Impulse Balance Equilibrium and Feedback in First Price Auctions

Axel Ockenfels*

Max Planck Institute for Research into Economic Systems, Germany

Reinhard Selten

University of Bonn, Germany

March 1, 2002

Abstract

Experimental sealed bid first price auctions with private values in which feedback on the losing bids is provided yield lower revenues than auctions where this feedback is not given. Furthermore, bids tend to be above the equilibrium predictions for risk neutral bidders. While the latter observation has been rationalized in terms of risk aversion, rational bidding is invariant to feedback on losing bids. We propose the concept of weighted impulse balance equilibrium that is based on learning direction theory and that incorporates a concern for social comparison. The one-parameter model captures both overbidding and the feedback effect. (*JEL* C7, C9)

* Correspondence: Axel Ockenfels, Max-Planck-Institute for Research into Economic Systems, Strategic Interaction Unit, Kahlaische Straße 10, D-07745 Jena, Germany; email: ockenfels@mpiew-jena.mpg.de. We thank Jeannette Brosig and Marco Bubke for their help in running the experiments at the Magdeburg Laboratory for Experimental Economics (MaxLab), and Bettina Büttner for helpful comments. Ockenfels gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft through the Emmy Noether-program.

I. Introduction

This paper presents an experimental investigation of two feedback treatments in the symmetric sealed bid first price auction with private values. The results throw light on different explanations of the bidding patterns reported in the literature. Cox, Roberson and Smith (1982) proposed the constant relative risk averse model (CRRAM) in order to account for the fact that bids tend to be higher than at risk neutral Nash equilibrium (RNNE) – a phenomenon called 'overbidding'. The model assumes that subjects maximize the expected value of a utility of the form $u(x) = x^r$ where u is the utility index, x is the amount of money gained, and r is the risk aversion parameter. Bidders are assumed to play Nash equilibrium strategies on this basis. The model is consistent with overbidding and fits the experimental data reasonably well (Cox, Smith and Walker, 1988). However, experiments by Selten and Buchta (1999) raise some doubts about the validity of an optimization approach such as CRRAM.[1] In these experiments, subjects did not make bids but had to construct bid functions on the basis of which bids were determined. The results showed that bid functions were not constant during the 50 rounds for which the auctions were played, but changed considerably in response to the experiences made in the game. In particular, the bid for the last observed value tends to change in accordance with learning direction theory. In the case that the price was higher than the bid, the theory indicates an increase of the bid, if it is changed at all. On the other hand, if the bid was successful, a decrease of the bid is indicated. Nevertheless, one might argue that a situation in which subjects have to specify a bid function rather than a bid is somewhat artificial and therefore does not necessarily exclude the possibility that CRRAM is the right explanation for auctions under a more natural response mode.[2] Therefore it seems to be important to examine whether learning direction theory is also confirmed by auctions in which subjects directly choose bids. It will be shown that this is the case.

The predictions of optimization models like CRRAM do not depend on the feedback given after the auction. It has been observed in the literature, however, that feedback can make a difference. Isaac and Walker (1985) compared sealed bid private value auctions, in which feedback was given only on the price, to other auctions, in which bidders were additionally informed about the losing bids after the auction. They observed that prices generated with limited

[1] Other studies also challenged CRRAM's explanation of overbidding (e.g., Kagel and Roth, 1992); see Kagel (1995) for a survey and discussion of the debate.

[2] Jim Cox argued this way in a public discussion.

feedback are higher than those generated under the full feedback condition. The experiments reported here are also run under these two conditions and replicate the effect. We will provide an integrated explanation of behavior in both types of auctions on the basis of learning direction theory.

The auction design underlying the experiments presented here tries to create a situation in which the predictions of learning direction theory can easily be examined. For this purpose the experiments were run with independent subject groups of six in which each bidder played five times with a constant own value in two-person auctions against all other subjects in the group. The interaction was anonymous and the order of opponents was randomly varied. This procedure was repeated 28 times for the same group of six subjects.

We derive quantitative predictions from the concept of weighted impulse balance. The model explains overbidding and the effect of the treatment variation quite well. It is based on learning direction theory (Selten and Stöcker, 1986). Furthermore, weighted impulse balance involves one parameter that may, in our auction context, be interpreted as a measure of a concern for relative standing. Bolton and Ockenfels (2000) and Fehr and Schmidt (1999) recently demonstrated that such social motivations appear to influence decision making in a wide variety of experimental games. However, neither CRRAM nor social preference models that are based on full rationality provide different predictions for each of our treatments.

