Group dynamics in experimental studies - the Bertrand Paradox revisited

Bruttel, Lisa V.

Empfohlene Zitierung / Suggested Citation:
Bruttel, L. V. (2008). Group dynamics in experimental studies - the Bertrand Paradox revisited. *Journal of Economic Behavior & Organization*, *69*(1), 51-63. https://doi.org/10.1016/j.jebo.2008.10.002

Accepted Manuscript

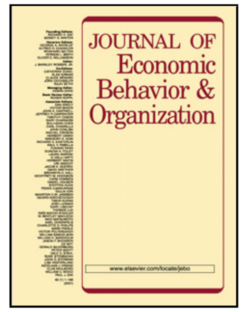

Title: Group Dynamics in Experimental Studies-The Bertrand Paradox Revisited

Author: Lisa V. Bruttel

PII: S0167-2681(08)00199-6
DOI: doi:10.1016/j.jebo.2008.10.002
Reference: JEBO 2265

To appear in: *Journal of Economic Behavior & Organization*

Received date: 19-3-2007
Revised date: 2-10-2008
Accepted date: 2-10-2008

Please cite this article as: Bruttel, L.V., Group Dynamics in Experimental Studies-The Bertrand Paradox Revisited, *Journal of Economic Behavior and Organization* (2008), doi:10.1016/j.jebo.2008.10.002

Group Dynamics in Experimental Studies - The Bertrand Paradox Revisited

Lisa V. Bruttel[†]

Abstract

Different information provision in experimental markets can drastically change subjects' behavior. Considering the repeated Bertrand duopoly game of Dufwenberg and Gneezy (2000), we find that population feedback about the prices in other markets outside a subjects' own current market causes group dynamics that prevent prices from convergence to Nash equilibrium. Limited information comprising only the decisions of a subject's own opponent, in contrast, leads to competitive behavior. When we extend the number of periods from 10 to 25 in the full information treatment we observe a very robust cyclical up and down movement of prices. We can explain tacit coordination in our experiment with an extended learning direction model and leadership by example.

JEL classification: C72, C92, D43, D83

Keywords: Bertrand duopoly, Tacit collusion, Learning, Leadership by example, Experiment

[†] University of Konstanz, Department of Economics, Box D-131, 78457 Konstanz, Germany, Phone +49-7531-883214, Fax +49-7531-882145, E-mail: lisa.bruttel@uni-konstanz.de. I would like to thank Ulrich Kamecke, Urs Fischbacher, Werner Güth and seminar and conference audiences at Magdeburg (GfeW), Berlin, Jena and Rome (ESA) for many helpful suggestions. The comments of the associate editor and two anonymous referees helped to improve the paper considerably. Stephanie Benke and Sebastian Schwenen provided excellent research support.

Group Dynamics in Experimental Studies - The Bertrand Paradox Revisited

Abstract

Different information provision in experimental markets can drastically change subjects' behavior. Considering the repeated Bertrand duopoly game of Dufwenberg and Gneezy (2000), we find that population feedback about the prices in other markets outside a subjects' own current market causes group dynamics that prevent prices from convergence to Nash equilibrium. Limited information comprising only the decisions of a subject's own opponent, in contrast, leads to competitive behavior. When we extend the number of periods from 10 to 25 in the full information treatment we observe a very robust cyclical up and down movement of prices. We can explain tacit coordination in our experiment with an extended learning direction model and leadership by example.

JEL classification: C72, C92, D43, D83

Keywords: Bertrand duopoly, Tacit collusion, Learning, Leadership by example, Experiment

1 Introduction

In experimental markets we often observe that the amount of information about their competitors that is provided to markets participants can drastically influence their behavior. Most authors argue that different information causes different learning possibilities that are crucial for the success or failure of coordination. The reported effect of detailed information in such studies, however, is ambiguous and depends on the market design. Huck et al. (2000), for example, find in an experimental oligopoly market that detailed information about rivals' actions leads to less cooperation than aggregate information (average quantities and prices), especially in the case of quantity competition. They explain this through imitation of successful firms. Similar results are found by Altavilla et al. (2006), who also study the dynamics of imitation and, furthermore, introduce aspiration rules as an explanation for behavior when only averaged information is known. In this view, the average profit in all markets in the previous period serves as an aspiration level. Offerman et al. (2002), on the contrary, demonstrate that detailed information about others' quantities leads to a higher frequency of collusive outcomes than aggregate information (sum of quantities). Their explanation for this observation also refers to imitation learning, but this time to imitation of exemplary firms that decrease their quantity to guide the other firms to cooperation. Selten and Apesteguia (2005) consider a location game with price competition on a circle. They explain the behavior of their subjects with a combination of cooperative attempts and imitation of successful players.

