

Last Minute Bidding on Tradera

David Nordström

Master Thesis I

Department of Economics, Lund University

Supervisor: Jerker Holm

August 2011

Keywords: Sniping, Last minute bidding, Online auction, Shill, Incremental bidding, Second-price auction, Bayesian game

Abstract

In online auctions bidding frequency tends to increase towards the last minutes of the auction. This thesis investigates why bidders might choose to engage in last minute bidding, *sniping*, in Tradera auctions (Tradera is a Swedish subsidiary of eBay). Tradera uses a hard-close rule and a second-price rule. In hard-close auctions the dead-line occurs at a specific time, after which new bids are not considered. Due to e.g. erratic network traffic, bids submitted close to the dead-line (snipe bids) might not be successfully transmitted. The second-price rule prescribes that the good is to be rewarded to the highest bidder, for a price equal to the second highest bid (see Vickrey, 1961).

Utilizing the hard-close and second-price rule we consider two different models to examine why bidders might snipe. In both models there is a positive probability, a , that a snipe bid is successfully transmitted. First we construct a static Bayesian game to evaluate how the presence of a shill affects bid timing. A shill is a seller that bid on his own item in order to boost up the final price thereby reducing the winner's surplus (see e.g. Bhargava et. al., 2005). In this model there is a probability p that a shill is present. We show that if a and p are sufficiently high; an equilibrium in which all players snipe may exist. Secondly, we review the discontinuous eBay-model proposed by Ockenfels and Roth (2005). eBay and Tradera auctions are by large identical, allowing us to directly apply their model for our purpose. In this model, sniping can be a best response against incremental bidding. An incremental bidder is interpreted as an inexperienced bidder that mistakes the second-price rule for a first-price rule.

Using data from Tradera consisting of 200 Iphone auctions and 200 art auctions we empirically test the theoretical predictions. The effects of a shill upon bid timing cannot be confirmed. Relevant coefficients exhibit the expected signs, but cannot be accepted on any relevant level of significance. The weak results are probably due to the restrictive environment of the game and inaccurate estimates of p . When testing for the effects of incremental bidding we observe a statistically significant and positive relationship between bidder rank and sniping. A bidder's rank is assumed to approximate his experience. This implies that rational (high rank) bidders snipe in order to avoid an early price war with incremental (low rank) bidders.

Table of contents

1. Introduction	5
1.1. Online auctions and last minute bidding	5-6
1.2. Purpose	6
1.3. Limitations	6-7
1.4. Structure	7
2. Tradera	8
3. Static Bayesian games	8
3.1. Games of incomplete information	8-10
3.2. Static Bayesian games and Bayesian Nash equilibrium	10-12
3.3. Dominant strategies	12
4. The model	12
4.1. Shill bidding	12-13
4.2. Setting and strategies	13-15
4.3. Pay-offs	15
4.3.1. Step 1, one player	15-16
4.3.2. Step 2, two players	16-17
4.3.3. Step 3, three players	17-20
4.4. Equilibrium of the model	20-22
5. Incremental bidding	22
5.1. Effects of incremental bidding on Tradera	22-26
6. Data analysis	26
6.1. Data set	26-27
6.2. The private and independent value paradigm	27-28
6.3. Results	28
6.3.1. The model	28-31
6.3.2. Incremental bidding	31-33
7. Conclusion	34
7.1. The effects of a shill	34-35
7.2. Incremental bidders	35-36

7.3 Difference between Iphone and art auctions	36-37
8. Final remarks and future research	37-38
Literature and other sources	39-40
Appendix I	41-44

List of figures, tables and graphs

<i>Figure 4.1. Normal-form representation of the limited game</i>	17
<i>Figure 4.2 Normal-form representation of the complete game</i>	18-19
<i>Graph 6.1. Amount of auctions with different number of snipe bids</i>	27
<i>Table 6.1 Commission rates on Tradera</i>	29
<i>Table 6.2 The effects of a shill – Regression output</i>	30
<i>Table 6.3 Incremental bidding - Regression output</i>	33

1. Introduction

1.1. Online auctions and last minute bidding

Auctions conducted online often display a relatively large amount of bids being submitted in the very last minutes. This poses the question if specific features of online auctions create incentives for such behavior. This thesis theoretically and empirically examines last minute bidding, *sniping*, in online auctions using auction data from the online auction house Tradera (www.tradera.se). Tradera is a Swedish subsidiary of eBay and their auctions are by large identical to eBay auctions. In particular, both Tradera and eBay employ a second-price rule which mean that the good being auctioned is rewarded to the highest bidder for a price equal to the second highest bid.

The commonly studied auction settings, such as the second-price (also called Vickrey auction from Vickrey (1961), who pioneered this type of auction setting) and the English auction, do not account for several important features present in online auctions. In this thesis we will consider some of these features and how they could affect bidders' behavior. More specifically we will look at why sniping could arise in online auctions. The main part of this thesis is devoted to developing and empirically testing a static game that examines how the possibility of a shill could affect the timing of a bid. A shill is a seller, disguised as a bidder, who bid on his own items to boost up the final price. By construction, second-price auctions create incentives for the seller to participate in the bidding procedure as otherwise the final price could be below the maximum willingness to pay amongst participating bidders. The model is completely developed by the author and relies on rather strong assumptions. Hence the implications of the model and the conclusions that can be drawn are limited.

To explore the data further, we also test the eBay-model developed by Ockenfels and Roth (2005) where the presence of incremental bidders can induce sniping. Incremental bidders submit a bid below their valuation early in the auction. Whenever he is outbid he increases his bid in increments, up to the point where he is again the winner, or he reaches his maximum willingness to pay. Incremental bidders are interpreted as naïve, or irrational, bidders that confuse the second-price auction for a first-price auction (open

out-cry auction to be specific). In an open out-cry auction it is rational for a bidder to behave in this manner.¹ As eBay auctions are identical to Tradera auctions with respect to properties considered here, we treat their description analogously with a Tradera auction.

Both models used in this thesis incorporate the possibility that a snipe bid might not be successfully transmitted. Tradera (and eBay) use a fixed deadline in every auction, known as a hard-close rule. Bids that are received after this deadline are not registered and are not considered as legitimate bids. If bidders wait till the last seconds of an auction to submit their bid then slow network traffic, or bad timing, could result in these bids not being successfully transmitted.² The hard-close rule is a crucial part of both models as a positive probability of a snipe bid not being successfully transmitted affect the timing of bids.

1.2. Purpose

The purpose of this thesis is to answer the two following questions:

- Does a positive probability of a snipe affect when bidders choose to submit their bids in Tradera auctions?
- Can incremental bidders affect when rational bidders choose to submit their bids in Tradera auctions?

The timing of bids is the relevant issue in this thesis. We investigate possible considerations to a bidder when choosing whether to bid early or snipe. When examining the timing of bids we look at whether they arrive in the very last minute of an auction (snipe bids), or before that (early bids). Besides determining whether a bid is submitted in the "early" part of the auction, or if it is sniped, we will not be more specific regarding bid timing.

1.3. Limitations

¹ The interpretation of an incremental bidder is analogous with the definition in Ockenfels and Roth (2005), and Ariley, Ockenfels and Roth (2005).

² Ockenfels and Roth (2005) report that roughly 20 % of snipe bids are not successfully transmitted on eBay.

Online auctions follow a strict set of rules that are set by the specific auction house. Though most of the rules and properties of these auctions are observable to a third-party such as a researcher, it would require advanced modeling techniques to fully account for these. We will not attempt to incorporate all such rules and properties in this thesis. Rather we focus on a few aspects of Tradera's auctions that I consider relevant for the purpose of this thesis. The results that stem from this work should not be seen as an exhausting description of bidder behavior in Tradera auctions. Instead it should be interpreted as an attempt to describe some aspects of these auctions and how those could affect sniping.

It might seem strange to use two rather different approaches to the issue of sniping. However shill bidding and incremental bidding are not mutually exclusive as they consider different participants of an auction (seller and bidders). It should be possible to incorporate both these features into one model but most likely we would need rather complicated formulations.

This thesis considers second-price auctions where bidders' valuations are private and independent. Though we let bidders' valuations to be common knowledge in the model by Ockenfels and Roth (2005), these are treated as idiosyncratic. In both models considered here, bidders have a partly, or complete, idiosyncratic valuation of the good being auctioned. For this reason, the connection to the private value auction setting seems suitable. In Section 6.2 we will discuss the properties of products with private value, and how these relate to our data.

1.4. Structure

The rest of this thesis is outlined as follows: In Section 2 we briefly describe how bidding is done on Tradera. In Section 3 we review static Bayesian games and the conditions for equilibrium. Section 4 is devoted to developing a model to test how a shill might affect bid timing. In Section 5 we review the model proposed by Ockenfels and Roth (2005) to examine how incremental bidding could affect sniping. Section 6 describe the data set collected from Tradera, and present the results from empirically testing both models. In

Section 7 we conclude the thesis and discuss the results. Section 8 briefly discusses future research related to online auctions.