II. Experimental Game and Design

The model underlying the experiment is a symmetric sealed bid first price auction with private values. The values v_1 and v_2 of the two bidders 1 and 2 are uniformly and independently distributed over the interval [0, 100]. Actually, the values were not really continuously distributed since values were decimal numbers with two digits behind the decimal point. Both bidders simultaneously and independently made bids b_1 and b_2. The higher bid b_j wins the auction, and the profit of the winner j is $v_j - b_j$. The loser receives nothing. In case of a tie, $b_1 = b_2$, the auction winner is selected randomly. Bids above the value were not permitted but otherwise no restrictions were imposed on bids. After each auction, each of both bidders was informed about whether he won the auction, the price, and his payment for that auction. In the treatment *NF* no additional feedback was given, but in the treatment *F* feedback on the opponent's bid was supplied to the winner of the auction. All together eight sessions with 12 subjects each were run for 140 rounds with one auction in each round. The subjects within one session belonged to two independent subject groups of six participants each.[3]

The 140 rounds were divided into 28 *weeks* with five rounds each. At the beginning of each week a value was drawn randomly for each subject. Then, each participant played against each of the five other participants in the same subject pool in a random order with the restriction that a subject could not be matched with the same opponent twice in a row. The interaction was anonymous and formal via computer terminals.[4] The computers were in three-sided cubicles and neither other subjects nor the experiment's monitor could watch someone make choices. The subjects did not know about the division into two independent subject pools. The impression was conveyed that in each week they play against opponents randomly selected among the other eleven participants (knowing that they cannot be matched twice in a row with the same opponent; see Appendix for the instructions given to the subjects). In total, 96 subjects participated in four sessions with treatment *NF* and four sessions with treatment *F*, yielding eight independent observations per treatment.

[3] Each money unit in the experiment was worth 2 German Pfennige (approximately 1 American cent). The total payoff of each subject was the sum of his winnings over all 140 auctions plus 10 German Marks show up fee. No session lasted longer than 1.5 hours and average payoffs were 49 German Marks ($23) with a standard deviation of 11 German Marks ($5).

[4] We used Fischbacher's (1998) z-Tree software tool.