In this paper, we consider a Bertrand price setting oligopoly, where theory predicts prices at the equilibrium to be equal to marginal costs. The Bertrand Paradox, however, contradicts this theoretical prediction with the observation that in real-life pricing situations, firms indeed set prices above marginal costs and face positive profits. Dufwenberg and Gneezy (2000) illustrate this contradiction in a well-known experimental study. They transfer the

Bertrand situation into a game where participants have to choose a number (price) from the interval [2:100]. The one who chooses the lowest number wins an amount of money equal to the price whereas all others receive zero. The game is repeated for ten rounds with a random matching of interacting groups in every round from a pool of overall 12 participants. In their paper, Dufwenberg and Gneezy find that with groups of three or four participants, the winning price converges to the lowest possible value while for the duopoly case the downward movement stops after some rounds and reverses into an increase of prices instead. The experimental setup of this game can also be translated into a first-price auction mechanism.[1] The first-price auction experiment of Ockenfels and Selten (2005) has a similar design, except for the fact that they consider two different private resale values. In their experiment they find that bidding is more aggressive when only winning (but no losing) bids of their auction are made known to the bidders. Neugebauer and Selten (2006) find similar results in an experiment where subjects played against computerized opponents.

In the experimental sessions of Dufwenberg and Gneezy (2000) participants were seated in a classroom. In every round they wrote their bid and their registration number on a coupon. All coupons were put into a box, and then groups were randomly drawn from the box by a neutral monitor. All bids and the related registration numbers were written on a blackboard to ensure public information and to demonstrate that no cheating by the experimenter was possible. This proceeding, however, has been questioned by other authors (see, for example, Huck 2004) because the experimental setup with public information and visual contact enables not only learning, as intended by the experimenters, but also unintended group dynamics. By placing high bids participants can demonstrate to others their willingness to collude. Indeed, Dufwenberg and Gneezy report a relatively high number of bids of 100 that can be interpreted as (successful) attempts to establish coordination between all 12

[1] Therefore, we will use "price" and "bid" as synonyms in our paper, always meaning the number between 2 and 100 a participant chooses.

participants. The ability of the experiment to demonstrate the Bertrand Paradox might therefore be questioned because leadership by example can also explain the observation of higher prices than in Nash equilibrium.

In their 2002 paper, Dufwenberg and Gneezy return to the information problem. They repeat the duopoly treatment, this time providing different levels of information history. In the full information treatment (FULL), data from the duopoly case of the 2000 experiment is used again. In a semi-information treatment (SEMI), only the winning bids are announced instead of making all bids publicly known. Finally, in a no-information treatment (ZERO), information is reduced further to the matching structure and the registration number of the winning participants without announcing the respective bids. Both reductions of the available information history lead to drastic decreases in cooperation. Thus, Dufwenberg and Gneezy conclude that high cooperativeness in the full information treatment of the 2000 experiment is explained by behavior that is similar to signaling. Following their argument, in every period players have the choice whether they want to play seriously (a best reply to the expected bids of the other participants) or to send the group a signal of their willingness to cooperate and therefore to set a high bid. If this action is well interpreted and accepted by other participants, it can pay off to abstain from today's chance to make a profit in order to gain larger profits in future periods.

In this paper, we follow some of the notions and extend the experiment of Dufwenberg and Gneezy in two ways. First, we modify the information that participants receive about surrounding markets to see whether this information is necessary for coordination. Second, we extend the duration of the experiment in order to consider group dynamics over a longer period of time.

In the first part of our series of experiments, we use a design very similar to Dufwenberg and

Gneezy, but focus on other aspects of information provision. Similar to the question in their 2002 paper, we ask whether information about the whole group of participants is crucial for successful tacit coordination. In contrast to their study, we provide participants with information about the decision of their direct opponent, but remove information about the rest of the group for the following reason: hiding all losing bids not only prevents coordination within the group, but also affects the participants' individual learning processes. Decreasing bids in the semi-information treatment can just as well be explained by imitation of successful strategies; participants only see the low bids of the auction winners but not the (possibly large) difference to the higher losing bids. They can only learn from successful bids that won an auction, but they are prevented from learning from lost opportunities when the own bid was much smaller than that of the opponent (see Selten and Buchta 1999). Thus, with our information set we allow for different learning dynamics than Dufwenberg and Gneezy (2002) because in our experiment, winners can learn from lost opportunities as they know the distance to their opponents' bids.

Additionally, we computerize most sessions to exclude potential effects of tacit face-to-face coordination. Though participants in a classroom experiment are not allowed to talk to each other, they might nevertheless be able to communicate in some visual way and, thus, establish a higher degree of cooperative-mindedness (see Bohnet and Frey 1999). Conducting the experiment in a computer lab allows us to diminish such tacit coordination. Table 1 summarises the treatments of our experiment.