2. Tradera

A user must have a registered ID on Tradera to be able to bid on items. To register an ID the user has to submit personal information such as social security number and address. Once registered, the user can both bid in auctions and sell items. It is not possible for a seller to bid on his own item when using the same ID. However, if a user has several registered accounts, e.g. by using social security numbers of family members, it is possible to sell the item using one ID and then bid on the item with another ID.

Bidding on Tradera is done via a proxy bidder. The bidder enters a maximum bid into the proxy bidder which then automatically does the bidding.³ Whenever the bidder is outbid the proxy automatically raises the bid to the minimum amount required to win. If the bid required to again be the highest bidder is above the maximum bid, the proxy submit the maximum bid and then drop out. If the bidder wishes to re-enter the auction, he can do this by submitting a new, and higher, maximum bid. When the auction has reached its deadline, the winner is the bidder that submitted the highest, successfully transmitted, bid. Auctions on Tradera normally run for 7 days. Though it is possible for a seller to choose his own closing-time, for a fee of 10 SEK, this is option is rarely used.⁴

The bidding history on Tradera lists every previously winning bid, the current winning bid, the time these bids were submitted and the ID of every bidder. As Tradera use a second-price rule, the current winning bid equals the second highest bid, plus some minimum increment. In Tradera auctions the minimum increment actually increases with the price, but we ignore this fact in our models (see Section 5). The maximum bid of the current highest bidder is not known by other bidders.

³ <http://www.tradera.com/help/HelpPage.aspx?NodeID=6741>

⁴ Most auctions in our data set run for 7 days. Some auctions run for e.g. 24 hours but since our models are not specified in absolute time, the strategic considerations mentioned here should still be valid.

3. Static Bayesian games

3.1. Games of incomplete information

The most simplistic setting of an economic game assumes perfect information. In such an environment, players are fully aware of their opponent's preferences and base their decisions on this information. The notion of perfect information might be applicable in some situations but real-life problems often involve some degree of uncertainty. When players are not fully informed of opponents' preferences we call this a game of incomplete information. In such games a player need to form beliefs about his opponents which must be based on opponents' beliefs about him. This reasoning extends in infinitum, making analytical work on such games rather complicated (Harsanyi, 1967). A widely used approach when analyzing such games is to instead let preferences follow a distribution that is common knowledge to all players. Preferences are then distributed amongst players according to some probability function that is common knowledge. The realization of each player's preference however, is private information. Players now have a common belief about each other, based on this distribution function. Every particular realization of players preferences constitute a state of nature, a subgame, in which the game will be played. Since the realization of opponents' preferences are private information, a player does not know in which subgame they are actually in. When the particular subgame is unknown, we say that players have imperfect information. Technically in a game of imperfect information, each player's decision node is a non-singleton. By following this procedure we have reinterpreted a game of incomplete information as a game of complete but imperfect information. Such a transformation is known as a Harsanyi transformation⁵ and we call these games *Bayesian games*. In Bayesian games we drop the requirement that players must form infinite contingent beliefs by instead letting preferences follow a probability distribution. This is the core of Bayesian games, and provides a useful analytical tool when studying situations (games) where participators have insecurity regarding their opponents' preferences.

⁵ The name is taken from Harsanyi (1967) who first proposed this technique.

Below we describe the equilibrium concept of static Bayesian games. The formulation of Bayesian games and the notion of equilibrium conditions in such games are to a large extent taken from the text-book "Microeconomic Theory" (Mas-Colell, Green and Whinston, 1995). The text-book is used as teaching material for advanced courses in microeconomics in academic institutions all over the world. For this reason I consider their work as recognized formulations of economic games.

Before we proceed a notational remark is needed: for the rest of this thesis we will index the opponents of player i as $-i$.

3.2. Static Bayesian games and Bayesian Nash equilibrium

We consider a static Bayesian game where all players move once and simultaneously. Since Bayesian games consists of several components this first part might seem somewhat technical. However, to explore equilibrium strategies later in Section 4 it becomes necessary to consider the full settings of these games.

A static Bayesian game consists of a set of $n = \{1, 2, 3, \dots, I\}$ players, indexed by $i = 1, 2, 3, \dots, I$, and *Nature*. Each player i has a pay-off function $u_i(s_i, s_{-i}, \theta_i)$, where $s_i \in S_i$ is the strategy for player i and $\theta_i \in \Theta_i$ is a random type distributed by *Nature*. A strategy s_i is a complete contingent plan that specifies an action for *any* possible information that the player could have. In this context, the relevant information for player i is his type θ_i , which is private information, and the distribution function $F(\theta_1, \dots, \theta_I)$ which is common knowledge. Each state of nature where every player is of a particular type is given by $\Theta = \Theta_1 \times \dots \times \Theta_I$. We can now summarize this game by

$$[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)] \quad (3.1)$$

Definition 3.1:

Any Bayesian game can be summarized by the data $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$.

A *pure* strategy for player i is given by the function $s_i(\theta_i) \in S_i$ that calls upon the player to take a specific action given his type θ_i . $\mathfrak{S}_i \subset S_i$ denotes the set of all such pure

strategy functions. The pay-off to player i given the profile of pure strategies for the I players, $(s_1(\cdot), \dots, s_I(\cdot))$, can then be written as

$$\tilde{u}_i(s_1(\cdot), \dots, s_I(\cdot)) = E_\theta [u_i(s_1(\theta_1), \dots, s_I(\theta_I), \theta_i)] \quad (3.2)$$

(3.2) is a von Neumann-Morgenstern (v.N-M) utility function as it is contingent on the random distribution of players' types $F(\cdot)$. The v.N-M utility function reflects that the pay-off for player i is a joint probability distribution over the strategies of all I players and the realization of his type. v.N-M utility functions require that we must make additional assumptions about players' attitudes towards risk.⁶ In this thesis we will only consider risk-neutral players. Pay-offs can then be directly compared in absolute terms and we do not have to consider different attitudes towards risk.

To formulate the equilibrium condition of a static Bayesian games we also need to define the normal form representation of a Bayesian game. It is given by

$$\Gamma_N = [I, \{\mathfrak{S}_i\}, \{\tilde{u}_i(\cdot)\}] \quad (3.3)$$

The normal form representation of the Bayesian game consists of I players, the set of pure strategies \mathfrak{S}_i and the pay-off function $\tilde{u}_i(\cdot)$ which is defined as in (3.2). Note that this representation includes the random variable, *Nature*, via \mathfrak{S}_i and $\tilde{u}_i(\cdot)$ so that no information is lost from Def 3.1. Once these settings are imposed it suffices to find a pure strategy Nash equilibrium (NE) for the game Γ_N , known in this context as a Bayesian Nash equilibrium (BNE).

Definition 3.2

A pure strategy Bayesian Nash equilibrium in the Bayesian game described by Def. (3.1), is a set of pure strategies $(s_1(\cdot), \dots, s_I(\cdot))$ that constitutes a Nash equilibrium of game Γ_N as defined in (3.3). For every $i = 1, \dots, I$ we must have

$$\tilde{u}_i(s_i(\cdot), s_{-i}(\cdot)) \geq \tilde{u}_i(s'_i(\cdot), s_{-i}(\cdot))$$

for all $s'_i(\cdot) \in \mathfrak{S}_i$ where $\tilde{u}_i(s_i(\cdot), s_{-i}(\cdot))$ is defined as in (3.2).

⁶ See e.g. Mas-Colell, Green and Whinston, Chapter. 6 (1995)

A BNE can be described as each player playing a *best response* strategy, given the conditional distribution of opponents' strategies, for *any type that he might get*. The notion of a best response is crucial to determine equilibrium in a Bayesian game. A best response strategy can only exist if it maximizes player i 's expected pay-off, as given by (3.2). The best response strategy of each player must coincide in equilibrium so that player i can not increase his pay-off by deviating to some other strategy s'_i . Also a best response strategy must exist *ex ante*. If a player, once he learns his type, change his strategy the model lack any predictable power. In a BNE we must be able to identify an equilibrium strategy before the game starts.

3.3. Dominant strategies

When equilibrium strategies exist in a Bayesian game these can either be weakly or strictly dominant strategies. In a set of strategies S_i for player i , the strategy $s_i \in S_i$ is *weakly dominated* if there exist another strategy $s'_i \in S_i$ that always yield at least as good pay-off given the strategy set S_{-i} of all opponents and strictly better pay-off for some strategy $s_{-i} \in S_{-i}$. A strategy is *weakly dominant* if it weakly dominates every other strategy in S_i . A strategy $s_i \in S_i$ is *strictly dominated* if there exists another strategy $s'_i \in S_i$ that, given the strategy set S_{-i} of opponents, always yield a better pay-off. A strategy s_i is *strictly dominant* if it dominates every other strategy in S_i .

4. The model

4.1. Shill bidding

In this section I develop a static model in an attempt to examine how the possibility of a shill affects players' bid timing. Before we formulate the model we will define the practice of shill bidding and briefly discuss shill bidding in the literature. To my knowledge there is no published work on the possible relation between shill bidding and sniping so the literature discussion will be confined to discussing the general purpose and actions of a shill.