III. Experimental results

III.1 Overbidding and feedback effect

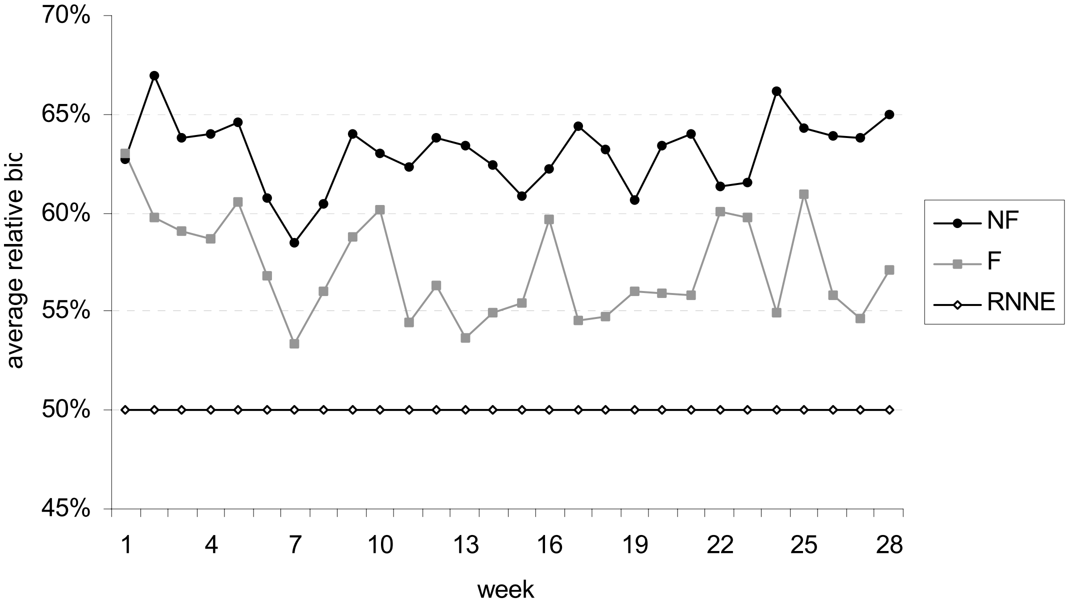

Figure 1. Average relative bids over weeks

Figure 1 shows the actual average relative bids in weeks 1 to 28 separately for the two treatments *F* (with feedback on losing bid) and *NF* (without feedback on losing bid), as well as the average relative bids predicted by the risk neutral Nash equilibrium (RNNE). The relative bid of a subject in a round is his bid as a percentage of his value. The average relative bid in a week is the average of all relative bids of all subjects in the same treatment *F* or *NF* in this week. Therefore, each data point in the graphic corresponds to 240 bids (and 48 values) of 8 groups of 6 subjects for 5 rounds. In the risk neutral Nash equilibrium, the relative bid in a two-person auction is 50 percent regardless of the feedback condition.[5]

Figure 1 shows that the curve for the feedback treatment *F* is entirely below the curve for treatment *NF*, and that the line for the risk neutral Nash equilibrium is entirely below the line for treatment *F*. The overall average relative bids are 63.1 percent for treatment *NF* and 57.2 percent for treatment *F*. The one-sided Mann Whitney *U*-test applied to average relative bids of subject groups rejects the null-hypothesis of equal average relative bids across treatments *F* and *NF* at the

[5] We have chosen a figure showing average relative bids rather than other data since RNNE, CRRAM and impulse balance theory all lead to constant predictions for average relative bids.

three-percent level. A similar comparison of average revenues or, in other words, of average high bids yields significance at the one-percent level. A comparison of actual relative bids and the predictions of the risk neutral Nash equilibrium yields significance levels of .00 for treatment *NF* and .04 for treatment *F*. These results confirm the picture obtained in the literature. Overbidding is the most common outcome in first price private value auctions (Kagel, 1995), and a comparison of analogous feedback conditions in the four-bidder case by Isaac and Walker (1985) yielded a similar feedback effect.[6]

III.2 Learning direction theory

In rounds 2 to 5 of a week in our experiment a bidder can look back on his experience from the previous round with the same value. There are three possible *experience conditions*:

- *Lost opportunity*: The price was higher than one's own bid but the value was higher than the price.
- *Overpayment*: The own bid won. In this case a smaller bid might have been sufficient in order to gain the object.
- *Outpriced value*: The price was higher than one's own value.

Learning direction theory predicts that the bid will tend to increase in the lost opportunity condition, and it will tend to decrease in the overpayment condition. In the outpriced value condition learning direction theory suggests that an unchanged bid is the most frequent response.

[6] Isaac and Walker (1985, p. 141) conclude from their experimental findings that "Central to the results of this study is the consistent evidence of the previous two studies [Walker, Smith and Cox 1983, and DeJong, Forsythe and Uecker 1984] which suggest that, in sealed bid (offer) auctions, the more limited the between market publicized information the lower the profits for those making bids to buy. […] If such evidence is found to be robust, this relationship has important implications for the evolution of market institutions in private economies, as well as policy implications for government implemented auctions."

experience condition	percent (number)	change of bid (percent) increase	decrease	unchanged
lost opportunity	22.6 (2415)	57.5	12.1	30.4
overpayment	50.0 (5359)	17.5	47.0	35.5
outpriced value	27.4 (2930)	26.6	22.6	50.8

Table 1. Experience conditions and bid changes

Table 1 shows the percentages of different responses to the three experience conditions in rounds 2 to 5 of each week.[7] The few cases in which both bidders bid the same amount are not included. The prediction of learning direction theory for these .02 percent of all observations are different for the treatments *NF* and *F*. In treatment *NF* such cases belong to the overpayment condition since as far as the bidder knows he could have won at a lower price. In treatment *F*, on the other hand, the case of equal bid must be regarded as a separate experience condition, which however is not explored explicitly in view of its very rare occurrence.