In the new treatment with no information about the rest of the group, we find that bids decrease constantly over the whole duration of the game. Thus, we can show that information about somebody's own lost opportunities is not sufficient to establish cooperation. Tacit coordination works solely with information about the whole group. Only in the sessions with

[2]See similar in the ultimatum game of Abbink et al. (2004)

Treatment	PP	FULL	LOW
Conduction	Paper & pencil, all bids on a blackboard	Computerized	Computerized
Information	All IDs & the respective bids	All IDs & the respective bids	Own bid & competitor's bid
Possible Interaction	Facial communication, tacit coordination	Tacit coordination within the group	Learning from own experience[2]
Expected Outcome	Cooperative behavior, high bids	Cooperative behavior, high bids	Competitive behavior, declining bids

Table 1: Overview treatments in the first series.

full information we observe (as Dufwenberg and Gneezy) a bid increase in the second half of the game after an initial decline of bids, no matter whether the session is conducted in the computer lab or with paper and pencil. We will illustrate that an extended variant of Selten and Buchta's learning direction theory can explain the differences between our treatments. In sections 3.1 and 3.2 we will describe the effects of different information provision and apply learning direction theory to our data.

In the second part of our experimental series, we extend the duration of some sessions to see how participants continue their play once they reached a high bid level. We might think that they just needed some time to coordinate and, once they have learned, remain at relatively high bids. However, it could also be argued that high average bids imply a stronger temptation to lower their bid because the possible profits from undercutting have increased. Since a bid decrease from a very high level is very tempting and since another bid increase from a very low level is cheap in comparison to the possible profits from serious play, we could expect that bids fluctuate between the two levels. In the second series of experiments we therefore analyze the behavior over a longer time horizon of 25 periods. Indeed, we find a very robust cyclical upward and downward movement of average bids. Each circle is triggered by an initial drastic bid increase of one participant who obviously tries to encourage the others to follow her example. The motivation of this individual initiative and the reaction of the others will be discussed in the second part of our analysis in sections 3.3 and 3.4. Section 4 concludes.

2 Experimental Design

In our experiment we focus on the importance of group dynamics and leadership by example for success of collusion in the Dufwenberg and Gneezy game. The primary goal of our experiment is to find out whether population feedback is necessary to enhance cooperation. In each session, 12 participants play the bidding game for 10, respectively 25, rounds in pairs of two with random re-matching of pairs after each round. To distinguish the Bertrand Paradox effect from possible tacit face-to-face coordination in the classroom, we compare the results of the original duopoly experiment in the classroom (treatment PP) with another one that is played in the computer laboratory (treatment FULL). All participants receive information about the bids of all other participants as well as the respective registration numbers and pairings. In the laboratory no interaction is possible, and visual communication is not only theoretically prohibited but also practically impossible.[3] In our new treatment participants receive only information about the decisions that are relevant to their own profit: their own bids and those of their opponents as well as their payoffs (treatment LOW). Registration numbers and all other participants' bids, however, remain unknown in this treatment so that group dynamics cannot occur. Thus, the participants receive all the information necessary to learn from successful bids as well as from lost opportunities. This is the key difference of this experiment from the SEMI and ZERO treatments of Dufwenberg and Gneezy (2002). In contrast to these treatments, even in our LOW treatment both participants in a market, winner and loser, know the distance to their opponent's bid, which is particularly important

[3]In contrast to the experiment of Bohnet and Frey, participants do not get to know their respective opponent.

in the case of winners who had a much lower bid than their opponent. They win much less than they could if they would have chosen a higher bid only a little lower than that of their opponent. At the same time, we can control for the effects that information may generate for the interaction of different markets. In our LOW treatment, participants have very limited opportunities to influence decision making within their group since only their respective opponent would see the distinctive action, not the group of all other participants. Table 2 summarises the number of participants in each treatment. All in all, 206 subjects participated in the experiment.

		PP	FULL	LOW
10 Rounds:	# Participants	26(52*)	60	60
	# Independent observations	**2(4*)**	**5**	**5**
25 Rounds:	# Participants	-	48	12
	# Independent observations	-	**4**	**1**
	Sum of independent observations	**2(4*)**	**9**	**6**

Table 2: Number of participants per treatment. *including the data from Dufwenberg and Gneezy

Participants were recruited from undergraduate classes of different subjects; they had no previous knowledge of similar games. In the computer treatments, we had twelve participants per session, whereas in the paper and pencil treatment, there was one additional participant, called "monitor", whose main task it was, as in the Dufwenberg and Gneezy experiments, to draw the bids from a box to ensure that we could not cheat with the matching procedure. After the experiment participants answered an ex-post questionnaire. On average, participants earned 8.04 Euros in the 10-period sessions and 14.15 Euros in the 25-period sessions, both including a show-up fee of 4 Euros. The 10-period sessions lasted 20-30 minutes and the 25-period sessions lasted 50-60 minutes, including the time for reading the instructions and answering a short post-experimental questionnaire. All experiments took place in the same room, the computer laboratory at Humboldt-University Berlin, in May 2006, July 2006, and June 2007. In the PP sessions we removed the walls between the computer workstations so that participants could see each other. The computer sessions were conducted using Fischbacher's z-tree software (2007).

