Experimental and theoretical essays in auctions and financial markets

Inaugural Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Sotirios Georganas

aus Athen, Griechenland

Bonn, 2010

Dekan: Professor Dr. Erik Theissen
Erstreferent: Professor Dr. Reinhard Selten
Zweitreferent: Professor Dr. Rosemarie Nagel
Mündliche Prüfung: 13.03.2009

στους γονεις μου (και χρηματοδοτες μου)

Abstract

This dissertation deals with auctions and other, financial market mechanisms, and how they work in real life situations. Such mechanisms have been extensively studied in economics and many elegant results have been gained. These results have been also used in the design of real life auctions, competitions and financial markets. In this dissertation, I show that there are relevant facts that should influence mechanism design for practical purposes, but are omitted in the main corpus of the literature.

One of these facts, explored in the first chapter of the dissertation, based on a single authored manuscript of this author, deals with resale possibilities. Whenever a durable good is sold, there is always a possibility that it will be resold. Even when such resale is forbidden, many imaginative ways have been found to overcome the restrictions. The possibility of resale alters many standard results regarding the efficiency of the commonly used mechanisms and their revenue ranking. While the importance of resale has been recognized early (see Milgrom 1987) there is a relative dearth of models including resale until quite recently (see the work of Philip Haile 1999, 2001, 2003). This is probably due to the widespread belief that resale possibilities are adequately covered by models with common values. However Haile (1999, 2001, 2003) has shown theoretically that this is often not true. The first chapter of the dissertation shows that resale also matters in the laboratory. Two treatments were designed and tested for this purpose. In both there is an English auction where bidders have private values. The initial auction is followed by a second stage where the winner is allowed to resell the good to the other players in an English auction with a reserve price. In one treatment, the winner can only base his reserve price on the signals the other players sent through their bids. In the other treatment

however, all private values become public after the initial auction and thus the reserve price is set equal to the highest private value among all players. In both treatments, bidding your value was an equilibrium strategy, exactly as in a simple English auction without resale. The results of the experiments show that resale possibilities, do indeed influence bidding even when the theory predicts no change. We attribute this deviation to the phenomenon of noisy behavior and employ a range of bounded rationality models like a QRE and a levels of reasoning model to formally test the hypothesis. Additionally we find that the exact structure of the resale market matters, as behavior in the two resale treatments differed significantly.

In the second chapter of the dissertation, based on a paper coauthored with Rosemarie Nagel, another kind of omitted characteristic of many auctions is treated, namely toeholds. In many auctions one or many of the potential bidders already own a part of the asset that is being sold. In these cases the equilibrium strategies can be very different. In Bulow et al. (1999) the authors predict that even very small toeholds can have an explosive effect. Strong players, ie the ones with the highest toehold, bid in equilibrium much more aggressively than in the simple case without toeholds and weak players bid much less aggressively. We decided again to test this prediction in the lab and failed to find any signs of an explosive equilibrium. However, bidders behave differently with respect to normal auctions without toeholds and the size of the toeholds matters. We attribute the failure of the Nash prediction to the extreme nature of this equilibrium. Payoff functions are unusually flat, in effect bidders can deviate as much as 50% from the equilibrium strategy with only minimal losses in payoffs. Such an equilibrium is highly implausible and unstable; if bidders have even small biases towards some particular action, in many cases there is no sufficient force, in the form of pecuniary incentives, that will move them towards equilibrium.

The results of these first two chapters not only advance knowledge in the application of the auctions and mechanism design theory. The behavioral results should add to the predictive power of game theoretic models in general. It is a widely applicable observation that in games where the equilibrium payoffs are flat, real players do not necessarily play strategies predicted by theory. Additionally, the papers focus on the behavioral effect of noisy behavior and the anticipation thereof. If a small amount of noise changes the players' best responses drastically, they will tend to deviate from equilibrium in order to insure themselves from unexpected behavior on behalf of their opponents. In most real life situations such noise is to be expected. It may also be a kind of experimentation, players try a range of strategies before they settle down. Moreover, noisy communication or false interpretation of instructions -when players interact through agents- can also have an effect. A final source for such noise can be just other considerations that are treated as exogenous in these models, reasons for which players might want to deviate from equilibrium such as focal points, altruism, fairness etc

The third chapter is based on joint work with Michael Zaehringer. It proposes a mechanism to sell an asset when there are many small players who have information about the private values of other players. In concrete terms, we think of the case where there is a seller of a large asset (say a company) and some strategic buyers who want to acquire the control of the asset. We divide the value of the asset in a common value part, that stems from the cash flow and a private value part that stems from private benefits of control. The first part is common knowledge. The private values however are not known to the seller, but are known to some small informed speculators, for example investment banks. The seller's problem is then to get this information in order to appropriate more rents from the potential buyer.

A mechanism that is often proposed when selling an asset to a few strategic buyers is an auction. However, such an auction will fail to use the information of the informed speculators, who are too small to participate. A mechanism with better information aggregation properties is a sale of shares of this company through an IPO. The downside of an IPO is that it removes power from the hands of the seller. The chapter proposes a two stage mechanism that combines the strengths of these two mechanisms. In the first stage a small part of the shares is sold through an IPO and a price is formed that, as we shall see, aggregates the information of the small speculators. The seller then uses this information to design an optimal auction in the second stage and maximize his revenue when selling the rest of the asset (including the controlling rights) to the large strategic investor.

A key detail that makes the two stage mechanism work is minority shareholder protection, in the form of the sell-out rule. This rule specifies minority shareholders can force majority shareholders to buy their shares for a fair price. Because of this rule, the speculators have an incentive to buy shares in the first stage and reveal their information. Thus the paper shows that a previously unstudied result of minority shareholder protection is that speculators are induced to participate and thus add to the informational content of market prices.

Acknowledgments

I would like to thank Reinhard Selten, Rosemarie Nagel, John Kagel and Benny Moldovanu for their support and fruitful conversations. I would also like to thank Charles Holt, Vincent Crawford, Thomas Palfrey, Jacob Goeree, Philip Haile and Paul Heidhues for useful comments. Finally I thank everyone at the BGSE, in particular Marc, Matthias, Michal, Michael, Lars, Steffen and Juergen, for being very capable peers in scientific times, supportive friends in the hard times and excellent buddies in the, not frequent enough, happy times. May every graduate school year have as nice a group as the one we had.

Contents

List of Tables

List of Figures

Chapter 1

English Auctions with Resale

1.1 Introduction

Auctions are very often followed by a resale opportunity. For instance, after virtually every durable good auction, the winner can choose to resell the good to the competing bidders or other third parties. Even when resale is explicitly prohibited, ways can be found to get around the prohibition. Consider the case of mobile phone and wireless spectrum licences where often "use-it-or-lose-it" conditions to prevent resale are imposed. Still these restrictions can be circumvented. The company holding the licence can be bought and there have even been cases where special companies were used to buy and resell the licence.

Until recently resale was not considered in auction models. However the work of Haile (1999, 2000, 2001, 2003) has shown that the existence of resale opportunities is important and can change many results of auction theory. In order for resale to be meaningful, the outcome of the initial auction must be inefficient with a positive probability. In Haile's paper (2003) the highest value player does not necessarily win in the initial auction because he does not perfectly know his value or because he is not participating. Other ways to induce resale in equilibrium can be asymmetries in the values of the bidders (Hafalir and Krishna 2008, Garratt and Troeger 2006) or new participants arriving in the resale stage (Haile 1999).

In this paper I test auctions with resale in the laboratory based on a simple model from Haile (1999) and find an alternative reason for resale that has not yet been considered in the abstract theoretical models but is plausible when human players are involved, namely noisy decision making. People make mistakes and anticipate others to make mistakes. This can lead players to deviate from theoretical predictions in systematic ways. To see how it can induce resale, consider the following example.

Suppose you are participating in an English auction for works of art, regarded often to be a textbook example of private values, in the absence of resale possibilities. Suppose that the current price for the Picasso painting under sale is 10 euros. Even if you do not have a taste for cubism and thus your private value is zero you might still want to participate in the auction, expecting to resell the painting for a higher price. Thus resale introduces a common value element to the valuations of bidders and can induce overbidding. We could use a similar argument in the markets for real estate, bonds, operating licences and more.

Thinking in line with standard models one could note that in such a simple setting, bidding one's value is still a symmetric equilibrium. A strategy of overbidding expecting to resell is not consistent with this equilibrium. If others bid their values, no profitable resale is possible. Thus winning with a bid higher than your value can only result in zero or negative payoffs. Crucially however, this is only true if bidders never make mistakes, as is usually assumed.

From our experience in the lab and the field, we know humans are prone to making mistakes. Expecting high value bidders to make mistakes can make it optimal for low value bidders to bid more than their values. This in response can lead to high value bidders bidding less and expecting to buy cheaper in the resale stage. Thus resale opportunities can be exploited even if standard theory predicts they will not and they can give rise to richer bidding strategies than theoretically expected. Let it be emphasized that this deviation from standard models is quite natural. There is no need for restrictive assumptions on the structure of markets or the private information of bidders to induce resale in real life situations. A sufficient condition, as will be shown, is the presence of a small amount of noise. Such noise exists in many markets, even in financial markets where stakes are very high (see Shleifer and Summers 1990). It can stem from experimentation, lack of experience or misunderstanding of the rules, false transmission of information and mistakes in the execution of orders, liquidity constraints or other exogenous reasons that are not adequately modeled in theory but whose presence in real markets cannot be easily dismissed.

To examine the importance of resale opportunities and the effect of noise in a controlled environment, I designed and ran two experimental treatments of English auctions with resale, with different informational backgrounds but with the same equilibrium bidding functions prescribing that players bid their values. Subjects exhibited significantly different behavior with respect to both the theory and previous auction experiments without resale. Instead of bidding their values in both treatments, they overbid relative to equilibrium when they can be certain they can reap all the rents in the resale markets, and they tend to underbid when the resale outcomes are uncertain. Moreover this result should not be attributed simply to irrational behavior in the laboratory, but seems to have a reasonable explanation. Subjects do try to maximize their profits. But while doing so, they anticipate the possibility of others making mistakes and they use this knowledge more or less optimally. In that sense, this paper presents a previously unstudied example of a more general class of games where the anticipation of noise drastically changes players' best responses[1]. In such cases, standard game theory loses much of its predictive power and concepts of bounded rationality, such as a Quantal Response Equilibrium (McKelvey and Palfrey 1995) and levels of reasoning (Nagel 1995, Stahl 1995, Camerer 2004, Crawford and Iriberri 2008), perform much better.

The experimental economics literature has not focused on auctions with resale yet, for the same reasons that there were precious few theoretical models of resale until recently. To my knowledge there exist three other experimental papers on auctions with resale. Two of these analyze symmetric auctions, Georganas (2003) on which part of the present paper is based and recent independent work by Lange, List and Price (2004). Their experimental treatments are similar to the ones in this paper, but they differ in important ways: first they used first-price sealed-bid auctions and

[1] Games with this property include the guessing game, the centipede game and the traveler's dilemma.

secondly they gave players noisy signals about their private values. They found deviations from equilibrium predictions, which they attribute to risk aversion. However risk aversion alone does not change the equilibrium in our games, so it cannot explain the data in the present study. Subsequent to the present paper, Georganas and Kagel (2008) analyze asymmetric first price auctions with resale and find support for the equilibrium that predicts weak players bidding more aggressively than without resale.

Even though the possibility of resale and its potential importance has been recognized in the theoretical literature (Milgrom and Weber, 1982 and Milgrom 1987 with the first models of auctions with resale) there has been a striking absence of formal models featuring resale until recently. A frequent argument has been that resale is covered by the assumption of common values. However, as shown in Haile (2003) players in the initial auction have common values when there is a possibility of resale but, importantly, valuations are endogenously determined and equilibrium strategies are not the same as in the simple common value model. Moreover Revenue Equivalence holds under some assumptions although it does not in the CV case. In Haile (2003) bidders have noisy signals about their values in the initial auction. Noisy signals work in a similar way as the noisy bids in this paper, as they lead to inefficient outcomes and profitable resale.

There are of course other possibilities to make resale potentially profitable. Haile (1999) assumes that an a priori known number of bidders is added to the bidder pool in the second period. These new subjects arriving in the resale auction can have higher private values than the winner of the first auction, opening up resale possibilities.

Alternatively, one can construct models with asymmetric equilibria. Intuitively, resale seems plausible in such an equilibrium, as the players will have asymmetric strategies and thus the highest value player will not necessarily be the highest bidder in the initial auction. This option is explored in Garratt, Troeger (2006). In a setup similar to ours they include speculators with zero valuations and find asymmetric equilibria where the speculator wins with positive probability. Gupta and Lebrun (1999) and Hafalir and Krishna (2008), on which the aforementioned Georganas and Kagel paper is based, have bidders with potentially positive use values, which however are asymmetrically distributed. This setup also gives rise to inefficient outcomes and subsequent resale.

Finally some other models including some flavor of resale are Ausubel and Cramton (1999), McAfee (1998) and Jehiel and Moldovanu (1999), although their setups are not directly related to the one in this paper.

This paper is structured as follows. The experimental procedure is introduced in section 1.2. Section 1.3 presents the equilibrium predictions and the results are presented in section 1.4. Models of bounded rationality involving some flavor of noisy decision making are presented in section 1.5. Ideas for future work are discussed in section 1.6 and section 1.7 concludes.

1.2 Experimental design

There are two stages in the game. In the first stage four bidders $i = 1, 2, 3, 4$ bid in an English auction[2] for one unit of an indivisible object. Each bidder has a use value v_i, which is identically independently drawn from a discrete uniform distribution with support [0,100]. The distribution of the use values is common knowledge, but the actual use values are private knowledge. We have to emphasize the distinction between a bidder's *use value*, i.e. the value a bidder places on owning the object ignoring any resale possibilities and the bidder's *valuation*, which is the value she places on winning the auction and which is determined endogenously, taking account of the resale opportunity.

For the auction we use an ascending clock design (see Kagel et al. 1987). There is a clock on each computer screen, starting simultaneously from zero and synchronously rising every second in steps of one unit. Each subject can exit the auction at any time by pressing a button. Once out of the auction no reentry is possible. The other bidders can observe the price at which one exits. After three bidders have left the auction, the last remaining bidder obtains the good and has to pay the price p_1 at

[2] To be more specific we use the Japanese variant of the English auction (see Milgrom and Weber 1982). There exist of course other, quite different variants which are not so practical to implement in a laboratory experiment.

which the last one left. This concludes the first stage.

Observe that the use of the English auction does not allow us to observe the intended bid of the winner, but only a lower bound. One could possibly argue that a second-price sealed-bid auction in the first stage would suit our purpose better. With this configuration the unobservable final bid problem is avoided. However behavior in sealed bid auctions usually presents large deviations from equilibrium, even without resale. Thus it is not sure one can separate the effect of resale from the other factors which are pushing behavior away from equilibrium[3]. On the other hand the English auction is widely studied and subjects seem to understand the Nash equilibrium and follow the predicted strategies quite closely.

In the second stage there is the possibility of resale. This is done through an English auction, where the seller can choose a reservation price. The difference between the two resale treatments, lies in the informational background of the second stage. As discussed there are many ways to model the resale stage. We chose two extremes with a big span between them, to test for a wide range of possibilities. In the first, incomplete information treatment (hence INC), the only information the bidders get about the others' values is through the bids in the initial auction. The seller decides about the reservation price r and then the remaining bidders can see the reserve price and decide simultaneously if they want to participate in the resale auction or not[4]. If no bidder chooses to participate then the ownership of the good and the payoffs remain the same as in the first stage. If only one bidder participates, then she obtains the good and pays the reservation price to the owner. Thus the final payoffs are $r - p_1$ for the first stage winner and $v_i - r$ for the second stage winner. If more than one bidder decides to participate we have an English auction like in the first stage, with the difference that this time the clock starts at the reserve price. Again when only one bidder remains, she obtains the good and has to pay the price p_2 where the last

[3] For results in sealed bid experiments see for example Kagel et al. (1987). Despite the mentioned problems, testing a second price sealed bid auction remains an interesting idea for the future. Additionally these experiments would afford a test of the theoretical result of revenue equivalence of the English and sealed-bid auctions under complete information in the resale stage (for a proof of revenue equivalence see Haile 2003).

[4] We did not give the seller the explicit choice not to put the good up for resale, however sellers were advised to set a reservation price of 100 if they did not want to resell.

bidder left the auction. The following payoffs are then communicated to the subjects: $p_2 - p_1$ for the first stage winner, $v_i - p_2$ for the winner in the second stage and zero for the others. In the same screen they can see the price of the initial auction, the reserve price, the number of participants and the price in the resale auction (zero if there was no resale auction), the highest private value and information about past periods. The information feedback was so rich in order to facilitate learning, as otherwise bidders would be getting too few experiences of winning and thus learning chances. Note that subjects win on average only 1/4 of the time, which means in the 30 periods they only get to win 7 or 8 times.

In the second treatment with complete information (COMP), after the first stage bidders get to know the use values of the others as in Gupta and Lebrun (1999). Thus in equilibrium they will ask for a reservation price equal to the highest private value. This amounts to a take-it-or-leave it offer to the person with the highest private value equal to his private value. We know however that subjects in experiments very often deviate from the equilibrium in the direction of a 50-50 split of the surplus, probably because of fairness considerations. As it is not the subject of this paper to treat bargaining games[5] we decided to avoid this problem by forcing the winner of the first auction to automatically resell the object in the second stage to the bidder with the highest value. She then received as payoff the highest private value minus the price she paid in the first auction. The rest of the players, including the person who obtains the good after resale, have a payoff of zero. After each auction players can see their payoffs, private values, the auction price, the private value of the winner, the highest private value and information about past periods.

A last treatment we ran is a standard English auction (ENG), with IPV drawn from a uniform distribution [0,100]. The experimental mechanism was in all other respects equal to the one used in COMP and INC, so we can use this as a reference treatment to check the robustness of our results. Thus, we will show that our unexpected results in COMP and INC are not due to some kind of framing effect. ENG will not be discussed on its own, but only in comparison with the other treatments,

[5]For a good review of this problem in bargaining games see Chapter 4, specially pages 258-274 in the Handbook of Experimental Economics by Kagel and Roth (1995).

Session	Treatment	Exchange rate	Paying Periods	Players	Location
1	COMP	20	30	16	UPF
2	COMP	20	30	16	UPF
3	INC	25	30	16	UPF
4	INC	25	35	16	UPF
5	ENG	20	30	16	UPF
6	COMP	20	30	16	Bonn
7	COMP	20	30	16	Bonn
8	INC	25	30	16	Bonn
9	INC	25	30	16	Bonn
10	ENG	20	30	16	Bonn

Table 1.1: Summary of sessions

as simple English auctions have been extensively discussed in the literature (see for example Kagel et al. 1987) and our results are quite similar to these previous studies.

We conducted 10 experimental sessions, with 16 participants in each. For the first 5 sessions, subjects were undergraduate students, mainly from the faculties of law and economics, at the Universitat Pompeu Fabra in Barcelona. The next five sessions were conducted at the University of Bonn, with subjects from many faculties. The analysis finds no consistent differences[6] between the two groups, so the data are pooled together.

At the beginning of the experiment the participants were divided in two subgroups[7] of 8 and then the players in every subgroup were randomly rematched every period in groups of 4. Subjects did not know what group they have been assigned to or who are the other members of the group. There were 31 periods in almost[8] every experimental session. The first period was a practice period that did not count for the players' payoffs and was not used in the statistical analysis of the data. After this period subjects received an initial capital of 150 units of our experimental currency,

[6] A Mann Whitney U test could not reject the hypothesis that behaviour was the same at the 0.05 level.

[7] This was done for statistical purposes, in order to have two independent observations in every session. Still the subjects did not know this and they thought they were being rematched with another 15 players. So the probability they will try to induce cooperating behaviour and the interperiod effect should remain small.

[8] There was one session with two practice periods, but they did not seem necessary so subsequent sessions had only one. It does not matter for the analysis, as we always use observations after the 9th paying period.

the drachma. In the following periods subjects were rewarded according to their success and their profits or losses were added to the initial capital. Despite the sometimes quite aggressive bidding, there were no bankruptcies, although two subjects came close. After the end of the game the experimental currency was transformed to euros in a ratio of 25 drachmas per euro in COMP and 20 drachmas per euro in INC. The reason for this difference is that INC is more complicated. Sessions lasted about thirty minutes longer than COMP and we wished to keep average profits per hour constant. Thus average profit in COMP was 10.56 euro and in INC 15.5 euro. Naturally, this difference is not only due to the different exchange rate but due to the different bidding behavior too.

1.3 Equilibrium predictions

1.3.1 Complete information - COMP

In this section I compute the symmetric equilibrium of COMP. An important difference with respect to usual auction models, is that in the presence of resale, players have a value v_i (their exogenous private value) for the good and a valuation u_i which is something else, the value she places on winning the auction. The valuation is determined endogenously, as it depends on the outcome in the resale market too. Use values v_i are drawn from a discrete uniform distribution with support $[0, 100]$. Consider the two-stage game COMP played by 4 risk-neutral players for a single indivisible object as described above. Let $y_1^i = \max\{v_j | j \neq i\}$, $i, j \in \{1, 2, 3, 4\}$ denote the highest use value among a given bidder's opponents and let v_{-i} denote the vector of the use values of all players, except i. Let f denote the final price of the game, which is equal to

$$f = \begin{cases} p_2 & \text{, if there was a resale auction} \\ r & \text{, if exactly one bidder participates in the resale auction} \\ p_1 & \text{, if no bidder participates in the resale auction} \end{cases}$$

We have following result for bidding in the first stage, similar to Theorem 1 in Haile (1999).

Proposition 1 *The symmetric bid your value equilibrium for an English auction without resale is also a Perfect Bayesian Equilibrium bidding strategy when the same auction is followed by a resale opportunity, where the private values of the bidders are publicly announced.*

Proof. Suppose bidder i with use value v_i deviates to a bid $\tilde{v} > v_i$, while all other bidders follow the proposed equilibrium strategy and bid their use values. This would change i's payoff only in the event that $\tilde{v} > y_1^i > v_i$. However if this is the case, i would have to pay y_1^i for the object but could only resell it for some price p_2 in the interval $[r, y_1^i]$. In equilibrium the reseller will set $r = y_1^i$ under complete information in the resale market, but this still leaves him with nonpositive expected profit. By bidding v_i, i would have received zero profit with certainty. A similar argument shows that bidder i would not profit by bidding less than v_i. ■

This proposition provides the risk-neutral symmetric Nash equilibrium[9] under complete information in the resale market, but as we shall see later the theorem remains valid under risk aversion and incomplete information. Therefore we will refer to this equilibrium as symmetric[10] Nash equilibrium (SNE). Also, this equilibrium covers the special case of the automated resale market that we actually used in the lab.

A characteristic of the equilibrium that should be noted is, that unlike the simple English auction, bidding your value is not a weakly dominant strategy in the presence of a resale possibility. If the person with the highest use value were to deviate from equilibrium and bid less than his value, it is clear that the best response for the others would be to bid up to this highest value (see Section 1.5 for an extensive discussion of this).

[9] Note that we treat the game as a second price sealed bid auction. The equilibrium we find is the equilibrium for the last stage of the English auction too, where only two bidders remain, as the previous 2 exits do not carry any important information that alters the players' strategies.

[10] It is worth noting here that in discrete value English auctions there exists an asymmetric equilibrium too where one player bids his value and the other bids her value minus one increment. In the non payoff relevant questions we asked after our experiments, some subjects actually reported playing such a strategy! However a few players emplying such a strategy does not significantly influence our analysis.

1.3.2 Incomplete Information - INC

In treatment two the theoretical prediction is the same as in COMP. The argumentation is similar to the one above. The only difference is that since private values do not become common knowledge in the resale stage, the reserve price has to be calculated in a more complicated way using the information that the signals (bids) in the initial auction have given us. However, independently of these signals the reserve price has to be higher than the private value of the first stage winner. This makes sure that in equilibrium we do not have resale and thus bidding one's value remains an equilibrium strategy in the initial auction. More formally:

Proposition 2 *Let bidders in INC have following pure strategy: i) In the first stage player i bids her value, ie* $b_i = v_i$ *ii) if i wins in first stage she sets a reserve price* $r_i \in [v_i, 100)$ *iii) if i loses in first stage she participates in the second iff* $v_i \geq r$ *iv) bidders in second stage bid their values, ie* $b_i = v_i$ *This is a continuum of strategies which constitute an equilibrium.*

Proof. The proof is similar to Proposition 1 and is omitted. ■

1.4 Experimental Results

In the following we present the general results and in the subsequent sections we offer explanations for the data. When making the statistical analysis of the results we will start with period 10 unless otherwise stated. The reason for this is that in almost all sessions there was a small number of subjects who indicated problems understanding the game, up to period 9 in the worst of these cases[11]. Additionally, this assumption of learning taking place before the 10th period is confirmed by the data.

The main question we are posing, is if resale possibilities alters behavior in auctions. The answer from our data seems to be a definite *yes*. In Figure 1 we compare

[11] See also Kagel et al (1987), p. 1286 where the authors claim "subjects' adjusting to experimental conditions argue for throwing out the first three auction periods" or Fehr/Schmidt (1999) who only use last period values.

the bidding in the three treatments, using simple boxplots[12] which include all but the winning bids[13].

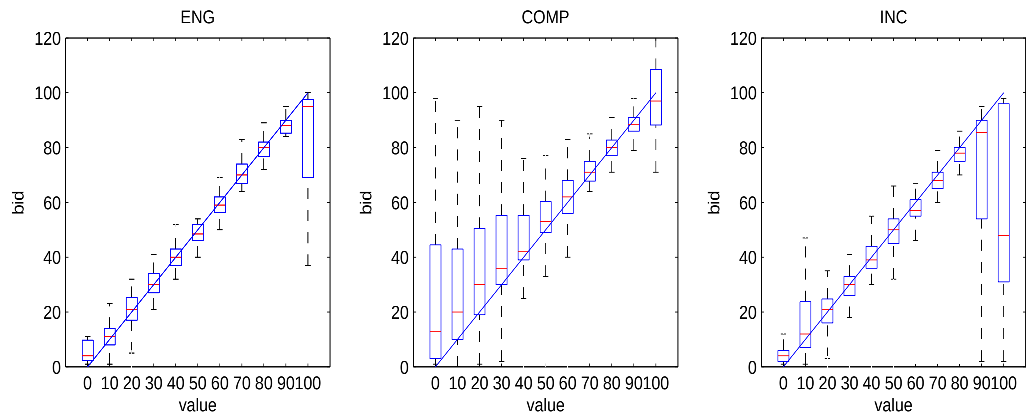

Figure 1.1: Series of boxplots of private values vs exits in the various treatments. Each box drawn represents the distribution of the bids for a block of values. The length of the box represents the interquartile range (IQR). The whisker extends from the box to the most extreme data value whithin 1,5 times the IQR.

We can see that although the three games have equilibria with the same bidding functions in the first stage, actual bidding behavior is quite different. The mere presence of a resale market makes subjects deviate much more from the equilibrium than in the simple English auction. The deviation is confirmed in Table 2, where we see that bidding in the simple English auction is significantly different to behavior in the treatments with resale (COMP, INC)[14]. This result shows that studies of auctions should take resale possibilities explicitly into account.

[12] More detailed ones follow in the individual analysis of each treatment.

[13] Keep in mind that the exit price of the last bidder is not equal to the maximum bid he was prepared to make, because he exits automatically once the last-but-one bidder exits the auction. As a consequence we only have a lower bound on the actual bidding strategy of these players. For the graphs and other statistics we exclude the winning bids. Although it leads to some bias, including them leads to an even higher bias. Techniques such as censored regressions do not completely eliminate this problem and they introduce new ones, e.g. they would rely on the restrictive assumptions that bidding is symmetric and follows some particular functional form.

[14] Bidding in COMP is significantly different from the bidding in the treatment ENG (without resale). Comparing INC with ENG we do not always have statistic significance. This can be attributed to some reasons. First, we do not have many observations for the simple English auction. Most previous experiments however, have found bidding which is very close to the bid-your-value equlibrium, and including these experiments we would get a significant difference between ENG and

	Values				
Treatments	0-20	21-40	41-60	61-80	81-100
COMP - SNE	-20.23	-13.80	–5.69	-1.12	1.91
	(0.000)	(0.000)	(0.006)	(0.083)	(1)
INC - SNE	-4.47	-0.99	2.16	7.97	20.25
	(0.000)	(0.404)	(0.0404)	(0.000)	(0.000)
ENG - SNE	-8.03	-2.70	-0.27	-0.40	5.73
	(0.004)	(0.004)	(0.150)	(1)	(0.000)
COMP - INC	-15.75	-12.81	-7.85	-9.09	-18.33
	(0.003)	(0.001)	(0.015)	(0.003)	(0.008)
COMP - ENG	-12.20	-11.10	-5.42	–0.71	–3.81
	(0.072)	(0.008)	(0.072)	(0.808)	(0.153)
INC - ENG	3.55	1.71	2.43	8,38	14.52
	(0.682)	(0.153)	(0.808)	(0.004)	(0.153)

Table 1.2: Differences in average deviations (private values minus bids), calculated excluding the censored observations. The numbers in parentheses are the p values of a Mann Whitney U test.