The most frequent responses to the different experience conditions are on the main diagonal. The numbers in Table 1 cannot be taken at face value, however. Suppose that the bid of a subject is determined by a constant bid function plus random error. If the value is constant a low bid is more likely to yield a lost opportunity than a high bid. At the same time, a low bid is more likely to be followed by a higher bid than a high bid. A similar argument applies to the overpayment condition and a subsequent decrease of the bid. We refer to this as the “regression effect” because it reminds us of the regression to the mean. The regression effect alone will already produce diagonal elements in the first two data rows of Table 1 that are higher than randomly expected. In the examination of predictive success of learning direction theory this has to be taken into account.

In order to control for the regression effect, we construct the following test. For every subject and every week we consider all possible permutations of the bids the subject made keeping the bids of the opponent constant. For each permutation of the actual bids we count how often we hit

[7] There is no indication that results with respect to learning direction theory differ across treatments so we present aggregate results only.

the main diagonal of Table 1. The *surplus* of a subject in a week is the difference of the actual number of cases on the diagonal minus the average number in all permutations. Learning direction theory predicts that this surplus tends to be positive, whereas the null-hypothesis is that the mean surplus is zero. For each of the 16 independent subject groups the average surplus over all six subjects and all 28 weeks is positive. By the binomial test this result is significant on the one-percent level for both treatments separately. Consequently, the hypothesis that bids change as a result of experience (as predicted by learning direction theory) is strongly supported, whereas the hypothesis that the bids of the subjects are a constant function of the value plus a random term (as assumed by CRRAM) is rejected.

The numbers on the main diagonal of Table 1 suggest that subjects do not respond equally strongly to the experience conditions. In particular, the probability that a lost opportunity is followed by a bid increase is by about 10 percentage points greater than the probability that an overpayment is followed by a bid decrease. In fact, we find that for 12 out of 16 independent subject groups the average surplus for the lost opportunity condition is greater than that for the overpayment condition.[8] By the binomial test this is significant at the 2.7-percent level. Evidently, the lost opportunity condition more strongly motivates subjects to change their bid in the expected direction than the overpayment condition does. This can be interpreted as due to social comparison processes. In the case of lost opportunity the forgone profit is not only connected to zero profits but also to an unfavorable relative standing, because a lost opportunity implies that the competitor won the auction (see e.g., Bolton and Ockenfels, 2000, and Fehr and Schmidt, 1999). For a bidder who wins and therefore experiences the overpayment condition, the forgone profit comes with a favorable relative standing. So, from the perspective of social comparison models, responses are conceivably stronger for the lost opportunity condition.

IV. Weighted impulse balance theory

IV.1 Basic idea

Impulse balance theory is an attempt to make quantitative predictions on the basis of learning direction theory without making use of a full-fledged learning model (impulse balance theory has been introduced by Selten, Abbink and Cox, 2001). Learning direction theory is applicable to situations in which a subject repeatedly has to make a decision on the same parameter, e.g. a bid in an auction. In addition, the feedback must permit at least qualitative conclusions about what

[8] Learning direction theory does not make any prediction about the difference of both average surpluses.

would have been a better choice of the parameter in the last period. It is postulated that the decision has a tendency to shift into the direction suggested by this counterfactual comparison.

We speak of an *upward impulse* if a higher profit could have been gained by a higher value of the parameter but not by a lower one. Analogously, a *downward impulse* is experienced if a higher profit could have been gained by a lower value of the parameter but not by a higher one. The strength of the impulse is the amount of forgone profit. The basic idea of impulse balance theory is that the size of the shift of the parameter tends to be proportional to the strength of the impulse. Impulse balance means that expected upward impulses equal expected downward impulses.[9]

In the following section we propose a modification of impulse balance theory that permits different weights for downward and upward impulses. We shall define a *weighted* impulse balance equilibrium that depends on the parameter λ called the downward impulse weight. This parameter is the number of upward impulse units equivalent to one downward impulse unit. The greater λ the more the average bid is drawn downward. Weighted impulse balance requires that the expected upward impulse is equal to λ times the expected downward impulse.

In the application of weighted impulse balance theory to our data we must distinguish between treatments *F* and *NF*. In treatment *F* the subjects see the amounts of upward and downward impulses. A comparison of expected amounts is behaviorally reasonable. In treatment *NF*, however, the amount of the downward impulse is not visible to the winner since feedback on the losing bid is not available. Therefore, in this treatment the principle of weighted impulse cannot be applied to expected amounts of impulses. Instead of this we shall apply it to expected numbers of impulses. In treatment *NF* subjects can still observe whether there is an upward or downward impulse even though amounts are not visible in the case of downward impulses.