3 Experimental results

Initial bids in the first period do not significantly differ across all sessions of our experiment. Here, we also include the data from Dufwenberg and Gneezy's (2000) sessions 2a and 2b, which were conducted in the same way as our PP treatment, into the analysis, calling their treatment D/G. Average bids in the first period are within the interval from 33.50 to 59.58. In a Kruskal-Wallis rank sum test we cannot reject the null hypothesis that all samples come from the same distribution at any reasonable significance level (d.f. = 18, p-value = 0.8184).

3.1 The effect of different information provision

The first question we want to answer is whether tacit coordination fails if we provide participants only with information about the behavior of their direct opponent but not about the rest of the group. Between the treatments with different information provision, the experiment shows a remarkable difference in bidding behavior. Figures 1 and 2 show the average (winning) bids in treatments with full information. Average bids and average winning bids follow very similar patterns. Therefore, the detailed description will concentrate only on the average bids with similar arguments holding for the average winning bids. In the 10-period sessions with full information we often (in 6 out of 9 sessions) find an initial decline of bids followed by an increase by the end of the game, whereas in sessions with no information the decline continues until the end of the game. For the treatments with full information (PP, D/G, and FULL 10) it does not seem to influence behavior whether they are conducted in classroom or in a computer lab. If at all, there is a little more cooperation in PP and

D/G than in FULL, but the basic down-and-up dynamic is mostly the same. Irrespective of the procedure, there obviously is interaction within the whole group, which is in contrast to the LOW treatment. Henceforth, we will mainly ignore the procedure of a session when we compare the effects of information. However, some sessions (PP 1, FULL 10-1, and FULL 10-5) do not fit into the down-and-up pattern. Understanding these three sessions will later become crucial for the interpretation of the experiment.

[Figure 1 & 2]

Considering the declining average bids in FULL 10-1 and FULL 10-5, we think that coordination at a higher bid level probably just takes longer than the time framework of only 10 periods. To answer this question a series of the computerized sessions is conducted, this time for a longer duration of 25 periods. Figure 3 shows the average bids in these sessions for treatment FULL. We find that in all these 25-period sessions, sooner or later the average bids start increasing again after an initial decline.

[Figure 3]

Furthermore, we now find a cyclical movement of bids. Initially, bids move down and at some point they move up again. Then, after some rounds with relatively high bids, they start lowering again. We think that this is because bid cuts become more and more tempting when the bid level is high. The lower bound of average bids that has to be reached before one of the participants initiates an increase of bids is different from session to session and even from circle to circle within one session. Thus, the underlying dynamics are not very easy to compare between sessions or even within a session. However, the cyclical movement itself seems remarkable robust since we find two to three circles within each of the full information sessions with 25 rounds.

In our LOW treatment with no information about the other markets, on the contrary, bids exhibit a downward sloping movement towards the Nash equilibrium from the beginning to the end (see Figures 4-6). After ten periods bids in the LOW sessions have decreased by approximately one-third. Compared to the full information sessions this observation clearly indicates that the removal of information about the decisions in the other markets has a drastic impact on the behavior of the subjects in a group. Population feedback enhances cooperative behavior which we cannot find in the sessions with limited information. We also compare the average bids of all LOW and FULL sessions (10 and 25 period sessions pooled) in the final period in a Wilcoxon-Mann-Whitney test and find that average final bids in LOW are significantly smaller than those in FULL (one-sided, p-value = 0.0360). Conducting the test only with the data from the 10 period sessions, the difference is statistically not significant (one-sided, p-value = 0.1199). Note that we have to be careful with the interpretation of this comparison of the average bids in the final period due to the fluctuating average bids in treatment FULL. Depending on whether the final period of a FULL session is by chance the maximum or minimum of such a bid circle, this might support or reject the hypothesis that the average bids in FULL are higher than in LOW.