This is not the only interesting result. The specific structure of the resale market makes a difference for the bidding strategies. We see in Figure 1 that in COMP, when subjects have common knowledge of the private values before the second stage, resale gives the low value types an incentive to overbid.[15] On the other hand, low value types bid close to their values in INC and ENG. High value types bid close to their values in COMP and ENG but not in INC. In general, bids in COMP are highly significantly different from bids in INC for all possible values.

Note that the underbidding of the high types in INC is not a spurious phenomenon due to the censoring of winning bids or to a presence of extreme observations driving the average down. A Mann Whitney U test shows that the percentage of bidders with use values greater 50 who bid less than 20 in INC are with 5.4% highly significantly more than the 1.2% in COMP. Comparing high value bidders who bid less than 50 we get an even higher difference with 17.97% versus 6.23% respectively.

INC. The second problem is that for high values, where bidding differs most from the equilibrium , in INC, we have a strong problem of unobservable bids. We tried to control for this by running a censored regression of bids on values including data from ENG and INC. The dummy and interaction terms were highly significant, which supports the hypothesis that bidding in ENG and INC was indeed different.

[15] In this paper we use the term overbidding/underbidding loosely, to describe bids higher/lower than a subject's use value, even when such bids are not necessarily irrational.

Theoretically the only difference between the two treatments was in the informational background of the resale stage. Naturally it is possible that differences in the bidding strategies of the subjects are not only due to the theoretical difference, but also due to the different mechanisms used in practice. In particular there is evidence from bargaining games where subjects do not behave "rationally" and split the surplus in ways that do not follow the Nash prediction. In COMP we did not allow subjects to deviate, enforcing on them exogenously the predicted outcome of the second stage. In INC this was not possible and as a consequence subjects were allowed to play the resale game themselves. This difference in the mechanisms used could be a problem, however the data about the rationality in the choice of reserve prices and in the choice of participation presented in Section 4.2 indicates that subject's behavior was fairly competitive. Any deviations from optimal behavior in the resale stage are probably not due to fairness concerns but can be attributed to other effects[16].

The difference in strategies between treatments translates into different prices in the auction. As we can see in Table 3, average prices in COMP were almost 18% higher than in INC and 15% higher than in ENG, whereas the average private values were very similar, as happened with the average equilibrium prices too. This difference is not only highly significant[17] but also quite large and economically important. Observe that the average highest value in every auction was ca. 80 so revenues in COMP were almost halfway between the Nash prediction and the maximum rents the seller could possibly appropriate. Prices in ENG were slightly higher than in INC but the difference is not so big and not significant. It should be noted that prices in both are a bit lower than the predicted ones though[18].

In the following we will analyze these results in more depth individually for every treatment and we will compare some behavioral models that could explain them.

[16]For instance, there is evidence that subjects cannot calculate difficult equilibria. As a matter of fact setting an optimal reserve price given your beliefs is a fairly complex task even for theorists.

[17]P values of a Mann Whitney U test using the 8 independent observations of COMP and ING and the four of ENG are 0.002 for the INC-COMP comparison and 0.008 in the case of COMP vs ENG.

[18]Our results regarding the simple English auction should be received a bit carefully as we did not have many observations. We did not run many experiments, as there exists already a very large literature on simple English auctions. Thus for our conclusions and comparison purposes we will use these results too.

	COMP	INC	ENG
Average Observed Price	67.04	56.85	57.97
Average Equilibrium Price	60.41	60.99	60.68
Average Pr. Value	50.37	51.34	50.68

Table 1.3: Average Prices, Equilibrium Predictions and Private Values. The Difference between COMP/INC and COMP/ENG are highly significant

1.4.1 Complete Information - COMP

Figure 1.2 graphs average[19] prices in the initial auction, average resale prices[20] p_2 and SNE predictions - which as shown in Proposition 1 are equal to second highest values in the group- over time, for the pooled data of all sessions of COMP. There were differences between individual sessions but the general tendency to overbid was the same in all of them, so it is not necessary to present individual session data. Table 3 reports median deviations from the SNE predictions pooled over all sessions of treatment COMP.[21]

In treatment COMP some underbidding is observed in the first few periods. As explained before, we can view these periods as adaptation periods. In the next periods mean prices in the initial auction lie always over the theoretical prediction, sometimes substantially so. Nonetheless, resale occurred in about 25.6% of the cases and mean resale prices are still higher than the initial auction prices, so the winners in the initial auction realize positive profits on average.

A closer view of individual bidding behavior, the box plot of values versus exits in Figure 1.3 can be very informative. Note that if the SNE prediction were valid, all bids should lie on the 45-degree line through the origin. In this plot the overbidding is even clearer than if we only look at auction prices, especially if we compare bids in

[19] In every period we take the average over the four groups that were formed in every experiment.

[20] Recall the resale price of the good is automatically equal to the highest value in group. Thus, the profit of the initial auction winner is just the difference between the highest value and the initial auction price.

[21] Following Harrison (1989) it seemed important to use two metrics to measure the deviation. Metric 1 is the usual metric, measuring the deviation in the message space of the auction, ie the deviation in the bids. Metric 2 measures the deviation in the payoff space. This measures the incentives the subjects have to play the equilibrium strategy, or alternatively how high is the cost of deviation. In this table only the results for Metric 1 are presented, as the statistical significances are quite similar using Metric 2.

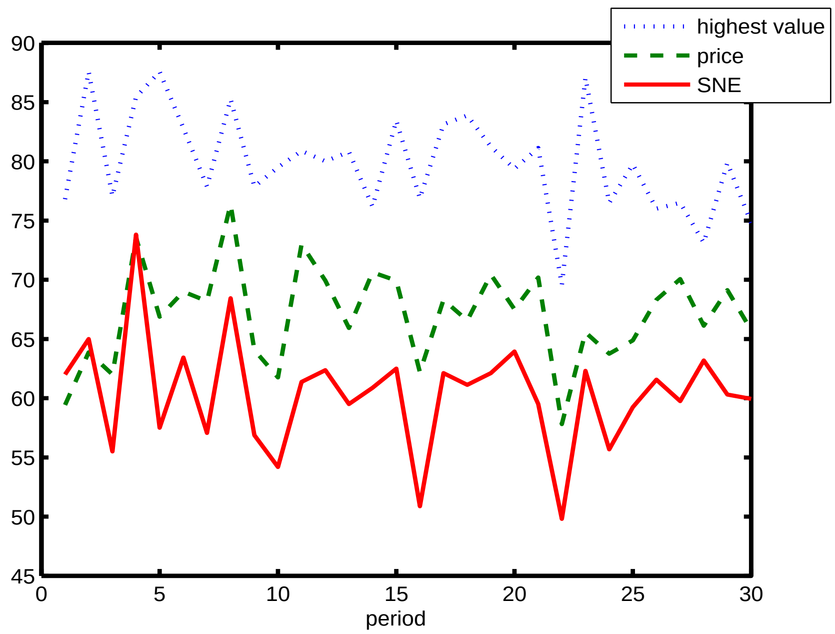

Figure 1.2: Prices, highest values and SNE predictions over time in first stage of treatment COMP.

this plot with the bidding in ENG or INC. We also see that the high auction prices come almost entirely from the overbidding of the low value players. In fact low value players overbid 40% of the time and in about 83% of the cases where they do so, the highest value bidder does not win the auction. But since there were 4 bidders in every group, overbidding did not necessarily lead to winning. Thus, players who did not have the highest value could keep overbidding without obvious punishment, as they did not win the auction very often and when they won their profits were not very low[22].

The persistent excess of bids and prices above the equilibrium predictions has to be compared with the results of ENG and the previous results in the literature, like the English auction experiment in Kagel et al. (1987). In that study the fast learning and eventual convergence to the equilibrium predictions was attributed partly to the

[22]Mean profit of bidders who did not have the highest value but won was 0.77 over all periods and -0.20 in the last 15.

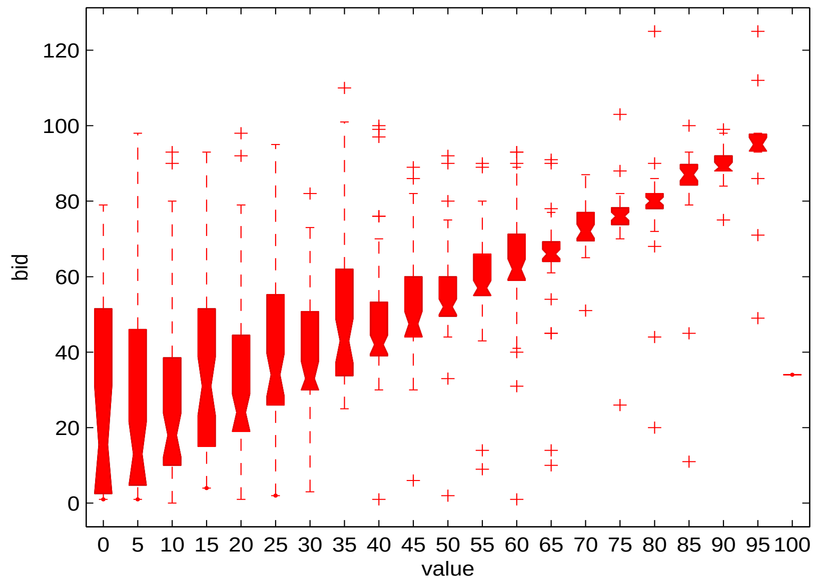

Figure 1.3: Box plots of values (x axis) versus exits (y axis) in all sessions of treatment COMP. The length of the box represents the interquartile range (IQR). The whisker extends from the box to the most extreme data value whithin 1,5 times the IQR. The crosses represent outliers outside this range. The middle of the notches is the mean, and the extent of the notches graphs a robust estimate of the uncertainty about the means for a box to box comparison.

negative profits of subjects who started by overbidding in the first periods. This effect, pushing subjects towards equilibrium behavior does not exist in sessions 1 and 6 and was very weak in sessions 2 and 7.

Thus subjects were not always best responding to the other bidders, but they were still choosing strategies that yielded payoffs close to their best response payoffs. In Section 1.5 I present models that allow for this kind of behavior, by assuming that subjects do not play pure best responses but have a mixed strategy, choosing a probability for an action depending on its expected payoff. I show that such a model rationalizes overbidding as a response to high value players underbidding with a positive probability and explains behavior better than the SNE.

1.4.2 Incomplete Information - INC

Figure 1.4 graphs average prices in the initial auction, average highest values and SNE predictions over time, for treatment INC. Table 3 reports mean deviations from the SNE model's predictions pooled for all sessions of treatment INC.

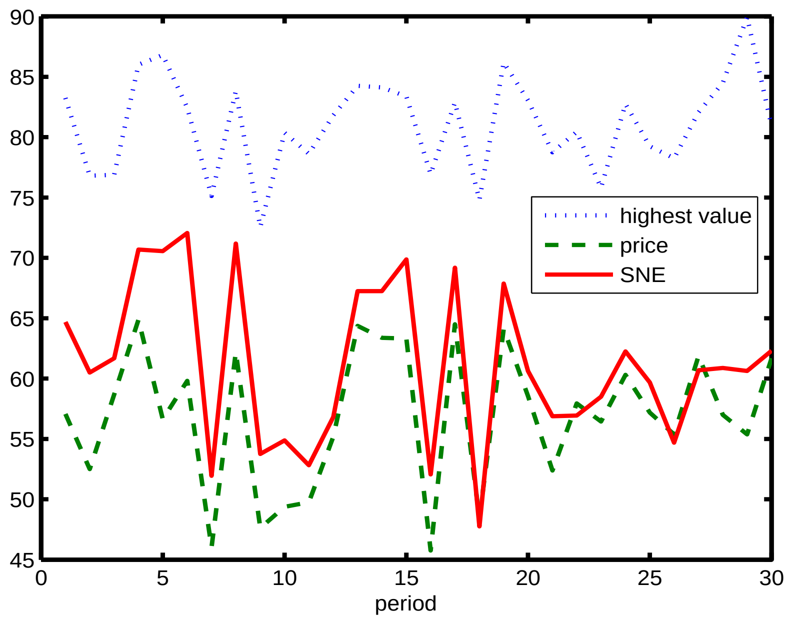

Figure 1.4: Prices, highest values and SNE predictions over time in first stafe of treatment INC.

As in treatment COMP we observe some learning in the first periods and then behavior tends to stabilize. We do not observe any big differences from session to session of treatment INC. It has to be noted that in this treatment the asymmetric information in the resale stage makes a richer strategic behavior possible. In particular signaling can be expected to play a significant role, so that looking just at prices or at aggregate values is less informative and a closer look on individual bidding behavior should be more revealing. Still there are some important facts to notice in Figures 1.4. The most obvious is that the overbidding from COMP has virtually disappeared. In most periods we even have underbidding. As we can see in Tables 2 and 3 this underbidding is small, but statistically significant. This underbidding for the high types means the highest value player loses the initial auction in about 13.4% of the cases. In the resale stage, prices are always higher than the prices in the initial auction, but still sometimes lower than the second highest value. In these cases it has to be that the subject with the second highest value does not participate in the resale auction or that she exits this auction before her use value has been reached. Measuring the rationality of subjects' behavior in the second stage is warranted.

To this end I have prepared two rationality indices. The first, RatR, is a measure of the optimality of the reserve prices in the resale auction. The optimal reserve price depends on the beliefs of the subjects and the beliefs depend on the signals from the initial auction, so it is impossible for us to calculate the optimal reserve price and deviations from it without knowing the subjects' beliefs. However we can expect that when all subjects are rational, the seller has to set a reserve price that is weakly higher than her use value for the good[23]. RatR measures the percentage of sellers who choose a reserve price $r > v_i - 3$. This index varied a bit in the four sessions. In the first session of INC it ranged from 0.5 to 1 with no trend to disappear over time. In session 2, RatR was higher and time had an effect. While in the first 10 periods it mostly ranges from 0.5 to 0.75 in the next periods it is always between 0.75 and 1 and the average is 0.88. In sessions 3 and 4 RatR was quite high, at 0.97 after the 10th period in both treatments. It is not clear why some subjects set reserve prices with such errors. The seller has her use value as an outside option and she should ask for this value at least, if others are rational. However as noted above there are subjects who think it is not interesting to participate in an auction where the starting price is below their value but very close to it, so maybe setting these reserve prices was a rational response to this behavior. A second explanation is that subjects just did not understand that when calculating the optimal reserve price they should think about their use values[24]. In an experiment such as this one, which was arguably more

[23] Note that if one of the subjects is not rational setting a reserve price becomes even more complex. The reasons for such irrationality vary. For instance it is possible and it was observed in some extreme cases that subjects do not participate in the auction when the reserve price is 10 units lower, or less, than their use value. This can be due to fairness considerations. Subjects may find an offer unfair if it gives them a small part of the available rent.

[24] In the questions subjects had to answer following the reading of the instructions, a number

complex than usual auction experiments, such mistakes could occur, so one might think more learning periods are necessary. However this argument is contradicted by the second index I calculated.

RatC, measures the percentage of the subjects who chose not to participate in the resale auction, despite the fact, that their use value was higher than the reserve price. Apart from very few mistakes in the early periods, subjects' behavior according to this index was 100% rational. This result is encouraging and suggests that probably the low RatR figures are also not due to miscalculation of the profits, but deliberate choices.

Figure 1.5 shows boxplots of values versus bids. The stark contrast to COMP becomes clear. Low value players bid close to their values with a small tendency to overbid, while high value players greatly underbid. Furthermore there are some cases, many more than in COMP, where subjects bid 0 or very close to it. These characteristics of the bids can also be explained with the anticipation of noisy bidding or signalling, as we shall see in the next section.

1.5 Bounded rationality and noisy decision making

In this section I present a variety of models of bounded rationality which are prominent in the literature, to explain subjects' behavior. Before proceeding I will first show that common explanations for overbidding in other auction experiments, like first-price auctions (see Cox et al. 1992), cannot explain the data. Consider "joy of winning", meaning that a player's utility is increased by a fixed amount if they manage to get the object and realize profits in the auction. A pure joy-of-winning model predicts the same absolute value of overbidding, for all private values, unless the joy of winning is

---

of subjects had answered the questions about second stage profits wrongly. Despite the efforts of the instructors to make these points clear after observing these mistakes, it could be the case that some players mistakenly thought the profit from the first period to be their outside option and set a reserve price that was just higher than this number but possibly lower than their value.

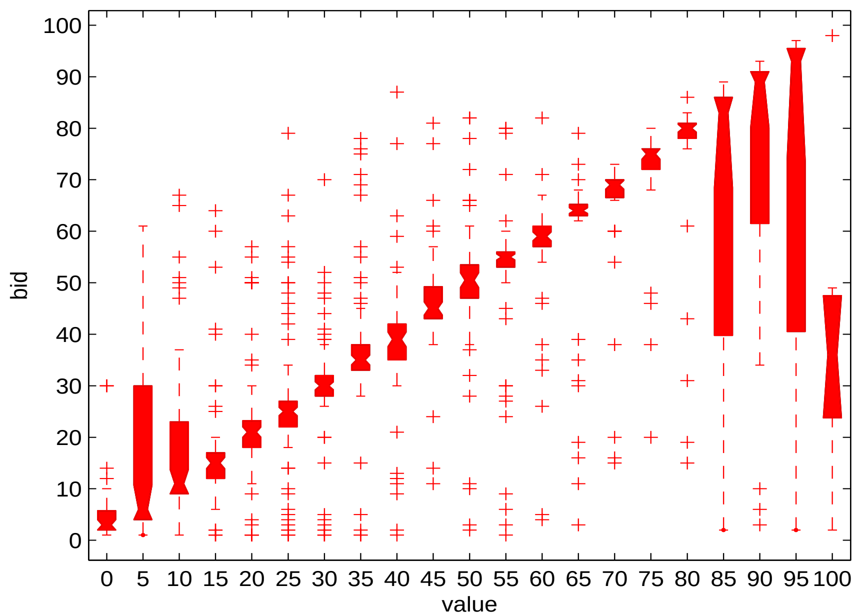

Figure 1.5: Box plot of values (x axis) versus exits (y axis) in all sessions of treatment INC. The length of the box represents the interquartile range (IQR). The whisker extends from the box to the most extreme data value whithin 1,5 times the IQr. The crosses represent outliers outside this range. The middle of the nothces is the median, and the extent graphs a robust estimate of the uncertainty about the means for a box to box comparison.

somehow correlated with use values. However, as we saw in the previous section low value bidders bid much higher than their values, whereas high value players' bids are very close to their values. More evidence against this hypothesis is that in the simple English auction, no overbidding is observed after the initial learning periods.

A second explanation, used for example in Kagel et al. (1987), is risk aversion. If subjects are risk averse they could value the higher probability of winning, when bidding above their values, more than the loss in their expected profit. In English auctions with no resale the equilibrium is, as noted before, in dominant strategies, so risk aversion does not induce different behavior. In COMP risk aversion *alone* does not change the equilibrium predictions. However risk aversion *combined* with some noise in the bidding (explained in the following section) could be a factor influencing the results.

Another motive for low value players overbidding is spite, as has been found for example in Andreoni et al. (2007). The authors gave subjects information about other bidders' values in second price private value auctions. When players have a low chance of winning due to a low value, but know that some other player has a high value they sometimes tend to overbid in order to lower the winner's earnings. The authors observe that when subjects get more information about others' value, this behavior becomes less risky and overbidding is more frequent. Note that if we model the auction as a series of stages (see Milgrom, Weber 1982), entering a new stage every time a bidder leaves the auction, this behavior is compatible with individual rationality. The SNE equilibrium described in Section 3 is unique only in the last stage of the auction when there are only two bidders left, while in the earlier stages other equilibria are possible. Of course any equilibrium arrives to the same outcome regarding the identity of the winner and the auction price.

While this explanation seems plausible it cannot account for the entire amount of overbidding observed. The first reason is that in simple English auctions with no resale, the extent of overbidding in early stages is much lower although the risk from overbidding is exactly the same as in INC and COMP. Secondly, in the last stage of the initial auction in COMP when only two bidders are left, overbidding is indeed less risky than in ENG, but unlike the previous stages it is still not part of any symmetric equilibrium strategy and can lead to negative profits[25]. It seems there is a need for more complex explanations, and in the following we propose some that consider noise.

1.5.1 Complete information in the resale market - COMP

Let us start from the basic observation that the Nash equilibrium of the games we tested is not robust to noisy behavior, as will be shown. That is, it is not robust to small perturbations of the bidding strategies. Human players make mistakes and

[25] In fact it can be part of an asymmetric equilibrium, where everyone bids his value except for one player who overbids. This strategy however is very risky. If for example there are two spiteful players with low values and they both overbid, the winner of the two will suffer a serious loss.

anticipate *others* to make mistakes. In general, as has been shown for example in Goeree et al. (2002), adding noise to the equilibrium bids can shift subjects' best responses quite radically. It remains to be seen if the same effect can be found in the present experiment. In the complete information treatment, if the other players use the SNE strategy, the expected payoff functions of a subject contemplating a deviation from this strategy are broadly the same as in a simple English auction with no resale. However, I will show that if there exists some kind of noisy behavior, which means that subjects make errors when choosing their bids, the payoff functions are quite different.

The following graphs in Figure 1.6 plot expected profits, depending on one's bid, in the case of a simple English auction (ENG) and an auction with resale (COMP). There are 4 curves plotted in every figure, representing expected profits calculated for use values of, 20, 30, 40 and 50. The upper left graph represents expected profits in ENG when three opponents bid their values without any noise, averaged over every possible value of the opponents. Notice that a bid is a best response given a use value, if it lies at the point where the expected profits reach their maximum value. In this case, payoff is maximized when a player bids her value. For example, the curve drawn for a use value of 50 reaches its highest value exactly for a bid of 50. In the upper right figure I calculate the expected profits, again given that other three subjects bid according to the Nash equilibrium but adding a normally distributed noise to these bids. This means that an opponent with a value of v is assumed to bid $v + \varepsilon$, where $\varepsilon\ \tilde{}N(0, \sigma^2)$ and $\sigma = 9$. I proceed to calculate by numerical simulation the expected profit functions of a player facing three opponents employing this noisy

Nash bidding[26]. Bid-your-value is still a best response. This is to be expected, as bid-your-value is a weakly dominant strategy in English auctions, ie. an optimal strategy regardless what others do.

The lower graphs depict the same for COMP. The utility functions without noise look similar to the ones of ENG. Expected utility now has a lower bound, equal to zero[27], but has a maximum at exactly the same points as in ENG. This corresponds to proposition 1, which says that bidding-your-value is the equilibrium strategy in COMP. However when we add noise as described previously, the curves change dramatically, as is evident in the lower right figure.[28] For every value we have a new maximum and its exact position depends on one's value. For all use values however, it is now optimal to bid very high (approximately 90).

The existence of the new hump[29], when opponents bid with some small errors,

[26] I calculate this by independently drawing 2 million sextuples of private values v and errors ε for the three opponents. For every opponent I obtain the noisy bid $\tilde{b} = v + \varepsilon$. I then calculate for every possible bid b_i of player i the winning frequency given this bid and the mean highest bid and highest value of her opponents, conditional on the highest bid $\max\{\tilde{b}_{-i}\}$ being lower than b_i. A player's expected profit is then caclulated for any given private value v_i as

$\Pi_i = prob\{b_i > max\{\tilde{b}_{-i}\}\}E[\max\{v_i, v_{-i}\} - \max\{b_{-i}\}|b_i > max\{\tilde{b}_{-i}\}]$

The numerical simulation is helpful because we can calculate these functions for any noise specification and the accuracy of the method is very high, as I have verified in the cases where an algebraic solution is straightforward to obtain.

[27] If I bid more than my value and I win, given that others bid their values, I can never pay more than the highest value among the other bidders. But this is exactly equal to what I receive in the second stage.

[28] The intuition of how mistakes can make overbidding with a low value profitable, is as follows. There are two cases of possible mistakes, opponents can (A) underbid or (B) overbid. In A there is a chance of winning and reselling at a profit. On the other hand in case B defeating opponents who have accidentally overbid is a bad idea, as there will be no profitable resale. Overbidding is a best response because, conditional on winning, case A is more likely than case B.

[29] It is interesting to note how the emergence of the new maximum is the result of the resale opportunity. The expected profit of a first stage bidder is a maximum of two values, expected utility if she consumes the good now and expected profit if she resells it in the second stage. The utility functions graphed are thus the maximum of these two utilities. In the right part of these curves, the resale effect dominates. In the left part, the usual utility enjoyed when she consumes the good herself is dominant. Without noise the utility from resale is zero, as the expected revenue in the second stage equals the expected price in the first auction (both equal the highest value among the other bidders). With noise however this not true anymore, as the expected price in the first stage becomes smaller than expected revenue in the second. This difference is maximized for a bid of around 90 (which is actually higher than the unconditional expected highest value, 75), depending on the size of the errors.

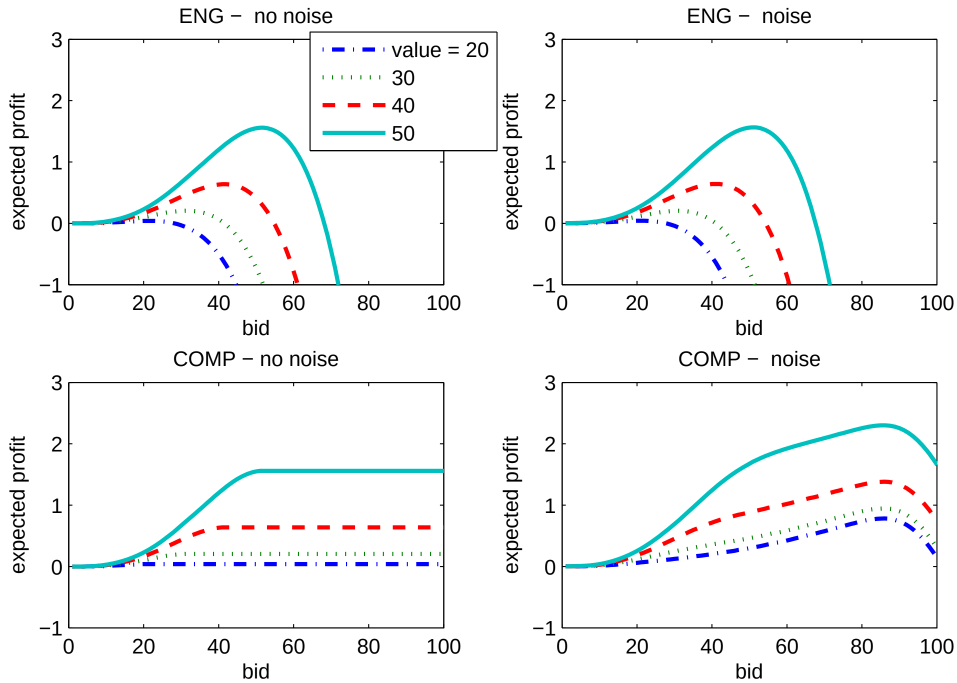

Figure 1.6: Expected utilites in ENG (upper two figures) and COMP (lower two) without and with noise (normally distriubuted with a σ of 9). The curves are drawn for private use value signals equal to 20, 30, 40 and 50. In the lower left panel utility is very flat but still maximized at a bid equal to value.

means that best responses change discontinuously with noise. Optimal bids jump from being equal to player's value to a bid that is significantly higher than a bidder's value, especially in the case of low values. Note that this discontinuity exists for many different specifications regarding the functional form of the noise distribution (e.g. triangle, logistic, uniform, Laplace) and even when the standard deviation is very small (a standard deviation of $\sigma = 1$ is enough in the case of the normal distribution). The intuition is that the hump rises higher the more noisy the bids, but its position on the x axis does not change much. Even when there is not much noise and the hump is very low, it will still be higher than the utility derived from a bid equal to or lower than one's value and its maximum will be positioned in the upper part of the interval $[0, 100]$. This means that even the lowest value players should overbid massively in the presence of small amounts of noise.