Definition 4.1:

A shill is a seller, or an accomplice of the seller, that participates in the auction disguised as a bidder. In a second-price auction with private, independent valuations the purpose of a shill is to capture a fraction between the highest and the second highest bid, thereby increasing the payment of the winner. Since the winner pay more when a shill is present, his surplus is reduced. The shill can only act, i.e. submit a bid, if an honest bidder submits an early bid. An early bid is any bid that is not sniped.

In the literature the definition of a shill is similar to Def. 4.1. The research focus on shill bidding has mainly been devoted to how deviations from the efficiency and revenue results of conservative auctions arise when a shill is present (see e.g. Porter and Shoham, 2005; and Bhargava et. al., 2005). The actions and preferences of a shill in second-price auctions with private values, are not specifically examined and is generally treated exogenously. The timing of a shill bid varies in different papers and either takes place at the same time as other bids (Chakraborty and Kosmopoulou, 2003) or in a second period after all honest bidders have submitted their bids (Harstad and Rothkopf, 1995).

Since I lack any previous work that is similar to the model I will develop we have to consider how a shill could reasonably choose to place a bid in an auction such as Tradera's. Using Def. 4.1 we will assume that a shill act only on *submitted bids*. Bidders must submit their bids before the shill might place his bid. As the shill wants to extract a fraction between the highest and the second highest bid, we require that these "reference points" must exist before the shill can act. Actually, we only need one bid to have been submitted since the second highest bid is then zero.⁷ Thus, the minimum information required for a shill to be able to submit a bid, is given when any bidder submits a bid. Below we will impose the rule that this only apply for early bids. When a bidder snipes, there exists no possibility for the shill to respond to such a bid.

4.2. Setting and strategies

⁷ If there is only one submitted bid, b , in a Tradera auction, the winning bid equals the minimum increment, m . This is true as long as the seller does not have a reservation price, r (and $b > m$). If $r > 0$, the winning bid is one increment above the reservation price, $r + m$ (assuming that $b > r + m$).

The model is formulated as a static game of imperfect information and we will be looking for a BNE for this game. The game is constructed to simulate some features of a real Tradera auction. In particular, we manipulate players pay-offs so as to account for a positive probability that a skill is present, and a positive probability that a snipe bid is not successfully transmitted. Rather than examining the actual size of bids, we focus on the timing of bids where we make a distinction between whether bids are submitted early, or sniped.

In this game:

- There is a set $n = \{1, 2, 3\}$ of risk-neutral players, indexed by $i = 1, 2, 3$.
- Each player's type $T_i \in \Theta_i = \{1, 2\}$ is private information, where both types are drawn with equal probability (=0.5). The distribution function $F(T_1, T_2, T_3)$ is common knowledge.
- Each player can choose two different actions: (*Snipe*), (*Don't*).
- First Nature distributes types according to $F(\cdot)$. After each player learns their type, they choose an action. After this the game ends and all players receive their pay-off.
- There is a positive probability a that a snipe bid is not successfully transmitted, and a positive probability p that a skill is present. Both a and p are exogenously given ex ante, and are common information to all players.

Though it is possible to allow for mixed strategies in this game, we will only focus on the set of pure strategies \mathcal{P}_i . Note that \mathcal{P}_i is actually a subset of some set S_i , where S_i contains both mixed and pure strategies available to player i . To denote the pure strategies in the set \mathcal{P}_i we add another index number $k = 1, 2, 3, 4$ for each strategy so that the k th strategy of player i is given by $s_{i,k}$. In this game, $\mathcal{P}_i =$

$\{(s_{i,1}: \textit{Snipe}$ if $T_i = 1$, \textit{Snipe} if $T_i = 2)$,

$(s_{i,2}: \textit{Snipe}$ if $T_i = 1$, $\textit{Don't}$ if $T_i = 2)$,

$(s_{i,3}: \textit{Don't}$ if $T_i = 1$, \textit{Snipe} if $T_i = 2)$,

$(s_{i,4}: \text{Don't if } T_i = 1, \text{Don't if } T_i = 2)$.

We have now specified a complete set of pure strategies for player i . Each strategy $s_{i,k}$ has a complete plan of action for any type T_i that player i could end up with. The pay-off to player i from playing strategy k is given by

$$u_{i,k}(s_{i,k}, s_{-i}, T_i) \quad (4.1)$$

Just like (3.2), (4.1) has the expected utility form where the expectation of player i 's type and the conditional distribution of opponents' strategies determine player i 's pay-off.

4.3. Pay-offs

Before we move on to considering possible BNE of this game we need to examine the pay-off functions more closely. The pay-offs are constructed so that some of the characteristics of Tradera's auctions are present. Instead of allowing players to submit a bid equal to some value we have limited the actions to (*Snipe*) and (*Don't*). The actions should be interpreted as players either choosing to snipe (*Snipe*), or placing an early bid (*Don't*). For simplicity the pay-offs are constructed in three steps. We first consider only one, then two, and then finally all three players.

4.3.1. Step 1, one player

Here we consider the pay-offs as if there was only one player, player i , in the game. Also, we ignore what specific type player i might be. If (*Snipe*) is chosen the pay-off is $u_i = aT_i$, where $0 < a < 1$ and u_i is defined as in (4.1). a should be interpreted as the probability that a late bid is *successfully transmitted*. If player i chooses to snipe he is exposed to a positive probability that his bid is not successfully transmitted. If (*Don't*) is chosen the pay-off is $u_i = (1 - p)T_i$, where $0 < p < 1$. p should be interpreted as the probability that a *skill is present*. From Def. 4.1 we know that when player i submits an early bid he is exposed to a positive probability that a skill will capture some of his surplus. The surplus for player i is given by his pay-off u_i . In this case we assume that if a skill is present, then player i loses his entire surplus ($u_i = 0$). The key-part of this game is that both actions are costly. By submitting an early bid (*Don't*), players risk that a skill

can observe their bid and act on it. If player submit a late bid (*Snipe*), there is no time for the shill to react. However, a snipe bid is costly since with probability a , the bid is successfully transmitted.

Now we move on to considering how a player's type affect his pay-off u_i . The specific type of a player and the associated pay-off is constructed so as to imitate the second-price rule of Tradera. For now, we let $a = 1$ and $p = 0$ so that the effect from playing different actions can be ignored. We assume that (when $a = 1$, $p = 0$) a player with a higher valuation ($T = 2$) receives a positive pay-off, and a player with a low valuation ($T = 1$) receives zero pay-off. If $T_i > T_{-i}$ only player i can get a non-zero pay-off. Player i receive a pay-off equal to his valuation, T_i minus the second highest valuation, T_{-i} . In this case since $T_i = \{1, 2\}$, the pay-off is equal to 1. If $T_i = T_{-i}$, all players would get zero payoff. When we ignore the effects of a and p , so that every bid is always submitted, and there is no shill present, the pay-off to the winner is determined by his valuation minus the highest valuation of his opponents. The pay-off for player i if he wins is given by $u_i = T_i - \max(T_{-i})$.

By manipulating players' pay-offs, we have made some advances towards a real Tradera auction. Allowing for a probability of a that a bid is successfully transmitted, and a probability of p that a shill is present, we have introduced two elements that could affect the outcome in such auctions. By only letting the player with the highest valuation receive a positive pay-off we have to some extent captured the second-price mechanism of Tradera.

4.3.2. Step 2, two players

Now we combine these properties to get a more complete description of players' pay-offs. For this part we only consider *two players*, player 1 and player 2. The reason for doing so is that we want to focus on how different strategies and valuations affect pay-offs. A third player is then trivially included by identical reasoning. In this part we describe the case $T_1 = 2, T_2 = 1$, depicted below in Fig. 4.1. Pay-off for player 1 and player 2 is given in the bottom-left and top-right respectively of each cell. It is important to note that Fig.4.1 is an incomplete representation of the game we are developing, so that

we cannot infer what strategies will actually be played from it. However, it is sufficient to exemplify the different features of the complete game.

		$T_2 = 1$	
		<i>Snipe</i>	<i>Don't</i>
$T_1 = 2$	<i>Snipe</i>	$2a - a^2$	a
	<i>Don't</i>	$(1 - p)(2 - a)$	$(1 - p)$

Figure 4.1. Normal-form representation of the limited game.

Consider any strategy profile which result in the cell (S, S) in Fig. 4.1. Since there is a positive probability that player 2's bid is *not* transmitted, player 1 can receive a pay-off of 2 if only his bid is transmitted. This occurs with probability $a(1 - a)$. If both bids are transmitted, which happens with probability a^2 , player i gets a pay-off of $1 (= T_1 - T_2)$. Adding these together we get $2a(1 - a) + a^2 = 2a - a^2$. Player 2 has a lower valuation, $T_2 = 1$, and can only receive a strictly positive pay-off if his bid is transmitted and player 1's is not. This occurs with probability $a(1 - a)$. Moving to cell (S, D) , where only player 1 plays (*Snipe*) and player 2's bid is always transmitted. Player 1 can only get a pay-off of 1, if his bid is transmitted. By playing (*Don't*), player 2 is exposed to the probability that a shill captures his surplus, T_2 , in which case he receives zero pay-off. This happens with probability p . If no shill is present player 2 can only receive a positive pay-off (=1) if player 1's bid is not transmitted. The joint probability of these events is given by $(1 - a)(1 - p)$. Adding up the pay-off for player 2 we get $p * 0 + (1 - a)(1 - p)1 = (1 - a)(1 - p)$. The pay-offs in cells (D, S) and (D, D) are solved in a similar manner. In (D, S) player 1 is exposed to the probability of a shill and the probability that player 2's bid is not transmitted. Player 1's pay-off is $p * 0 + (1 - p)a + 2(1 - p)(1 - a) = (1 - p)(2 - a)$. Since player 1's bid is always transmitted, and $T_1 > T_2$, player 2 has zero pay-off. In (D, D) both players bids are

transmitted and only player 1 can get a pay-off of 1 which occurs with probability $(1 - p)$.

