.509	0.502
Number of groups		3942
Number of observations	8316	8316

Notes: The dependent variable is log(sale price$_i$); Robust standard errors in parantheses; Observations and errors are clustered by seller ID in FE model; ***-statistical significance at 1%; We do not report tractor brand dummies; Tractor brand dummies are jointly statistically significant at 1% with p-value<0.0001;

To get a better idea about the impact of secret reserve prices on sale prices, we run a simulation exercise where we implement different strategies of an uninformed seller. To introduce dynamics into seller actions, we impose a cost of an additional relisting. We introduce a relisting cost, because by Proposition 2, an uninformed seller who does not face any costs of relisting and cares only about learning the distribution function of bidders' valuations does not have an incentive to adjust a secret reserve price. In the presence of relisting costs, the seller has to balance the expected payoff from keeping a secret reserve high and learning bidders' valuations against the expected payoff from selling an item earlier and avoiding the relisting costs.

In our simulation exercise we impose three kinds of relisting costs. The first cost is the listing fee of $25. The second one is a future discount factor of 0.99. The third cost is that the seller's payoff becomes zero if she does not sell an item after the fifth relisting. To make our simulation exercise specific, we assume that a seller wants to sell a John Deere tractor with a 95 HP engine and which is about 37 years old. There are 25 of these tractors with a public reserve price less than US 100 dollars. We use 104 bids from these 25 tractors to non-parametrically estimate a cumulative distribution function of bidders' valuations and its associated inverse cumulative distribution function. We draw five valuations from the estimated inverse cdf and calculate the seller's revenue in an ascending second-price auction under different regimes.

We assume that the seller is uninformed about the estimated distribution function and only knows the maximum sale price of John Deere tractors with ages above 30 and horse power above 40. The maximum sale price of these tractors in our sample is $18988. Under the first regime, the seller sells an item in an ascending second-price auction without a public or a secret reserve price. As a result, she does not need to invoke her beliefs about the parameters of the distribution function and always sells at the second highest valuation at the first listing.

Under the second regime, the seller uses a secret reserve price. Since the seller is uninformed, it is optimal to set a secret reserve price equal to the expected maximum of the highest order statistic. In our case it is $18988 - the maximum sale price of John Deere tractors with ages above 30 and HP above 40. Once the seller is able to observe the highest unsuccessful bid from the an auction for her specific John Deere tractor, the seller can adjust the secret reserve price accordingly. We assume that the seller follows a simple updating rule of the following form, $S_t = \alpha * S_{t-1} + (1-\alpha) * Y^1_{t-1}$, where S_t is a secret reserve at time t, S_{t-1} is the previous period secret reserve, Y_{t-1}^1 is the highest unsuccessful bid in the previous auction stage, and α, $\alpha \in [0,1]$, is the weight the seller attaches to the previous period secret reserve price.

After the seller sets the initial secret reserve price, we calculate the seller revenue in an ascending second-price auction given the listing cost of \$25 and the draw of 5 independent valuations from the estimated distribution of valuations for the 37 year old John Deere tractors with 95 HP. If the highest valuation is below the initial secret reserve price, we proceed to the next stage, where we draw another set of 5 valuations, subtract \$25 listing fee, and discount the seller revenue at the rate of 0.99. We continue drawing new sets of 5 valuations, charging \$25 listing fee, and discounting seller payoff if there is no sale up to 5 times. If there is no sale at the fifth draw, the seller obtains zero payoff. We run 10000 iterations to obtain an average seller revenue for 100 weights α, where α ranges from 0 to 1 with an increment of 0.01. The simulation results are presented in Table 13.

Table 13. Simulated mean price of a 37 year old John Deere tractor with 95 HP.