To test in a systematic way if the characteristics of the game discussed above are indeed influencing the bidders' behavior we can consider following model. Suppose a player believes her 3 opponents want to bid their values but make small errors, distributed normally[30] with $\sigma = 15$. Then this player's best response[31] given such beliefs is approximated by:

$$BR_{naive} = 0.000057v^3 - 0.0046v^2 + 0.098v + 79$$

This is a concave bidding function that starts at 79 for a value of zero and reaches approximately 100 for a value of 100. An alternative to this model is to calculate the best response to the *actual bidding distribution* (and not to the one predicted by the theory). It predicts serious overbidding of approximately[32] the following form:

$$BR_{act} = \begin{cases} 47, v \leq 39 \\ 0.87(v - 39) + 47, v > 39 \end{cases}$$

The hypothesis that subjects were responding optimally to the actual play of the others can be tested using the BR_{act} model. The fit of both these models is presented in table 1.4.

Levels of reasoning A model that has been found to explain many anomalies in experiments is a level of reasoning model (see for example Nagel 1995, Stahl 1995, Camerer 2004). In specific I will use the level-k version (Crawford and Iriberri, 2008).

[30] I choose here a σ that is lower than the minimum of the actual estimated standard deviation σ of players' bids in the various sessions of COMP, assuming errors are distributed normally, as can be seen in Table 4. I also tried other distributions and the result was robust to these variations.

[31] The BR and other alternative models we present will be under the assumption that bidders do not update their beliefs after they observe the exits of other players. It does not change the results by much but it greatly simplifies the calculations. Additionally it is confirmed by the data, the main determinant of a player's bid was her use value and the unconditional distribution of the other player's values. Actual observed exits were not a significant factor.

[32] To calculate this best response I first estimated a joint bid-value distribution using the data from the experiments. Then I find by numerical simulation the bidding function that maximises a player's expected payoff when playing against 3 opponents who are employing this empirical bidding strategy. The best response is not exactly piecewise linear, but well approximated by the given function.

The idea is simple and rather intuitive. There exist k types of players, varying in their degree of sophistication. Level 0 (L0) players bid randomly with a uniform distribution. Their bids can be interpreted as pure noise, given that values do not correlate with bids at all. In this way this version of the levels of reasoning model used in the literature has built-in the idea of noisy behavior. Level 1 players believe they are playing against L0 players and play a best response, Level 2 players play a best response to Level 1 and so on. I first derive the strategy for a Level 1 player best responding to N players who bid randomly. Her expected profit will equal the maximum of her value and the expected highest value among the opponents minus the expected highest bid, given that the latter is lower than her own bid.

$$\Pi_i = \int_0^{b_i} (\max\{v_i, E[\max\{v_{-i}\}]\} - x) N F(x)^{N-1} dx$$

Note that opponents' values and bids are not correlated. Rearranging and taking first order conditions (see appendix) leads us to following strategy for a level 1 player when we have 4 bidders in total and values are uniform in [0,100]

$$b_{L1} = 25(v_i/100)^4 + 75$$

Thus, Level 1 players bid an increasing concave function of their values, from a bid of 75 for a zero value type to a bid of 100 for a bidder with value 100. Level 2 types best respond to Level 1. It is simple to show that this results in a bid your value strategy. Given that the opponents all bid at least equal to their values there is no opportunity for profitable resale, thus the game reduces to a simple English auction with no resale and bid your value is a best response[33]. However the expected profit curves are not the same as in the Nash equilibrium, as the probability of winning with a bid lower than 75 is zero[34]. L3 is then exactly equal to the Nash equilibrium, with the same expected payoff functions.

[33] The best response to a level one player is actually a correspondence and not a function, for private values under 75. Any bid is in principle part of the best response, however bid your value is the obvious focal part of this best response. I accordingly expect L2 players to bid their values. Note this is only important for the calculation of L3 as a response to L2. For the actual fitting of L2 the errors are logistic, which means that every action yielding the same payoff is treated as equally likely, thus every action in the best response correspondence is treated the same.

[34] This will be important when fitting the model to the data using logistic errors, as they depend on the exact structure of the expected payoff curves.

The model is fit assuming that each observed bid is a draw from a common distribution over the three types. The frequency of L1 players in the population is χ_1, χ_2 is the frequency of L2 players and the remaining $1 - \chi_1 - \chi_2$ is the frequency of L3. L0 just exists in the mind of L1, as has been found in Crawford and Iriberri's work[35]. For each type we can calculate expected utility for every possible action, given the beliefs of this type. I assume that a player of a certain type makes errors with a frequency that depends on the expected utility of each action, according to a logit specification. This means the probability for a subject i playing a particular action j out of all actions J is calculated in the following way:

$$p_{ij}(\lambda) = \frac{e^{\lambda U_{ij}}}{\sum_{k=1}^{J} e^{\lambda U_{ik}}}$$

The numerator is the utility from each action transformed by an exponential function, in the denominator we have the sum of all these exponential weights as a scaling factor, so the probabilities add up to one. The parameter λ determines how sensitive errors are to payoff differences. Bids become uniform as $\lambda \to 0$ and errors are eliminated when $\lambda \to \infty$.

QRE The last model I calculate is a Quantal Response Equilibrium which captures the idea of noisy behavior but predicts that players' deviations will be systematic. Similar to the level k model, I use a logit specification that has been found to give intuitive theoretical predictions in auctions (see Anderson et al. 1998) and to fit experimental data well (see Goeree et al. 2002). Bidders with a given use value have a probability distribution over every possible bid which depends on their payoff sensitivity parameter λ and the actual payoffs.

Players correctly anticipate the bidding distributions of their opponents and all choose the probabilities according to the rule above. Thus, a best response will be played with a higher probability, but not with certainty. The equilibrium is a

[35] In this paper there is also an alternative specification of the model with truthful bidding as a L0 starting point, but it is not useful here, as it would obviously lead to the L1 type bidding her value and thus no difference with the SNE.

fixed point of a mapping from choice probabilities to choice probabilities. Note, that although a QRE approaches a Nash equilibrium in the limit when the noise parameter tends to infinity, it can be far away from it for intermediate values of the parameter.

Calculating a QRE with such a large strategy space is a daunting task. With the usual differential equations approach (used for example by Goeree et al. 2002) it is even considered numerically impossible, to the best of my knowledge, as it involves solving a system of 101 simultaneous non-linear differential equations. Therefore I use a different method to calculate the QRE, namely a Cournot process. Starting with a random bidding function[36], the expected utility of a player facing three bidders employing such a bidding function is calculated. I proceed to calculate a quantal response by weighing the utility, to get choice probabilities according to the formula above. This process is then iterated until the quantal responses converge to a stable state.[37] Convergence is usually reached after about 15 periods and does not depend much on the initial bid function.

A question when calculating the QRE is the choice of the parameter λ, which is a measure of sensitivity to payoff differences. A different λ can lead to radically different predictions. Because of this I decided to restrict λ to values found in earlier research with auction data[38], that is to values around 1 and not to include a risk aversion parameter[39]. It so happens that the restriction of λ was not binding and the values estimated in other experiments, including auctions, gave a very good fit

[36] Note that the way we calculate the QRE reveals an interesting relationship to the levels of reasoning model. Every iteration of the Cournot process to find the QRE, corresponds to a level of reasoning. A somewhat similar analysis is done in Goeree and Holt (2004), where they propose a model of noisy introspection using logit response functions, but relaxing the equilibrium assumption. Their model can be seen as an alternative to the level k model (which it includes as a special case), where there is more noise associated with players' beliefs about higher levels of iterated expectations.

[37] An iterative method has been used to calculate a QRE in coordination games and the traveller's dilemma using a simple spreadsheet (see chapters 25-26 in Holt 1996).

[38] Note that λ depends on the payoff space and we adjust it accordinlgy. For example, in Goeree et al (2002), where the private values have a support of [0,11], λ is found to be on average 10 (actually they use a parameter μ which is equivalent to $1/\lambda$ and they find $\mu = 0.1$). Then for this auction where values are in $[0, 100]$, the values of λ are restricted to be close to 1.

[39] Haile et. al (2006) note that a QRE with two parameters, suitably chosen, can be used to fit any data. This is an additional reason why I did not want to include a risk aversion parameter. This critique does not apply to this analysis, as with the logit structure of the errors I have assumed, the payoff perturbations are i.i.d. See Goeree, Holt and Palfrey (2005) for a discussion.

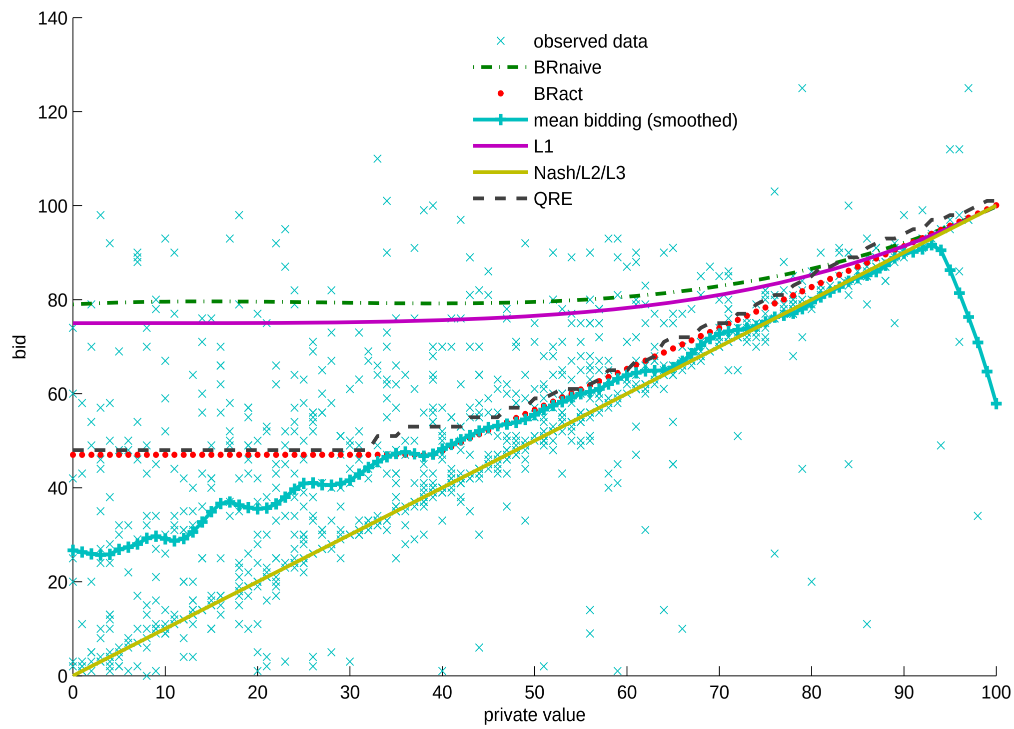

Figure 1.7: Comparison of the different models for treatment COMP. The QRE predicts a distribution of bids for every use value, so the mean of these bids is presented. Keep in mind however, the model with the best fit is not the one closest to the mean actual behaviour but the one where the whole predicted distribution is closest to the actual one.

for our data too, which indicates that the QRE is an appropriate model to explain behavior in a wide range of auction experiments.

Comparison of models The different models calculated above for treatment COMP are fit to the data in this section and compared with the same models for treatment ENG. The predictions of the various models in the simple English auction are straightforward. The strategy is actually equal to bid your value for all models except QRE, as this is the weakly dominant strategy in simple English auctions with IPV. Note that unlike in treatment COMP, L0 is exactly equivalent to bid your value, as the L1 type's payoff is not influenced any more by the values of his opponents but just by their bids, which are uniform in both L0 and Nash. The QRE predictions are calculated by simulation, as in the case of COMP.

Results are shown in Figure 1.7 and the goodness of fit can be found in Table 4. Maximized log likelihood values for each model are presented in the first row. In the case of the pure strategy models (Nash and BR), where no dispersion is predicted by theory, I allowed for normally distributed errors and the estimated σ is shown in brackets. The QRE predicts a dispersion according to the logit formula presented above, so there was no need for additional errors. For the levels of reasoning models I posit logistic errors as described above and calculate them numerically. I assume that λ is independent of subject or type.[40] Thus in total the model has three parameters, the common precision λ, χ_1 and χ_2.

Lastly I estimate a Nash model but with logistic instead of normal errors, calculated in the same way[41]. This gives us a fairer comparison to the QRE and level k. Since the models are not nested I use the Bayesian Info Criterion (BIC) for model selection. Recall that we have more observations for treatment COMP than ENG so the respective likelihoods are not directly comparable. Note that all models except the level k have just one free parameter.

Model	**Nash**	**BR**$_{naive}$	**BR**$_{act}$	**L1**	**L2**	**L1+L2+L3**	**QRE**	**L3/Nash+logit**
LL COMP	-4469.2	-5044.5	-4391.5	-4555.8	-4492.8	-4312.9	-4207.7	-4367.5
BIC	8945.3	10095.9	8789.9	9118.6	8992.6	8646.6	8422.3	8741.9
Est. λ or σ	20.9	37.1	19.4	0.1	0.22	1.84	1.23	0.87
LL ENG	-2706.2	-2706.2	-2706.2	-2732.4	-2732.4	-2732.4	-2715.1	-2732.4
BIC	5419	5419	5419	5471.4	5471.4	5471.4	5436.8	5471.4
Est. λ or σ	10.38	10.38	10.38	1.13	1.13	1.13	1.1	1.13

Table 1.4: Goodness of fit of different models for treatments COMP and ENG. LL is the maximised log likelihood.

The model that performs best in explaining the results in COMP is the QRE model, followed by the mixed levels of reasoning model. The estimated frequencies

[40] In Crawford and Iriberri's study they compared such a model to models where precisions can be type-specific or subject-specific. Forcing subjects to be of a single type in COMP leads to a very large increase in the number of parameters, without adding much to the fit. Also the estimated population frequencies do not change much. If we assume type-specific precisions we get a LL of 4199.6, BIC=8433.8, a frequency of 4.3% L1 types with precision $\lambda = 0.039$, 33.3% L2 with $\lambda > 18$ and 62.4% L3 with $\lambda = 12.7$.

[41] Actually note that Nash+logit is exactly the same as the L3 model fitted with logistic errors.

of the types was 1% for L1, 24.8% for L2 and 74.2% for L3. Nash with normal errors does not fit the data well, while BR_{act} (the best response to actual behavior model) performs better. There is a leap in the likelihood when we estimate the Nash model with logistic errors instead of normal. This means that the assumption of symmetric errors is not plausible. Still, the QRE with logistic errors performs even better than Nash+logit which shows that logistic errors are not enough to explain the subjects' behavior. Players not only make errors in a systematic way as is modelled through the logistic distribution, but anticipate others to make errors too, which leads them away from the Nash prediction and towards a quantal response equilibrium. These results are reinforced by the fact that the estimated error parameter λ was quite similar for both treatments and for both models, Nash and QRE.

All the models except Nash predict overbidding in the resale treatment COMP. Thus there is no way to separate them based on the qualitative predictions. They differ however in their predictions when we vary the number of bidders in the auction. The levels of reasoning model clearly predicts a monotonic rise in the total amount of overbidding while the QRE predicts an initial rise up to 4 bidders and then a slight fall in parts of the bidding function[42]. In specific, for middle-of-the-range use values, the QRE prediction falls when there are many bidders. Thus an experiment with COMP and 5 or more bidders would help to separate the models.

In ENG all models exhibit a similar performance, which is not surprising given that their predictions are very similar and the greatest difference stems from the different distributions of the errors (logit vs normal). The Nash and BR models have the lowest log likelihood with the QRE a bit worse and the mixed LOR and Nash+logit with still a bit higher LL. Overall however the great improvement in fit given by the last three models in COMP means that the total predictive power of these models is higher. If we use the average performance in the two treatments as a

[42] The reason for the fall has to do with the feature of the QRE, where strategies with a payoff of zero are played with a positive probability. As the number of players grows the probability of winning with a low to middle bid falls dramatically. Thus the part of the bidding function that gives an expected payoff of zero grows and the bidding distribution for a given value comes close to uniform. Thus while for $n = 3$ the QRE predicts a bidder with a value of 70 to bid on average close to 60, for $n = 8$ she will bid close to the average of the uniform distribution in [0,100] which equals 50.

selection criterion, the levels of reasoning model and QRE emerge as clear winners[43].

1.5.2 Incomplete information in the resale market - INC

As in treatment COMP, the anticipation of noisy behavior can be used to explain the data in INC. Very high value players know they will win with a very high probability. But if they try to win in the first stage they would possibly have to pay a price higher than the second highest value in the group because some low value players can be (relatively costlessly) overbidding. They prefer to signal low values[44] and wait for the second stage auction where they know that overbidding for the low value players is exactly as costly as in a simple English auction and will thus be avoided. Given actual behavior such a strategy would be more profitable than the Nash prediction.

Note that this logic is exactly captured by the logistic errors which allow for the fact that players do not always best respond, but still try to avoid the most costly mistakes. Low value types can costlessly overbid in the first stage but avoid overbidding in the second stage. High value types will not avoid underbidding in INC as much as in COMP or ENG, since in case of losing in the initial auction they can still make some profit in the second stage. The question that arises is which of the previously discussed models employing logistic errors will fit the actual behavior better. The QRE assumes that subjects correctly anticipate the logistic errors of their opponents and arrive at an equilibrium where subjects play noisy best responses to each other. On the other hand the level k model assumes bidders do not think past a limited number of iterated best responses. They intend to play a best response to their opponents, given the beliefs that correspond to their level of reasoning, but make logistic errors. Finally the Nash+logit model just assumes players intend to

[43]Since Nash+logit is equivalent to L3, it is nested in the level k model. A likelihood ratio test rejects the hypothesis that the two models are equivalent at the 0.001 level.

[44]Suppose there are two biders. Imagine a bidder with value 50 believes the other player is playing the bid your value equilibrium with small symmetric mistakes of maximum magnitude 10, as described in the previous section. He will then bid more than his value, say 60, expecting to resell. In that case, his opponent will have an incentive to bid much less, if he has a high value, say 90. For example if he bids 40 he will lead the winner to believe that he has a value of maximally 50 and he will thus get the good in the second stage for a price of 50. See example 1 in Hafalir and Krishna (2008) for a similar argumentation.

play the Nash strategies, but make logistic errors.

Due to the interdependence of the two stages and the additional complexity, it is not straightforward to calculate the logistic errors and the QRE and level k models for treatment INC. One has to use a shortcut, building a reduced form of the game to make it tractable. I therefore assume that the reserve price in the second stage is equal to the use value of the seller and that players who have a higher value than the reserve, do participate in the auction. I also assume that players who have decided to participate in the second stage auction proceed to play exactly as in a simple English auction, bidding their values. These assumptions are largely consistent with actual behavior and partly with theoretical arguments too[45]. I then plug the expected continuation payoffs from the second stage subgame in the first stage payoffs and this results in a game where the various models' predictions can be calculated[46] as described in treatment COMP. The main difference of this reduced version of INC with COMP is that the winner of the first stage can only expect to get the second highest value among the other bidders in the second, but the losers now have a chance to win in the second stage and appropriate a part of the rent (for example in case they have the highest value, they will get the difference between this value and the second highest). Note that in this reduced game, bid your value is still a Nash equilibrium.

Level 1 play, meaning a best response against opponents who are bidding randomly, results in a bidding function that starts at (N-1)/(N+1) for a value of zero (see appendix). It rises monotonically to N/(N+1) for a value of 100. Level 2 players bid their values[47], up to a maximum bid of N/(N+1) and level 3 do the same, but have different expected payoff functions and thus different error distributions.

[45] The bidding behaviour prescribed for the second stage bidders is rational. On the other hand, for the second stage seller setting a reserve price equal to her use value is not an optimal choice. However, when the number of bidders is high enough, the reserve price becomes irrelevant. For example when selling to 3 bidders as in our experiments with values uniformly distributed in [0,1] the expected revenue under the optimal reserve price is around 0.53 and the second highest value (which is the expected revenue without a reserve price) is 1/2. This means the reserve price enhances revenues by not more than 6%.

[46] $\mathbf{BR}_{act}$ is not caclulated as we did not have enough observations in the second stage to estimate the empirical distribution of reserve prices, participation strategies etc.

[47] As with treatment COMP, the best response to level 1 is a correspondence. Where multiple bids are possible for a given value, bid your value is chosen as the obvious focal one.

Model	Nash	L1+L2+L3	QRE	Nash+logit
LL INC	-4659.8	-4559.3	-4216	-45863
BIC	9326.6	9152	8466	9179.6
Est. λ or σ	18.1	0.65	1.1	0.73

Table 1.5: Goodness of fit of different models for treatment INC. LL is the maximised log likelihood, BIC is the Bayesian Information Criterion for model selection.

The estimated frequencies for the mixed levels of reasoning mode are 31.5% for L1, 1% for L2 and 67.5% for L3, relatively close to the values estimated in the previous section for COMP. The simple Nash model with symmetric normal errors performs once more very badly. All the models using logistic errors fit the data better. As in treatment COMP, the QRE outperforms the mixed levels of reasoning model.

These models improve upon Nash mainly by predicting some underbidding for the high types. An alternative reasoning for very low bids is reported in Kamecke (1994). In this study it has been found that some subjects tended to bid very low when they thought they did not have a good chance of winning in order to raise the profits of the winner. In Cox et al. (1982) this tendency for low value holders to throw away bids was argued to make economic sense, once one accounts for subjective costs of calculating a more meaningful bid under the circumstances. However this can not explain why so many high value players were underbidding. Also, all these arguments can not explain the level of very low bids in INC. After all in the similar COMP the magnitude of low bids was significantly lower as noted in the beginning of this section.

1.6 Discussion and ideas for future work

The preceding results show, that when we are interested in real bidding behavior, Nash equilibrium analysis is not adequate. It does not suffice to study the best responses *in* equilibrium to arrive at an equilibrium bidding strategy and work with this as a prediction, as is commonly done. The exact shape of the expected payoff functions is important, as it will influence the errors of players. We find these errors to be asymmetric and systematically depending on the expected payoff of each action.

The stability of the equilibrium is important too, that is, we have to study what happens to expected payoffs and best responses when some player *deviates a bit* from equilibrium. In games such as the present, where as we have seen the best responses change dramatically when opponents tremble a bit around the equilibrium, we should not expect subjects to play according to the Nash prediction.

As an extension of this work, in order to test our hypothesis of subjects anticipating mistakes, one could run an experiment where human players face computerised opponents. As computerised opponents do not make mistakes, we should expect very similar behavior in all three treatments. On the other hand it is questionable whether players' behavior when playing against machines allows us useful predictions of how they will play against real humans.

A promising idea for future research is the explicit inclusion of a speculator in the game as in the Garratt and Troeger (2006) paper. This experiment will be very useful to compare with INC and will give us valuable insights to the source of the asymmetric behavior in our data. Other models where resale happens in equilibrium are also interesting, in particular the model of Hafalir and Krishna with asymmetric values seems to be promising and our results hint that the weak players will indeed bid much higher with resale than without, if the resale market is appropriately designed.

Finally, as discussed in the design section, experiments with sealed bid auctions can also be interesting. It would be additionally useful to design these experiments in a way that makes the results comparable with the results of the empirical study in Haile (2001), which has found evidence of the effect of resale markets on US Forest Service timber auctions. As already mentioned, independent work of List et al. (2004) has run first-price sealed bid experiments and compared them with these timber auctions. They seem to have found a significant presence of risk aversion in the data. While this seems like a plausible explanation, it is very likely that the combination of risk aversion with noisy behavior can enhance their results.

1.7 Conclusions

In the resale treatment under complete information it seems we have a case similar to the "ten little treasures" in Goeree and Holt (2001). The simple English auction represents the "treasure treatment", the case where Nash theory seems to work perfectly, predicting subjects' behavior with a very high accuracy. When we change the game a bit, adding the resale opportunity, the Nash equilibrium remains the same, prescribing that players should bid their values. Nonetheless, subjects seem to see a difference where theory does not see one. Players significantly overbid in the presence of a resale opportunity, under complete information in the resale market, and that this overbidding does not tend to fade away with the passage of time and the effect of learning.

However when there is no complete information in the resale market, the results are quite different. Subjects with low values tend to bid a bit more than their values, whereas high value bidders bid much less than their values. This indicates that instead of the usual separating equilibria there is pooling[48], high value players pretend to have smaller values and expect to get a better offer in the resale market.

In both cases the addition of the resale opportunity alters the strategic behavior of the subjects significantly in comparison to common results in simple English auctions. In most cases these changes in the bidding behavior lead to substantial differences in the revenues that accrue to the initial seller. For policy prescription purposes these findings should be taken carefully into account. While some features of a laboratory experiment will probably not apply in real markets (for instance we do not expect real-life investors to have fairness concerns or to display altruistic behavior), others like noisy behavior and the anticipation thereof are surely present and of significant importance. Thus we believe our results to have some external validity.

The second and more general result of this paper is the importance of thinking about noisy decision making and the exact form of expected payoff functions. The three treatments I tested had exactly the same Nash equilibrium, but subjects' behavior was quite different in each one of them. I argued that the reason for this is the

[48] For a similar pooling effect see Haile (2000).

presence of errors (even small ones suffice) on behalf of some players. These errors can be attributed to experimentation with different strategies, trembling, idiosyncratic preferences and moreover, not adequately modelled liquidity constraints in the case of auctions in the field. Errors and noise are present even in the most important financial markets where the stakes are very high (see Shleifer and Summers 1990). In cases where the anticipation of such errors on behalf of some players does not alter best responses by much, the Nash prediction can be valid. However, in cases as the present experiments where best responses are sensitive even to small amounts of noise, we should not expect rational human subjects to follow the Nash equilibrium strategies. Additionally, I find that whatever the reason for subjects making errors, they systematically try to avoid the most costly ones, thus the shape of the payoff functions is a good indicator for the empirical distribution of players' errors.

1.8 Appendix

Derivation of level 1 bids for treatment COMP (all values in the interval [0,1]).

All values are scaled to be in the interval [0,1], N is the number of opponents.

$$\Pi_i = \int_0^{b_i} (E[\max\{v_i, v_{-i}\}] - x) N F(x)^{N-1} dx$$

$$= \int_0^{b_i} (E[\max\{v_i, v_{-i}\}] - x) N x^{N-1} dx = \int_0^{b_i} N(E[\max\{v_i, v_{-i}\}] x^{N-1} - x^N) dx$$

$$= [N(E[\max\{v_i, v_{-i}\}] \tfrac{1}{N} x^N - \tfrac{1}{N+1} x^{N+1})]_0^{b_i} = N(E[\max\{v_i, v_{-i}\}] \tfrac{1}{N} b_i^N - \tfrac{1}{N+1} b_i^{N+1})$$

Taking first order conditions:

$$N(E[\max\{v_i, v_{-i}\}] b_i^{N-1} - b_i^N) = 0 \rightarrow b_i = E[\max\{v_i, v_{-i}\}]$$

$$E[\max\{v_i, v_{-i}\}] = prob(v_i > \max\{v_{-i}\}) v_i + (1 - prob(v_i > \max\{v_{-i}\})) E[\max\{v_{-i}\} | v_i < \max\{v_{-i}\}]$$

$$= v_i^N v_i + (1 - v_i^N) \int_{v_i}^{1} x \tfrac{N x^{N-1}}{1 - v_i^N} dx = v_i^{N+1} + (1 - v_i^N) \int_{v_i}^{1} \tfrac{N x^N}{1 - v_i^N} dx$$

$$= v_i^{N+1} + (1 - v_i^N) [\tfrac{N x^{N+1}}{(N+1)(1 - v_i^N)}]_{v_i}^1 = v_i^{N+1} + \tfrac{N(1 - v_i^{N+1})}{N+1} = \tfrac{N + v_i^{N+1}}{N+1}$$

In the case of $N = 3$, as in the experiments, we have $b_i = \frac{1}{4} v_i^4 + \frac{3}{4}$

Chapter 2

Auctions with Toeholds

2.1 Introduction

Competition for the control of a company can be essentially viewed as an ascending auction. The bidders in such an auction have more or less similar valuations for the contested company. This leads to the literature often viewing such takeover battles as common value auctions. While there is a strong common value element in these auctions there very often exist small asymmetries which can radically change the strategic interplay between the bidders and the outcome of the contest.

If the asymmetries are due to some private control benefits or idiosyncratic synergies then we can speak of *almost common value auctions* (Klemperer 1998), auctions where one of the bidders has a small payoff advantage, a value that is slightly higher than the common value. The asymmetries can also arise when some bidders already own a part of the company that is being sold. Ownership of such a part is called a *toehold* and is quite common in takeover battles (Betton and Eckbo 2000). This paper presents results from experiments on auctions with toeholds and compares these results with the theory and other experimental results in almost common value auctions.