4.3.3. Step 3, three players

We now have sufficient information to impose a complete description of our game.

Definition 4.1: We can now formulate the setting outlined above as a Bayesian game summarized by the data

$$\mathcal{B} = [3, \{\mathcal{S}_i\}, \{u_{i,k}\}, \Theta, F(T_1, T_2, T_3)] \quad (4.2)$$

Definition 4.2:

Using Def. 4.1 we can describe the normal form representation of the game as

$$\Gamma_N = [3, \{\mathcal{P}_i\}, \{u_{i,k}\}] \quad (4.3)$$

A graphical representation of Γ_N is given in Fig.4.2 below. Fig 4.2 consists of two pay-off matrices, treating player 3's choice as fixed at either (*Snipe*) or (*Don't*). A profile of strategies $(s_{1,k}, s_{2,k}, s_{3,k})$, tell us what cell $((S, S, S), (S, S, D), \dots, (D, D, D))$ we end up in. Remember from (3.2) that the pay-off for player i is taken over the expectation of his valuation and the distribution of opponents strategies $F(\cdot)$. Thus each pay-off cell must consider every possible state of nature, given by the possible configurations of valuations (T_1, T_2, T_3) , and the outcome associated with that configuration. Each configuration occurs with probability $0.125 (= (\frac{1}{2})^3)$. Complete calculations for Fig.4.2 can be found in Appendix I.

P. 3 = <i>Snipe</i>		P. 2	
		<i>Snipe</i>	<i>Don't</i>
P. 1	<i>Snip</i>	$u_1 = \frac{1}{8}(a(11(1-a)^2 + 4a(1-a))),$	$u_1 = \frac{1}{8}(a + a(1-a)),$
	<i>e</i>	$u_2 = \frac{1}{8}(a(11(1-a)^2 + 4a(1-a))),$	$u_2 = \frac{1}{8}(12(1-a)^2(1-p) + 4a(1-a)(1-p)),$

		$u_3 = \frac{1}{8}(a(11(1-a)^2 + 4a(1-a)))$	$u_3 = \frac{1}{8}(a + a(1-a)),$
	<i>Don't</i>	$u_1 = \frac{1}{8}(12(1-a)^2(1-p) + 4a(1-a)(1-p))$ $u_2 = \frac{1}{8}(a + a(1-a)),$ $u_3 = \frac{1}{8}(a + a(1-a)),$	$u_1 = \frac{1}{8}((1-p) + (1-a)(1-p)),$ $u_2 = \frac{1}{8}((1-p) + (1-a)(1-p)),$ $u_3 = \frac{1}{8}a$
P. 3 =Don't		P. 2	
		<i>Snipe</i>	<i>Don't</i>
P. 1	<i>Snipe</i>	$u_1 = \frac{1}{8}(a + a(1-a)),$ $u_2 = \frac{1}{8}(a + a(1-a)),$ $u_3 = \frac{1}{8}(12(1-a)^2(1-p) + 4a(1-a)(1-p)),$	$u_1 = \frac{1}{8}a,$ $u_2 = \frac{1}{8}((1-p) + (1-a)(1-p)),,$ $u_3 = \frac{1}{8}((1-p) + (1-a)(1-p)),,$
	<i>Don't</i>	$u_1 = \frac{1}{8}((1-p) + (1-a)(1-p)),$ $u_2 = \frac{1}{8}a,$ $u_3 = \frac{1}{8}((1-p) + (1-a)(1-p)),,$	$u_1 = \frac{1}{8}(1-p),$ $u_2 = \frac{1}{8}(1-p),$ $u_3 = \frac{1}{8}(1-p)$

Figure 4.2. Normal-form representation of the complete game

Now we have a complete description of our game, given by Fig. 4.2, containing all relevant information. Before we move on, let us recap: we are analyzing a static Bayesian game where players know their own type (private information) and the distribution of opponents' types (common knowledge). The game is formulated as to simulate some of the features of a Tradera auction. For this purpose we explicitly analyzed the pay-offs. By submitting an early bid players were exposed to a positive probability that a shill reduced their surplus to zero. If a player snipes there is a positive probability that his bid is not successfully transmitted. Given that his bid was transmitted, the player with the highest valuation always wins.

We have imposed some rather strong assumptions into this model. First we treat the auction as a static game where players can only choose one action. This limit our analysis as we do not allow players to update their information, and actions, as the auction

proceeds. Secondly, we assume that as in conservative auction theory, the player with the highest valuation win the good (as long as his bid is successfully transmitted). Again, the dynamics of a real Tradera auction might result in bidders *not bidding* their true valuation. For example, Ockenfels and Roth (2005) show that in the eBay-model with two rational bidders with private valuations, there is generally no dominant bidding strategy.

It should be noted that we did not formally include the variables a and p in any definition of our game. This was deliberately done in order to treat our model in a comparable way to the outline of Bayesian games given in Section 3. The pay-off corresponding to some strategy is not always affected by both variables. Also the strategies of opponents affects how player i 's pay-off is determined with respect to a and p . If, for example, all players play (*Snipe*) then p does not determine the pay-off function for player i . If player i play (*Don't*) then his opponents are also affected by p no matter what strategy they play. As we will see the realized values of both variables affect equilibrium strategies which we do not fully capture in Def. 4.1-4.2. I was not able to find any good suggestions on this problem so I will not attempt to incorporate a and p into these definitions. Instead we let the graphical representation of Γ_N given in Fig. 4.2 serve as an indicator as to how the pay-off for player i is affected by a and p .

4.4. Equilibrium of the model

From our settings described above we can now look for a BNE in this game. To do this we will use Fig. 4.2 and delete strategies that are never a best response given some play by opponents. The proof is carried out similarly to the process of iterated deletion of strictly dominated strategies (see Mas-Colell, Green and Whinston, 1995, pp. 238).

Proposition 4.1:

In the Bayesian game described by \mathcal{B} , $(s_{1,1}, s_{2,1}, s_{3,1})$ is a profile of strategies that may constitute a pure strategy Bayesian Nash equilibrium, if and only if

$$a > (1 - p) * y(a).$$

$$\text{where } y(a) = \frac{12-8a}{11-7a}.$$

Proof: We need to establish that if $a > (1-p) * y(a)$, player i always find it a best response to play (*Snipe*) given any strategy of his opponents. In order to do this we need to check that given any play by opponents, player i will not find it profitable to deviate from (*Snipe*). Practically this amounts to checking that three different, and strict, inequalities hold for player $i = 1$.

Consider the cell (S, S, S). Player 1 will strictly prefer to play (*Snipe*) if and only if (*iff*)

$$\frac{1}{8} \left(a(11(1-a)^2 + 4a(1-a)) \right) > \frac{1}{8} \left(12(1-a)^2(1-p) + 4a(1-a)(1-p) \right)$$

after removing common terms on both sides, we obtain

$$a > (1-p) \left(\frac{12-8a}{11-7a} \right) \rightarrow$$

$$a > (1-p) * y(a) \quad (4.5)$$

Now we move to cell (S, D, S). Player 1 strictly prefers (*Snipe*) *iff*

$$\frac{1}{8} (a + a(1-a)) > \frac{1}{8} ((1-p) + (1-a)(1-p)) \quad (4.6)$$

by removing common terms we get

$$a > (1-p).$$

From the perspective of player 1, (S, D, S) and (S, S, D) are equivalent so the same condition has to hold in both cases.

Moving to cell (S, D, D), player 1 prefers (*Snipe*) *iff*

$$\frac{1}{8} a > \frac{1}{8} (1-p) \quad (4.7)$$

which obviously only holds if

$$a > (1-p).$$

We have now established that (4.6) and (4.7) holds whenever $a > (1 - p)$. For (4.5) to hold we need the additional requirement $a > (1 - p) * y(a)$.

We will now show that whenever (4.5) holds, (4.6) and (4.7) always hold. To do this, consider the first derivative of $x(a)$ which is given by

$$y(a)' = \frac{-4}{(11 - 7a)^2} < 0 \quad (4.8)$$

Though our model only allow values of a that are strictly below 1, we check the limit value of $y(a)$ when $a = 1$. This is given by

$$y(1) = \frac{12 - 8}{11 - 7} = 1 \quad (4.9)$$

Thus, in the limit, $y(a)$ approach 1 "from above" as a approach 1. Since $y(a) > 1$ for $0 < a < 1$, this prove that $(1 - p) * y(a) > (1 - p)$ for every possible value of a . We have now established that whenever (4.5) holds, (4.6) and (4.7) always hold.