Auction Format and Data Type	Mean Price in dollars	Standard Deviation	Mean Number of Relistings
Actual data	7947.1	3357.7	1.08
Simulated without Public or Secret Reserve	8333.3	2421.6	1.00
Simulated with Secret Reserve and $\alpha = 0$	9881.4	1795.5	2.70
Simulated with Secret Reserve and $\alpha = 0.15$	9979.5	2141.8	2.90
Simulated with Secret Reserve and $\alpha = 0.5$	8311.2	4773.5	4.50
Simulated with Secret Reserve and $\alpha = 0.75$	1361.3	3998.5	>5
Simulated with Secret Reserve and $\alpha = 1$	-125.0	0.0	>5

The results in the table show that depending on the updating rule, the introduction of a secret reserve price can raise the simulated mean price quite significantly, however, at the cost of an additional relisting. For example, the simulated mean seller revenue with $\alpha = 0$, where the seller adopts the highest unsuccessful bid as a secret reserve price and completely ignores the prior secret reserve price, is \$1547.7 more than the benchmark case without a public or a secret reserve price.

In addition, note that when $\alpha = 0$ and the seller sets the secret reserve price equal to the highest unsuccessful bid from the first stage, the seller does not need to relist the unsuccessfully sold item in an auction in the second stage. Instead, the seller can switch to a posted price sale with a price equal to the highest bid from the unsuccessful first-stage auction. Under this strategy that the seller sets a secret reserve price at the maximum of the highest order statistic, the seller's revenue from a posted price sale is exactly equal to the seller revenue from a repeated auction with a secret reserve price and $\alpha = 0$. Hence, if sellers in general have weight $\alpha = 0$ and the cost of a posted price sale is lower than that of an auction, it is optimal for sellers to switch to a posted price sale after running just one unsuccessful ascending second-price auction[36]. This can explain a large number of delisted tractors after one or two unsuccessful auction rounds in the data.

The results in our analysis crucially depend on the assumption that entry is exogenous and that bidders' entry decisions do not depend on the presence of a secret reserve price. However, since the higher revenue in an auction with a secret reserve price is driven by the ability of a seller to sell to the highest valuation bidder at the highest possible price and entry decisions of highest valuation bidders are least affected by the presence of a secret reserve price, we expect that the lack of exogeneity of the number of bidders in auctions with secret reserve prices should not bias our results too much.

Another critical assumption is that we are assuming that bidders do not act strategically in repeated auctions and do not shade their bids. This assumption is satisfied for large sale platforms such as eBay, where many sales are conducted

[36] Recall that eBay offers a posted price sale format in addition to an auction-type sale format.

at the same time and where it is difficult to coordinate participation in repeated auctions[37]. In live auctions, or in electronic auctions with costly or restricted entry, the assumption of non-strategic bidders is harder to justify. However, if it is the case that bidders are strategic, then we should expect the convergence of sale prices in repeated auctions to sale prices in static auctions as was shown by McAffee and Vincent (1997).

Nevertheless, the results of the simulation exercise show that the use of a secret reserve price can potentially raise sale prices if sellers are sufficiently patient and do not face high costs of relisting, while bidders bid according to strategies of a single-stage ascending second-price auction and do not take into account the possibility of a future resale. The cost of this increase in seller revenue is the necessity to run at least one unsuccessful ascending second-price auction.

7 Conclusion

In this paper we attempt to explain several stylized facts found in the eBay data such as the widespread use of secret reserve price, the presence of relisted items and a large number of unsuccessful sales. We believe that sellers relist their items because they want to learn bidders' valuations. We also believe that there are many unsuccessful auctions on eBay, because uninformed sellers switch to other selling formats after acquiring enough information by running unsuccessful auctions[38]. We cannot demonstrate this argument empirically, because to test this argument, we need data from alternative sales venues used by sellers who delist their items from eBay or about eBay posted price sales.

In the paper we give an explanation to why sellers may use a secret reserve price. We argue that a secret reserve price serves as a cheap and an efficient tool

[37] We do not consider sellers' response to bidders' strategic behavior in repeated auctions, because in eBay auction sellers cannot differentiate bidders' identities and, therefore, most likely treat all bidders as non-strategic.

[38] In particular, sellers can switch to posted prices sales on eBay after running unsuccessful eBay auction sales.

for uninformed patient sellers to run unsuccessful auctions to learn parameters of the distribution function of bidders' valuations. Our findings depends on two key assumptions: that sellers' do not know parameters of the distribution function of bidders' valuations and that sellers can observe truthful revelation of bidders' valuations in unsuccessful ascending second-price auctions.