In theory ownership of a toehold can deter competitors from bidding for the company and can give its owner a strong strategic advantage. Bulow et al. (1999) give a good illustration of how toeholds can be useful in takeover battles. The authors use an English auction framework, where bidders for a company have similar restructuring plans but differing estimates of the expected returns. Under this setup, the buyers have common values but imperfect signals. The analysis proceeds to find that with common values, toeholds can have a profound effect on players' optimal strategies. Players with a toehold bid more aggressively as they know they will not have to pay the full price and in the case they lose they will get part of this payment. On the other hand players facing an opponent who owns a toehold, have to play less aggressively than in the case the playing field were level. In equilibrium, even with a small toehold of 5% or 10% the bidder who owns it will get the company for a much lower price than without toeholds. Thus, theory gives strong reasons for bidders to acquire toeholds. The empirical findings however are not in full support of this idea. Betton and Eckbo (2000) find that only about half of the bidders acquire toeholds before trying to buy a majority stake.

Our paper addresses the conflict between this observation and theoretical results. Although theory predicts that the toeholds should have a big effect on the players' predicted strategies, the effect could be much smaller when human players participate in this game, for reasons that will become clear in the analysis. Thus we designed and ran a series of experiments to test this idea. We choose an English auction with two players and common values, similar to the Bulow et al. (1999) setup. The major simplification is that we let the total value simply be the sum of the signals the players receive. This is to keep the setup simple and to avoid understanding problems on behalf of the players. What we found is indeed that although toeholds give bidders an advantage, it is not nearly as strong as theory predicts. Thus, under some circumstances it is not advisable for an agent planning a takeover to acquire toeholds. Moreover, we find that the players' deviation from the theoretical prediction is not unreasonable, but rather has deep roots in the structure of the equilibrium proposed by Bulow et al (1999). The equilibrium payoff functions are in some cases extremely flat, meaning that large deviations from equilibrium are practically costless. In particular, we find that when the ratio of the two players toeholds is larger than 10 (e.g. 1% and 10%), the strong bidder can deviate almost 50% from his optimal bid with a negligible loss in expected payoff. Consequently, there is no reason to believe that human agents – be it in the lab or in real markets – would play their exact best responses and thus convergence to the theoretical equilibrium is very unlikely. We show that a levels-of-reasoning model with bounded rationality of the players generates more intuitive predictions and fits the observed behavior more precisely.

To our knowledge there are no other experimental studies focusing on toeholds. There are however models with almost common values that as mentioned above lead to similar theoretical results (see for example Kagel and Levin 2003). When a player is known to enjoy a payoff advantage in a common value auction, theory predicts an explosive effect in the bidding strategies, similarly to the effect of toeholds. The player with the advantage bids more aggressively, his opponents less, which leads to the strong player winning almost all the time. Avery and Kagel (1997) have sought to test this theory and they found that the differences in common values have a linear and not explosive effect. Moreover, they find advantaged bidders' behavior resembles a best response to the behavior of disadvantaged bidders. The latter bid much more aggressively than in equilibrium, which leads to negative average profits. Experienced players bid consistently closer to the Nash equilibrium than inexperienced bidders, although these adjustments towards equilibrium are small.

In a recent paper with a similar setup, Kagel and Rose (2006) again find that the Nash prediction fails to prognose the subjects' behavior. They find rather that behavior is characterized by a behavioral model where the advantaged bidders simply add their private value to their private information signal about the common value, and proceed to bid as if in a pure common value auction. The model they chose is actually, as we shall see later, a special case of the more general toehold framework. The main theoretical difference between their model and ours is that the high types should win the auction with probability one in the almost common value setting, while in our experiments the effect is predicted to be much weaker.

While our paper finds no explosive effect of small asymmetries, similarly to the above papers, our design has the advantage of varying toehold differences which allow us to see if the comparative statics predicted by theory hold, even when subjects are not following exactly the equilibrium strategies. Our finding is that in general weak types tend to bid less aggressively the higher the toehold difference, which is only partially in accordance with the theory but much more consistent with the predictions of the levels of reasoning model.

Section 2.2 introduces the model. Section 2.3 presents the experimental setup and Section 2.4 analyzes the data. Section 2.5 concludes.

2.2 The model

Two risk neutral bidders i and j bid in an English auction for one unit of an indivisible good. Bidders' signals t_k are independently drawn from the uniform distribution in [0,1]. The value of the good to every bidder is then just the sum of these signals. Additionally the bidders already own a share of the company θ_k, which we will call a *toehold*. Ownership of a toehold means that in case the company is sold the owner will get θ_k times the sale price, thus if she wins she only pays $1 - \theta_k$. Bidder's shares are exogenous and common knowledge.

The unique symmetric equilibrium is calculated in Bulow et al. (1999).

Proposition 1 *The equilibrium bidding functions of the game are given by* $b_i(t_i) = 2 - \frac{1}{1+\theta_j}(1 - t_i) - \frac{1}{1+\theta_j}(1 - t_i)^{\frac{\theta_u}{\theta_j}}$ *A discussion and the proofs can be found in the aforementioned paper.*

The proposition is true for all $\theta > 0$. For $\theta = 0$ we would have a usual English auction with common values, with the well known equilibria. That is, in the absence of toeholds the equilibrium bidding functions would be just symmetric, straight lines[1] through the origin with slope 2. Even when players have toeholds, if they are symmetric, the bidding functions are still symmetric straight lines with a slope that depends on θ.

Now, when the toeholds are asymmetric there is the explosive effect described in the introduction. The bidding functions of the two players grow apart very rapidly.

[1]This can be seen by the standard methods used in the literature. There is however a more straightforward way to see what happens for very small toeholds, by taking the limit of the bidding function in proposition 1 with the toeholds being equal and tending to zero. The function then reduces to just $b(t) = 2t$.

In Figure 2.1 you can see the shapes of the equilibrium bidding functions, separately for the low and high types. It can be observed that for toehold differences greater than 10 percentage points, the functions have parts with extremely high slopes. For signals close to 0 the high types' bids rise very steeply and similarly for signals close to 100 the low types' functions are rising very fast.

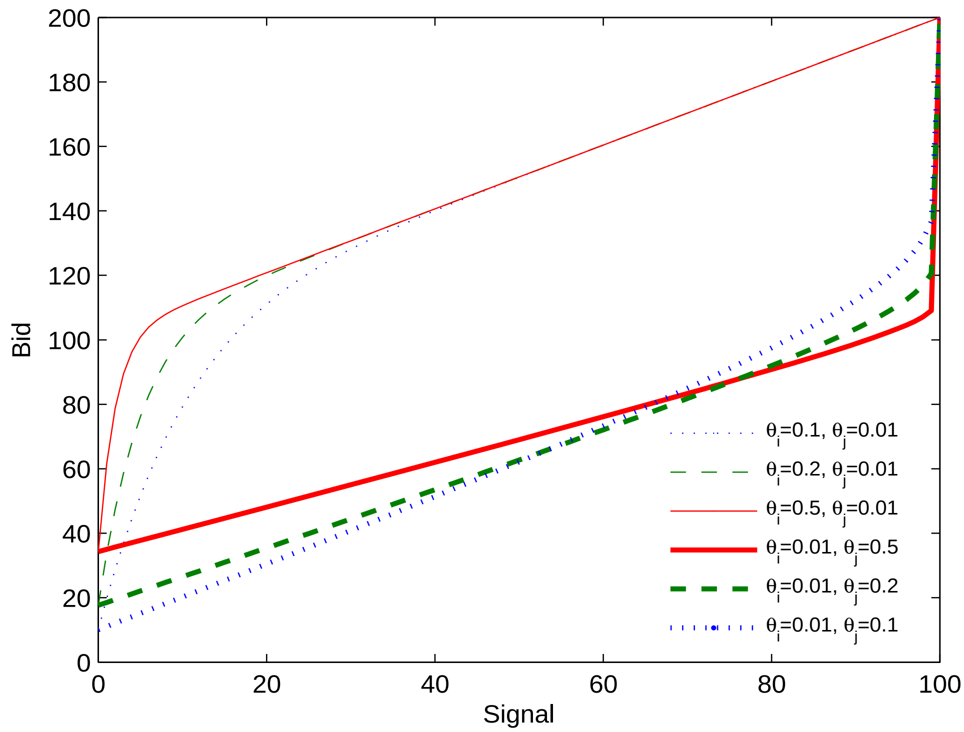

Figure 2.1: The equilibrium bidding functions for $\theta_i = 0.01$ and $\theta_j = 0.05$, 0.2 and 0.5. The lower thick lines represent the bids of the low toehold type, the upper thin lines are the bids of the high type.

Observe that the bidder with the large toehold bids for every possible signal more than in the symmetric case where no bidder has a toehold. On the other hand, the bidder with the smaller toehold bids lower than in the symmetric case for almost all but the smallest values of her signal. Finally it is obvious from the figure that when the difference between the toeholds becomes larger, the high type tends to become more aggressive for all signals he can get. The low type tends to bid less aggressively for almost all of her possible signals.

Results from the theoretical paper that will be useful for our analysis are

a) the probability of winning the auction for agent i is just $\theta_i/(\theta_i+\theta_j)$

b) increasing a bidder's toehold always makes the bidder more aggressive.

c) increasing a bidder's toehold increases her profits regardless of her signal.

2.3 The experimental setup

The experiments were run with undergraduates of all faculties in the LeeX of the Universitat Pompeu Fabra, in Barcelona. No subject could participate in more than one session. Upon arrival students were randomly assigned to their seats. One of the instructors read the instructions aloud and questions were answered in private. Sessions lasted about 1 hour including the reading of the instructions. All sessions presented here were run by computer using z-tree tools (Fischbacher 2007).

Our design consisted of three treatments with two players, one owning a low toehold and the other owning a high toehold. The low toehold was always equal to 1%, the high toeholds were equal to 5%, 20% and 50% respectively. We had one session of the combination 1%-5% (hence treatment 1-5), two sessions of 1%-20% (treatment 1-20) and three sessions of 1%-50% (treatment 1-50). Players alternated roles every turn[2] and the assignment of the toeholds was common knowledge. Note that the treatments we chose are representative of all cases where the toeholds have a ratio of 1/5, 1/20 and 1/50. This means treatments 1-20 and 1-50 should not be dismissed as extreme cases that have no practical relevance.[3]

[2] We had the players alternate roles because of the big asymmetry induced by the toeholds. Theoretically the low toehold types were predicted to make close to zero profits in treatments 1-20 and 1-50!

[3] For the bidding strategies the ratio of the toeholds is of big significance, but the absolute size of the toeholds plays a much smaller role. It is easy to see that the predicted bidding functions are virtually identical between the case of 1-20 and other cases with the same toehold ratio. This includes for example cases that are more frequently found in the field, such as toeholds of 0.1% and 2%. The toehold configuration of 1-20 and 1-50 was chosen in order to make computations easier for the subjects.

Each session consisted of 16 subjects, which were divided into 2 independent subgroups of 8 subjects. This way we obtain two independent observations for each session. Each session consisted of 50 rounds. In each round or period, a signal between 1 and 100 was drawn randomly and independently for every bidder. Subsequently the players participate in an English auction. This means they had in their screen a clock that was constantly ticking upwards. Bidders were considered to be actively bidding until they pressed a key to drop out of the auction. Once they dropped out, they could not re-enter the auction. As usual in English auctions, when all but one players have exited the auction stops. Since we had only two players, once one of them dropped out, the auction ended and the other player was assigned the good. The winner was paid the common value (sum of the two values of the two players) and had to pay the price shown in the clock. Additionally every player received her portion of the price according to her toehold. The information feedback the players received after every round was the value of the asset, the selling price, whether she was the buyer of the asset or not, the gain/loss that she made if she was the buyer of the asset or the gain/loss that she made if she was not the buyer of the asset. Players were given some time to review this information before going to the next round. After every round subjects were randomly matched with the next opponent.

During the experiment, subjects were always able to check the History of the last six rounds they played, with all the relevant information. The rest of the rounds were viewable by using a scroll bar.

The currency of the experiment were Thalers. At the outset of the experiment, each of the subjects received a capital balance of 1000 Thalers. Total gain from participating in this experiment was equal to the sum of all the player's gains and her capital balance minus her losses. If ever the player's gains fell below 0, she would not be allowed to participate any more. Fortunately this did not happen. At the end of the experiment the gains were converted to pesetas at the rate of 1.5 pesetas per Thaler[4].

[4]The peseta has meanwhile given its place to the euro. One euro corresponds to approximately 166 pesetas.

2.4 Experimental results

The main question we are trying to answer is to what extent owning a toehold alters the strategic behavior of a bidder in an English auction. Then we want to see if this change in behavior is translated into a difference in prices.

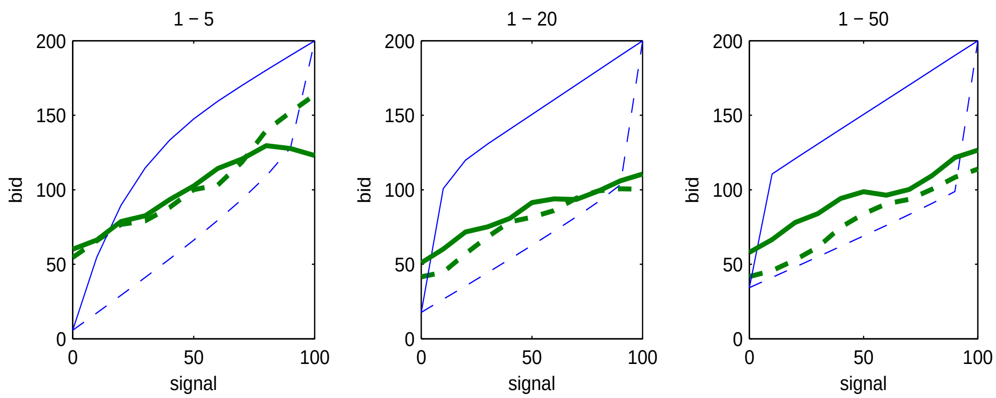

Figure 2.2: Actual (thick lines) vs theoretical bid functions (thin lines) in the three treatments. The dotted lines represent bids of the low type, solid lines are bids of the high type.

We start with the strategies. In Figure 2.2 we have plotted the average exits for the three treatments for given signals[5]. Furthermore, we plot the equilibrium bids of all players. Clearly for the treatment 1-5 and there seems to be no difference in behavior between the two types. For treatments 1-20 and 1-50 high types bid more than low types. Players in general do not follow the shape of the equilibrium bidding functions ie, bidding seems to be linear instead of the highly convex and concave shapes of the equilibrium bids.

Players do not even seem to be influenced by the toehold, when it is low, as can be seen from the fact that in Treatment 1-5 the low toehold type wins approximately half the time, when theoretically she should win only 17% of the time. These results are presented in Table 1. Note that although the signals were drawn at random,

[5]For graphs with details for every individual experiment see Appendix.

	Observation	**should win**	**won**
Treatment 1-5	1	0.2	0.48
	2	0.135	0.54
treatment average		**0.167**	**0.51**
Treatment 1-20	3	0.055	0.495
	4	0.045	0.39
	5	0.015	0.495
	6	0.04	0.425
treatment average		**0.038**	**0.451**
Treatment 1-50	7	0.01	0.37
	8	0.02	0.36
	9	0.03	0.255
	10	0.015	0.345
	11	0	0.385
	12	0.02	0.31
treatment average		**0.0158**	**0.3375**

Table 2.1: Win Frequency of the low type: Theoretic vs Actual

the theoretical ex post winning possibilities are close to the ex ante ones[6] of 1/6 for treatment 1-5, 1/21 for 1-20 and 1/51 for treatment 1-50. In treatment 1-20 the low toehold type still wins more often than she should, and the discrepancy between the theoretical frequency and the predicted one is slightly bigger. In treatment 1-50 the discrepancy between the theoretical winning frequency and the empirical one is smaller. However it has to be noted that the low type should win only about 1.5% of the time, while actually she won in 33.8% of the cases!

In total, there seems to be a tendency for the low toehold type to win less often, the higher the toehold of her opponent. This means naturally that a higher toehold, brings a higher chance of winning, both theoretically and in the experiments. However this effect of the toehold on bidding behavior is not very clear, so we try to estimate its statistical significance. Note, that it is an inherent characteristic of an English auction that we cannot observe the intended bids of the winners, as the winner exits the auction automatically once the one but last bidder leaves. To overcome this we use tobit techniques, or censored regressions (see Kirchkamp, Moldovanu 2004) to

[6] Recall from section 2 that the ex ante probability of player i winning is $\theta_i/(\theta_i + \theta_j)$.

estimate these unobserved bids. The regression we estimated was

$$\text{Bid} = \text{constant} + \alpha^{*}\text{value} + \beta^{*}\text{toehold} + \varepsilon$$

We run this regression for each independent observation. Note the toehold variable is not a binary dummy, but equals the value of the toehold (1, 5, 20 or 50). We add a dummy for the period variable, to control for learning effects. There seemed to be some learning in the first 5 to 10 periods. We always excluded these first 10 periods from the subsequent analysis. Other factors we tried in the analysis, like cubic or interaction terms were not significant and thus are not presented. The results of the regressions for the various treatments are summarized in Table 2.

Treatment	**constant**	**value**	**toehold**	**mean $\mathbf{R}^2$**
1-5	55.876	1.08	0.74	0.67
σ	3.82	0.06	0.79	
	(2/2)	(2/2)	(1*/2)	
1-20	49.53	0.84	0.365	0.51
σ	4.15	0.07	0.19	
	(4/4)	(4/4)	(2**/4)	
1-50	43.92	0.887	0.51	0.60
σ	3.88	0.068	0.08	
	(6/6)	(6/6)	(6***/6)	

Table 2.2: Results of the tobit regressions. There was one regression for each independent session. Numbers in parentheses are significant cases out of total. In treatment 1-5 the one asterisk means that one observation was significant at the 0.1 level. In 1-20 there were two significant observations, both at 0.05. In 1-50 all cases were significant at the 0.01 level.

In parentheses is the number of observations where the coefficient was significant and the asterisks denote the level of significance. Note that the toehold dummy is equal to 5, 20 and 50 in the relevant cases. We observe that in treatment 1-5 the possession of a higher toehold makes almost no difference for the subjects' bidding behavior. However, in 1-20 the toehold sometimes has a significant effect. On average, the high toehold type bid 0.365*20=6.94 more than the low toehold type. In 1-50 the effect of the toehold is always significant and quite high. The high toehold type will bid on average 24.99 more than the low toehold type.

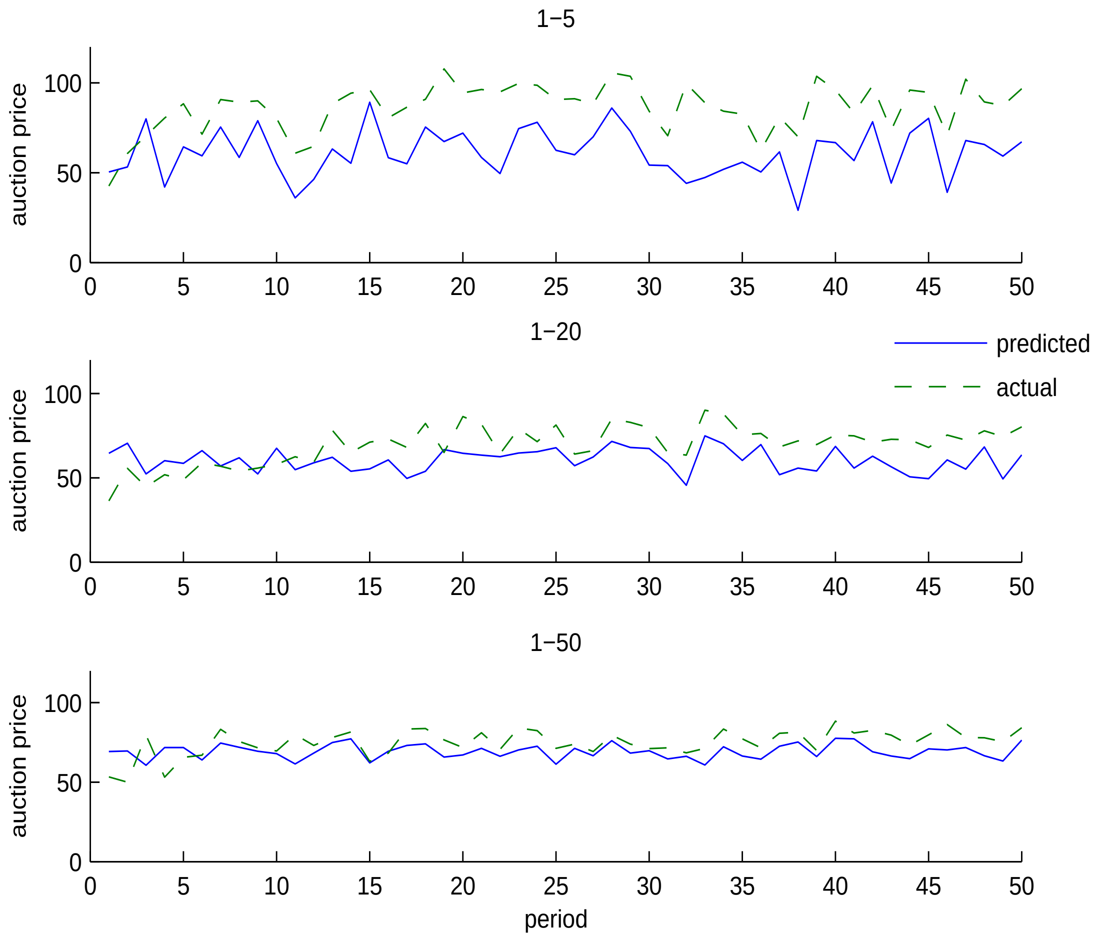

Figure 2.3: Predicted and actual prices over time in the three treatments.

The unique equilibrium predicts[7] prices should fall slightly with the high type getting a toehold 20 instead of 5. This is reflected in our data. The mean price in treatment 1-5 was 89.7, in treatment 1-20 it was much lower at 73.8. Going from a high type with toehold 20 to the high type having 50, the prices were expected to rise by more than 10%, but they only rose to 76.9 which is a 4.2% rise. In general our results mean that *ceteris paribus* the seller's revenues will tend to fall when there

[7]The a priori expected price is $\frac{\theta_j(2\theta_j+\theta_i+1)}{(\theta_j+1)(2\theta_j+\theta_i)} + \frac{\theta_i(2\theta_i+\theta_j+1)}{(\theta_i+1)(2\theta_i+\theta_j)}$, which gives us 63.1, 62.8 and 69.2 for treatments 1-5, 1-20 and 1-50 respectively. For our purposes however we use the theoretical prices given the actual values that the players had, so there is a small difference.

exist players with larger toeholds.

Interestingly the deviation of actual prices from the theoretical ones tended to fall the higher the toehold. Actual mean deviations over all periods were 28 for treatment 1-5, 14 for treatment 1-20 and 8 for treatment 1-50. Note of course that when calculating the mean, positive and negative deviations tend to cancel out. This is why it seems useful to have a look at Figure 4, where we present the evolution of the deviation of observed prices from the equilibrium prices, over time and for the different treatments.

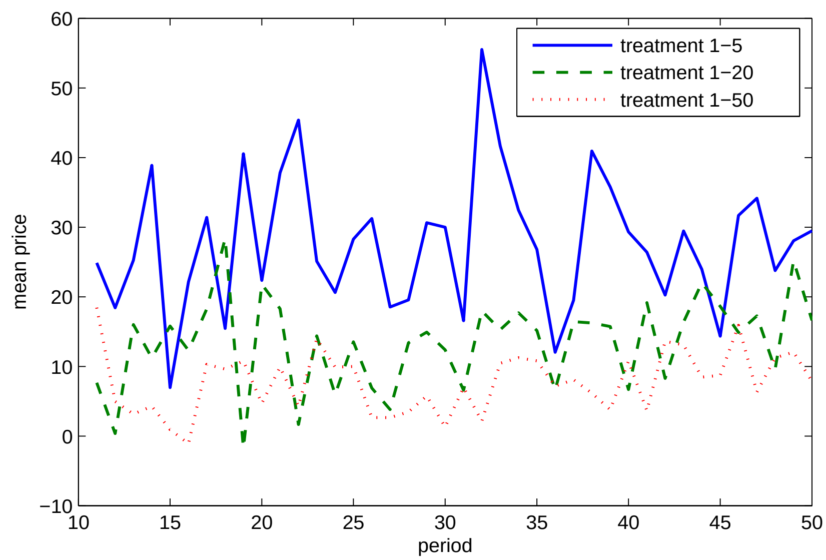

Figure 2.4: Deviation in average prices (actual minus predicted) over time for the three treatments.

The deviation in prices seems to be highest in treatment 1-5, where prices were usually quite a bit higher than predicted by the theory. This is due to the fact that the low types bid more aggressively than they should. In treatments 1-20 and 1-50 the deviation becomes smaller, with a tendency for the deviation to be higher in treatment 1-20. This again can be explained by the fact that low toehold bidders in treatment 1-20 were a bit more aggressive. Table III summarizes these results.

Treatment	Mean actual price	Mean predicted price	Mean deviation
1-5	88.7	61	27.7
1-20	73.8	60.3	13.5
1-50	76.9	69.2	7.7

Table 2.3: Mean prices and deviations from the theoretical predictions, in the various treatments

2.4.1 Theoretical analysis

As we have seen in Figure 2, the subjects' behavior constitutes a deviation with respect to the equilibrium prediction. Does this deviation evade any systematic rational analysis or are subjects responding to a feature of the game that was not obvious from the previous theoretical analysis? Our paper claims that the latter is the case.

There is some literature showing that we should not expect subjects to play the equilibrium strategies if a deviation from these does not cost very much (Harrison 1989). Players will make some small errors when bidding, which produces noise and this noise will be in some way indirectly proportional to the cost of a deviation (see for example McKelvey, Palfrey 1995). To examine this, we will calculate the equilibrium expected payoff functions for each type in every treatment. To be precise, the equilibrium expected payoff functions are the functions which depict one player's expected payoff depending on her bid. The expectation is taken over all possible signals of the opponent, given that this opponent will play the strategy predicted by the Nash equilibrium in Section 2.2.

A closer look at these functions in our experiments, reveals that payoffs are very flat around the maximum. This means that a player anticipating the others to be in equilibrium, will not expect a big punishment for deviating from his equilibrium bid. Figure 2.5 visualizes the concept. The different lines in each of the graphs in figure 5 are drawn for selected signals (0, 25, 50, 75, 100) of a player with toehold 1 (graphs on the left) and those of a player with a high toehold (graphs on the right). The x-axis depicts a players bid and the y-axis the expected payoff given the behavior of the other type, and given the private signal (0, 25, ... ,100). As we can see for the low toehold type the expected payoff is near 0 in treatments 1-20 and 1-50 as theoretically

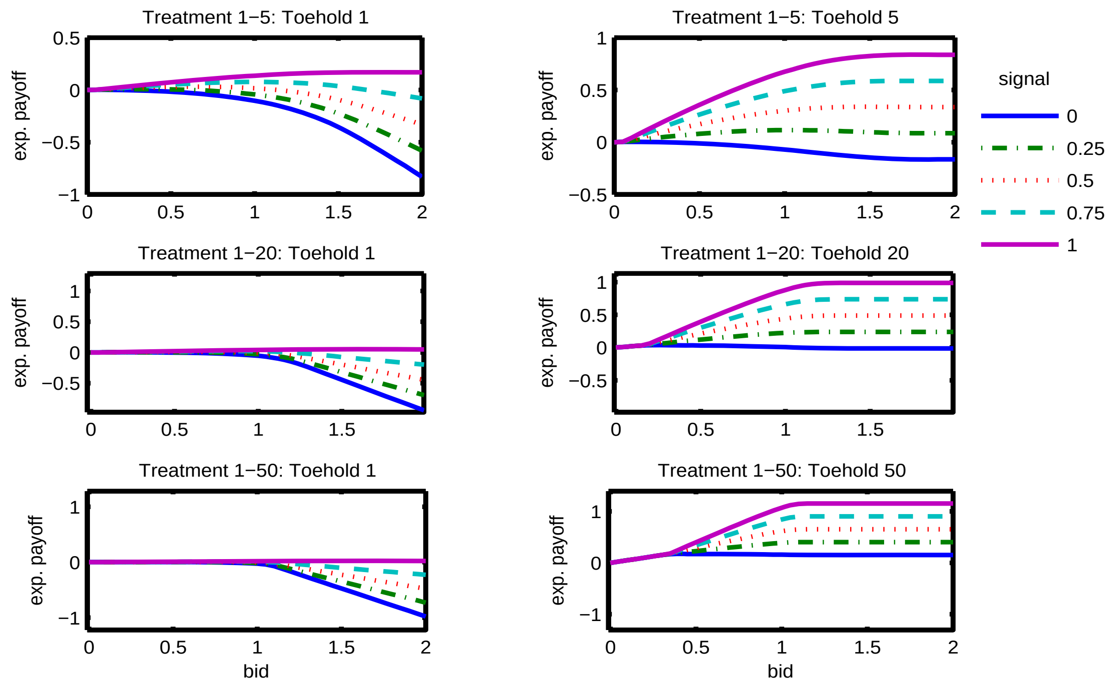

Figure 2.5: Payoff flatness in the various treatments. The various curves depict expected profits depending on bids (both scaled by 100) for signals 0, 25, 50, 75 and 100 given that the opponents play their equilibrium strategies.

the low type never wins. Additionally this flatness is growing with the difference in the toehold sizes[8]. This means the punishment for deviations is smallest in treatment 1-50, where it makes virtually no difference for the high toehold type if she bids even 50% less than the theoretical best response.