We now know that whenever $a > (1 - p) * y(a)$, the best response for player 1 against any opponents' strategy is to play (*Snipe*). In other words, given any valuation that player 1 end up with, and any possible strategy $s_{-1,k} \in \mathcal{P}_{-1}$; player 1 weakly prefer $s_{1,1}(\cdot)$, iff $a > (1 - p) * y(a)$. $s_{1,1}(\cdot)$ is a weakly preferred strategy since $s_{1,2}(\cdot)$ and $s_{1,3}(\cdot)$ yield the same pay-off when player 1 draw type $T_1 = 1$ or $T_1 = 2$ respectively. Since player 1 choose to play (*Snipe*) in those cases, the pay-off is the same as for $s_{1,1}(\cdot)$.

Since the game is symmetric in strategies and valuations we can apply above argument for player 2 and player 3. We have now established a best response strategy for player i in the normal-form representation of \mathcal{B} given in Fig. 4.2. The best response strategy $s_{i,1}(\cdot)$ thus constitute a NE of Γ_N . Using Def. (3.2) we know that any NE Γ_N is a BNE of \mathcal{B} . ■

Proposition 4.1 tell us that if the probability that a snipe bid is transmitted, and the probability of a skill, are sufficiently high, then $s_{i,1}$ is a best response for player i against any strategy $s_{-i,k} \in \mathcal{P}_{-i}$. This result seems intuitive. Players will find it profitable to

always snipe if the risk of doing so is outweighed by the risk that a snipe captures all surplus.

5. Incremental bidding

5.1. Effects of incremental bidding on Tradera

In this section we present a second approach to why sniping can occur in auctions such as Tradera's. Below we will review the auction setting proposed by Ockenfels and Roth (2005) (the authors henceforth). The authors consider a dynamic second-price auction with a continuous "early" and a static "late" stage. Their model is built so as to simulate eBay auctions where bidders have the possibility to continuously update their bids, based on new information. In this auction, bidders compete for a single indivisible good. Compared to Section 4 we will be somewhat loose considering the initial settings of the game where we describe the path of play and rules. The reason for this is that the authors consider several different modifications of the auction regarding types of valuation, information and strategy sets. The path of play and rules are however identical for each modification. As we consider incremental bidding later we will be more specific regarding these issues.

In the early stage bidders can submit a bid at any time $t \in [0,1) \cup \{1\}$. Every bid submitted in the early stage is successfully transmitted with probability $a = 1$. A bidder can always react to a bid placed at time t' by an opponent. However, the reaction can only take place at some $t_n > t'$ where t_n is chosen from a countably infinite subset $\{t_n\} \cup [0,1)$ where $\{t_n\}$ converges to 1. If for some t we have $t_{n-1} < t < t_n$ then the information set for all bidders is the history up to t_{n-1} . This technicality creates a half-open interval so that bidders can make their reaction contingent on opponents' actions within this interval. Thus bidders always have the possibility to react in the early part of the auction.

The history at t_n consists of the current winning bid and the submission time and ID of each bidder's last submitted reservation price. The current winning bid at time t is $b_2 + s$

where b_2 is the second highest submitted bid, and s is some minimum increment. If the highest submitted bid is $b_1 < b_2 + s$, the current winning bid is b_1 . Whenever $b_1 \geq b_2 + s$, b_1 is private information. Any submitted bid by bidder i must exceed the current winning bid and any previous bid made by bidder i . If two or more identical bids are submitted then the current winning bid is that which was submitted first. If these were submitted at the same time the winning bid is randomly assigned with equal probability.

These rules imitates a Tradera (and eBay) auction rather well as bidders can only submit new bids of ascending values with a minimum increment rule. Also, in line with auctions on Tradera, bidders can observe the current winning bid at any time t , but not the value of the highest bid (as long as $b_1 \geq b_2 + s$).

At $t = 1$, the auction enters the late stage. At this point every bidder know the complete history of the early stage prior to t . The game now move into a static stage where every bidder can submit one more bid which has a probability a , where of $0 < a < 1$, of being successfully transmitted. a is exogenously given and is not affected by events within the game. Bids submitted at this stage are considered as snipe bids. Bidders cannot observe opponents actions at $t = 1$. After the late stage the auction ends and the good is rewarded to the bidder that submitted the highest, successfully transmitted, bid.

With this setting the authors also capture the hard-close rule of Tradera where snipe bids have a positive probability of not being transmitted. When bidders snipe there exists no possibility for their opponents to react to this. Under this assumption the strategy of bidders concerns a trade-off between the risk of inducing a price-war in the early stage and the risk of not having their bid successfully transmitted if they snipe.

We have now outlined the settings of the auction and can now proceed to consider the effect of incremental bidding in this auction. To do so we first need to clarify the behavior of an incremental bidder:

Definition 5.1:

The incremental bidder i only plays one strategy, s_i . The strategy s_i is to bid $m + s$ where m is the sellers reserve at $t = 0$. Whenever he is outbid at some $t' < 1$ he raises

his bid in increments of s for every t'' where $t' < t'' < 1$. He will continue to do so until he is again the highest bidder, or he reaches his maximum willingness to pay θ_i , whichever occurs first. The incremental bidder can be thought of as a program which activates whenever another bid is submitted by bidder j .

Proposition 5.1:

Bidding at $t = 1$ can be a best response against an incremental bidder. The strategy of the incremental bidder follows Def. 5.1.

Sketch of proof:

I will attempt to provide a similar proof as that of the authors below. As the authors only provide a sketch of the proof, I leave the details of the proof to them. In the proof provided by the authors there is one rational bidder j and one incremental bidder i . The seller has a reservation value of m . The valuation of bidder i and j is $\theta_i > m + s$ and $\theta_j > m + s$ respectively. Both θ_i and θ_j is common knowledge.

We will show that the strategy s_j where j only submits a bid $b = \theta_i$ at $t = 1$, is a best response to s_i when $\theta_i > \theta_j + s$.⁸ This strategy yields an expected pay-off of $a(\theta_i - m - s)$. Any other bid $b \leq \theta_j$ by j at some $t < 1$ would be outbid by i at some t' where $t < t' < 1$. Thus if bidder j wants to re-gain the position as the highest bidder whenever he has submitted a bid at t' , this would require a bid above his valuation θ_j . Any such bid would give j a negative pay-off. Also any bid b at $t < 1$ would raise the minimum achievable price at $t = 1$ above $m + s$. ■

The proof provided above is limited to only account for the case when bidders' valuations are common knowledge. As the authors do not discuss the case of private values we do not cover such a case. Relating this to a real Tradera auction, it seems unlikely that bidders in a Tradera auction would be aware of the valuation of opponents. Ariely, Ockenfels and Roth (2005) provide a similar proof, based on the same game setting as above. They show that, as long as bidder j knows he faces an incremental bidder, and

⁸ Though we will not examine it further here, the authors claim that other bids at $t = 1$ qualify as a best-response. This seems intuitive since some bid b' where $(m + s) < b' < \theta_i$ would yield the same expected pay-off as b

$a = 0.8$, bidding at $t = 1$ is a best response, for any value θ_j . The rationale for bidding at $t = 1$ is the same as above. Even though there is a positive probability equal to a that a snipe bid gets lost, this is outweighed by the benefits of avoiding an early price-war. Since this holds for all values of θ_j , we can assume that sniping as a best response to incremental bidding holds under quite broad circumstances with regards to bidders' information.

In this section we have discussed the model developed by Ockenfels and Roth (2005) and how incremental bidding behavior in such a model affects the timing of bids. It should be noted that the authors also show that in any auction with a hard-close, it is possible for a sniping equilibrium to exist, even without incremental bidders. The hard-close rule used on Tradera is in itself sufficient to result in a sniping equilibrium (though this equilibrium is not unique). The authors show that with two bidders where valuations follow a degenerate distribution with all its mass on $H(> m + 2s)$, they can credibly commit to the threat of an early price-war. By sniping, bidders can avoid a price-war and get a positive expected pay-off.⁹ Thus we should not expect incremental bidding by itself to be responsible for all observed sniping in these auctions. Rather it is one effect out of several, present on Tradera that can induce sniping. As we are interested in investigating how sniping varies between different Tradera auctions (all with a hard-close rule), the effects of incremental bidding is suitable for our purpose. By detecting incremental bidders in an auction, we can compare the timing of bids to that of rational bidders.