There are two directions of future empirical and theoretical research. The first direction is related to the analysis of seller behavior under non-commitment, or in the presence of strategic bidders, who anticipate that unsold items will be relisted later and who take this fact into consideration. The empirical analysis of such a model involves a structural estimation of a dynamic game with evolving state variables.

The second direction of future research involves an analysis of hybrid sales, where an uninformed seller makes a choice for an optimal selling mechanism as she acquires more and more information about the distribution function of bidders' valuations. The analysis of such hybrid sales requires the data on different selling formats and the degree of seller informativeness.

8 References:

1. Bajari, P. and Hortacsu, A. (2003). "The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions." *Rand Journal of Economics*, Vol. 34(2), pp. 329–355.

2. Bergemann, D. and Said M. (2011). "Dynamic auctions," in *Wiley Encyclopedia of Operations Research and Management Science*, edited by J.J. Cochran. Wiley: New York, 1511–1522;

3. Cai, H., Riley, J., and Ye, L. (2007). "Reserve price signaling," *Journal of Economic Theory*, 135 (1), 253-268;

4. Carare, O. (2012). "Reserve prices in repeated auctions," *Review of Industrial Organization*, 40, 225–247;

5. Einav, E., Kuchler, T., Levin, J. and Sundaresan, N. (2012). "Learning from Seller Experiments in Online Markets," *working paper*;

6. Hasker, K. and Sickles, R. (2010). "eBay in the economic literature: analysis of an auction marketplace," *Review of Industrial Organization*, Vol. 37, 3–42;

7. Hidvegi, Z., Wang, W., and Whinston, A. (2006). "Buy-price English auction," *Journal of Economic Theory*, Vol. 129(1), 31-56;

8. Kagel, J. and Levin D. (2002). *Common value auctions and the winner's curse*, (Princeton University Press);

9. Katkar, R., and Reiley, D. (2006). "Public versus secret reserve prices in eBay auctions: results from a Pokemon field experiment," *Advances in Economic Analysis and Policy*, 6(2), Article 7;

10. Levin, D. and Smith, J. (1996). "Optimal reservation prices in auctions," *The Economic Journal*, 106(438), 1271-1283;

11. Lucking-Reiley, D., Bryan, D., Prasad, N., and Reeves, D. (2007). "Pennies from eBay: the determinants of price in online auctions," *Journal of Industrial Economics*, Vol. 55(2), 223–233;

12. Mathews, T., (2003). "A risk averse seller in a continuous time auction with a buyout option," *Brazilian Electronic Journal of Economics,* Vol. 5(2);

13. McAffee, P. and Vincent, D., (1997). "Sequentially Optimal Auctions," *Games and Economic Behavior,* 18, 246-276;

14. Milgrom, P., and Weber, R. J., (1982). "A theory of auctions and competitive bidding," *Econometrica*, 50(5), 1089-1122;

15. Myerson, R. B., (1981), "Optimal auction design," *Mathematics of Operations Research,* 6, 58–73;

16. Ockenfels, A. and Roth, A. (2006). "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction," *Games and Economic Behavior*, 55, 297–320;

17. Riley, J., and Samuelson,W., (1981), "Optimal auctions," *American Economic Review*, 71(3), 381–392;

18. Rosenkranz, S., and Schmitz, P. W., (2007). "Reserve prices in auctions as reference points," *The Economic Journal*, 117, 637–653;

19. Vincent, D., (1995), "Bidding off the wall: why reserve prices may be kept secret," *Journal of Economic Theory,* 65(2), 575–584;

9 Appendix: Tests of Directions of the Bias

9.1 Repeated versus non-repeated bidders

The most important assumption of Proposition 1 is that bidders follow a single-stage ascending second-price auction bidding strategy in the presence of a secret reserve price. Since a bidder may expect that an item listed in an auction with a secret reserve price is very likely to be relisted again, the highest unsuccessful bid of this bidder does not necessarily reveal her valuation. We call bidders who take into account the possibility of participation in a future listing due to a secret reserve price as repeated bidders. The change in the bidding behavior of such bidders is driven by a number of factors.