The flatness we observe has a quite intuitive explanation and is a general feature of other auction models too, whenever parts of the bidding function are very steep. In the "explosive" equilibria predicted by theory the low types bid very defensively up to a very steep last part. In treatment 1-50 the low type bids less than 140 for almost all signals he gets. This means the high type has no big incentive to bid more than

[8]Recall here that as explained in the design, our treatments are representative of a much wider class of possible configurations. This means that payoffs are flat not only in 1-20 and 1-50 but in all cases where the toehold ratio is greater than 20.

this value, as the probability of winning remains virtually unchanged. This flatness in the payoff functions can explain the difference between the results in our experiments and the usual results in common value English auctions, where bidders tend to follow their equilibrium strategies more closely. In common value English auctions, payoffs are not flat and the payoff maxima are quite pronounced. Thus bidders get stronger incentives to play the equilibrium strategies.

Now, given the flatness of the payoff functions it is interesting to investigate, how big was the deviation of our subjects in the *payoff* space[9]? The reason is that although bid differences might be significant, they could lead to insignificant differences in payoffs, which is what really motivates subjects. Figure 2.6 illustrates the difference between actual and theoretical payoffs in all treatments.

If subjects' payoffs were close to the equilibrium payoffs all dots should lie close to the 45 degree line. We see however that this is not the case. Only in some observations in Treatment 1-20 and in almost all observations in treatment 1-50 are the payoffs of the high type close to equilibrium. The payoffs of the low type are very often away from equilibrium. This is due to the fact that in Treatments 1-20 and 1-50 the low type sometimes wins the auction, although, as we have seen in table I, theoretically she should virtually never win!

So we see, bidders in our experiment had no incentives to play the equilibrium strategy. But their behavior doesn't seem to be completely irrational. Their strategies were approximate best responses[10] to the actual bidding behavior of the others in treatments 1-5 and 1-20, at least qualitatively as we can see in Figure 7. The low toehold types bid more than predicted and the high types less, thus they converge to a middle ground.

It is worth noting that the best responses given actual behavior are not very different between the low and the high type in the first two treatments and even the inter treatment difference is not high. Only in treatment 1-50 do we have a clear separation of the two types. Note that unlike in Kagel and Rose (2006), the high

[9] According to many authors (eg Harrisson 91) this is the naturally relevant space to study.

[10] We get the best responses by calculating the expected payoff given actual bids, and then maximising it. Actually, we calculated the average payoff for each bid in the sample when matched up with every other bid and signal value in the distribution, including that players other bids.

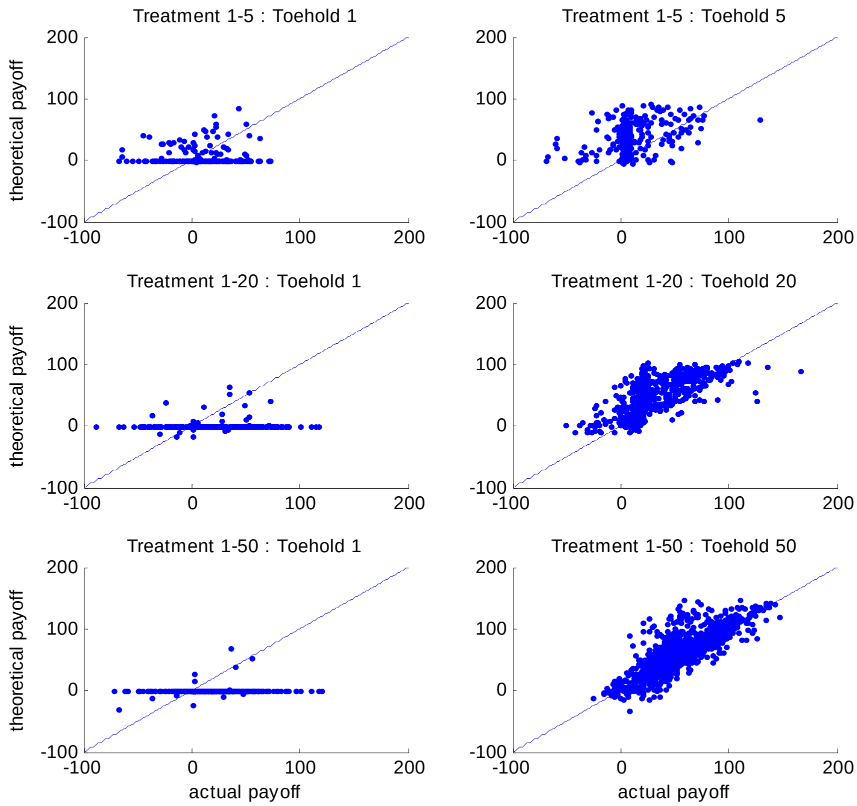

Figure 2.6: Actual vs Theoretical payoffs. The straight lines describe the equilibrium relationship.

type would not have made a much higher profit in expectation, had he chosen the equilibrium bids instead of the actual ones. This is due to the substantial overbidding of the low types, which makes the option of winning less attractive to the high type than predicted by the equilibrium.

We also observe that the expected payoff functions are not so flat, if we calculate them this time assuming that the opponents' strategies follow the actual empirical distribution of the bids. This means that subjects now have higher incentives to play strategies that resemble their best responses. This is visualized in Figure 2.8.

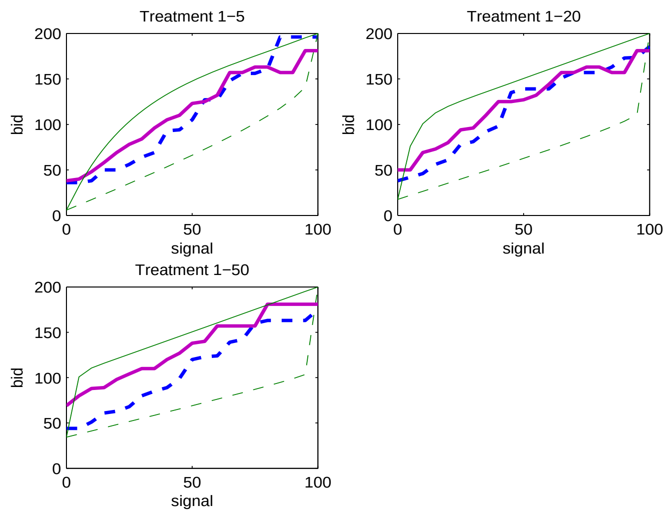

Figure 2.7: Best responses to actual bidding behaviour of the opponents. The dashed lines are the bids of the low type type, solid lines are bids of the high type. The thin lines depict the equilibrium best responses, while the thicker lines depict the best responses to the actual bid distributions.

Bounded rationality

As the shape of the payoff functions is leading to deviations from equilibrium, one could use an equilibrium concept that incorporates the ideas of subjects being influenced by the exact shape of payoff functions. In particular, we could calculate a quantal response equilibrium (McKelvey and Palfrey 1995), where players put weights on their strategies that are proportional in some way to the expected payoff from each action. Unfortunately the calculation of a QRE in auctions with continuous strategy spaces is to date generically impossible. An approximation using a discrete version of the game with a 10x10 bidding space, shows that the QRE would go in the direction we observed.

Given that payoffs are very flat, any kind of learning model would predict very slow

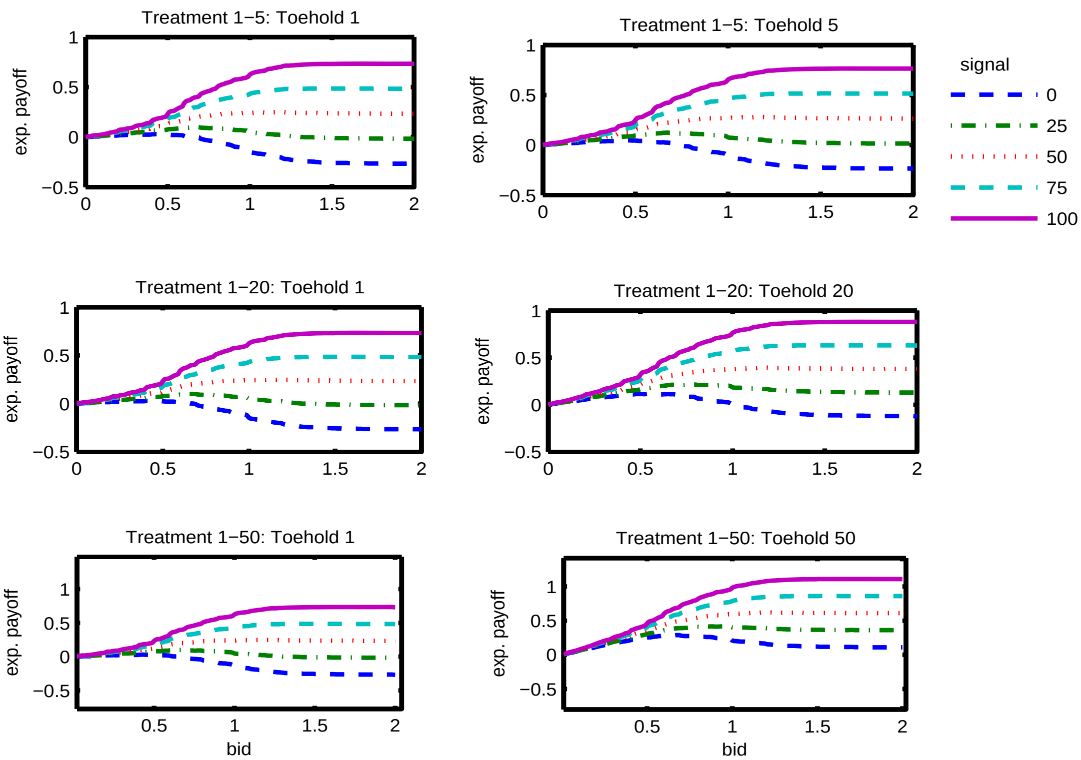

Figure 2.8: Payoff functions given the actual behaviour in the various treatments. The various curves depict expected profits depending on bids (both scaled by 100) for signals 0, 25, 50, 75 and 100.

convergence to the equilibrium. So instead of an equilibrium concept it is interesting to use an explanation that assumes bounded rationality and does not expect subjects to reach an equilibrium, such as a levels of reasoning model (see Nagel 1995, Stahl 1995, Camerer 2004, Crawford 2008). Suppose there exist some Level 0 players who are completely irrational and play randomly. Then the expected payoff of a Level 1 (L1) player who anticipates this behavior is:

$$\Pi_i(b_i) = \Pr\{b_i > b_j\}E[t_i + t_j - (1-\theta_i)p|b_i > b_j] + \Pr\{b_i \leq b_j\}E[\theta_i p|b_i \leq b_j]$$

Since Level 0 bids randomly with a uniform distribution

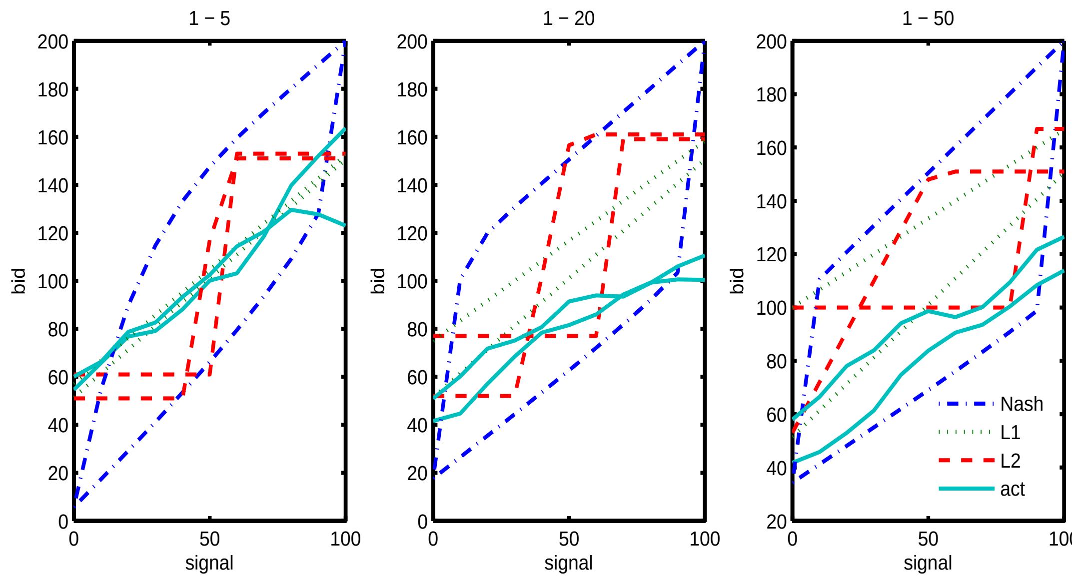

Figure 2.9: Levels of reasoning: actual behaviour vs the Nash prediction and the Level 1 and 2 models.

$$\Pi_i(b_i) = 0.5b_i[t_i + 50 - 0.5(1-\theta_i)b_i] + (1 - 0.5b_i)\theta_i b_i$$

Maximization for a Level 1 player leads to following best response bidding function:

$$b_{L1}(t,\theta) = \tfrac{1}{1+\theta_i}t_i + \tfrac{50+2\theta_i}{1+\theta_i}$$

Note that for a toehold of zero, L1 means the player bids the expectation of the other type's signal (50) plus her own signal, that is just her unconditional expectation for the total value of the company. As toeholds become larger the constant part of the bidding function rises above 50 and the slope falls. L1 bidding is identical to the behavior of a naive player who does not know how the opponents will bid and thus assumes they will bid randomly. Alternatively it describes a player who does not realize that in the event of winning with a given bid his expectation of his opponent's signal has to be updated. In both these cases the player is willing to bid up to the unconditional expectation of his opponent's signal plus his own, known, signal.

Another interesting feature is that for L1 players the size of the opponent's toehold is irrelevant. This is quite intuitive as L1 players do not follow the chain of reasoning that leads to a Nash equilibrium, where bids are usually dependent on the best responses of the others (except if there exists a dominant strategy). In Figure 9 we observe that L1 fits our experimental results rather well in treatment 1-5, much better than the Nash prediction. For treatments 1-20 and 1-50, recall that the bids of the winner are censored. As the high type tends to win more often in these treatments the *observed* exits tend to be more downwardly biased than the underlying bidding strategy. If instead of the observed exits we use the results of the censored regression from table 2, L1 describes the high type's strategies better than the Nash prediction. Also, in general, unlike the Nash equilibrium, L1 bidding predicts a linear, non explosive, effect of toeholds. This is qualitatively in line with actual bidding.

What is missing however is an explanation of the fact that some low toehold types tended to bid a bit less aggressively in treatment 1-50 and 1-20 than in 1-5. Such an effect can be explained when we examine the bidding strategy of level 2 players, who are more sophisticated than L1 players. They anticipate the bidding strategies of L1 and realize that winning against players who use this strategy has implications for their estimate of the asset value. They thus best respond to L1 players. The calculation of these Level 2 strategies is not so simple as above and the players do not use a linear strategy like L1. However, as expected, they do respond to the L1 players in a way that makes low toehold types bid less the higher the toehold of their opponent.

We fit the levels of reasoning model to the data, assuming that the population consists of a mixture of L1 and L2 types, as is found in most experiments in the literature. The model has two parameters, the frequency of the L1 types which is μ and the SD of the normally distributed errors σ which we assume is equal for both types[11]. We also fit the unique Nash equilibrium model assuming normally distributed errors with a SD of σ. A comparison of the models follows in table 2.4.

Overall the mixed L1+L2 model performs better than the Nash prediction and the estimation of the mixture parameter μ is similar across types and treatments. A serious outlier is found in the case of toehold 50 in treatment 1-50. We think the explanation is to be found within the fact that this case suffers most from the aforementioned unobservable final bid problem.

[11] In the presented estimations we forced the individual mixture of levels to be equal to the overall frequency in the population for the same type in the same treatment. We have done calculation with individual estimation of the level and the fit was not enhanced by much, but the number of free parameters grows by the number of subjects. Thus we preferred the more parsimonious model. However it is of interest that the type frequencies found with individual estimation where quite close to previous results at ca. 0.05 for L0, 0.6 for L1 and 0.35 for L2.

	1-5	**5-1**	**1-20**	**20-1**	**1-50**	**50-1**
Nash -LL	798.87	864.32	1744.0	1581.1	2999.8	1734.3
σ	44.81	42.8	32.91	62.11	26.07	55.56
mixed L1+L2 -LL	709.08	758.49	1742.2	1437.5	3125.5	1610.5
σ	24.92	22.71	32.74	37.52	31.72	37.67
μ	0.9316	0.9781	1	0.8591	1	0

Table 2.4: Maximized log likelihoods for the Nash and LOR models.

2.4.2 Does a toehold grant its holder a real advantage?

We can now answer the question if a toehold is beneficial for its holder, at least in the lab. There are two ways to view this, from the *ex ante* or from the *ex post* viewpoint.

In the interim stage, where the company has bought the toehold and is preparing for the acquisition, all the investment the company initially made to buy the toehold is a sunk cost. So the only important questions is: does a toehold raise my chances to win in the auction? Does the expected price fall? As we have seen, the answer is positive in both cases. The bidding in the various experiments depends on the size of the available toeholds. Although the high toehold type does not always win (especially not in treatment 1-5), the auction prices fall monotonically in the size of the high type's toehold. This means the presence of a bidder with a high toehold benefits both bidders, usually asymmetrically, and lowers the revenue that the seller can expect. The choice of what toehold to have is fairly clear cut. As we see in Figure 10, the bidders with a toehold of 50 fared better than the others for almost any private value they had.

The ex ante discussion is a bit more complicated. In particular, it is not generally known how the bidder acquired the toehold in the first case. Let us assume that the price per share paid by the prospective owner of the toehold was reflecting the true value of the company[12], so that for example a 50% toehold of a company of value 100 would have cost exactly 50. Assume additionally that each bidder got a signal of 50. Then we find that buying this toehold was a wise choice for this bidder in case he

[12] This assumption can be justified, if we suppose that shares of the object under sale were floated in financial markets. Then in these markets some informed investors would drive the price to the true value of the company, as in Kyle 1989.

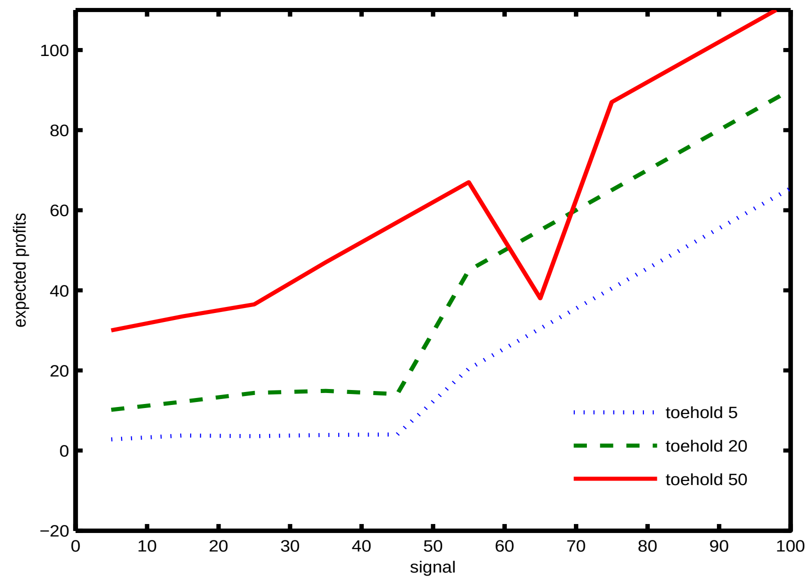

Figure 2.10: Average profits of holders of toeholds 5 (low curve), 20 and 50 (highest curve) in our experiments for different signals.

wins the auction as he gets the rest of the company for only ca. 40 (as we can see in Figure 2), but a suboptimal choice in case he loses, as he would just get ca. 40 for his share of the company, leaving him with a loss of 10.

In general given the average behavior of subjects in our experiments we can calculate the expected profit for a bidder with a signal X if he buys a toehold of 5, 20 or 50 and given that the other bidder has a toehold of 1. The results are depicted in Figure 11. The difference between the ex ante and ex post cases is just the inclusion of the payment for the toehold[13].

We now see the results are now reversed! Acquiring a toehold of 50 is almost never a good strategy. For low signals all toeholds are quite close, but for signals higher than 50 a toehold of 20 is always the best choice.

We could assume a different setup. Imagine the players acquire the toeholds *before* the private signals are drawn. The player does not know his own private signal and

[13] We calculated these expected payoffs assuming that the signal of the other bidder is unknown. Thus we just take its expectation which is equal to 50. It it is of course conceivable that a bidder knows the signal of the other bidder (or has an estimate thereof), but this would completely change the game.

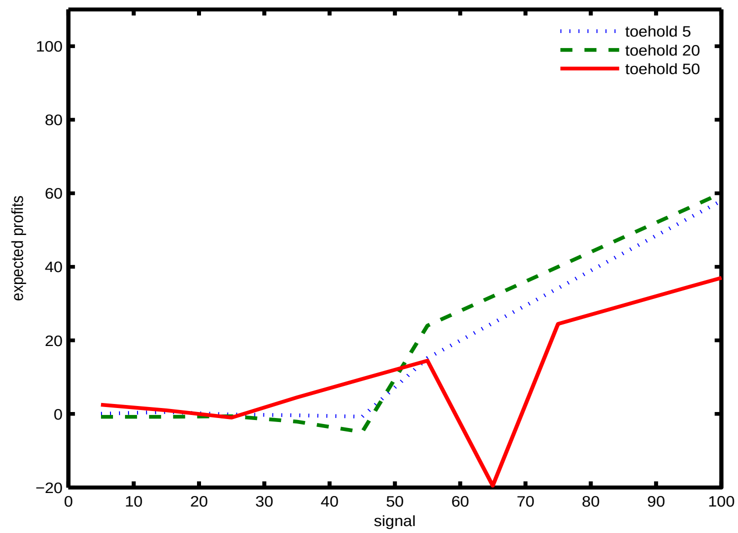

Figure 2.11: Average profits of bidders holding a toehold of 5, 20 and 50, including the expenditure to acquire the toehold, calculated with method 1 (see text).

thus the only information available is the expected value of the company, which equals 100. Then a 5% toehold would cost exactly 5, a 20% would cost 20 and a 50% toehold would cost 50. The results with the toehold prices calculated with this method, are illustrated in Figure 12.

The image is similar to the one above in Figure 11, with the difference that toehold 50 becomes more attractive for high signals and less attractive for low signals. Still in each case we conclude that acquiring a high toehold can sometimes be too costly.

2.4.3 Toeholds and almost common values

Almost common values can be seen as a limit case of the more general toehold framework. In an almost common value auction all but one subjects have the same common value, that is they possess a toehold of zero. The last person has an advantage over the common value, that is, a positive toehold. As the probability of winning in the two person toehold game is equal to $\theta_i/(\theta_i+\theta_j)$ in the limiting case of almost common values the strong type wins with probability one!

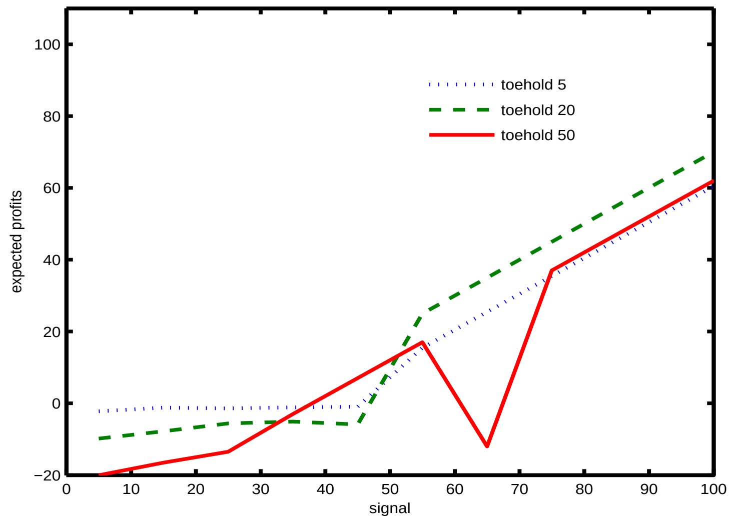

Figure 2.12: Average profits of bidders holding a toehold of 5, 20 and 50, including the expenditure to acquire the toehold, calculated with method 2 (see text).

The size of the private advantage of the strong type, that is the size of his toehold, does not influence her probability of winning theoretically. However Rose and Kagel (2006) find that bidders do not follow the strategies predicted by the explosive equilibrium. The authors find that advantaged bidders won only 27% of the auctions, where 25% would be predicted by chance factors alone. Additionally there was no significant change in average revenue compared to a series of pure common value English auctions.

Combining our results with these findings leads to the following hypothesis: the explosive equilibria are not to be found in real markets. At and close to these equilibria, payoffs are extremely flat, which means subjects have no pressure to play the predicted strategies. Instead they seem to be playing a naive linear strategy. The explanation of Rose and Kagel that the strong type just adds her private advantage to her signal and proceeds to bid like in a pure common value auction, seem to be a plausible first explanation, similar to our L1 model. There is however a feature that remains unexamined: how is the low type playing, how does he respond to a variation of the high type's private advantage? We claim the low type will bid lower the higher the toehold of the opponent, as predicted by the L2 model. Thus, we can make a testable prediction for almost common value auctions. The winning probability of the high type should not be independent of her private advantage as predicted by theory. This is the case because as Rose and Kagel predict, the high type will be *more* aggressive but critically the probability will also rise because the low type as in our experiments will become *less* aggressive in his bidding behavior. This effect however does not converge to the explosive bidding as predicted because due to the flat payoffs subjects do not have sufficient monetary incentives to follow such a counterintuitive strategy.

2.5 Conclusions

We have found that higher toeholds do raise the probability of winning and the profits of their owners. Moreover the seller's revenue tends to fall the higher the discrepancy between the two players' toeholds. However, this fall is not linear, which means that the revenues fall faster when the toeholds are small than when they are greater. We additionally find that these results are not as strong as predicted by theory, although they are broadly in the right direction. Importantly, we show that the high deviations from equilibrium bids are not reflected in high differences of payoffs between actual and equilibrium payoffs, which could thus be an explanation of the subjects' behavior. Our results have some implications for the seller. When one player has a small toehold, it might be of benefit to the seller to award the other buyer some shares to level the playing field.

In general we conclude that small toeholds are not very effective when we observe real human players, in contrast to the theory which predicts a very high effect of even the smallest toeholds. On the other hand, we have seen that big toeholds give their owners a significant advantage in the laboratory. Our result is in support of the empirical literature which finds acquiring companies owning sometimes quite large toeholds. This observation is contrary to the theory which predicts a small advantage would do as well and contrary to the strategic thought which says potential buyers should avoid signalling their intentions by prematurely buying too big shares of the company. Finally, although we find big toeholds to be effective, we show that, under some circumstances, acquiring such large toeholds might be too costly and their cost might not be justified by the advantage one gets in the subsequent bidding for the control of the company.

Chapter 3

Information Revealing Speculation

3.1 Introduction

The owner of a highly valuable and divisible asset contemplating to sell it, is faced traditionally with two options: an auction with a reserve price and an initial public offering (IPO). The typical setting of an IPO is a large market where small players have pieces of relevant information. The information aggregation properties of these markets are the main points of interest. On the other hand, an auction is used to sell an asset to relatively few, big players, who act strategically. Revenue is maximized by manipulating the incentives of the buyers to reveal their true valuations. While the seller is imperfectly informed, the buyers' valuations are not necessarily private information. There may exist third parties with relevant information and inducing them to share it is of great interest to the seller, as it influences her ability to extract surplus from a transaction.

In this paper is that we analyze an alternative two-stage mechanism which brings forth a higher revenue by revealing the information of the third parties and allowing the seller to appropriate more rents from the buyers. It consists of

1. an emission of a minority part of the shares (partial IPO)

2. an auction with optimal reserve price for the rest of the shares.