[Figure 4 & 5]

If we repeat LOW for 25 periods the bid level continues to fall slowly (see Figure 6). After 25 periods, average bids have reduced by about 70 percent, compared to their initial distribution. Towards the end of the game, in round 21, one of the subjects makes an attempt to increase bids by choosing number 100 for some rounds. This action, however, can only be seen by the direct opponents of this subject, not by all others, and therefore does not lead to more cooperative choices.[4]

[4]This observation supports the argument of Dufwenberg and Gneezy (2000) that high bids might be

[Figure 6]

3.2 Explaining the data with learning direction theory and best replies

Different behavior in sessions with more and less information can be explained qualitatively by Selten and Buchta's learning direction theory, taking into account that participants have different learning opportunities under different information sets. For the analysis of our LOW treatment, we can directly apply their original directional learning model. When analyzing sessions with full information, we extend their model by an additional component to include learning that adapts to the group as a whole. Subsequently, we perform a short quantitative analysis with respect to the size of bid changes and their fit to the best reply prediction.

Learning direction theory distinguishes participants' reactions to situations where they won and lost, respectively. According to the theory, we expect that subjects react to their experiences by adjusting their bid downwards when they lost and upwards when they won. The data from our experiment confirms the predictions of the learning direction theory (see Table 3). Losers tend to lower their bid in the following period because they previously learned that their bid was higher than that of their opponent and, thus, they lost the opportunity to make a profit. In our data, losers decrease their bids in all treatments in two-thirds of the cases, but increase their bid only in 8(15)% of the cases in LOW (FULL). On average,

reasonable if a participant expects that another participant for some reason might choose a bid of 100. They show that even for a very small probability of meeting someone with a bid of 100, a bid of 99 is a profit-maximising action in the duopoly case. Here we find that such behavior is reasonable since there is indeed a bid of 100 in some periods. In our case, however, no subject thinks one step further and chooses a bid of 99.

losers decrease their bid by 8.14 units in LOW and by 7.15 units in FULL. Winners should, in contrast, tend to increase (or at least not lower) their bid in the following period because they just experienced that a higher bid would have resulted in a higher profit, given that there was some distance to the bid of their opponent. The data from our LOW treatment confirms this prediction since 80% of the winners increase their bid or keep it constant. For the FULL treatment the picture is less obvious. There are nearly equally as many bid increases (42%) and decreases (37%) after winning situations. On average, winners increase their bid by 5.81 units in FULL and by 3.91 units in LOW.

We can also show that in treatment LOW (but not in FULL) the absolute reaction of losers is larger than the reaction of winners (Wilcoxon signed rank test, one-sided, p-value(LOW) = 0.0156, p-value(FULL) = 0.1602). This result is in line with the results of Ockenfels and Selten (2005) and Neugebauer and Selten (2006), who also found that the downward impulse to losers is stronger than the upward impulse to winners. It is reasonable due to the consideration that losers have to reduce their bid by more than the distance to their opponent if they want to win. Winners, on the contrary, could still have won if they reduce their bid by less than the distance.

		FULL			LOW		
		Increase	Constant	Decrease	Increase	Constant	Decrease
Loser reaction	#	131	169	598	32	76	282
	%	15%	19%	66%	8%	19%	73%
Winner reaction	#	373	194	331	175	137	78
	%	42%	22%	37%	45%	35%	20%

Table 3: Loser and Winner behavior.[5]

[5]Ties are excluded from the analysis. We included all data that was generated by our own experiments. In the Dufwenberg and Gneezy data there is no information contained about the matching structure in each round, and therefore we could not use their data here because ties cannot be specified in neither of their experiments. However, we conducted our tests also with their FULL, SEMI, and ZERO data, treating ties as winning bids. We find the results from their FULL and ZERO data being relatively similar to our FULL and LOW data, whereas the reactions in SEMI are more extreme in such a way that there are generally

Although Selten and Buchta's learning direction argumentation holds for all our treatments, in sessions with full information an additional component influences the decision about the change in direction. When full information about all others' decisions is provided, participants can react not only to their immediate experience, but they can also utilise the information about what is happening around. Thus, they additionally take into consideration what would have happened if they had met one of the other opponents in the previous period. More precisely, they could determine a best reply to the distribution of bids within their matching group.[6] If their own bid was lower (higher) than this best reply bid, than we could expect that they increase or keep constant (decrease) their bid in the following period to maximise their expected profit. Moreover, the size of the adjustment might depend on the size of the distance, similar to the strength of impulses in Ockenfels and Selten's experiment. We therefore run a clustered linear regression to see how the distance to the opponent's bid and to the best reply to the group matter for the bid adjustment $b_t^i - b_{t-1}^i$ in both treatments.[7] In the regression, *win* and *below_BR* are dummy variables indicating whether a subject's bid b_{t-1}^i in period $t-1$ won against its opponent's bid b_{t-1}^j and whether b_{t-1}^i was lower than the best reply to the group BR_{t-1}^i, whereas $b_{t-1}^i - b_{t-1}^j$ and $b_{t-1}^i - BR_{t-1}^i$ contain the size of the distance to the opponent's bid and to the best reply to the group in $t-1$. C is a constant.