IV.2 The model

We look at a symmetric sealed bid private value first price auction with n bidders. The values v are uniformly and independently distributed over [0, 1]. Let x be the value of bidder 1 and y be the maximum of the values of bidders 2 to n. Then, the density of y is $f(y) = (n-1)y^{n-2}$. Moreover, let p be the price or, in other words, the highest bid of all players. Our analysis is

[9] In the definitions of Selten, Abbink and Cox (2001) losses are counted double in the impulse balance equation. However, in the auction situation of this paper losses cannot occur. Therefore, we shall neglect this feature of impulse balance theory here.

based on the assumption that each bidder has the same homogenously linear bidding function $b(v) = av$ with $0 \leq a \leq 1$ where $b(v)$ is the bid at the value v.[10]

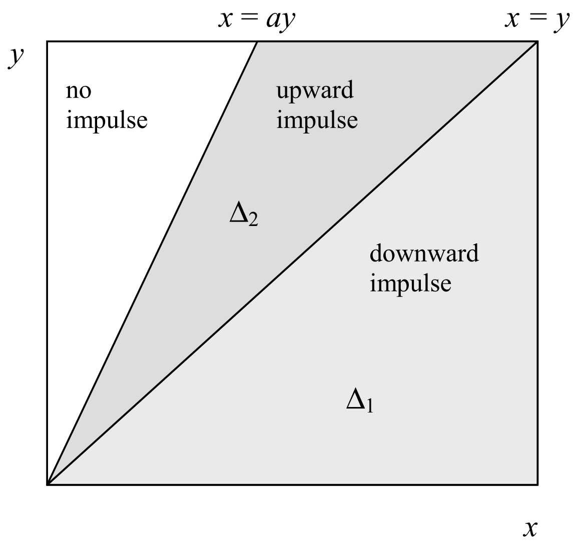

regions	*characterization*	*amount of impulse* F	NF
downward impulse: Δ_1 (overpayment)	$p = ax > ay$	$ax - ay$	not known
upward impulse: Δ_2 (lost opportunity)	$p = ay < x$	$x - ay$	$x - ay$
no impulse	$p = ay > x$	0	0

Figure 2. Regions and amounts of impulses

The diagram in Figure 2 shows the value x of bidder 1 with $0 \leq x \leq 1$ on the abscissa and the maximum value y of the other bidders with $0 \leq y \leq 1$ at the ordinate. We distinguish three regions in this diagram. In region Δ_1, x is greater than y so that bidder 1 wins. In this region bidder 1 observes a downward impulse. In region Δ_2 the maximum value of the others is greater than bidder 1's value but the highest bid ay of the others is still smaller than x. Here, an upward

[10] RNNE and CRRAM lead to bidding functions of this kind. In particular, for RNNE we have $a = (n - 1)/n$, and CRRAM yields $a = (n - 1)/(n - 1 + r)$. Our own data and earlier empirical results suggest that bidding functions are approximately linear (e.g., Cox et al., 1988), even if some deviations like non-monotonicities can be observed (Selten and Buchta, 1999).

impulse is observed. The remaining region corresponds to the outpriced value condition in which there is no impulse. The table below the diagram in Figure 2 describes the three regions. The difference between the treatments *F* and *NF* is that the amount of a downward impulse is not known.

IV. 3 Weighted impulse balance equilibrium in treatment F

The expected downward impulse E_- is computed as follows:

$$E_- = a\int_{\Delta_1}(x-y)f(x)f(y)dxdy$$
$$= a\int_0^1\int_0^x (x-y)(n-1)y^{n-2}dydx$$
$$= \frac{a}{n(n+1)}.$$

For the expected upward impulse E_+ we obtain

$$E_+ = \int_{\Delta_2}(x-ay)f(x)f(y)dxdy$$
$$= \int_0^1\int_{ay}^y (x-ay)(n-1)y^{n-2}dxdy$$
$$= \frac{(1-a)^2(n-1)}{2(n+1)}.$$

Weighted impulse balance equilibrium requires

$$E_+ = \lambda E_- .$$

This impulse balance equation is a quadratic equation for the slope *a* of the bid function. We obtain

$$a = 1 + \frac{\lambda}{n(n-1)} - \sqrt{(1+\frac{\lambda}{n(n-1)})^2 - 1} .$$

The solution for *a* is in the interval (0,1) for all positive values of the downward impulse weight λ. The other root of the quadratic equation is greater than 1 and therefore outside the interval of admissible values for *a*.