The two stage mechanism is a mixture of the one stage alternatives and combines their respective advantages. It exploits the presence of the informed agents to gather information about the buyers' valuations and uses this information to set an optimal reserve price in the second stage.

Both one-stage options mentioned above, are often encountered empirically. Russia for example has privatized many companies (Yukos, Sibneft) by direct bargaining with the prospective owners. Germany has organized large auctions for former East German assets (through the Treuhandanstalt). On the other hand, the Japanese monopolistic power utility was recently privatized through an IPO for all shares. But there are cases, where due to a variety of reasons, first a part of the company is sold through an IPO and then the majority is sold to a strategic partner. Such a two-stage method is widely followed in Europe[1] in countries like France (EdF), Greece (Hellenic Telecom, Emporiki Bank), the Scandinavian countries or during the privatization programmes in Central Europe (e.g. Czech Telecom). Gaz de France is another prominent example. After an initial offer of ca. 20% of its shares to the public, is it now being effectively privatized through bargaining with Suez. And lastly there are cases where the public unloads its shares in a series of public offerings, as in the case of Deutsche Telekom, TeliaSonera, National Bank of Greece and others. Thus the mechanisms we examine represent options encountered in real markets. We will argue that the two stage mechanism is the most effective one under some conditions, based on three basic features: the existence of big buyers with control benefits, the existence of small agents with information about these and minority shareholder protection rules in financial markets. We will also argue that the same forces that make our mechanism effective are at play whenever a listed company is considered a possible takeover target and thus influence the market price of shares.

Small informed agents are ubiquitous in financial markets. These can be investment banks, pension funds or even individual analysts, who acquire this information in the course of their everyday business. Before we proceed to a further analysis of their information, it will be of use to dissect the value of the company in two parts:

[1] One can plausibly explain the use of such methods by the need for public revenues without the backlash privatizations tend to bring in these countries. However we do not think this method would be continuously used if it brought consistently suboptimal results.

a cash flow part and a corporate control part (see for example Zingales, 1995). Cash flow rights are enjoyed by all shareholders, in proportion to their equity stake. Therefore we assume the cash flow part is the same for every shareholder, and commonly known. The corporate control part however, depends on who controls the management of the company. Every possible owner of the company has different benefits she can derive from controlling the company, which are known only to herself. These benefits accrue only to the owner and can range from the purely psychological value of being in control (Aghion, Bolton 1992) to perks enjoyed by top executives[2]. An additional reason for private control benefits is that the ownership of some share might affect other shareholdings of an individual or company. For example Porsche recently acquired a 20% stake in fellow carmaker Volkswagen. This control gives it strategic benefits that the other shareholders of Volkswagen do not enjoy.

Control benefits can be quite large.[3] There are empirical studies estimating them, based mainly on the different prices paid for individual shares and for packages of shares carrying the control of the company. The size of this discrepancy is found to be quite significant, for example by Dyck and Zingales (2004) it is estimated at 14% on average. Thus the control benefits are a private value and constitute a sizeable portion of the possible total value of the company to any particular majority owner.

The informed agents have some knowledge of these private values, but no control benefit value of their own. This is due to the fact that the banks and analysts we mentioned above do not have the intention or the capacity to manage the company. They are just buying shares with the speculative motive to resell them at a higher price. Usually these financial investors are liquidity constrained[4] and thus unable to

[2]Perks can be the use of corporate assets and infrastructure, club memberships, special discounts etc but also importantly other indirect benefits. For example, the suitability if a new subcontractor or partner in a new project is not always clear. The person who has the power to choose a partner can expect personal benefits from this choice, without any anticipated damage to, or reaction from, the shareholders.

[3]A spectacular example was observed in the recent takeover of TXU, where KKR and Texas Pacific offered a 25% premium over the average closing price in the 20 days before the offer. The New York Times actually reported the control benefits must be even higher, as the markets (some informed agents?) responded by raising the stock price even higher than the offered by KKR.

[4]Most financial investors are small relative to the size of the companies they study. In the case of large investment banks their size is not that insignificant, but the investors who have the relevant information will belong to a division of this bank, which surely can not use all resources of the

influence the outcome of an eventual sale of 100% of the company through e.g. an auction. Under some circumstances it can be of benefit to the seller if these financial investors somehow revealed their information or participated in the sale. However, the identity of the informed financial investors is unknown, so the use of a direct mechanism to elicit their information is impossible. And a simple auction is not a solution either, as the small financial investors cannot possibly influence the outcome, due to their liquidity constraints. A big impersonal market, e.g. the stock market, is a natural alternative to these mechanisms.

In most important financial markets, small investors are protected by minority shareholder protection regulation, in particular by a *sell out rule*. This rule states that when any investor buys more than a certain percentage of the shares of a company (ranging from 30% to 50%) she has to offer a *fair price*[5] for the shares of all other remaining minority shareholders. Such regulation has very important consequences, as it allows the small investors to acquire stakes in the company using their information regarding its value to a potential buyer, without fearing they will be bypassed in the takeover agreement. The presence of these speculating investors, who just buy to resell, could be a factor raising the revenue of sellers conducting IPOs.[6]

Our model examines the role of these speculators and provides insights in the way minority protection influences the investors' behavior and how it enhances the information aggregation properties of financial markets. We claim this is an unintended and not much studied effect of minority shareholder protection. The usual analysis deals with this protection on the basis of its effects on the efficiency of takeovers (see the seminal paper of Grossman and Hart, 1980) or perceives it as a rule to protect

company. Additionally, regular banks almost never buy majority stakes in a whole corporation from a non-related field. This can be due to regulation and/or diversification reasons.

[5] According to a recent EU directive, a fair offer is an offer equal to the maximum price the acquiring investor has paid for shares of the company under sale, in the recent months. The length of the period considered is allowed to vary in the member states between 6 and 12 months.

[6] From this hypothesis it follows we should expect these investors to be more active, the higher the possibility of an eventual takeover after the IPO. Actually, empirical evidence suggests that many IPOs are followed by an eventual merger or acquisition by another company. Pagano, Panetta and Zingales (1998) find that IPOs are followed by a much higher turnover of control than that of similar privately held companies.

small shareholders from exploitation by the private benefit seeking majority owners. In this paper we show how, further to these effects, minority protection makes markets informationally more efficient, which can benefit all types of investors, be they in the minority or majority. Our model also applies to cases where an agent attempts a takeover of an already listed company. Our results can thus be used to answer questions regarding the reaction of the share price and its information content after the announcement of the takeover attempt. We find that under some conditions the target company can plausibly claim that its share is undervalued, even after the takeover is announced[7].

Our model differs from most IPO papers in the techniques applied, as our focus is on the special informational structure outlined above and how a strategic player (the seller) can use it to his advantage. We want to abstract from other phenomena like the strategic behavior of the underwriting banks and consequent underpricing which are often discussed in this literature on the role and design of IPOs when outsiders can generate information about the firm (see Rock 1986, Benveniste and Spindt 1989). Due to this, we build mainly upon the theoretical literature on financial markets, among others the seminal paper of Grossman and Stiglitz (1980). They use a simple model, where agents can either be fully informed or uninformed. Assuming traders have CARA utility function and that the return of the asset is normally distributed, they are able to find linear equilibria. The authors proceed to analyze how information is conveyed from the informed to the uninformed through the price.[8] We additionally include the standard assumption of noise traders (see for example Hellwig 1980) who bring a stochastic element to the models and allows an only partial information revelation in the markets. The assumptions of CARA utility and normally distributed random values, are crucial in these and most other papers in the literature (e.g. Verrechia 1982, Admati 1985) for the existence of a tractable model with a linear solution. Unfortunately, as we shall see, in our two-stage mechanism it is guaranteed that the posteriors will not be normally distributed, which precludes the use of standard techniques. Due to this, we follow Barlevy, Veronesi (2000) which is one of

[7] Yahoo has recently claimed its share was undervalued, while bargaining over a merger with Microsoft.

[8] In our model information flows from the informed traders to the uninformed seller.

the few tractable models which do not make use of these assumptions. The authors construct a model with a binomial state space and risk neutral informed/uninformed traders.

There are few theoretical models asking similar questions to our paper. Boone and Goeree (2005) explore the sale of an asset when there is a single insider bidder who possesses better information about the asset's risky value and bidders differ in their costs of exploiting the asset. The insider's presence results in a strong winner's curse for the uninformed bidders and devastates expected revenue. The authors show that the optimal mechanism discriminates against the informationally advantaged bidder to ensure truthful information revelation by employing a two stage mechanism. In the qualifying auction, non-binding bids are submitted to determine who enters the second stage, which consists of a standard optimal auction (i.e. second-price auction with an optimal reserve price).

Zingales (1995) focuses on the role of an IPO when there is perfect information about the buyer's impact on cash flow and the control premium. He shows that direct bargaining maximizes the proceeds from the sale of the control right. On the other hand, an IPO is more appropriate to extract rents from cash-flow rights to dispersed shareholders. The decision whether to go public and which fraction to issue depends on the trade off between the two effects. Biais et al (2002) discuss optimal IPO mechanisms when there exist professional investors with private information and liquidity constrained retail investors. However private control benefits and a possible takeover of the company are not considered in this paper. Thus, there is no role for speculating small agents who are covered by minority shareholder protection, which is crucial in our model. In Subrahmanyam and Titman (1999) firms do IPOs because the price revealed in secondary market trading can be useful. This paper shares with our model the market microstructure approach to how information gets reflected in the firm's price. However the analysis focuses on the way that information in the stock market can help entrepreneurs make better production choices. A possible sale of the company and agents' information about the values of potential buyers is not considered.

Section 3.2 introduces the model, section 3.3 presents the results followed by remarks and extensions in section 3.4. Section 3.5 concludes. Omitted proofs and a numerical example including comparative statics can be found in the appendix.

3.2 The model

There are two assets. One is a riskless asset, with return R, scaled without loss of generality to zero. The other asset is a firm with a total value θ, which is the sum of the common value created by the cash flow part plus the private control benefit, which can differ depending on who owns the company. To simplify the setup, we assume there is only one *strategic investor* B interested in acquiring the company and her control benefits are binomially distributed as in Barlevy and Veronesi (2000). We further assume the cash flow part is equal to zero.[9] This gives us following distribution for the total value θ:

$$\widetilde{\theta} = \begin{cases} \overline{\theta} & \text{with prob} \quad \sigma \\ \underline{\theta} & \text{with prob} \quad 1-\sigma \end{cases} .$$

From the discussion of the control benefits, it follows we can assume they are always positive. We have $0 < \underline{\theta} < \overline{\theta}$.

The prior probability σ of θ being high is assumed to be low:

$$0 < \sigma < \underline{\theta}/\overline{\theta} \tag{A.1}$$

As we shall see later, this assumption means that without additional information the optimal take-it-or-leave-it (tioli) offer to a single buyer is $\underline{\theta}$.

There is a continuum of financial investors $i \in [0, 1]$ whose valuation of the firm is zero. These agents all have the same endowment of money, which we set equal to 1.[10] We assume they are risk neutral, so they invest all their endowment in the asset with the highest expected return. This allows us to avoid the usual problems of investors

[9] Any value of the cash flow part, as long as it is deterministic, results in a binomial distribution of the total value.

[10] Setting the endowment equal to one is not restrictive, as for our results only the relative size to w matters.

having a nonlinear demand, as described in the introduction. Additionally, we assume that the financial investors are liquidity constrained and short selling is prohibited. This precludes spending more than their endowment. These assumptions represent the idea that there are many small financial investors, with no market power. All these investors are assumed to be informed[11] of the true value of θ.

The original owner chooses a mass of shares λ of the company to sell in an IPO and $1-\lambda$ to offer subsequently to the strategic investor. He enjoys no private control benefits[12], thus it is always efficient to sell the control of the company to the strategic investor. In the setup we have described so far, the revenue maximization problem of the seller can be quite trivial. Given that the total wealth of the investors is high enough to clear the market, she can offer any mass of shares (though less than 0.5 to avoid ceding management control) through an IPO in the first stage, announcing she will use the price of the IPO as a reserve price for the rest of the shares in the second stage. For all prices less than θ, aggregate demand will exceed the supply as informed investors buy all shares up to a price equal to θ. Note that we have as many possible realizations of the market clearing price as states of the world, in this case two. This invertible price function leads to full information revelation. The seller uses this information to extract all rents in the second stage, by charging the strategic investor her full value to transfer control. The informed, financial investors, subsequently sell all their shares to the strategic investor at a price equal to θ, due to the fair price rule described in the introduction.

Of course this example is highly stylized. The model becomes more plausible when we introduce some noise into the system, which precludes prices from revealing all available information. According with standard practice we include so called *noise traders*, who possess total wealth $w > 0$. Due to exogenous reasons which will not

[11] We will later explain what happens when this is not true and how the key insights of our model transfer to the more general case, where we allow for the presence of uninformed, rational players. Also, the assumption that speculators are perfectly informed can be replaced by the assumption that every speculator gets a noisy signal, with the noise having a zero mean and cancelling out on aggregate. It is straightforward to extend the results of our model to this case, however the analysis would be unnecessarily complicated.

[12] Actually he might have control benefits which we suppose are lower than the potential buyer's. Especially when talking about privatizations, we could speak about the higher efficiency private ownership brings.

be motivated strategically in the current study (for a nice discussion of noise trade see Shleifer and Summers 1990), they spend a random share $\widetilde{x}$ of their wealth buying stock.[13] Let p denote the price of the total asset; total noise trade becomes

$$x_0(p) = \widetilde{x}\ \frac{w}{p} \tag{A.2}$$

Since this model is describing an IPO, we do not allow the noise trade to become negative, i.e. there can be no short sales.

As is usually found in the literature, the seller cannot distinguish between demand coming from the noise traders or from the informed investors, else she could just invert the price function to reveal the state of the world. We furthermore assume that w is large enough to keep the market liquid for a given part of the shares λ that are offered and for any reasonable price below the maximum value of θ,

$$w > \frac{1}{2}\overline{\theta} \tag{A.3}$$

The game proceeds as a sequence of seven steps.

Step 1: Random draw of θ out of a binomial distribution with $prob(\theta = \overline{\theta}) = \sigma$.

Step 2: Choice of λ

The seller S selects a portion $\lambda \in [0, \frac{1}{2}]$ of the company to be sold through an initial public offer. The fraction λ is publicly announced. In the case $\lambda = 0$ steps 3 to 5 are omitted.

Step 3: Random draw of the noise trader wealth investment share $\widetilde{x} \in (0, 1]$. The share $\widetilde{x}$ has a twice continuously differentiable and logarithmically concave density f, which is positive on the whole interval $[0, 1]$. No information about the realization of $\widetilde{x}$ is given to any player.

Step 4: Given a θ, each investor $i \in [0, 1]$ chooses a piecewise continuous demand schedule $x_i(p)$. This schedule assigns a set of demands $x_i(p)$ to every $p \geqslant 0$, with

[13] To allow for prices below $\underline{\theta}$ we need an additional assumption that informed traders are relatively poor, i. e. $w_I = 1 < \lambda\underline{\theta}$

$\sup\{x_i(p)\}p \leqslant 1$ (due to the liquidity constraint). No other player than i receives any information about $x_i(p)$.

Step 5: Market for the stock of the initial public offer

The price p of the stock is determined using the equilibrium in demand functions concept (see Kyle 1989) as follows. Define the aggregate demand of the informed investors

$$x_I(p) = \int_0^1 x_i(p)\, di$$

for values of p such that the integral on the right hand side exists[14].

If the market equation

$$\frac{\widetilde{x}\ w}{p} + x_I(p) = \lambda \tag{3.1}$$

has a smallest solution p_o for p, then p_o is the publicly announced market price. We speak of market failure[15] if no smallest solution p_o exists. In this case the price is set at $+\infty$ and no shares are sold in the IPO.

Step 6:

The seller makes a take it or leave it offer $r \geqslant 0$ to the buyer. This means that S is willing to sell the fraction $1-\lambda$ of the company for $r(1-\lambda)$ money units to B. The offer r is then made public.

Step 7:

[14] Note the demand schedule is a correspondence, as we allow the investors to be indifferent between many demands for a given price. We use here the integral of a correspondence, for a definition see Handbook of Mathematical Economics, p. 206.

[15] There are three possible reasons for a market failure. (a) The integral defining f(p) may not exist for any price. (b) The market equation has no solution. (c) The market equation has no smallest solution (the set of all solutions is an open interval). In real markets such market failure can happen, if for example the computerized systems overload or the software is confronted with unforeseen contingencies. We specify that in the case of market failure no orders are executed. Note that in our equilibrium there will never be a market failure.

The buyer can accept ($\psi = 1$) or reject ($\psi = 0$) the offer r of the seller. The game ends with step 7, unless it has already ended in step 5.

3.2.1 Equilibrium

We focus on pure strategies. A strategy combination will always be a combination of pure strategies and an equilibrium will be an equilibrium in pure strategies. A strategy of a player is defined as a function which assigns a choice at every information set u of the player.

Information sets Player S has one information set u_2 at step 2 and an information set $u_6(\lambda, p_0)$ for every pair (λ, p_0) with $\lambda \in (0, \frac{1}{2})$ and $p_0 > 0$ at step 6.

An investor i has an information set $u_i(\lambda, \theta)$ for every pair (λ, θ) with $\lambda \in (0, \frac{1}{2})$ and $\theta \in \{\underline{\theta}, \overline{\theta}\}$.

Player B has one information set $u_7(\lambda, p_0, r)$ for every triple (λ, p_0, r) with $\lambda \in (0, \frac{1}{2}]$, $p_0 > 0$ and $r \geqslant 0$ at step 7. Player B also has an information set $u_7(0, r)$ for $\lambda = 0$ and every $r \geqslant 0$.

Strategies A strategy φ_S of S assigns a $\lambda \in [0, \frac{1}{2}]$ to u_2 and an offer $r(\lambda, p_0) \geqslant 0$ to every u_6.

A strategy φ_i of an investor i assigns a demand schedule $x_i(\theta, p, \lambda) = \varphi_i(u_i(\lambda))$ to every one of his information sets $u_i(\lambda)$. This schedule must have the properties mentioned in the description of step 4.

A strategy φ_B of B assigns $\varphi_B(u_7) \in \{0, 1\}$ to every information set $u_7(\lambda, p_0, r)$ or $u_7(0, r)$ of player B.

A strategy combination φ is a collection of exactly one strategy φ_S for S, an φ_B for B as well as exactly one strategy for every investor $i \in [0, 1]$. A strategy combination is symmetric if for every λ and θ all investors i gave the same demand schedule $x_i(\theta, p, \lambda) = \varphi_i(\lambda)$. A combination φ' is a deviation from φ, if the strategy of exactly one player is different in φ and φ'. This player is called the deviator from

φ in φ'. A strategy combination is an equilibrium if no deviation φ' from φ yields a higher payoff to the deviator.

Table 3.1 shows the payoffs. Payoffs are calculated by assuming that in case of market failure or rejection of the reserve price, the company is liquidated. Then the payoff of the seller S and the buyer B is zero.

			$\mathbf{A}_S$	$\mathbf{A}_i$	$\mathbf{A}_B$
$\lambda \in (0, \frac{1}{2})$	market price p_0	$\psi = 1$	$\lambda p_0 + (1-\lambda) r$	$(r - p_0) x_i (p_0)$	$\theta - r$
	market price p_0	$\psi = 0$	0	$-p_0 x_i (p_0)$	0
$\lambda = 0$ or market failure	$\psi = 1$		r	0	$\theta - r$
	$\psi = 0$		0	0	0

Table 3.1: Player payoffs. A_S is the payoff of the seller, A_B is the payoff of the buyer, A_i is the payoff of the informed investors.

In the following we focus on the substructure of the game where some $\lambda \in (0, \frac{1}{2})$ has already been chosen[16], that is we treat λ as exogenous and the strategies of the buyer B and the investors i do not depend on λ[17]. We then solve by backward induction. The seller knows the equilibrium strategies of the players in the first stage, this means he knows the (stochastic) equilibrium relationship of the price with the unknown variable θ and thus can build a price rule. He uses it in the second stage to determine his posterior beliefs about the value of the asset, after observing the stock price. We then characterize the optimal take-it-or-leave-it (tioli) offer the seller will make to the strategic investor, focusing on subgame perfect Nash equilibria where the seller is allowed to use a cutoff rule. Using the outcome of the second stage to calculate the returns for the buyers in the first stage, we derive the optimal demand schedules.

3.2.2 Second stage

In the second stage the seller knows the price rule, which she uses to update her beliefs about the state of the world. Given the posterior probability $\widehat{\sigma}$[18] (which is a

[16] For an actual calculation of an optimal λ we refer to the numerical example.

[17] This is true because the best responses of B and the investors i are the same for all $\lambda \in (0, \frac{1}{2})$.

[18] We explicitely determine the seller's posterior beliefs in the next subsection.

function of the observed price p) she offers the asset to the investor for the price r. If the investor rejects her offer, the asset will be liquidated which results in zero payoffs for all parties. The seller's expected revenue, given she observes p and offers r, is:

$$E[\widetilde{v}|\, p, r] = \begin{cases} \widehat{\sigma}(p) \cdot \overline{\theta} & \text{if} \quad r > \underline{\theta} \\ \underline{\theta} & \text{if} \quad r \leq \underline{\theta} \end{cases}$$

It is obvious that all reserve prices other than one of the two realizations of theta are dominated.[19] If the seller charges more than $\overline{\theta}$ the company is liquidated and her payoff is zero. This offer is dominated by $r = \overline{\theta}$ which results in positive revenue if the buyer's control premium is high. The converse holds for reserve prices below $\underline{\theta}$ that are always accepted. A transaction price between the two realizations only occurs if $\widetilde{\theta} = \overline{\theta}$ and is therefore dominated by $\overline{\theta}$. In equilibrium the seller offers

$$r^*(p) = \begin{cases} \overline{\theta} & \text{if} \quad \widehat{\sigma}(p) > \underline{\theta}/\overline{\theta} \\ \underline{\theta} & \text{if} \quad \widehat{\sigma}(p) \leq \underline{\theta}/\overline{\theta} \end{cases} \tag{3.2}$$

and the buyer accepts whenever the reserve price r^* does not exceed his value θ.

The seller's second-stage behavior feeds back to the valuation of financial investors. Thus, to determine the first stage outcome, we have to make a conjecture about the optimal reserve price which in turn depends on the IPO outcome itself. By assumption, only aggregate demand can be observed. Relevant information can therefore solely be revealed by the market-clearing price. Suppose that the price rises weakly monotonically in the true value θ, as we shall show to be true in equilibrium. This monotonicity means that a high IPO price signals a high value of the company. Then, due to the binary state space it seems sensible that the seller's offer will have a single discrete jump in the IPO price. Thus we focus on strategies where the seller uses the

[19] The simple optimization problem the seller faces in the second stage is one major advantage of our binary setup, as it allows for tractability of the model. In most other cases reserve prices can only be defined implicitly.

following cut-off rule p^* in the price interval $(\underline{\theta}, \bar{\theta})$:

$$r^* = \begin{cases} \overline{\theta} & \text{if} \quad p \in (p^*, \overline{\theta}] \\ \underline{\theta} & \text{if} \quad p \in [\underline{\theta}, p^*] \end{cases} \tag{3.3}$$

If the financial investors anticipate this cut-off rule correctly, their demand has to be zero for any price $p \in (\underline{\theta}, p^*]$, independently of their information, since the price will then exceed the proceeds from the second stage ($\underline{\theta}$ according to (3.3)). On the other hand, if the price lies in the interval $(p^*, \overline{\theta}]$ investors' demand does depend on their information. If the actual value and the reserve price (resulting from the IPO price) coincide, they invest all their wealth in the risky asset. In the contrary case however, their demand is zero.

Let us now consider prices outside the interval $(\underline{\theta}, \bar{\theta})$. Financial investors act rationally, i.e. they never buy stock for a price above $\overline{\theta}$. For prices below $\underline{\theta}$ we have to proceed one step further: the seller could now potentially choose a high reserve price, which would give the informed agents a zero payoff in case the company value is in fact low. In such a case the investors' demand schedule x_I should become zero, even for very low IPO prices. On the other hand, if the reserve price is indeed low, the informed agents demand becomes positive for every price lower than $\underline{\theta}$. As we will argue later, the only reserve price consistent with equilibrium behavior for prices below $\underline{\theta}$ is $\underline{\theta}$. Thus the informed investor's equilibrium demand will be strictly positive for $p \in (0, \underline{\theta}]$.

In contrast to the financial investors, noise traders' demand is by construction inelastic to the seller's second-stage decision. Their demand is simply a hyperbolic function in p for any $x > 0$ and contains no information about the state of the world.

To determine the optimal reserve price, the seller is interested in the realization of $\widetilde{\theta}$ but not in the amount of noise $\widetilde{x}$. Both random variables affect the price though. She knows which parties exhibit a positive demand in equilibrium, given a certain price. To be able to update her beliefs about the probability of a high θ, we need that at least one party's demand is elastic with respect to $\widetilde{\theta}$. If the seller for example observes a price below the cut-off point, no information is revealed since the noise traders' demand does not depend on the state of the world and informed traders do not buy in any case. Prices outside the interval $[\underline{\theta}, \overline{\theta}]$ reveal no information either. If p is smaller than the lower bound, informed agents demand in both states of the world while for prices exceeding the upper bound, their demand is always zero. For prices in $(p^*, \overline{\theta})$ the seller has to deliberate about which probability mass to put on combinations of (x, θ) which are consistent with the observed outcome. To be able to run through this procedure we first have to determine the equilibrium relationship of the price with the unknown variables $P(x, \theta)$.

In the next section we derive the equilibrium in the financial market given that the seller uses the cut off rule in the second stage.

3.2.3 First stage

In the first stage informed investors submit demand schedules $x_i(\theta, p)$ given the expected value of $\widetilde{v}$ from the second stage. Recall that these demand schedules can be any piecewise continuous correspondence mapping prices p into non empty subsets of the interval $[0, +\infty)$. An auctioneer receives the demand schedules and calculates the set of market clearing prices and corresponding allocations as described in Section 2. This procedure gives a well defined price for any pair (x, θ), which will be denoted as $P(x, \theta)$.[20]

Notice that in our setup a positive and finite equilibrium price always exists due to the following. Noise trade is always positive but monotonically continuously falling in the price, asymptotically reaching zero as price goes to infinity[21] and going to infinity as p goes to zero. The informed buyer's demand is an upper hemi continuous correspondence in $(0, +\infty)$ except for the upward jump in p^*, going to zero as price goes to infinity. Also recall that short selling is not allowed. Given these facts it is easy to verify that aggregate demand is surjective in $(0, +\infty)$, thus if supply is constant and non trivial or infinite, there is always a positive, finite price at which

[20] Note that due to the liquidity constraints there is no need to account for the case of infinite demands as in the Kyle model.

[21] Note that the liquidity constraints actually imply demand becomes zero for some large enough price.

demand equals supply.

In equilibrium all informed agents maximize utility given the demand functions of the others and the information revealed by the resulting price. Formally, we have

Definition 1 *A symmetric Nash equilibrium in trading strategies is defined as a function $x_i(\theta, p)$ such that x_i solves the maximization problem of the agents conditional upon their information:*

$$\max_{x_i} E_{x,\theta}[\widetilde{v}]x_i + (1 - px_i)$$

We have assumed there is a continuum of financial investors who are price takers. In contrast to the one-stage model the asset's return to the financial investors is determined endogenously. It depends on the realization of the random variable $\widetilde{\theta}$ but also on the resulting price in the aftermarket, $\widetilde{v}$. Recall that in case the asset is not sold, it has to be liquidated[22].

In equilibrium, the informed financial investors correctly anticipate the seller's reserve price decision when the market clears at price p. We have assumed that these agents are perfectly informed about the value to the investor, therefore they can foresee whether a transaction will take place in the aftermarket. If the seller asks for a high price $\bar{\theta}$, no transaction will take place when the true value is low. This leads to a liquidation and zero payoff. On the other hand, a reserve price of $\underline{\theta}$ ensures an efficient transaction but does not extract the full surplus if the control premium is high. The reduced form value function to the informed financial investors is:

$$v(p, \theta) = \begin{cases} \overline{\theta} & \text{if} \quad r^*(p) = \overline{\theta} \text{ and } \theta = \overline{\theta} \\ \underline{\theta} & \text{if} \quad r^*(p) = \underline{\theta} \\ 0 & \text{else} \end{cases} \tag{3.4}$$

We are now equipped with all the ingredients to solve stage 1. Recall that the financial investors are risk neutral and liquidity constrained. Optimal behavior – as defined above– requires that they invest all available funds in the asset with the

[22] At this point one has to mention the Coase Conjecture. To ensure the credibility of the reserve price, the seller could delegate the sale in stage 2 to an agent who is committed to the strategy of selling for the ex-ante optimal reserve price or liquidating the asset otherwise.