$$(b_t^i - b_{t-1}^i) = \beta_1 \cdot C + \beta_2 \cdot win + \beta_3 \cdot (b_{t-1}^i - b_{t-1}^j) + \beta_4 \cdot below\_BR + \beta_5 \cdot (b_{t-1}^i - BR_{t-1}^i)$$

many more bid decreases than in any other treatments. On average, winners here decrease their bid in 57% and losers even in 81% of the cases.

[6] In our analysis, we do that by comparing the expected profits from undercutting each of the rivals by exactly one unit. We multiply each of those 11 candidate 'undercutting bids' with the share of participants, against which a certain subject would have won with the respective 'undercutting bid' and consider the 'undercutting bid' that yields the highest expected profit in this comparison as the best reply.

[7] Of course, participants in LOW do not know the best reply to the group because they are informed only about their direct opponent's bid, not about the whole bid distribution after each round. We nevertheless include a hypothetical best reply to the group into the regression for LOW in order to have a reference point to compare the FULL data with. If participants in LOW actually apply best reply strategies, they can only best reply to their immediate experience. Thus, the directions of change predicted by learning direction and best reply coincide. The direction of those reactions was already presented in table 3.

The results from this regression are presented in table 4. First, we see that the dummy whether or not a subject won its own match is only significant in LOW but not in FULL. Second, the distance to the opponents' bid is significant in both treatments, but the size of its influence is larger in LOW than in FULL. Both results suggest that in treatment FULL there are other strong influences on subjects' behavior than only the opponents' bid. The mere fact whether a subjects' bid was lower or higher than the best reply to the group is not significant in neither of the two treatments (a little confusing here is the negative sign of the - although not significant - variable *below_BR* in treatment FULL). The distance to the best reply, however, has a strong negative impact on the bid change in FULL but not in LOW. This again indicates that the reaction to the distribution of bids within the matching group in FULL is intentional and not just random.

	LOW	**FULL**
C	-3.9366***	-1.3099
	.(6937)	(1.0273)
win	3.0290**	1.5827
	(1.0613)	(1.1379)
$b^i_{t-1} - b^j_{t-1}$	-.2158***	-.1411***
	(.0319)	(.0337)
$below\_BR$	1.9228	-1.1082
	(1.0760)	(1.2192)
$b^i_{t-1} - BR^i_{t-1}$	-.0864	-.2548***
	(.0885)	(.0465)

Table 4: Parameter values in the regression. Standard error in brackets. *** denotes significance at the 1% level and ** at the 5% level.

The most informative case occurs in situations in treatment FULL in which learning direction theory and directional best replies to the group predict bid changes in different directions. If a participant's bid was higher than that of her respective opponent but nevertheless below the best reply to the distribution of all other bids, it might be rational to increase the bid in the following period. The same holds in the opposite case. If she won in a certain round but sees that she was really lucky because most of the other participants around actually had an even lower bid, then she might tend to decrease her bid although she won in her own match. Exactly this can be observed in the data.

In FULL, there are clearly more adjustments into the opposite direction than predicted by pure learning direction theory than in LOW. In cases where a losing bid was smaller than the best reply to the group, the bid is increased in the following period in only 13 percent of the cases in LOW. In FULL, this holds in contrast for 33 percent of the cases (see Table 5). This difference of loser behavior is statistically significant at the 1-percent level in a Chi-squared test (d.f. = 2, p-value = 0.0078). Losers, whose bids were smaller than the best reply, increase their bid on average by 3.61 units under FULL information and by 2.13 units under the LOW information condition.

	FULL			**LOW**		
	Increase	Constant	Decrease	Increase	Constant	Decrease
#	55	33	80	4	14	13
%	**33%**	20%	47%	**13%**	45%	42%

Table 5: Loser behavior when their bid in t-1 was smaller than the best reply to the group.

For winning situations, the pattern is even more pronounced. If a subject won in a certain round, but the bid was higher than the best reply action, 58% of the reactions in this case do not follow learning direction but rather the directional best reply theory in FULL. These subjects decrease their bid although pure learning direction theory would predict a bid raise. On the contrary, this holds only for 29% of the cases in LOW (see Table 6). The difference of winner behavior is statistically significant at the 1-percent level in a Chi-squared test (d.f. = 2, p-value < 0.001). Winners whose bids were above the best reply increase their bid on average by 0.28 units under full information. In LOW, in contrast, they increase their bid on average by 1.88 units.