First of all, a repeated bidder realizes that by participating in an ascending second-price auction she reveals information about her valuation to other repeated bidders and that other repeated bidders may exploit this information in future stages. As a result, a repeated bidder in the current stage auction may not bid truthfully. We do not intend to go into more details about this aspect of repeated bidders; for more details, see Bergemann and Said (2011).

Second, a repeated bidder realizes that her bid reveals information about her valuation to the seller, and the seller can exploit this information in the future as well. Under the assumption that the seller has imperfect information about the distribution function of bidders' valuations, a repeated bidder acting strategically has incentives to under-bid to convince the seller of a different distribution function. By under-bidding and forcing the seller to have a wrong update about the distribution function of valuations, a repeated bidder may obtain substantial gains if the seller with a wrong updated belief chooses a selling format with a higher trade surplus allocated to bidders or changes auction parameters benefitting participating bidders.

Of course, if bidders in auctions with a secret reserve price do not intend to participate in future sales, then they have no incentive to act strategically and their bidding behavior in auctions with a secret reserve price coincides with their bidding behavior in single-stage auctions with a secret reserve price. To make sure that there is no strategic component in bidders' behavior in our sample and the identification assumption is satisfied, we remove all auctions were there are two or more identical bidders participating in two or more stages of a multi-stage auction. We present the distribution of auctions by the number of identical bidders across relistings in the next table.

The results in Table A1 show that there is a sizeable share of auctions in our data, where there are at least two identical bidders participating in at least two listings of the same tractor. To satisfy the identification assumption that there are no repeated bidders in our data, we omit observations from the bottom half of Table A1 when estimating distribution functions of bidders' valuations and when testing for determinants of secret and public reserve prices. In other words in our sample we retain only those auctions, where either all bidders in all listings are different or there is at most one identical bidder participating in at most two listings.

Table A1. Share of identical bidders in different stages in repeated auctions.

Type of Auctions	2-stage auctions	3-stage auctions	4-stage auctions	5-stage auctions
Auctions with at most 1 identical bidder in at most 2 listings (number)	709	76	20	5
Auctions with at most 1 identical bidder in at most 2 listings (percent)	69.04%	46.34%	38.46%	23.81%
Auctions with multiple identical bidders in at most 2 listings (number)	318	88	32	16
Auctions with multiple identical bidders in at most 2 listings (percent)	30.96%	53.66%	61.54%	76.19%

9.2 Entry rates

To test the assumption that entry decisions of bidders do not depend on the presence of a secret reserve price and/or BINs, we regress the number of bidders on a number of determinants including the presence of a secret reserve price, the size of a public reserve price, and the size of a BIN.

The results in Table A2 show that the presence of a secret reserve price has a strong negative impact on entry. For instance, the presence if a secret reserve price lowers the average number of bidders by more than 0.8 bidders. In addition, note that although the simultaneous presence of a BIN and a secret reserve price has a statistically significant negative impact on entry, the size of the impact is quite small. Hence, we find evidence to support our claim that the simultaneous use of a BIN and a secret reserve price does not excessively reduce entry.

Table A2. Poisson regression of the number of bidders.

Variable	Coefficient	Standard Error
Public reserve$_i$	-0.0001***	0.000001
Secret reserve dummy$_i$	-0.206***	0.004
BIN$_i$	0.000002***	0.0000007
BIN$_i$*Secret reserve dummy$_i$	-0.00001***	0.0000008
Public reserve$_i$*Secret reserve dummy$_i$	0.00008***	0.000001
John Deere$_i$	0.080***	0.003
Age$_i$	-0.008***	0.0001
Engine HP$_i$	0.002***	0.00005
ln(Seller feedback$_i$)	0.022***	0.0008
ln(Buyer feedback$_i$)	0.008***	0.0008
Constant	2.589***	0.006
Number of observations		43138

Dependent variable is the number of bidders; ***-statistical significance at 1%;

The negative impact of a secret reserve price on entry suggests that both the non-parametrically estimated distribution function of valuations and the empirical distribution of bids, which uninformed sellers use to form their beliefs about the distribution functions of valuations, are likely to over-represent higher valuations and under-represent lower valuations. It is natural to expect that a bidder with a lower valuation is less likely to enter an auction with a secret reserve price, since the probability of winning in such an auction given the equilibrium strategy of an uninformed seller is virtually zero. To test whether the presence of a secret reserve price biases the distribution of valuations toward the right tail, we regress average bids on a secret reserve price dummy. The results in Table A3 show that the presence of a secret reserve price raises the average bid. Hence, the distribution of bids observed by an uninformed seller and our estimated distribution functions are skewed to the right.