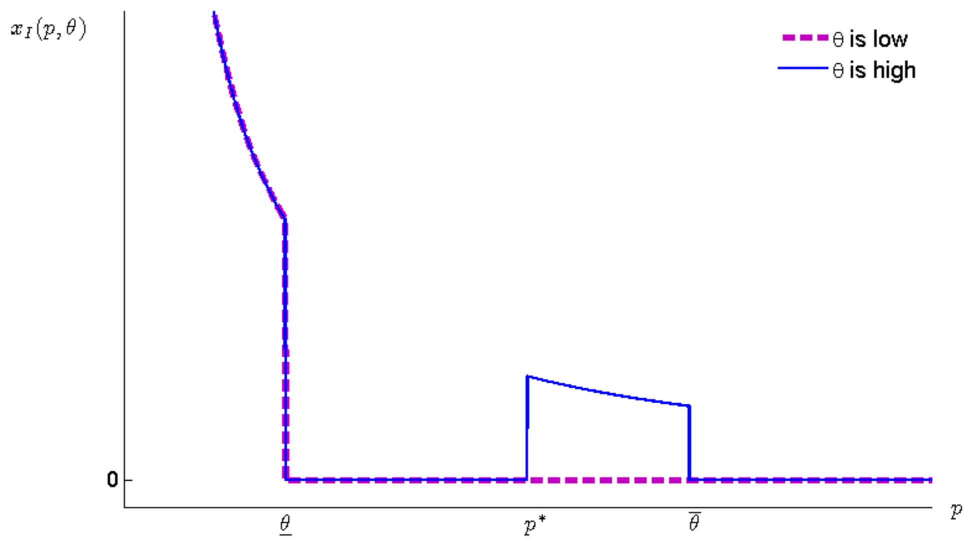

Figure 3.1: Demand of the informed investors depending on price, when the value of the company is high/low.

highest return. Due to the riskless bond having zero yield, demand for stock will be positive as long its value exceeds its price. Thus, the aggregate demand of the informed sector, denoted by x_I, becomes:

$$x_I(p,\theta) = \begin{cases} \frac{1}{p} & \text{if} \quad p < v(p,\theta) \\ \left[0, \frac{1}{p}\right] & \text{if} \quad p = v(p,\theta) \\ 0 & \text{if} \quad p > v(p,\theta) \end{cases} \tag{3.5}$$

Note that no individual informed agent has an incentive to deviate from this strategy as she can not influence the equilibrium price.

Let $P(x,\theta;p^*)$ denote the market-clearing price for a pair (x,θ) given a cut-off point p^*. Using the market clearing condition leads to the following price function in equilibrium[23]:

[23] We limit our analysis to $\underline{\theta} > 1/\lambda$ in order to avoid additional case distinctions.

Proposition 1 *Given that the seller issues a fraction λ of the asset and determines her optimal offer to the buyer by a cut-off rule $p^* \in (\underline{\theta}, \overline{\theta})$. Then there exists an equilibrium in the asset market, where*

1. *Aggregate demand is given by (A.2) and (3.5).*

2. *The market-clearing price is*

$$P(x, \underline{\theta}; p^*) = \begin{cases} \frac{wx}{\lambda} & \text{if} \quad x > \frac{\lambda\underline{\theta}}{w} \\ \underline{\theta} & \text{if} \quad x \in \left[\frac{\lambda\underline{\theta}-1}{w}, \frac{\lambda\underline{\theta}}{w}\right] \\ \frac{wx+1}{\lambda} & \text{if} \quad x < \frac{\lambda\underline{\theta}-1}{w} \end{cases} \tag{3.6}$$

and

$$P(x, \overline{\theta}; p^*) = \begin{cases} \frac{wx}{\lambda} & \text{if} \quad x > \frac{\lambda\overline{\theta}}{w} \\ \overline{\theta} & \text{if} \quad x \in \left[\frac{\lambda\overline{\theta}-1}{w}, \frac{\lambda\overline{\theta}}{w}\right] \\ \frac{wx+1}{\lambda} & \text{if} \quad x \in \left(\frac{\lambda p^*-1}{w}, \frac{\lambda\overline{\theta}-1}{w}\right) \\ \frac{wx}{\lambda} & \text{if} \quad x \in \left(\min\{\frac{\lambda\underline{\theta}}{w}, \frac{\lambda p^*-1}{w}\}, \frac{\lambda p^*-1}{w}\right] \\ \underline{\theta} & \text{if} \quad x \in \left[\frac{\lambda\underline{\theta}-1}{w}, \min\{\frac{\lambda\underline{\theta}}{w}, \frac{\lambda p^*-1}{w}\}\right] \\ \frac{wx+1}{\lambda} & \text{if} \quad x < \frac{\lambda\underline{\theta}-1}{w} \end{cases} \tag{3.7}$$

Proof. The price function is obtained by inserting the informed demand x_I and the noise trade x_0 into the market clearing condition 3.1 and solving for the market clearing price. ■

If the cut-off price p^* is sufficiently close to $\underline{\theta}$ and the realized premium is high then there is an interval of x, each consistent with two different prices: in the low-price case, only noise traders demand while at the higher price also informed traders participate.[24] We have constructed the function in proposition 1 by selecting the higher of the two prices.

Using (3.6) and (3.7), we can calculate the posterior probability $\Pr(\overline{\theta}|p, p^*)$ of the true value being high after the seller has observed a price p. The determination of these beliefs is illustrated in Figure 3.2.

[24] This is a consequence of the non-monotonic demand function of the informed financial investors.

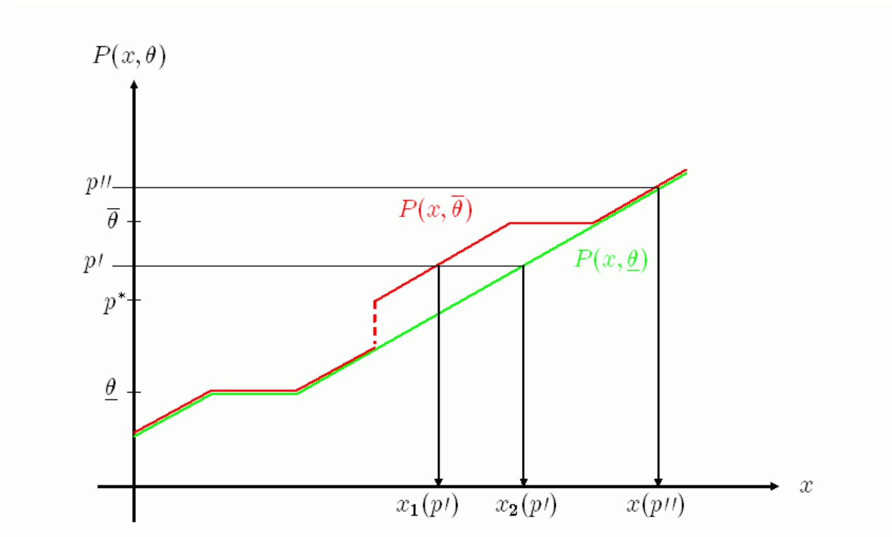

Figure 3.2: Determination of posterior beliefs.

If the seller observes a market-clearing price $p\prime$, then inverting the family of price functions $P(\theta)$ gives her two realizations of $\widetilde{x}$ which are consistent with equilibrium, $x_1(p\prime)$ and $x_2(p\prime)$. From the ex ante distribution of the noise she can infer how likely it is that this price was generated by a high θ or high noise trade. On the contrary, $p\prime\prime$ contains no such information since it is associated with a single realization of $\widetilde{x}$. Full information revelation is only possible for prices which correspond to just one state of the world. Such a price can be found in Figure 2 as a point on the y axis from which a line parallel to the x axis intersects with only one of the two depicted curves.

Algebraically, the application of Bayes' rule[25] leads to the following posterior probability for a high control premium:

25

$$\begin{aligned}\Pr(\overline{\theta}|p) &= \frac{\Pr(\overline{\theta}\cap p)}{\Pr(p)} = \frac{\Pr(\overline{\theta})\Pr(p|\overline{\theta})}{\Pr(p)} \\ &= \frac{\sigma\Pr(p|\overline{\theta})}{\sigma\Pr(p|\overline{\theta}) + (1-\sigma)\Pr(p|\underline{\theta})} = \frac{\sigma}{\sigma + (1-\sigma)\frac{f(\lambda p/w)}{f((\lambda p-1)/w)}}\end{aligned}$$

$$\Pr(\overline{\theta}|p,p^*) = \begin{cases} \sigma & \text{if} \quad p > \overline{\theta} \\ 1 & \text{if} \quad p = \overline{\theta} \\ \xi(p) & \text{if} \quad p \in (p^*, \overline{\theta}) \\ 0 & \text{if} \quad p \in (p^* - \frac{1}{\lambda}, p^*] \\ \sigma & \text{if} \quad p \leq p^* - \frac{1}{\lambda} \end{cases} \tag{3.8}$$

with $\xi(p) \equiv \sigma \left[\sigma + (1-\sigma)\frac{f(\lambda p/w)}{f((\lambda p-1)/w)}\right]^{-1}$.

Prices outside $(p^* - \frac{1}{\lambda}, \overline{\theta}]$ contain no information. Their occurrence stems from high and low realizations of the noise trade component $\widetilde{x}$ respectively. In contrast, if the price hits the upper bound, information is fully revealed due to the indifference of the informed traders. At this price they will demand any amount of shares, which leads to a range of realizations of $\widetilde{x}$ that support $P = \overline{\theta}$. However, when the state of the world is $\underline{\theta}$ the value of the asset is strictly lower than its price. Thus demand is solely driven by noise traders which leads to exactly one x where demand equals supply[26]. Only for prices in $(p^* - \frac{1}{\lambda}, \overline{\theta}]$ can the seller actually update her beliefs.

Now we can proceed to characterize the equilibrium of the game.

3.2.4 Existence of equilibrium

In the next proposition we show which conditions ensure the existence of an equilibrium.

Proposition 2 *Suppose the seller floats a fraction λ of the asset in the financial market. If the distribution f is log-concave there exists a cut-off equilibrium $p^* \in [\underline{\theta}, \overline{\theta}]$ such that the optimal tioli-offer to the buyer is*

$$r^* = \begin{cases} \overline{\theta} & \textit{if} \quad p \in [p^*, \overline{\theta}] \\ \underline{\theta} & \textit{else} \end{cases}$$

[26]This results hinges upon the distribution of $\widetilde{x}$. Since we modeled it as a continuous random variable, every realization of $\widetilde{x}$ which could cause the price to correspond to $\overline{\theta}$ when this is not the true state of the world is a zero-probability event.

Proof. Recall from (3.2) that the seller sets a reserve price $r = \overline{\theta}$ whenever the posterior probability of a high value exceeds $\underline{\theta}/\overline{\theta}$ and $r = \underline{\theta}$ otherwise. According to (3.8), the posterior probability can only exceed the ratio of the two realizations for prices in $[\underline{\theta}, \overline{\theta}]$. Thus, an optimal cut-off has to lie in this interval. Choosing any cut-off point, the resulting posterior will always be, by construction, consistent with optimal behavior for prices below the cut-off: the resulting posterior is either σ or 0, and by assumption A.1 lower than $\underline{\theta}/\overline{\theta}$. This cut-off is suboptimal if there are prices above p^* for which the posterior $\xi(p)$ induces a low instead of a high offer, i.e. $\xi(p) < \underline{\theta}/\overline{\theta}$. Suppose $\xi(p)$ is monotonically increasing. Then there are three possible cases:

1. $\xi(p) > \underline{\theta}/\overline{\theta}\ \forall p \in [\underline{\theta}, \overline{\theta}]$: the posterior always exceeds the ratio and the optimal cut-off is $\underline{\theta}$ (high offer at all prices)

2. $\xi(p) < \underline{\theta}/\overline{\theta}\ \forall p \in [\underline{\theta}, \overline{\theta}]$: the posterior lies always below the ratio and the optimal cut-off is $\overline{\theta}$ (low offer at all prices)

3. $\exists! p^* \in [\underline{\theta}, \overline{\theta}] : \xi(p^*) = \underline{\theta}/\overline{\theta}$: the posterior and the horizontal line at $\underline{\theta}/\overline{\theta}$ intersect just once. Uniqueness and existence are guaranteed by strict monotonicity and continuity of ξ in p. The latter property follows from the differentiability of f.

What remains to be shown is that the log-concavity of f is sufficient for the monotonicity of ξ.

$$\frac{\partial \xi(p)}{\partial p} = -\sigma\left(1-\sigma\right)\frac{\lambda}{w}\frac{f'(\lambda p/w) f\left((\lambda p-1)/w\right) - f'\left((\lambda p-1)/w\right) f\left(\lambda p/w\right) f\left((\lambda p-1)/w\right)^{-2}}{\left(\sigma + \left(1-\sigma\right)\frac{f(\lambda p/w)}{f((\lambda p-1)/w)}\right)^2}$$

This derivative is strictly positive if $\forall p \in \left(\underline{\theta}, \overline{\theta}\right)$:

$$\frac{f'(\lambda p/w)}{f(\lambda p/w)} < \frac{f'\left(\left(\lambda p-1\right)/w\right)}{f\left(\left(\lambda p-1\right)/w\right)}$$

This is true if $\ln(f(.))'' < 0$. The log-concavity of f is therefore sufficient for the monotonicity of $\xi(p)$. ■

The interpretation of the conditions is straightforward. The interior solution requires the seller's decision to change from a low to a high reserve price in $\left[\underline{\theta}, \overline{\theta}\right]$. If she observes a low market-clearing price, the conditional probability of a high control premium has to be sufficiently low to choose $\underline{\theta}$. Or, in other words, the probability of a realization x associated with a low premium has to be sufficiently high, if prices approach $\underline{\theta}$. The contrary has to hold for prices above the cut-off. Note that for the existence of a cut-off equilibrium log-concavity is sufficient but not necessary.[27] Common cases that fulfill log-concavity include the Normal, Poisson, and triangle distributions (see [8] for a detailed survey). When this monotonous ratio condition is violated, any number of cutoff points is possible. Cases 1 and 2 in the proof are corner solutions. To ensure an inner cut-off point we need two additional endpoint conditions (see appendix).

In the next section we compare the expected revenue generated by the two-stage mechanism with two obvious alternatives: a sealed-bid auction with optimal reserve price and a public offering of 100% of the shares.

3.2.5 Revenue comparison

Let us start with the calculation of the expected revenue in an auction for the whole asset. Recall that the financial investors would never plausibly participate, as they are financially constrained and cannot influence the outcome of a second price sealed bid auction with a reserve price. Thus the optimal indivisible good auction with just one strategic investor reduces to an optimal take-it-or-leave-it offer.

The prior distribution of the control premium is such that the seller offers $\underline{\theta}$ and the buyer always accepts. Therefore, the expected revenue from the one-stage auction yields $\underline{\theta}$ in equilibrium.

$$\Pi_A = \underline{\theta} \tag{3.9}$$

In the case of an IPO for 100% of the shares, as we have explained in the previous section, the strategic investor will not participate as she can purchase a majority

[27] Necessary is that the density function is such that once the posterior probability exceeds the indifference point $\underline{\theta}/\overline{\theta}$ it remains above it.

stake later by making a minimal offer to the small investors. In contrast to the noise traders, informed traders anticipate this behavior correctly and demand no stocks at any positive price. The only demand component which drives the price above zero is the noise trade. Therefore, the expected revenue for the seller equals the total expected wealth of the noise traders

$$\Pi_{IPO} = w \int_0^1 x f(x) dx \tag{3.10}$$

The advantage of going public over an optimal tioli-offer is obvious: the seller can fleece noise traders. If sufficiently high probability mass is on realizations below $\underline{\theta}/w$ then an optimal auction outperforms the wholesale IPO.[28]

If the seller chooses the two-stage mechanism instead and issues a fraction λ of the asset, her expected revenue consists of the expected IPO price (Π_{PIPO}) and the revenue from the subsequent bargaining (Π_{RP}):

$$\Pi_{TS} = \lambda \Pi_{PIPO} + (1 - \lambda)\, \Pi_{RP} \tag{3.11}$$

The IPO price serves as a signal for the seller to extract information from the informed financial investors and thus to update her beliefs. She will only switch from $\underline{\theta}$ to $\overline{\theta}$ if a higher reserve price generates a higher expected revenue. It immediately follows that in the two-stage mechanism she will be better off than in the optimal auction regarding the non-issued fraction $(1 - \lambda)$. If an interior cut-off point exists then the expected revenue from the second stage can be written as

$$\Pi_{r^*} = \int_{[0,p^*]\cup(\overline{\theta},\infty)} \underline{\theta} g(p) dp + \int_{(p^*,\overline{\theta}]} \overline{\theta} \Pr(\overline{\theta}|p) g(p) dp$$

with $g(p)$ as the distribution of prices in equilibrium.

[28] If an auction with optimal reserve price is chosen by the seller, this results in a market with a monopolist facing a monopsonist. Such a case should leave the buyer worse off compared to bargaining with a continuum of agents who all possess the same outside option: liquidation of the asset. We claim that any other reasonable bargaining specification should not alter our results qualitatively but would make the analysis more cumbersome. See section 4 for further discussion of this issue.

Let us compare this to the expected revenue of the wholesale auction, which is the lower realization according to (A.1):

$$\Pi_{RP} > \Pi_{OA}$$
$$\Leftrightarrow \int_{[0,p^*]\cup(\overline{\theta},\infty)} \underline{\theta} g(p)dp + \int_{(p^*,\overline{\theta}]} \overline{\theta} \Pr(\overline{\theta}|p)g(p)dp > \underline{\theta}$$

Since $[0, \infty)$ covers the whole support of prices in equilibrium, the first integral can be rewritten as one minus $\underline{\theta}$ times the probability of prices in $[p^*, \overline{\theta}]$. Therefore we get

$$\int_{(p^*,\overline{\theta}]} \overline{\theta} \Pr(\overline{\theta}|p)g(p)dp > \int_{(p^*,\overline{\theta}]} \underline{\theta} g(p)dp$$

Sufficient for this inequality to hold is that the integrand in the left part is pointwise bigger than the integrand to the right, i.e. for all $p \in (p^*, \overline{\theta}]$

$$\overline{\theta} \Pr(\overline{\theta}|p)g(p) > \underline{\theta} g(p) \Leftrightarrow \Pr(\overline{\theta}|p) > \underline{\theta}/\overline{\theta}$$

which follows from Proposition 3. Thus we have shown the revenue per share in the second stage is always higher than in the auction.

Now, it is easy to show that the revenue in the first stage $\lambda\Pi_{PIPO}$ is always greater than the revenue in the full IPO Π_{IPO}. Observe that the revenue in an IPO without a second stage, where only the noise traders participate will always be the same, independently of λ.This is due to the fact that the noise traders always spend their whole wealth so that the price elasticity of demand is always -1, a higher supply λ leads to a one to one reduction in the price and vice versa. This implies $\lambda P_{IPO}(\lambda) = \Pi_{IPO}$.Given that the only difference between the first stage of the two stage mechanism and an IPO without a second stage is the possible extra demand coming from the informed traders, the revenue $\lambda\Pi_{PIPO}$ will always be higher than λP_{IPO} which in turn is equal to the revenue in the full IPO. Thus we see the revenue in the first stage is always higher than the revenue in a full IPO. Combining these two observations we have following result:

Remark 1 *The two stage mechanism performs always better than the full IPO and the optimal auction. The ranking between the optimal auction and the full IPO is ambiguous and depends on the size of the noise trade $E[x]w$ and the prior σ.The higher the noise trade, the more attractive the IPO becomes while the opposite is true when the prior σ becomes higher.*

The first part follows from the discussion above and the fact that λ is chosen optimally.[29] The revenue in the two stage mechanism is a convex combination of two elements that are always respectively higher than the revenues in the two other mechanisms. Thus an optimal λ leads to always higher revenues than each of the two other mechanisms. Actually, as we have shown, the two stage mechanism is better than the the full IPO for any possible $\lambda < 0.5$. In both mechanisms the seller extracts all the noise traders' wealth, but in the two stage one she can also extract a part of the strategic investor's revenue.

A natural question arises as to why a full IPO is not a good alternative, especially if we think that a two stage mechanism is more complicated and probably more costly in reality. The answer lies in the incentives of the financial investors. In the absence of a second stage the financial investors cannot expect a resale and with it a sell-out rule to apply[30]. They cannot try to buy a majority stake themselves in order to resell it, as they are liquidity constrained. Thus their valuation of the shares equal just the cash flow part. The strategic buyer has no incentive to participate in a full IPO either, as he can always wait and make an offer to the small investors after the IPO. Recall there is a continuum of them so they have no bargaining power and the strategic investors will just pay their reserve price, which equals the common part of the valuation of the company.

Another remark is in order here. In real markets IPOs bring along significant underwriting and marketing[31] costs. Assuming a fixed cost of underwriting, the auction

[29] We do not derive the optimal λ explicitely but there is at least one solution in $[0, 0.5)$. For an explicit determination be referred to the numerical example.

[30] In all IPOs no investor gets more than a limited percentage of the company. Buying a majority stake is usually impossible. Examining the optimality of these rules is outside the scope of this paper. However it should be noted that our framework can be useful in analyzing such regulations.

[31] Actually part of the job of the underwriting banks is to raise the amount of noise trade!

can actually be more interesting to the seller than the two stage mechanism. Additionally, having several potential acquirers improves the performance of the auction. To illustrate our results we calculate a simple numerical example and present some comparative statics in appendix B.

3.3 Discussion and Extensions

In this section we will discuss some characteristics of the model and present possible extensions.

An unusual result of our model is that the informed investors benefit from more informative market prices and thus have a preference for low amounts of noise trade. This is in contrast to models like Kyle (1989) where the noise traders are exploited by the informed, who as a consequence prefer markets with plenty of noise trade. The intuition is that in our model informed traders want to signal the value of the premium to the seller accurately, in order for him to set a more beneficial reserve price. Noise traders are only hindering this task.

Also worth noting is the seemingly paradoxical result that in our model more (correct) information can lead to less efficiency. The most efficient reserve price is one set at the lowest value of the premium, where the company is sold for sure. If the seller chooses direct bargaining she sets such a reserve price, given our assumption on the prior. However when the seller uses the two stage mechanism she gets more information and updates her prior. The updated probability of a high premium can induce her to raise the reserve price even if the premium is not actually high! There is an intuitive parallel we can draw with models of imperfect competition where market power in general lowers efficiency. In our model information gives the seller an advantage in the bargaining game which raises his expected revenue but is possibly detrimental for overall efficiency.

Another interesting property of the equilibrium is that it involves a demand function which is discontinuous and non-monotonic in the price. This is due to the cutoff strategy of the seller. Even when the value of the premium is high, informed traders are interested in buying only if the price is above some limit (in our model p^*). For prices under this limit they expect the seller to set a low reserve price in the second stage and thus do not want to pay more than $\underline{\theta}$ for the shares. This characteristic of the demand function is a robust feature of the two stage mechanism and will not fade away if we have more states of the world.

Lastly, a natural extension would be a deeper modelling of the information gathering process. Endogenous information acquisition, where the investors can buy degrees of precision would add generality to our model. Naturally, this addition will make the model very complicated, as we can see in the following section on strategic uninformed traders, but is a promising avenue for future research.

3.3.1 Strategic uninformed traders

So far we have assumed there is a continuum of informed traders and an exogenous amount of noise trade, thus there are no real strategic players in the first stage market. In the other extreme case where all agents act strategically there would actually be no uncertainty, once a market-clearing price is observed, as we have seen in Section 3.2. The non-strategic noise trade component in aggregated demand is responsible for shading the true θ.

Now, what happens if there is still some noise trade but only a fraction z informed about the control premium? Such a structure is consistent with the majority of financial market models.

Starting from the setup we analyzed before we add uninformed investors with total wealth w_U. Recall that total initial endowment of the informed agents was normalized to one. Since all atomic agents are identical except for the information they hold we can express the fraction z of the informed investors in the following way:

$$z = \frac{1}{1 + w_U}$$

Uninformed investors face a similar problem than the seller when they observe a market-clearing price: to which extent is the price driven by noise traders and by informed agents respectively? To determine how their presence alters the outcome we proceed in the following way: we take the original equilibrium and analyze optimal behavior of the uninformed when they enter this market. Demand schedules are submitted simultaneously. Thus in the second step we have to check whether this behavior is consistent with an equilibrium and how it affects the other market participants' strategies. A priori we do not know if the seller's cut-off point is a function of z and how it affects informed investors' optimal demand.

Recall that the inner optimal cut-off point was implicitly defined as the price p^* such that

$$\Pr\left(\overline{\theta}|P = p^*; p^*\right) = \frac{\underline{\theta}}{\overline{\theta}}$$

given the conditions stated in Proposition 3 are met.

Informed agents invest all their wealth in the asset with the highest yield, i.e. they buy stocks whenever the price is lower than their value. Uninformed traders conduct a similar calculation. Since they cannot observe θ, they form conditional expectation of $\widetilde{v}$ (which is a function of θ) based on the price.

Let x_u denote the optimal demand schedule, represented by

$$x_U(p) = \begin{cases} \frac{w_U}{p} & \text{if} \quad p < E[\widetilde{v}|P = p] \\ \left[0, \frac{w_U}{p}\right] & \text{if} \quad p = E[\widetilde{v}|P = p] \\ 0 & \text{if} \quad p > E[\widetilde{v}|P = p] \end{cases} \tag{3.12}$$

Let us ignore the impact of this additional demand component on the market-clearing price for a moment. It is therefore very easy to determine the net benefits as a function of the price. If the price is smaller greater than $\underline{\theta}$ we know that no info is revealed but net benefits are positive. Unless the price exceeds the cut-off point p^*, the seller's action results in a second-stage price of $\underline{\theta}$ and demand will thus be zero in $(\underline{\theta},\ p^*)$. As the price approaches the upper realization net benefits converge to

$$\Pr(\overline{\theta}|p; p^*)\overline{\theta} - \overline{\theta} \leq 0$$

There are two possible cases how net benefits evolve between the cut-off point and the high realization as demonstrated by the next figure.

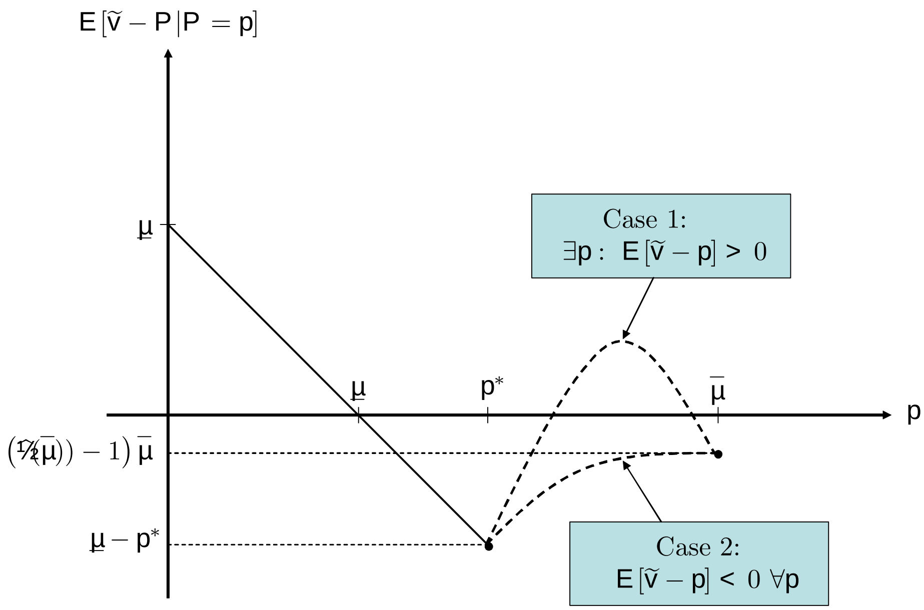

Figure 4: Net value for uninformed investors.

We know from the previous analysis that the conditional probability of a high premium is a continuous function in the price in the upper interval. Therefore, net benefits are also a continuous function and we can distinguish two generic cases:

1. There exists a non-degenerate price interval in the interior of $(p^*, \overline{\theta})$ such that net benefits are strictly positive.

2. Net benefits are negative for all prices greater than $\underline{\theta}$.

In the latter case the only change occurs to the price function for the interval below $\underline{\theta}$. Proposition 3 remains valid. On the other hand in the first case a general analytical result is not attainable.