After qualitative directional learning we now apply a quantitative best reply mechanism to our data, still assuming that participants adjust their bid b^i as a best response to the infor-

	FULL			LOW		
	Increase	Constant	Decrease	Increase	Constant	Decrease
#	95	59	209	58	68	51
%	26%	16%	**58%**	33%	38%	**29%**

Table 6: Winner behavior when their bid in t-1 was larger than the best reply to the group.

mation they receive about the previous period. In LOW, this information is just the bid of their opponent, whereas in FULL they receive information about the whole bid distribution. We compare the size of the bid change predicted by the best reply, $b_t^i - BR_{t-1}^i$, with the actual bid change $b_t^i - b_{t-1}^i$ in every period. If b_t^i is closer to BR_{t-1}^i than b_{t-1}^i, the 'best reply fit' BR_fit takes a positive value according to the percentage of how close the actual change comes to the predicted change:

$$BR\_fit = 1 - \frac{\left|b_t^i - BR_{t-1}^i\right|}{\left|b_{t-1}^i - BR_{t-1}^i\right|} \text{ if } \left|b_t^i - BR_{t-1}^i\right| < \left|b_{t-1}^i - BR_{t-1}^i\right|.$$

In cases in which $\left|b_t^i - BR_{t-1}^i\right|$ is larger than (or equal to) $\left|b_{t-1}^i - BR_{t-1}^i\right|$ or when b_t^i is farther from BR_{t-1}^i than b_{t-1}^i, BR_fit is defined to be zero. Again, we distinguish between cases where the best reply predicts upward or downward adjustments. Taking the above definition, best reply theory captures on average 21% (18%) of the movements in which it predicts an upward adjustment in treatment FULL (LOW) and 35% (40%) of the movements in which it predicts downward adjustment. When best reply predicts no change of the bid, this prediction is correct in 27% (54%) of the cases in treatment FULL (LOW). We see that this quantitative best reply analysis leads to a lower fit to the data than the simpler qualitative directional analysis.

3.3 Dynamic leadership by example

In section 3.2 we have seen that participants in the full information treatment not only learn from their own experience according to learning direction theory, but also consider the distribution of bids within their group when making their decision. In other words, the participants' beliefs about what will happen in the next round depend on the behavior of the population in the previous round. If a participant anticipates this consideration of the other group members, this person might think one step further and recognize that by varying her own bid, she can change the distribution of bids that the others see and, thus, their beliefs. Therefore, this certain participant could influence the behavior of the others assuming that their decisions depend on their observation of the bids within the whole group. If, for example, one subject increases her bid by a large number of points, this action increases the best reply of the others. From what we have seen in section 3.1 we could expect that such an individual drastic bid increase leads on average to an increase of the other bids.[8]

At this point, understanding the behavior in sessions PP 1, FULL 10-1, and FULL 10-5 becomes crucial for the interpretation of the experiment. Remember that our treatments PP and FULL provide the same information as the full information treatment in the papers of Dufwenberg and Gneezy, D/G. In 6 out of those overall 9 sessions with full information, bids typically first move down and then up again. We have argued that the increase of average bids is a collective reaction to an initial increase by one participant. In session PP 1 one of our subjects selects the number 100 all the time from the beginning on. Thus, we do not find a downward movement of average bids during the first half of the experiment because many other participants constantly choose relatively high numbers from the beginning, too, hoping to be matched with this certain subject. In FULL 10-1 and FULL 10-5 basically the opposite happens. None of the subjects seriously tries to initiate a price increase, and therefore the overall bid level declines constantly until the end of the session. In all other sessions, sooner or later one of the subjects drastically raises her own bid to initiate an overall bid increase of the group. We think that the lack of leadership in the two FULL 10 sessions is because

[8]Selten and Apesteguia already suggested such a connection between leadership and cooperation. In some of their experimental sessions they find cyclical pricing similar to the behavioral patterns of our participants. In contrast to our data, however, the price increase in their experiment occurs slowly step by step.

leadership by example requires a higher cognitive effort (see also the analysis in Offerman et al.), and the short duration of these sessions does not always suffice for at least one subject developing the idea of giving a good example.

To capture the dynamic effects that lead to the cyclical bid movements, we need to define what should be treated as leading behavior in the analysis. Obviously, we can speak of leadership if one subject *initially* increases her bid drastically with the intention to make the others follow this increase. Thus, subsequent bid increases are not contained in the definition. To be sure that such an initial bid raise is really intended to make the others follow this step we require a leading bid to fulfill two conditions.

Definition: A bid b_t^i of player i in period t is called *"leading bid"* if the following conditions are satisfied:

(1) $b_t^i > b_{t-1}^j \quad \forall\, j \;\in\; [1;12]$

(2) $b_t^i > b_{t-1}^i + 30$

The player i who places the leading bid is called a *"leader"*.