6. Data analysis

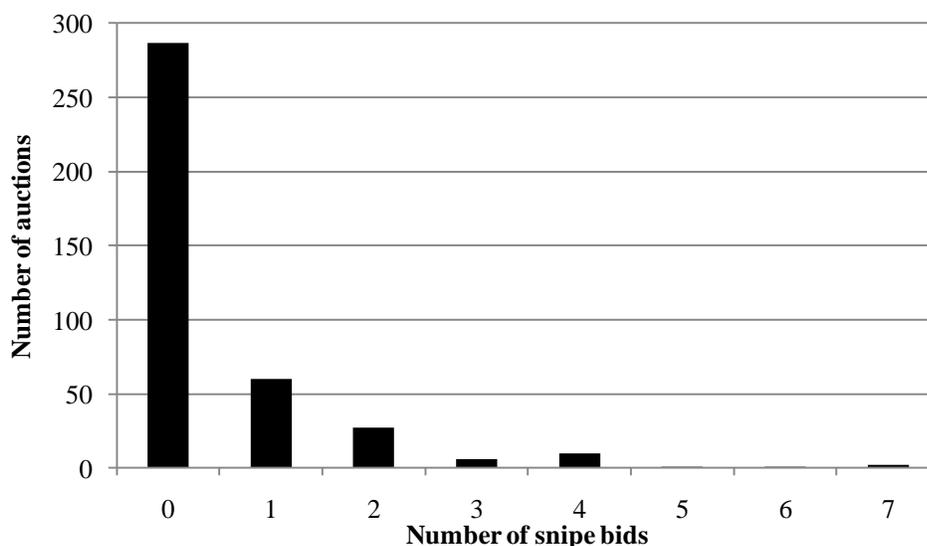
6.1. Data set

The data consists of 200 auctions for art and 200 auctions for Apple iPhones (various versions) gathered from Tradera between April-May 2011. There are totally 7291 bids submitted in the 400 auctions where 5436 of these bids are in iPhone auctions and 1855 are in art auctions. All bids that arrive in the last full minute of auction i are considered as

⁹ For details on this see Ockenfels and Roth (2005).

snipe bids. In Iphone auctions there are a total of 146 snipe bids and in art auctions there are 57 snipe bids. Sniping occurs in 113 auctions where 71 out of these are for Iphones and 42 for art. 2251 active bidders participated in Iphone auctions and 819 in art auctions. We do not account for the fact that one bidder might bid in several auctions. Thus the number of unique bidders is most likely lower in both cases.

In graph 6.1 below depict how many auctions that had a particular number of snipe bids. The maximum amount of snipe bids in an auction is 7 (two auctions), and the minimum amount is 0 (287 auctions). As is apparent from graph 6.1, the data is skewed towards lower amounts of snipe bids.



Graph 6.1. Amount of auctions with different number of snipe bids.

6.2. The private and independent value paradigm

In our theoretical models we have assumed that the value of a good to a bidder is independent in relation to another bidders' valuation. The distinction between independent, private value goods and other types is theoretically straight forward, but often hard to apply to real products.¹⁰ As is the case here, goods can display both private and common value "components". We thus need to make a choice as to what valuation

¹⁰ Actually in the incremental model we assumed that though bidders valuations were independent of each other, they could be common information. However the distinction being made here is between private and common valuation of a good (though one bidder might be informed about the private value of another bidder).

model better suits the data. Below we present considerations as to why the data fit the private value paradigm.

Different goods can have private value if there is no prestige in owning the good or if the possibilities to resell the good at similar prices are limited (Milgrom and Weber, 1982). A Picasso painting or another "famous" piece of art, would possess both these qualities. The goods of art in our sample do not come from famous artists, so the valuation of a good is rather related to bidders' taste for different techniques, colors etc.¹¹ Assuming that bidders' tastes are idiosyncratic, the private value description seems appropriate for art auctions.

Iphones have cited a market price which is of common value to all bidders. However, the valuation of a second-hand good can differ between bidders (Ockenfels and Roth, 2005). Also, Iphones have limited possibilities of resale as new versions become available on the market and rather quickly deteriorate in value. The possibility to resell an iPhone at a price equal to what it was originally bought for should thus be limited. These properties should validate the private value framework for iPhone auctions.

6.3. Results

In this section we present the results from data analysis. Using GLS and probit regressions we test if the data fit with the theoretical results. The results from each model are presented individually below.

6.3.1. The model

The model developed in this thesis considers how the possibility of a snipe could affect the timing of a bid. If the probability of a snipe, p , and the probability of a snipe bid being successfully transmitted, a , was sufficiently high bidders always chose to snipe. More specifically if $a > (1 - p) * y(a)$ an equilibrium where bidder always snipe exist. Now, how can we infer the probability that a snipe is present in a specific auction? To determine how p might vary between auctions, we utilize the different commission rates imposed

¹¹ I base this statement from a brief survey of the different art objects in the sample and on the fact that the final price of art auctions are rather low compared to i.e. prices in a gallery or other art auctions such as Sotheby's.

on the seller. Tradera's commission rate consists of a fixed part (r) and a floating part (k) which generally decreases in the price. Fig. 7.1 below displays the commission rates charged on Tradera.

Winning bid, SEK (P)	Fixed Rate (r)	Moving rate (k)
0-3	-	-
4-200	7,5 %	3 SEK
201-1000	4 %	18 SEK
1001-3499	2 %	50 SEK
≥ 3500	-	100 SEK

Table 6.1. Commission-rates on Tradera, (<http://www.tradera.com/help/HelpPage.aspx?NodeID=6790>).

As a potential snipe bid is not informed about the highest bid in a Tradera auction there is a risk that a snipe bid is "too high" so that the snipe ends up winning the auction. The cost for the seller if he wins is determined by the commission rate that he has to pay to Tradera. We will assume that a higher commission rate decreases p . Determining a is more complicated since we have no data on the probability that a snipe bid is not successfully transmitted.¹² We will assume that in every auction in our sample, a is identical so that snipe bids always have the same probability of being transmitted in any Tradera auction. By keeping a fixed we let p vary to test if the data can confirm our theoretical predictions.

Earlier we noted that the numbers of auctions with respect to the total number of snipe bids were strongly skewed to the right (graph 6.1). When this is the case, OLS is not a very useful tool since it assumes homogeneous residual variance over the sample. Running an OLS regression would thus result in inconsistency problems due to heteroscedasticity. To account for this problem we run an OLS regression with robust standard errors. This method relaxes the assumption that the residuals are identically distributed across the sample. Though OLS with robust standard errors do not change the

¹² A survey amongst experienced bidders in Ockenfels and Roth (2005) report that 20 % of snipe bids are not transmitted on eBay. However we should not use this value as we treat a different online auction which might have different network resources.

coefficients, it provides consistent p-values compared to OLS. We run the following regression:

$$S_i = C_1 + C_2K_i + C_3B_i + C_4I \quad (7.1)$$

for every auction $i = 1, \dots, 400$. S_i is the dependent variable and is equal to the number of snipe bids in auction i . The independent variables are commission rate, K_i , the total number of bids, B_i , and a dummy for Iphone auctions, I_i (=1 if Iphone, =0 if not).¹³ Commission rate is calculated as $C_i = r_i + k_i/P_i$, (see Fig. 7.1). We add B_i as the distribution of timing of submitted bids in auction i might be correlated with S_i . From the theoretical model we expect that K_i is negative in S_i .

The result from (7.1) is depicted in the table below.

Regression output - Dependent variable S_i	
Independent variables	Coefficient estimates (robust standard errors)
C_1	0.334*** (0.0920)
K_i	-0.00206 (0.00165)
B_i	0.00224 (0.00425)
I_i (=1 if Iphone, 0 otherwise)	0.356*** (0.135)
Observations: 400	
R-squared: 0.048	
*** p<0.01, ** p<0.05, * p<0.1	

Table 6.2 The effects of a snipe –Regression output

It is unclear if there is actually any effect on sniping from different commission rates. The coefficient for K_i cannot be accepted on any conventional level of significance (p-value =

¹³ As discussed in 6.2., there could be a difference between how well the private value paradigm fit into each product category. Also there were more Iphone auctions where sniping occurred than in art auctions. The dummy-variable is included as to account for these effects.

0.21). However it exhibits the expected sign which implies that our theoretical model has some relevance. The model treat p as exogenously given and does not offer an explanation as to why it could differ between auctions. The probability of a snipe could be dependent on variables that were omitted here, in which case K_i is an insufficient proxy for p .

The only relevant coefficient that is statistically significant at any relevant level is I_i (p-value < 0.01). The coefficient for I is positive, implying that on average there are 0.36 more snipe bids in Iphone auctions than in art auctions. In light of our theoretical model that in Iphone auctions, either, or both, p and a are systematically higher. Thus it is possible to we have to drop the assumption a is identical across auctions, or that p is higher in Iphone auctions. There are generally more bidders in Iphone auctions so it seems unreasonable to assume that a would be higher. If anything, more bidders would intensify network traffic, and we would expect a to be lower. The other explanation is that the probability of a snipe is higher in Iphone auctions. As mentioned above our model does not offer guidance to any discrepancy in p between art and Iphone auctions.¹⁴

6.3.2. Incremental bidding

The theoretical model developed by Ockenfels and Roth (2005) (the authors) examined how an incremental bidder affects the timing of bids. We showed that it can be a best response for a rational bidder to snipe against incremental bidding behavior. In line with the authors we will use the ranking of bidders as a determinant of their experience. Every time a user on Tradera has completed a transaction (either as a winning bidder, or a seller) he has the possibility to leave a feedback to the counterpart of the transaction. The ranking of a user on Tradera is equal to the number of positive feedbacks minus the number of negative feedbacks.¹⁵ Though it is possible that positive and negative feedback cancel out, empirical research shows that a vast majority of user feedback is positive. "...almost all of the feedback on eBay is positive. Resnick and Zeckhauser (2001) report that only 0.6 percent of feedback comments left on eBay by buyers about sellers was negative or neutral" (Bajari and Hortacsu, 2004). As the reputation system is constructed

¹⁴ We will go outside our model in Section 7.3 to further discuss this observation.