If a cut-off exists then prices where the uninformed exhibit positive demand are always a strict subset of prices for which the informed investors buy. This means that being informed is truly better independent of the fraction z. If the cost of information acquisition is small enough all investors will be informed in equilibrium. This resembles the situation we analyzed throughout the paper.

3.3.2 More than two states of the world and zero information equilibria

The equilibrium in Proposition 1 is informative, in the sense that the speculators reveal their private information and the noise traders are the only hindrance to full information revelation. In case we have more states of the world such an equilibrium continues to exist. Assume for example there are three states of the world θ_l, θ_m and θ_h with a vector of prior probabilities σ. Let ξ^* be a vector of posterior probabilities for each state of the world (ξ_l, ξ_m, ξ_h) such that, for all $\xi^{'}$ where θ_m and/or θ_h are more likely than in ξ^*, the seller chooses a reserve price of θ_m and conversely for all $\xi^{'}$ that assign a lower probability weight on θ_m or θ_h than in ξ^*, θ_l is chosen. Define ξ^{**} similarly as the vector of probabilities so that the seller switches from θ_m to θ_h. A sufficient condition for the informative equilibrium is that the posterior ξ "crosses" ξ^* and ξ^{**} once and only once, meaning that there is a price such that $\xi_m + \xi_h < \xi^*_m + \xi^*_h$ for all lower prices and for all higher prices $\xi_m + \xi_h > \xi^*_m + \xi^*_h$. The analogous must be true for ξ_h and ξ^*_h. There are parameters for which these conditions hold.

However there also exist zero information equilibria. Suppose for example the informed speculators believe that the company will be sold for θ_l and no other speculator will buy shares. Independently of the number of possible states of the world, if the prior is such that without additional information the seller chooses a reserve price equal to θ_l, such beliefs are self fulfilling. The speculators buy no shares and the seller does not update the prior. This actually leads to a reserve price of θ_l and in this equilibrium the market price reveals no information whatsoever.

Note there are also partially informative equilibria of this type. Assume the speculators believe the company will never be sold for more than θ_m. Then, if the above conditions hold, the speculators will sometimes be able -depending on the noise- to signal that θ is not low by buying shares at the appropriate prices. However they will never buy above θ_m and the seller will never ask for a reserve price higher than this[32]. Thus this equilibrium is also self fulfilling and not all private information is included in the market price.

A way to discard these equilibria is market power. If one informed speculator has enough market power to move the market price and reveal his information, subgame perfection requires that the seller responds by choosing the appropriate reserve price. Since the speculator always wants the reserve price to reflect his information it is optimal for him to indeed buy shares and signal his information. Thus with market power zero information equilibria do not exist.

3.3.3 Bargaining Power

A point that should be discussed is the extreme bargaining power we attribute to the seller, by allowing him to make a take it or leave it offer. Obviously, if we move to the other extreme and the expected allocation does not depend on the seller's information, e.g. if the seller has no bargaining power, the seller will not want to acquire information and the two stage mechanism is rendered useless. Arguably such a case will be very rare.

A more plausible configuration is a bargaining model where the seller has no full bargaining power but the expected allocation still depends on the information the agents possess. For example in Rubinstein and Wollinsky (1985) one of the two parties is selected randomly to propose a split of the gains from trade. In our framework this means the seller will want to acquire information, to choose a better proposal in case she is selected to make an offer. In the simplest case where the bargaining game is played just once and rejection of the offer leads to liquidation the basic results of our model hold. There is no qualitative change, just a quantitative shift in the parameter constellations where a two-stage mechanism dominates the others.

[32] There is an additional condition on the prior probabilities and on the $\theta's$ for this equilibrium to exist. It must be that the seller wants to set a medium reserve price whenever he cannot distinguish between demand coming from a high state of the world and demand coming from a medium state.

3.3.4 Information structure

As a last note, we would like to point out that the special information structure of our model can have an alternative interpretation. Assume there is no control premium, but the value of a company stems just from its discounted stream of dividends. Investors, however, differ in their prognosis of events that can have an industry-wide effect on all firms in the relevant market. Insiders, such as companies in the field or industry-specific analysts, can be assumed to have a superior prediction of the future. On the other hand, governments or large non-focused corporations do not have access to such information. When these large agents try to sell a firm, they can benefit from the information of the small players in a very similar way to the one described above[33].

3.4 Conclusions

Our analysis shows that small financial investors can help a seller extract higher rents from the potential buyers in an IPO. The use of a two-stage mechanism for this purpose yields a higher revenue, under some conditions, than a simple IPO or an optimal auction. Important parameters for a seller contemplating a decision between the three alternatives, are the amount of noise trading in the market, the number of financial investors that can be informed of the value of the company and the number of strategic investors who are possibly interested in acquiring it. With a large number of strategic investors the advantage that a reserve price can give, becomes quite small. On the other hand a large number of informed investors makes the information aggregation through the IPO stronger and the two-stage mechanism more attractive. A great amount of noise trade can have ambivalent effects. It will make the information aggregation in the IPO worse. On the other hand, it will raise the demand for the shares and thus raise the seller's revenues. When there is a lot of noise trade, as might be the case in a bull market, the best option for the seller is to

[33]The only siginificant difference in this case, is that the value of the company to potential owners would be an affiliated/common value and for reasons that become clear when one tries to solve the model, we would like to abstract from problems coming with such a specification.

sell the entire company through a full IPO.

In general we think our framework is a useful tool to discuss the ongoing privatization programs all around the World and the sale of divisions by big corporations. This paper demonstrates that the presence of informed but liquidity constrained financial investors, can explain in part the broad usage of IPOs in the field. Additionally we show that sell-out rules are important for the presence of small informed investors and thus for the informational content of prices in a stockmarket. Our analysis should thus interest regulators contemplating plans to impose sell-out rules in financial markets, such as the recently voted EU takeover directive.

Finally, our results can be useful to analyze the informational content in the price of a company's shares in the stock market once a takeover attempt has been announced. Under some conditions the share price will be an accurate signal for the valuation of the target company. On the other hand, an unfavorable result is also possible. The market can be stuck in a zero information equilibrium and the market prices only reflect noise. However, when some informed agents have market power these zero information equilibria cease to exist.

3.5 Appendix A: Proofs

Let us define the following parameter for the *ex ante* profitability of a high offer in stage 2:

$$\zeta \equiv \frac{\overline{\theta} - \underline{\theta}}{\underline{\theta}} \frac{\sigma}{1-\sigma}.$$

By assumption A.1, $\zeta < 1$, that is without additional information the optimal tioli-offer is $\underline{\theta}$. The following proposition contains technical conditions when the cut-off equilibrium is an interior solution:

Proposition 3 *If f is log-concave and*

$$\frac{f(\lambda\underline{\theta}/w)}{f\left(\left(\lambda\underline{\theta}-1\right)/w\right)} > \zeta \tag{C.1}$$

$$\frac{f(\lambda\overline{\theta}/w)}{f\left(\left(\lambda\overline{\theta}-1\right)/w\right)} < \zeta \tag{C.2}$$

then there exists a unique cut-off p^**, implicitly defined by* $\Pr(\overline{\theta}|P = p^*, p^*) = \underline{\theta}/\overline{\theta}$*, in the interior of* $[\underline{\theta}, \overline{\theta}]$*.*

Proof. In proposition 2 we have shown that log-concavity of the density function is sufficient for the existence of a unique cut-off. It lies in the interior if the posterior starts below $\underline{\theta}/\overline{\theta}$ and eventually exceeds it in the relevant interval. This is fulfilled if

$$\xi(\underline{\theta}) < \underline{\theta}/\overline{\theta} \tag{C.1'}$$

$$\xi(\overline{\theta}) > \underline{\theta}/\overline{\theta} \tag{C.2'}$$

It remains to see that conditions (C.1) and (C.2) are sufficient for (C.1') and (C.2') to be true. C1' can be rewritten as

$$\left(1 + \frac{1-\sigma}{\sigma}\frac{f(\lambda\underline{\theta}/w)}{f\left(\left(\lambda\underline{\theta}-1\right)/w\right)}\right)\frac{\underline{\theta}}{\overline{\theta}} > 1$$

Rearranging terms leads us immediately to the expression in (C.1). Note that necessary for this condition to hold is

$$1 + \frac{1-\sigma}{\sigma}\frac{f(\lambda\underline{\theta}/w)}{f\left(\left(\lambda\underline{\theta}-1\right)/w\right)} > 1$$

since $\underline{\theta}/\overline{\theta} < 1$ by assumption A.1. This is always true as long as the density is strictly positive. If we proceed in the same manner with (C.2') we get

$$\left(1 + \frac{1-\sigma}{\sigma}\frac{f(\lambda\overline{\theta}/w)}{f\left(\left(\lambda\overline{\theta}-1\right)/w\right)}\right)\frac{\underline{\theta}}{\overline{\theta}} < 1$$

which is equivalent to (C.2.). The term in brackets cannot get smaller than one while $\underline{\theta}/\overline{\theta}$ is always smaller than one. This means that the first term must be sufficiently close to one, given the realizations of $\widetilde{\theta}$. In other words, the density must have a sufficiently lower value at $x = \lambda\overline{\theta}/w$ than $x = \left(\lambda\overline{\theta}-1\right)/w$. Note that the lower σ

the more difficult it becomes to meet this condition. The reverse holds for condition (C.1.). ■

Corollary 4 (of Lemma 3) *There is no density function f which ensures the existence of an interior cut-off independent of the other parameters of the model.*

Proof. Suppose that f is log-concave (i.e. the posterior for prices above p^* is monotonically increasing in p) and the parameters $\omega = (\underline{\theta}, \overline{\theta}, \sigma, w, \lambda)$ are such that the conditions of Lemma 3 are fulfilled. Then there exists a interior cut-off p^* defined by $\Pr(\overline{\theta}|P = p^*, \hat{p} = p^*) = \underline{\theta}/\overline{\theta}$. Take a second parameter vector ω' which differs from ω in the first two components such that their ratio does not change:

$$\overline{\theta}' = (1-\epsilon)p^* \text{ and } \underline{\theta}' = \frac{\underline{\theta}}{\overline{\theta}}(1-\epsilon)p^* \text{ for } \epsilon > 0$$

The indifference condition remains unchanged by this transformation of parameters and is fulfilled at a same price p^* which lies now outside the interval. Since such a transformation can be conducted for any density f, the existence of an interior solution has to depend on the parameters of the model. ■

3.6 Appendix B: Numerical Example

In this section we will use a specific density to illustrate our results and conduct some comparative statics. As we have noted, there is a wide range of distributions that fulfill the log-concavity assumption. We choose a simple bounded distribution with enough versatility. *Kumaraswamy's* double bounded distribution [39] has a simple closed form for both its PDF and CDF. In the simplest case, we can take the bounds to be $x \in [0, 1]$ in which case the probability density function is:

$$f(x) = abx^{a-1}\left(1 - x^a\right)^{b-1}$$

The mode for $a, b > 1$ is given by

$$\left(\frac{a-1}{ab-1}\right)^{1/a}$$

With parameters $a = 1.5$ and $b = 8$ we have following density.

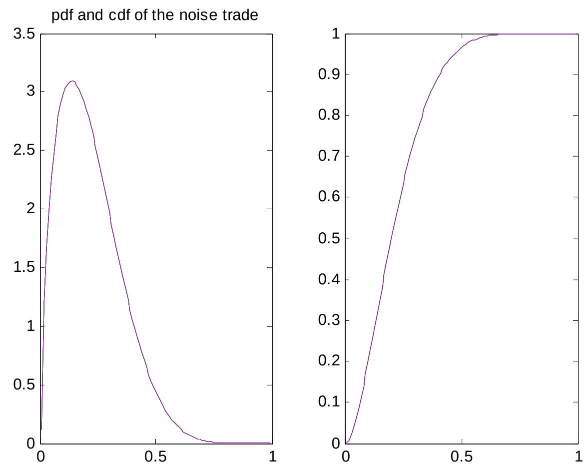

For a given σ we can now find the optimal cutoff p^* depending on the emission size λ. The optimal cutoff is going to fall in λ, as the increased percentage of shares offered makes the noise trader component of the demand less important and thus the price more informative.

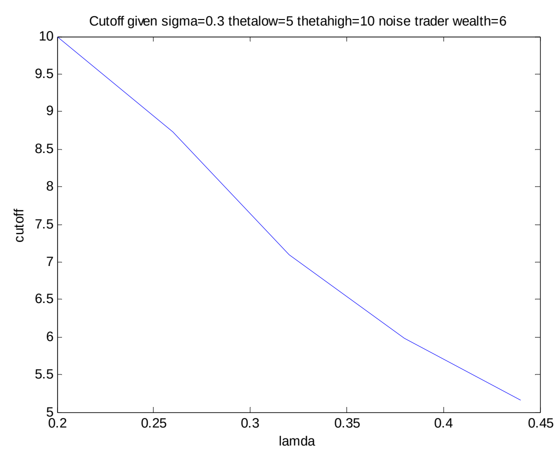

Given this we can calculate the optimal λ by maximizing expected revenues. In following plot we draw the expected revenue for a series of $\sigma \in [0.2, 0.5]$. Observe for high σ the revenue has an inverse U shape so λ^* is an interior solution, while for low values we get a corner solution; the seller tries to issue an infinitesimally small share in the first stage IPO.

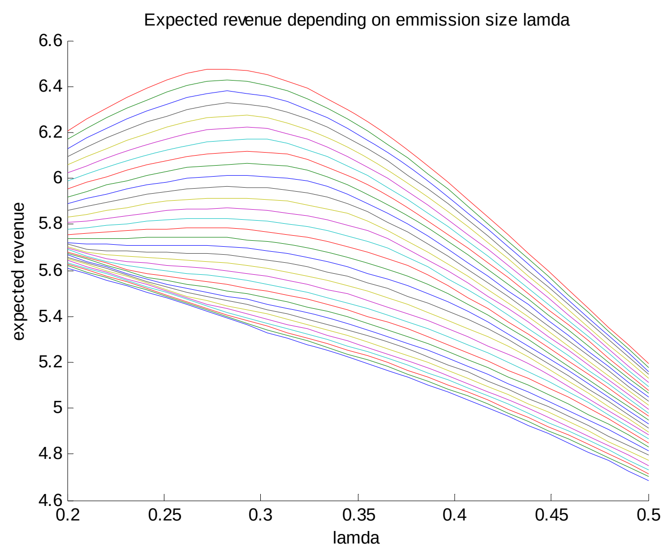

The revenue comparison, as we described in the previous section, is strict. The revenue of the two stage mechanism depends on σ but is always higher than the revenue in two other mechanisms which are independent of σ. The auction is better than the full IPO only when there is not much noise (which means when $E[x]$ is low, or when w is low).

3.6.1 Comparative statics

We will now evaluate our model under different parameterizations of the noise trade distribution. In particular we keep $a = 1.5$ and let $b = 1.1, 2, 4, 8$ and 15.

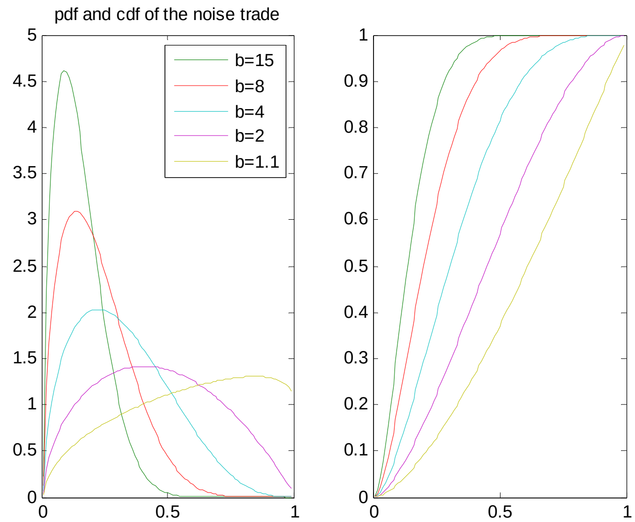

Now we find the optimal IPO issue size λ^* plotted against the prior σ for the different distributions. We will assume there is an exogenously given minimum $\lambda_{\min} = 0.2$. Observe that for $b = 1.1$ and distributions that are even more left skewed (that is have a lower b) we have a corner solution, λ is always low. In general there is no clear relationship between the shape of the distribution and the optimal λ, independently of σ. For low values of σ right skewed distributions lead to a higher λ, while for high values of σ the opposite is true.

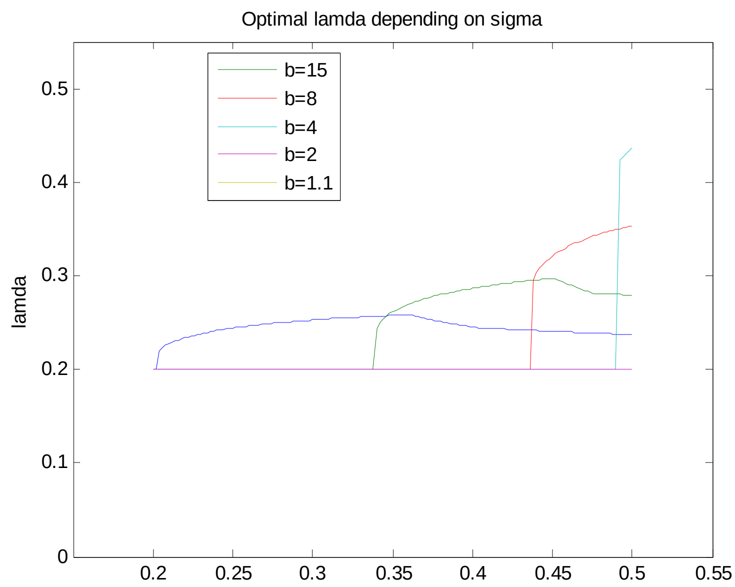

In the next graph we plot the optimal cutoffs depending on sigma for the 5 different parameterizations. Here a clear relationship can be seen. Again for $b = 1.1$ we have a corner solution, but for higher b (more right skewed distributions) the cutoff is falling.

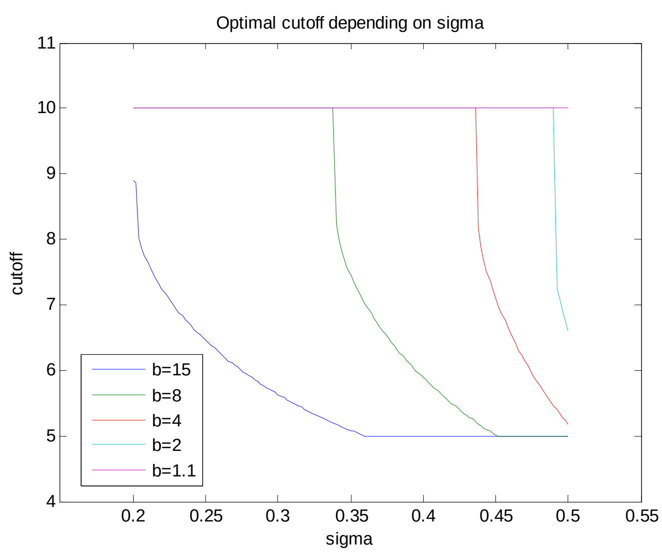

It is also worth investigating what happens for different sizes of the noise trade, that is when we vary the noise trader wealth w. We find that the effect of a change in w affects the optimal cutoff and optimal λ in a very similar way as a change in the distribution parameters. That is, a higher w makes the cutoff p^* always smaller, but the effect on λ is ambiguous and it depends on σ.

Bibliography

[1] Andreoni, James, Yeon-Koo Che and Jinwoo Kim, (2007). "Asymmetric information about rivals' types in standard auctions: An experiment," *Games and Economic Behavior*, Elsevier, vol. 59(2), pages 240-259, May.

[2] Anderson, Simon P., Jacob K. Goeree and Charles A. Holt, (1998). "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction," *Journal of Political Economy*, University of Chicago Press, vol. 106(4), pages 828-853, August.

[3] Avery, Christopher & John H. Kagel, (1997). "Second-Price Auctions with Asymmetric Payoffs: An Experimental Investigation," *Journal of Economics & Management Strategy*, Blackwell Publishing, vol. 6(3), pages 573-603, 09.

[4] Ausubel, Lawrence M. and Peter Cramton (1999) "The Optimality of Being Efficient," Working Paper, University of Maryland

[5] Admati, Anat R, (1985). "A Noisy Rational Expectations Equilibrium for Multi-asset Securities Markets," *Econometrica*, Econometric Society, 53(3), May, 629-57.

[6] Aghion, Philippe, and Patrick Bolton, (1992), "An incomplete contract approach to financial contracting", *Review of Economic Studies*, 59, 473-494.

[7] Arrow, Kenneth J & M.D. Intriligator (ed.), (2000). "Handbook of Mathematical Economics," *Handbook of Mathematical Economics*, Elsevier, edition 4, volume 1

[8] Bagnoli, Mark and Ted Bergstrom, (2005). "Log-concave probability and its applications", *Economic Theory*, 26, 445-469.

[9] Barlevy, Gadi, and Pietro Veronesi, (2000), "Information Acquisition in Financial Markets", *Review of Economic Studies*, 67, 79-90.

[10] Boone, Jan & Jacob K. Goeree, (2005), "Optimal Privatization Using Qualifying Auctions," CEPR Discussion Papers 5118, C.E.P.R. Discussion Papers.

[11] Bulow, Jeremy & Ming Huang & Paul Klemperer, (1999). "Toeholds and Takeovers", *Journal of Political Economy*, University of Chicago Press, vol. 107(3), pages 427-454

[12] Benveniste, L. and P.A. Spindt, (1989), "How Investment Bankers Determine the Offer Price and Allocation of New Issues," *Journal of Financial Economics*, 24, 343-362.

[13] Betton, Sandra & B. Espen Eckbo, (2000), "Toeholds, Bid Jumps, and Expected Payoffs in Takeovers" *Review of Financial Studies*, vol. 13 (4), pages 841-882

[14] Biais, Bruno, Bossaerts, P. and J.-C. Rochet, (2002), "An Optimal IPO Mechanism," *Review of Economic Studies*, 69, 117-146.

[15] Camerer, Colin; Teck-Hua HO; and Juin-Kuan Chong (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119, 861-898.

[16] Calzolari, Giacomo and A. Pavan (2002) "Monopoly with resale", mimeo University of Toulouse

[17] Camerer, Colin; Teck-Hua HO; and Juin-Kuan Chong (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119, 861-898.

[18] Crawford, Vincent and Nagore Iriberri (2008) "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, forthcoming

[19] Cox, James C.; Vernon L. Smith, and James M.Walker (1992), Theory and Misbehavior of First-Price Auctions: Comment, *American Economic Review*, 82, 1392-1412

[20] Dyck, Alexander & Luigi Zingales, (2004). "Private Benefits of Control: An International Comparison", *Journal of Finance*, American Finance Association, 59(2), 537-600.

[21] Fehr, Ernst and K. Schmidt (1999), "A Theory of Fairness, Competition, and Cooperation", *Quarterly Journal of Economics*, 817-868

[22] Garratt, Rod and T.Troeger (2006) "Speculation in Standard Auctions with Resale" *Econometrica* 74, 753-769.

[23] Georganas, Sotirios (2003) "Auctions with Resale: An Experimental Study", masters thesis, University of Bonn

[24] Goeree, Jakob K. and C. Holt (2001) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402-1422

[25] Goeree, Jacob K., Charles A. Holt and Thomas R. Palfrey (2002). "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory*, Elsevier, vol. 104(1), pages 247-272, May.

[26] Gupta, M. and B. Lebrun (1999) First Price Auctions with Resale, *Economics Letters*, 64, 181-185

[27] Grossman, Sanford J & Oliver D. Hart, 1980 "Takeover bids, the free-rider problem and the theory of the corporation" *The Bell Journal of Economics*, vol. 11 (1), pages 42-64, Spring

[28] Grossman, Sanford J & Stiglitz, Joseph E, (1980). "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, American Economic Association, vol. 70(3), pages 393-408, June.

[29] Harrison, Glenn (1989), "Theory and Misbehavior of First-Price Auctions", *American Economic Review*, 79, 749-762

[30] Hafalir, Isa Emin and Vijay Krishna, "Asymmetric Auctions with Resale", *American Economic Review*, forthcoming

[31] Haile, Philip A. (1999) "Auctions with Resale", working paper

[32] Haile, Philip A. (2000) "Partial Pooling at the Reserve Price in Auctions with Resale Opportunities," *Games and Economic Behavior*, Elsevier, vol. 33(2), pages 231-248, November.

[33] Haile, Philip A. (2001) "Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales" *American Economic Review*, vol. 91, issue 3, pages 399-427

[34] Haile, Philip A. (2003) "Auctions with private uncertainty and resale opportunities", *Journal of Economic Theory* 108, 72-110

[35] Haile, Philip A., Ali Hortacsu and Grigory Kosenok (2006) "On the Empirical Content of Quantal Response Equilibrium", working paper

[36] Harrison, Glenn (1989) "Theory and Misbehavior of First-Price Auctions", *American Economic Review*, 79, 749-762

[37] Hellwig, Martin, (1980): "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory*, 22(3), 477-498.

[38] Jehiel, Philip and Benny Moldovanu (1999) "Resale Markets and the assignment of property rights", *Review of Economic Studies* 66, 971-999

[39] Kumaraswamy, P. (1980): "A generalized probability density function for double-bounded random processes," *Journal of Hydrology*, 46, 79–88.

[40] Kagel, John H. ; Ronald M. Harstad and Dan Levin (1987) "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study", *Econometrica*, 55, 1275-1304

[41] Kagel, John H. and Alvin E. Roth (1995) The Handbook of Experimental Economics, Princeton University Press, 1995

[42] Kagel, John H & Dan Levin (2003), "Almost Common-Value auctions revisited", *European Economic Review*

[43] Kagel, John H & Susan Rose (2006), "Bidding in Almost Common Value Auctions: An Experiment" submitted to *Review of Economics and Statistics.*

[44] Kagel, John H. (1995) Auctions: A survey of experimental results In: Kagel, John H. and Alvin Roth (eds.) The Handbook of Experimental Economics, Princeton University Press, Princeton NJ

[45] Kamecke, Ulrich (1994) Competitive Bidding in an English Multi-object Auction: An Experimental Study, Discussion paper, University of Bonn

[46] Kirchkamp, Oliver and Benny Moldovanu, (2004), "An experimental analysis of auctions with interdependent valuations", *Games and Economic Behavior,* Vol. 48 (1), July, pp. 54-85

[47] McKelvey, Richard R. and Thomas R. Palfrey (1995), "Quantal Response Equilibrium for Normal Form Games", *Games and Economic Behavior* 10, 6-38

[48] Lange, Andreas, John A. List, Michael K. Price (2004) Auctions with Resale When Private Values Are Uncertain: Theory and Empirical Evidence, NBER Working Paper No. w10639

[49] McAfee, R.Preston (1998) "Four issues of market design", *Revista de analisis economico* 13, 7-24

[50] Milgrom, Paul R (1987) Auction Theory in T.F.Bewley ed., Advances in Economic Theory: Fifth World Congress, Cambridge: Cambridge University Press, 1-32

[51] Milgrom, Paul R. and Robert J. Weber (1982) "A Theory of Auctions and Competitive Bidding", *Econometrica* 50, 1089-1122

[52] Nagel, Rosemarie; Antoni Bosch-Domènech; Albert Satorra and José García Montalvo (2002) "One, Two, (Three), Infinity: Newspaper and Lab Beauty-Contest Experiments", *American Economic Review*, 92, 5,

[53] Nagel Rosemarie (1995) "Unraveling in Guessing Games: An Experimental Study", *American Economic Review* 85(5), 1313-1326

[54] Ockenfels, Axel and A. Roth (2002) "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet", *American Economic Review*, 92 (4)

[55] Pagano, Marco & Fabio Panetta & and Luigi Zingales, (1998). "Why Do Companies Go Public? An Empirical Analysis," *Journal of Finance*, American Finance Association, 53(1), 27-64.

[56] Rock, Kevin, 1986, "Why New Issues Are Underpriced," *Journal of Financial Economics*, 15, 187-212.

[57] Shleifer, Andrei and Lawrence H. Summers (1990) "The Noise Trader Approach to Finance," Journal of Economic Perspectives, American Economic Association, vol. 4(2), pages 19-33, Spring.

[58] Stahl, Dale and Paul Wilson (1995): "On Players' Models of Other Players: Theory and Experimental Evidence", *Games and Economic Behavior*, 10, 218-254.

[59] Subrahmanyam, A. and S. Titman, (1999), "The Going-Public Decision and the Development of Financial Markets," *Journal of Finance*, 54, 1045-1082.

[60] Zingales, Luigi, (1995), "Insider Ownership and the Decision to Go Public", *Review of Economic Studies*, 62, 425-448.