Table A3. OLS regression of the average bid on a secret reserve dummy.

Variable	Coefficient	Robust Standard Error
Number of bidders$_i$	26.349***	5.109
John Deere$_i$	1214.455***	91.461
Age$_i$	-81.396***	1.963
Engine HP$_i$	27.612***	1.186
Public reserve$_i$	0.785***	0.023
Secret reserve dummy$_i$	400.334***	64.826
BIN$_i$	.095***	0.012
Constant	2783.676***	95.752
Number of observations		13057
R^2		0.592

The dependent variable is the average bid; ***-statistical significance at 1%;

9.3 Independent versus inter-dependent valuations

The identification of the highest valuation from the highest observable bid in unsuccessful auctions is based on the assumption that bidders' valuations are independently distributed. If bidders' valuations have a commonly distributed component, then a bidder with the highest bid over-bids relative to her valuation[39].

We argue that bidders' valuations are independently distributed, although a

[39] In the literature this phenomenon is known as the "winner's curse," and it occurs because a bidder with the highest bid takes the fact that she is the winner as a negative signal of the quality of the item on sale. Hence, in the presence of a commonly distributed component in bidders' valuations the highest observable bid exceeds the true highest valuation. For an overview, see Kagel and Levin (2002).

commonly distributed component is likely to be present as well. The main justification for the independence assumption comes from the nature of the items on sale in our sample. In section 2 we show that the average age of the tractors in our sample is about 22 years with a large share of tractors exceeding 30 years. Since most of the tractors in the sample are old, they are probably purchased for own use rather than for a later resale. By "own use" we mean that the bidders most likely intend to use the purchased tractors for farming purposes or for spare parts. In either case, the bidders' valuations of the tractors on sale are mostly guided by bidders' individual preferences and less by valuations of other bidders. Based on this consideration we argue that the independence assumption is likely to fit the data more accurately.

We follow the approach of Bajari and Hortacsu (2003) to empirically test whether bidders' valuations are independently distributed or have a commonly distributed component. The empirical test of Bajari and Hortacsu is based on the argument that if there is a commonly-distributed component in bidders' valuations then bidders should rationally lower their bids in larger auctions. Hence, Bajari and Hortacsu argue that the number of bidders should be negatively correlated with realized bids if there is a commonly-distributed component in bidders' valuations. In contrast, if bidders' valuations are completely independently distributed, then the number of participating bidders should not affect realized bids.

To control for affiliation among bidders' valuations we include dummies for 18 categories of tractors. In addition, since the number of bidders is endogenous to realized bids, we use the minimum bid as an instrument for the number of bidders. Since we regress individual realized bids, we also include bidders' feedback scores among determinants. We present results of our test in Table A4.

The IV regression results in Table A4 show that the number of bidders is statistically significant and negatively correlated with realized bids. According to Bajari and Hortacsu (2003), this indicates that bidder's valuations have a commonly distributed component, and the highest bids in unsuccessful auctions exceed their underlying valuations.

Table A4. IV regression of realized bids on the number of bidders.

Variable	Coefficient	Robust Standard Error
Number of bidders$_i$	-1222.166***	198.562
Secret reserve dummy$_i$	-1689.940***	386.458
BIN$_i$*Secret reserve dummy$_i$	-3796.384***	619.197
BIN$_i$	0.073**	0.033
ln(Seller feedback$_i$)	404.160***	142.116
ln(Buyer feedback$_i$)	-146.798**	68.294
Age$_i$	-209.094***	23.806
Engine HP$_i$	56.059***	13.494
John Deere$_i$	-942.256	942.939
Number of observations		43138

The dependent variable is realized bids; ***,**-statistical significance at 1%,5%;

In the table we do not report coefficients for category dummies;

Category dummies are jointly statistically significant at 1% with p-value<0.0001