¹⁵ <http://www.tradera.com/help/HelpPage.aspx?NodeID=1326>

in a similar way on eBay (though the rank can only vary between -1 and +1) we will assume that this observation can be applied to Tradera. Another consideration to the use of ranking as a proxy for experience is that providing feedback is voluntary action. Thus it is possible that a user actually has participated in more transactions than indicated by his rank. A third issue that arises is that user rank does not admit for the possibility that experience could have been gained in some other way. If for example a user had a previous or a different account on Tradera, then this will not be included in our proxy. There are obviously some considerations of how good user rank is as a proxy for experience. However, bidder rank is easily obtained from Tradera and should at least suffice as an imperfect proxy for experience.

We will use a probit model to test if bidders with a higher rank choose to snipe, and if bidders with a low rank submit bids early. This is done by assigning every bid that arrived in the last minute a value of 1, and all other bids a value of 0. We then use this variable as the dependent variable in the probit regression. As independent variable we use the bidder rank corresponding to every bid. Since bidders might submit several bids there is a possibility that they both submit an early bid and snipe. In such a case the rank of the particular bidder will enter the dependent variable for both 1 and 0. This should not pose any particular problem in our regression as we expect rational bidders to only submit one bid in the end of the auction, and incremental bidders to submit one bid in the beginning of the auction. It seems reasonable that if several incremental bidders compete in an auction they will engage in an early price war, where they submit several bids each (the number of bids compared to the number of bidders in 6.1 suggest that this is the case). This should, however, not create any bias in our probit model as we expect this behavior. We specify the following probit regression model

$$\Pr(S_i = 1) = C_1 + C_2 R_i \quad (7.2)$$

for every bid $i = 1, \dots, 7291$. The dependent variable S_i is binary and is equal to 1 if bid i arrived in the very last minute of an auction, and 0 otherwise. R_i is the rank of the bidder submitting bid i . We expect C_2 to be positive in $\Pr(S_i = 1)$. The results from the regression are given in Table 6.3 below.

Regression output - Dependent variable $\Pr(S_i=1)$	
Independent variables	Coefficient estimates (standard error)
C_1	-1.941*** (0.0314)
R_i	0.000155** (6.74e-05)
Observations: 7291	
*** p<0.01, ** p<0.05, * p<0.1	

Table 6.3 Incremental bidding – Regression output

As expected, R_i has a positive coefficient. We can accept this result on a 5 % level of significance, implying that a bidders rank affect the timing of a bid. It seems as if bidders with a low rank generally submit bids early in the auction, which corresponds to our interpretation of low rank bidders as inexperienced incremental bidders. Bidders with a high rank, which we interpreted as a sign of experience, choose to snipe in order to avoid an early bidding war. As previously mentioned we could not be certain that ranking was a sufficient proxy for experience. The obtained results however, imply that there is a positive correlation between experience and rank.

Our results from (7.2) is similar to that reported in Ockenfels and Roth (2005) (the authors) regarding incremental bidding on eBay. They use a somewhat different approach when estimating their probit model. Instead of considering every bid as an observation as we did here, the authors only considered the last bid of every bidder. First they check if the last bid of every bidder arrived in the last ten minutes of the auction. If so, that observation was assigned a value of 1, if not, it was assigned a value of 0. Every bidder represented one observation, where their rank was used as independent variable. The dependent, binary variable was whether the last bid of every bidder arrived within the last ten minutes of the auction. They report estimates for both all bidders' last bid, and the last bid in every auction. Similarly to our estimates they report a significant positive effect between rank and sniping (below 1 % level of significance). As the authors only consider the last bid of every bidder, their method will differ from ours the more bids that every bidder submits. If the amount of early bids per bidder increases, their method, in comparison to ours, will exaggerate the amount of sniping observations.

7. Conclusion

The purpose of this thesis was to investigate whether the presence of a shill and the presence of incremental bidders could explain why bidders snipe in Tradera auctions. To empirically test both models we collected data from 400 Tradera auctions. 200 of these were Iphone auctions, and 200 were art auctions. Though it is probable that both of these product categories partly have a common value to bidders, we provided some considerations as to why they could at least to some extent be considered to fit the private value paradigm.

7.1. The effects of a shill

To investigate the effects of a shill we developed a static Bayesian game with three bidders. In this game, bidders faced a positive probability that a shill was present, and a positive probability that a snipe bid was not successfully transmitted. These effects induced a strategic consideration where bidders faced a trade-off between bidding early or sniping. Bidding early assured that the bid was being transmitted but exposed the bidder to the risk of having his surplus reduced to zero by a shill. By sniping, bidders avoided the negative effects of a shill but instead faced a risk of their bid not being successfully transmitted. In this game we showed that when the probability of a shill and the probability of a bid not being successfully transmitted were above a certain level, an equilibrium where all bidders snipe may exist.

To apply the data to our theoretical model of the effects of a shill, we needed to determine how the probability of a shill, p , could vary between auctions. As a proxy for p we used the commission rate on Tradera. In Tradera auctions a seller choosing to submit a shill bid face a risk that he might outbid every honest bidder. The commission rate thus provides a measure of the cost of submitting a shill bid. When the commission rate is higher we should expect that p is lower. The results from our regression did not yield statistically significant estimates. However, the coefficient had the expected sign which provided some credibility to our model. The generally weak results from data analysis could be due to several reasons. The commission rate used in our regression is not completely determined until the auction is finished. As the commission rate changes

whenever there is a new highest bid, this has implications for how bidders adjust to the probability of a snipe. Thus initially when the auction begins, p would generally be higher than when the auction finishes. In a Tradera auction p changes dynamically which we do not account for in our model. Furthermore we lacked sufficient data to determine the probability that a snipe bid was successfully transmitted, a . We assumed that a was constant across auctions so that only differences in p across auctions would yield different outcomes. As there are a different number of bidders and different amounts of snipe bids in auctions it is possible that a could vary. In an auction with many bidders submitting late bids the impact on Tradera's network might be higher, thereby reducing a . A second reason that the results of our data analysis were weak could be because the nature of the snipe was not explicitly analyzed. In our model we treated the snipe exogenously and did not model the strategic considerations that he could face. We assumed that the cost of submitting a snipe bid could be determined by the commission rate, however this notion might require further inspection of the role of a snipe.

7.2. Incremental bidders

To investigate the effects of incremental bidding we used the model developed by Ockenfels and Roth (2005). The authors use a discontinuous game with one early and one late part. Due to the continuous property of the game, bidders always have the possibility to respond to another bid in the early part. In the late part of the game, bidders could only submit one more bid, and it was not possible to observe the bids made by others. In this static, late part of the game, bidders faced a positive probability that their bid was not successfully transmitted. The definition of an incremental bidder was that of a naïve bidder, that mistook a second-price auction for a first-price auction. In a second-price auction, the winner's payment is independent of his bid, whereas under a first-price rule it would be rational for a bidder to submit incremental bids as the winner pays his bid. Using the proof given by the authors we showed that when bidders face an incremental bidder they might choose to snipe. Though sniping is costly this is, most often, outweighed by the fact that any early bid will result in a bidding war, which increases the payment for the winner.

When testing whether the presence of incremental bidders would result in experienced bidders sniping we estimated a probit model. The dependent variable was 1 if a bid was submitted in the last minute of an auction and 0 otherwise. As a proxy for experience, we used the rank corresponding to each bid as independent variable. We presented some considerations as to the suitability of a bidder's rank as a signal of experience. We obtained significant estimates that were in line with our theoretical predictions. If a bid belonged to a low rank bidder then this arrived early in the auction. If on the other hand a bid belonged to a high rank bidder, it was submitted in the last minute. This suggests that a bidder's experience can be somewhat determined by his rank. The results suggest that the model developed by Ockenfels and Roth (2005) and the effects of incremental bidding within this model, can partly explain sniping on Tradera.

Though there was not time to do it here, it should be possible to augment (7.2) to account for auction specific conditions in our data. As we did not include in what particular auction a bid was being made in, it is possible that a rational bidder that participated in several auctions engaged in different strategies depending on the competition. If, for example, he participated in an auction with no incremental bidders he might choose to bid early, whereas in an auction with many incremental bidders he might choose to snipe. By re-estimating (7.2) to account for specific factors of an auction, such as the relative amount of incremental bidders, we might be able to obtain stronger results.