IV. 4 Weighted impulse balance equilibrium in treatment NF

Under the treatment *NF* the expected number of downward impulses P_- is

$$P_- = \frac{1}{n}.$$

This is due to the fact that all bidders use the same bid function and therefore have the same chance $1/n$ of winning the auction. The expected number of upward impulses P_+ is the area of the region Δ_2:

$$\begin{aligned} P_+ &= \int_{\Delta_2} f(x) f(y) dxdy \\ &= \int_0^1 \int_{ay}^{y} (n-1) y^{n-2} dxdy \\ &= \frac{(1-a)(n-1)}{n}. \end{aligned}$$

Weighted impulse balance equilibrium requires

$$P_+ = \lambda P_-.$$

This is a linear equation for a. We obtain

$$a = 1 - \frac{\lambda}{n-1}.$$

a is in the interval (0,1) if $0 < \lambda < n - 1$. In our experiments we have $n = 2$ which means that we must have $\lambda < 1$ if a is positive.

IV.5 Weighted impulse balance equilibrium and overbidding

In the risk neutral Nash equilibrium, the bidding function is $b(v) = (n - 1)v/n$. Overbidding therefore requires $a > (n-1)/n$. Depending on the value of the downward impulse weight λ the impulse balance equilibrium predicts overbidding. Let us turn first to treatment F and assume $a > (n-1)/n$. Then we have

$$E_- = \frac{a}{n(n+1)} > \frac{n-1}{n^2(n+1)}, \text{ and}$$

$$E_+ = \frac{(1-a)^2(n-1)}{2(n+1)} < \frac{n-1}{2n^2(n+1)}.$$

This yields

$$\lambda = \frac{E_+}{E_-} < \frac{1}{2}.$$

So, in treatment *F* overbidding results in weighted impulse balance equilibrium if and only if λ is smaller than 1/2. In treatment *NF*, on the other hand, $a > (n-1)/n$ is equivalent to $1 - \lambda/(n-1) > (n-1)/n$ and therefore to $\lambda < (n-1)/n$. This condition is necessary and sufficient for overbidding in impulse balance equilibrium under treatment *NF*.

Summing up, overbidding requires $\lambda < 1$ regardless of whether we look at treatment *F* or *NF* and for all number of bidders *n*. This means that upward impulses must get more weight than downward ones. The root of this asymmetry – and therefore of overbidding – may be the social comparison process described in Section III.2 that explains why paying 'too much' and being the winner yields a weaker response than bidding 'too little' and being the loser of an auction.

IV.6 Weighted impulse balance equilibrium and the treatment effect

We now turn to the observation that in treatment *F* bids tend to be lower than in treatment *NF*. We show that this effect follows from impulse balance theory if the downward impulse weight is in the region connected to overbidding. Let S_- be the conditional expectation of the amount of the downward impulse under the condition that such an impulse occurs. Similarly, let S_+ be the conditional expectation of the amount of an upward impulse if it occurs. We have

$$S_- = E_-/P_- = \frac{a}{n+1}, \text{ and}$$

$$S_+ = E_+/P_+ = \frac{(1-a)n}{2(n+1)}.$$

Thus, $S_- > S_+ \Leftrightarrow \frac{a}{1-a} > \frac{n}{2}$, which holds if $a > \frac{n}{n+2}$. Since $\frac{n}{n+2} \le \frac{n-1}{n}$ for all $n > 1$, overbidding implies $S_- > S_+$. This explains the treatment effect. The conditional average amount of a downward impulse is larger than the conditional average amount of an upward impulse. So taking into account both amount and number of impulses yields on average a stronger downward impulse than taking into account only numbers. As a consequence, bids in treatment *NF* are higher than bids in treatment *F*.