7.3 Difference between Iphone and art auctions

We were not able to explain the discrepancy in sniping between Iphone and art auctions. In Iphone auctions there was a statistically significant higher amount of snipe bids than in the other product category. Ockenfels and Roth (2005) shows that in a common value auction, bidders might choose to snipe to avoid revealing information about the value of the good to others. Their theoretical predictions are confirmed in their data where they report a higher amount of sniping in antique auctions than in computer auctions. This contrasts our results, suggesting that in our sample, Iphones could exhibit more common value "components" than art. Several Iphones in our data were relatively new, so that private second-hand values might only have marginal effects. On the other hand, market prices for Iphones are, unlike common value goods, public information. The price

information that a bidder could signal to other bidders by bidding early should thus be limited. A possible explanation to our observation is that bidders' participation signals "trust-worthiness" of the seller. Bidders might then choose to snipe in order to avoid signaling that the seller is reliable, which could attract other bidders and increase competition.

8. Final remarks and future research

Online auctions provide a great possibility to investigate and develop the field of auction theory. As the rules of these auctions are fixed it is possible to construct models that could closely resemble the "real world" and the strategic considerations taking place. The models presented in this thesis attempt to move towards a more complete description of online auction but there is still a lot of work to be done. Modeling online auctions in a more realistic perspective should involve dynamic, continuous models where bidders and the seller can update their information as the auction proceeds. The effects of a skill and the probability of snipe bids not being transmitted would be more realistically interpreted if they were continuously changing. Also, by explicitly modeling the considerations of a skill we could move towards a more complete description of these auctions.

Some interesting work towards some of these issues has been undertaken by Bose and Daripa (working paper, currently under review at a journal). Bose and Daripa develop a continuous auction game with private values and a hard-close rule. They analyze the timing of bid submission and bid arrival to explicitly determine when a bidder might choose to submit his bid (unlike the models in this thesis that considers if a bid arrives in the early or late part of the auction). Also, they include the strategic considerations of a skill so that his actions are treated endogenously in the model. Within this auction model they show that in equilibrium bidders choose to snipe as to avoid the possibility that a seller can successfully submit a skill bid.

The examination of online auctions poses several interesting questions for academic research. For example, it would be interesting to see how efficiency and revenue is affected by properties such as those discussed in this thesis. An experiment conducted by

Ariley, Ockenfels and Roth (2005) shows that both revenue and efficiency is higher in auctions with a soft-close compared to auctions with a hard-close rule. This imposes the interesting question as to why sites such as eBay and Tradera keep this rule. One explanation could be that in hard-close auctions strategic effects of late bidding create additional incentives for bidder participation as there exists possibilities to minimize the final payment. Another interesting extension to the one good auction would be to model online auctions as multi-good auctions with perfect and/or imperfect substitutes. A typical example would be Iphone auctions where bidders often can choose between a wide range of auctions for the same type of phone. The substitutability of one phone for another could depend on factors such as bidders' valuation of second-hand goods, other bidders' actions and the rank of the seller. Lastly, it could be fruitful to further explore the rationale behind incremental bidding. Though the definition by Ockenfels and Roth (2005) seems to translate rather well into empirical data and reality, incremental bidding could for example be a response to shill bidding. By just raising the bid incrementally the possibility for a shill to extract the winner's surplus would be limited in comparison to bidders immediately bidding their maximum willingness to pay.

Literature

Ariely D., Ockenfels A., Roth A.E., 2005, "An experimental analysis of ending rules in Internet auctions", *RAND Journal of Economics*, Vol. 36, pp. 891-908.

Bajari P., Hortacsu A., 2003, "The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions", *RAND Journal of Economic Literature*, Vol. III.II, pp. 329-355.

Bajari P., Hortacsu A., 2004, "Economic Insights from Internet Auctions", *Journal of Economic Literature*, Vol. XLII, pp. 457-486.

Barbaro S., Bracht B., 2006, "Shilling, Squeezing, Sniping: Explaining bidding in online second-price auctions", Working paper.

Beckman M., "Why do sellers at auctions bid for their own goods – Theory and evidence", German economic association of business administration, Discussion paper no. 03-11.

Bhargava B., Jenamani M., Zhing Y., 2005, "Cheating in online auction – Towards explaining the popularity of English auction", *Electronic commerce research and applications* 6, pp. 53-62.

Bose S., Daripa A., "Shills and Snipes", working paper (currently under review at a journal)

Chakraborty I., Kosmopoulou G., 2003, "Auctions with shill bidding", *Economic theory* 24, pp. 271-287.

Harstad R., Rothkopf M., 1995, "Two models of bid-taker cheating in Vickrey auctions", *The Journal of Business*, Vol. 68, pp. 257-267.

Hlasny V., 2007, "Shill-bidding in private values auctions", *Journal of information, communication and ethics in society*, Vol. 5, pp. 307-320.

Kaufmann R., Wood C., 2003, "Running up the bid: detecting, predicting, and preventing reserve price shilling in online auctions", Working paper.

Milgrom P., Weber R., 1982, "A theory of auctions and competitive bidding", *Econometrica*, Vol. 50, pp. 1089-1122.

Ockenfels A., Roth A.E., 2005, "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction", *Games and Economic Behavior*, No. 55, pp. 297-320.

Porter R., Shoham Y., 2005, "On cheating in sealed-bid auctions", *Journal of decision support systems: the fourth ACM conference on electronic commerce*, Vol. 39.

Rasmusen, E.B., 2006, "Strategic Implications of Uncertainty over One's Own Private Value in Auctions", *Advances in Theoretical Economics*, Vol. 6, Art. 7.

Riley J.G., Samuelson W.F., 1981, "Optimal auctions", *The American Economic Review*, Vol. 71, No. 3, pp. 381-392.

Vickrey W., 1961, "Counterspeculation, Auctions, and Competitive Sealed Tenders", *Journal of Finance*, Vol. 16, pp. 8-37.

Other sources

Tradera (www.tradera.se)

Mas-Colell A., Green J.R., Whinston M.D, "Microeconomic Theory", 1995, Oxford University Press

Lecture notes, "NEKM31 Microeconomics – Theory for Individual choice and Game Theory", Hans Carlsson/Alexander Reffgen, Department of Economics, Lund University.

Appendix I

For every pay-off cell in Fig. 4.2, we must consider every possible state of nature. The distribution function $F(T_1, T_2, T_3)$ determine the particular outcome. Since each player i can observe two different signals $T_i = \{1,2\}$ there are eight possible outcomes $((1,1,1), (1,1,2), \dots, (2,2,2))$ which we will index by $h = 1, \dots, 8$. Each of one of these outcomes occur with an equal probability of $\frac{1}{8}$.

Below we calculate the pay-off for player 1 for every cell in Fig. 4.2.

(S, S, S) .

h	Pay-off in state h
1	$a(1 - a)^2$
2	$a(1 - a)^2$
3	$a(1 - a)^2$
4	$a(1 - a)^2$
5	$2a(1 - a)^2 + 2a^2(1 - a)$
6	$2a(1 - a)^2 + a^2(1 - a)$
7	$2a(1 - a)^2 + a^2(1 - a)$
8	$a(1 - a)^2$
Expected pay-off	$\frac{1}{8} [11a(1 - a)^2 + 4a^2(1 - a)] =$ $\frac{1}{8} a [11(1 - a)^2 + 4a(1 - a)]$

(S, D, S) .

h	Pay-off in state h
1	0
2	0
3	0

4	0
5	a
6	$a(1 - a)$
7	0
8	0
Expected pay-off	$\frac{1}{8} [6 * 0 + a(1 - a) + a] =$ $\frac{1}{8} [a(1 - a) + a]$

From the perspective of player 1, the pay-off in cell (S, D, S) is identical to the pay-off in cell (S, S, D) .

(S, D, D)

h	Pay-off in state h
1	0
2	0
3	0
4	0
5	a
6	0
7	0
8	0
Expected pay-off	$\frac{1}{8} [a]$

(D, S, S)

h	Pay-off in state h
1	$(1 - a)^2(1 - p)$
2	$(1 - a)^2(1 - p)$
3	$(1 - a)^2(1 - p)$

4	$(1 - a)^2(1 - p)$
5	$2(1 - a)^2(1 - p) + 2a(1 - a)(1 - p)$
6	$2(1 - a)^2(1 - p) + a(1 - a)(1 - p)$
7	$2(1 - a)^2(1 - p) + a(1 - a)(1 - p)$
8	$2(1 - a)^2(1 - p)$
Expected pay-off	$\frac{1}{8} [12(1 - a)^2(1 - p) + 4a(1 - a)(1 - p)] =$ $\frac{1}{8} (1 - p) [12(1 - a)^2 + 4a(1 - a)]$

(D, S, D)

h	Pay-off in state h
1	0
2	0
3	0
4	0
5	$(1 - p)$
6	0
7	$(1 - a)(1 - p)$
8	0
Expected pay-off	$\frac{1}{8} [(1 - a)(1 - p) + (1 - p)]$

From the perspective of player 1, the pay-off in (D, S, D) is equivalent to the pay-off in (D, D, S) .

(D, D, D)

h	Pay-off in state h
1	0
2	0
3	0
4	0
5	$(1 - p)$
6	0
7	0
8	0
Expected pay-off	$\frac{1}{8}[(1 - p)]$