IV.7 Quantitative estimation and prediction

On the basis of our data we close our discussion of the weighted impulse balance equilibrium with a rough estimate of the value of the parameter λ. We use the relationship between the slope a and λ and insert the average slope $\bar{a}$ in our data. This is done for each treatment separately. We obtain $\lambda = (1 - \bar{a})/(2\bar{a}) = .32$ for treatment *F* (where $\bar{a} = 0.572$), and $\lambda = 1 - \bar{a} = 0.37$ for treatment *NF* (where $\bar{a} = .631$). Both treatments yield slightly different estimates of λ which however seem to be near enough to each other in order to justify the assumption that the two values of λ are the same for both treatments. Accordingly, we estimate λ as the average $\lambda = .34$ of both estimates.

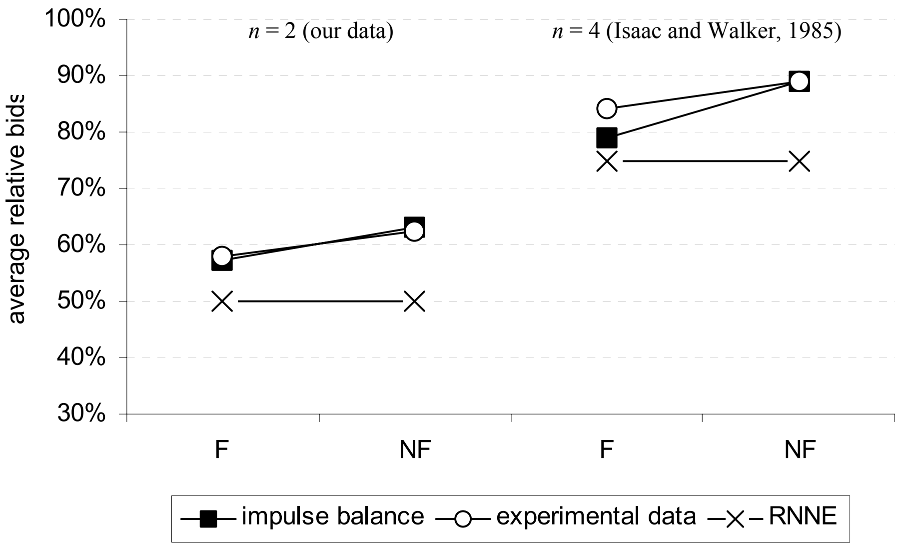

Figure 3. Average relative bids and theoretical predictions

Based on this estimate of λ Figure 3 shows empirical and theoretical values of average relative bids for both treatments in our experiments and those of Isaac and Walker (1985).[11] Overbidding and lower slopes for treatment *F* are correctly reproduced for both experiments. Furthermore, we obtain a good quantitative fit of our data. Actual average bids are 57 and 63

[11] We thank Mark Isaac and Jim Walker for making their data available to us.

percent for treatments *F* and *NF*, respectively, while predicted average bids are 58 and 62 percent, respectively. Using the estimate of the downward impulse weight λ from our data, we also correctly predict the average relative bid in Isaac and Walker's treatment *NF*, and only somewhat underestimate the average relative bid in their treatment *F*. Actual values are 84 and 89 percent for treatments *F* and *NF*, respectively, while predicted values are 79 and 89 percent, respectively.

V. Conclusions

The concept of weighted impulse balance equilibrium has been introduced as a quantitative version of learning direction theory that incorporates a concern for social comparison and that avoids the specification of a full-fledged learning model. The weighted impulse balance equilibrium does not only yield a possible explanation of overbidding but it also predicts the treatment effect of lower average relative bids in treatment *F* under the assumption that the downward impulse weight is in the range where overbidding occurs. This treatment effect cannot be explained by optimization approaches since the additional feedback on the losing bid supplied in treatment *F* is irrelevant for the maximization of profits on the basis of the rules of the game. This feedback is also unimportant for pure reinforcement theories in which probabilities of bids depend only on experienced past profits. An adequate explanation of bidding behavior in our simple auction environments seems to require a picture of human learning that is neither completely mechanistic nor hyper-rationalistic. Learning direction theory, which postulates simple processes of cause and reasoning, seems to be a more adequate approach. The theory of weighted impulse balance equilibrium shows that quantitative behavioral equilibrium concepts can be devised which may serve as benchmarks in the evaluation of experiments.

References

Bolton, Gary, and Axel Ockenfels (2000). "ERC – A Theory of Equity, Reciprocity and Competition. *American Economic Review*, 90(1), 166-193.

Cox, James C., Bruce Roberson, and Vernon L. Smith (1982). "Theory and Behavior of Single Object Auctions." In: Vernon L. Smith (ed.), *Research in Experimental Economics*, Vol. 2, Greenwich, CT: JAI Press.

Cox, James C., Vernon L. Smith, and James M. Walker (1988). "Theory and Individual Behavior of First Price Auctions." *Journal of Risk and Uncertainty*, 1, 61- 99.

DeJong, Douglas V., Robert Forsythe, and Wilfred C. Uecker (1984). "The Effects of Alternative Liability Fees on the Price and Quality of Audit Services: A Laboratory Market Study. Working paper, University of Iowa.

Fehr, E., and Schmidt, K. (1999). "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 114, 817-868.

Fischbacher, Urs (1998). "z-Tree. Zurich Toolbox for Readymade Economic Experiments." Working Paper, University of Zurich.

Isaac, R.M., and James M. Walker (1985). "Information and Conspiracy in Sealed Bid Auctions", *Journal of Economic Behavior and Organization*, 6, 139-159.

Kagel, John (1995). "Auctions." In: J.H. Kagel and A.E. Roth (eds.), *The Handbook of Experimental Economics*. Princeton University Press.

Selten, Reinhard, Klaus Abbink, and Ricarda Cox (2001). "Learning Direction Theory and the Winner's Curse." Working Paper, University of Bonn.

Selten, Reinhard, and Joachim Buchta (1999). "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions." In: D. Budescu, I. Erev., and R. Zwick (eds.), *Games and Human Behavior: Essays in the Honor of Amnon Rapoport*. Lawrenz Associates Mahwah NJ.

Walker, James M., Vernon L. Smith, and James C. Cox (1983). "Bidding Behavior in Sealed Bid Discriminative Auctions: An Experimental Analysis of Multiple Sequence Auctions with Variations in Auction Information." Working paper, University of Arizona.

Appendix: Instructions to the subjects (translation from German)

Instructions

General Information

The purpose of this session is to study how people make decisions. If at any time you have questions, raise your hand and a monitor will assist you. From now until the end of the session, unauthorized communication of any nature with other participants is prohibited.

During the session you will participate in auctions that give you an opportunity to make money. You will be paid your earnings plus a DM 10 show-up fee at the end of the session. Decisions and payments are confidential: No other participant will be told your actions during the game or the amount of money you make.

Description of the game

Each participant will have the opportunity to submit one bid in each of 140 subsequent auctions. In each auction, there will be exactly two bidders. The higher bid wins. The winner receives a payoff, the loser receives nothing in that auction.

What is the value of the item auctioned off to me?

The precise value of the (fictitious) item differs across bidders and between auctions. However, in each auction, before you submit your bid, you will be told exactly what the commodity is worth to you, i.e. what DM-amount we pay you if you win the auction. Specifically, your value is randomly drawn such that each value between 0 and 100 money units has the same probability of being drawn. The values of all bidders are independently determined, so that generally each bidder has a different value. (But you won't know the values of the other bidders.) The value does not change for 5 auctions. After each 5 auctions, new values are randomly drawn for all bidders.

What are my earnings?

If you submit the higher bid in an auction, you win the auction. Your payment from this game is then your value minus your bid. (Therefore, bids higher than your value may lead to losses and are therefore not allowed.) In the case that your bid is smaller than the bid of your opponent, you lose the auction and won't receive any payment for this auction. In the case of a tie, a chance move randomly determines the winner. Your earnings paid at the end of the session is the sum of payments over all 140 auctions. Each money unit earned in the auction is worth 2 Pfennige.

Who is the other bidder?

All pairings are anonymous. Your identity will not be revealed to the persons you are playing with either before, during or after the game. In each auction, your opponent will be selected randomly. However, you won't be paired with the same person twice in a row. Note that by this matching scheme, the values of your (changing) opponents generally also change from round to round.

[TREATMENT *F*] *What is the feedback after each auction?*

After each round you will be informed whether you won the auction, the price (that is, the winning bid), the opponent's bid, and your payment from that auction.

[TREATMENT *NF*] *What is the feedback after each auction?*

After each round you will be informed whether you won the auction, the price (that is, the winning bid), and your payment from that auction.

Good luck!