First, the leading bid has to lie outside the support of the distribution of bids in the previous period. This criterion makes sure that we capture only those bid increases that cannot be explained by other theories like the impulse balance theory of Ockenfels and Selten, learning direction of Selten and Buchta, or imitation learning according to Huck et al. (1999). Second, we require the bid change of the leader from $t-1$ to t to be considerably large. Otherwise, we would include too many relatively small bid increases from a very low overall bid level that are obviously not intended to lead the group to more cooperation but that are nevertheless a little larger than the highest bid in the previous period. The second criterion, that a drastic bid increase should be of at least 30 units, takes two considerations into account. On the one hand, the increase should be large enough not to comprise other bid increases resulting from win experiences, but on the other hand, it should also be small enough so that all real leadership attempts are included. From the picture of individual data we decided to define initial bid increases of more than 30 units as drastic. The border of 30 units is, of course, arbitrary to some extent. However, it is selected to give more weight to the second and more important criterion not to exclude any potential leading bids. Thus, if we consider the reaction to a drastic bid increase, we will at most diminish its observed effects and therefore decrease the statistical significance in a comparison with regular rounds, but we can be sure not to amplify the difference in the comparison artificially by counting only very evident bid increases. Table 7 shows how many leading bids can be determined according to our definition in each session of the experiment.

Treatment	**PP**		**D/G**				**FULL 10**					**FULL 25**			
Session	1	2	2a	2b	2a"	2b"	1	2	3	4	5	1	2	3	4
# Leading bids	0	1	1	0	1	0	1	1	1	1	0	4	4	3	3

Table 7: Number of leading bids per session.

The reaction of the other participants to such a leading bid often starts with some delay. In most situations it seems as if they first wait and see what the others do. Usually only one or two of them follow the leader immediately, but the majority reacts in the second period after the leading bid was made or even later, which can exemplarily be seen from the average bids before and after a leading bid in the FULL 25 sessions in table 8. We also see that the variance in the bid distribution increases after a leading bid. We excluded the leading bid from the calculation of the average bids for two reasons: first, because this might artificially increase the observed variance, and second, because in this data we are mainly interested in the reaction of the other participants.

Thus, a statistical test of the response to a leading bid should concentrate on the bid changes from the initial bid increase in t to the main reaction in $t + 2$. Within all sessions with full

Session	t - 1	t	t + 1	t + 2
FULL 25-1	19.66 (8.20)	16.84 (7.14)	20.34 (13.61)	22.34 (15.64)
FULL 25-2	39.41 (15.73)	33.47 (12.43)	38.47 (28.18)	49.47 (35.30)
FULL 25-3	12.97 (6.17)	13.12 (8.22)	15.91 (12.66)	17.52 (12.99)
FULL 25-4	42.62 (17.48)	39.19 (13.63)	43.62 (19.41)	46.86 (22.86)

Table 8: Average bids (excluding the leading bid) in the FULL 25 sessions before and after a leading bid. Standard deviations in brackets.

information of our experiment and the data from Dufwenberg and Gneezy (2000), we count how many participants on average increase their bid after a leading bid increase in comparison to the period where the initial bid increase took place. Subsequently, we compare this number with the average number of subjects who increase their bid from t to $t + 2$ if there was no drastic bid increase in the two previous periods. If people react to a leading bid in the way we assume, then we would expect that there are significantly more bid raises (and fewer bid reductions) after a bid increase than in the remaining periods of the game.[9] In fact, we find exactly this in all sessions except FULL 10-3 (see Table 9).

In a Wilcoxon signed rank test we can separately show at the 1-percent-level that after a drastic bid increase more participants increase (one-sided, p-value < 0.001) and fewer decrease (one-sided, p-value $= 0.0019$) their bid than in other periods.[10] The same holds if we compare the size of changes quantitatively (one-sided, p-value < 0.001): on average, participants have increased their bid by 11.38 in the second period after a leading bid compared to an average decrease of bids of -2.15 during the remaining periods of the game.

[9]Some sessions are excluded from the analysis because it is not possible to identify clearly periods in which signals were sent. In PP 1 one subject chooses a bid of 100 from the beginning so that our criterion of a drastic bid increase is not matched. If we would include the observations from this session concerning period 1 as the signaling period, this would not change our results in any aspect. In the other cases we cannot detect any clear intention to signal according to our criteria.

[10]The share of constant bids is not significantly different in both cases.

	Leading bid in t-2			No leading bid		
	Increase	**Constant**	**Decrease**	**Increase**	**Constant**	**Decrease**
PP 2	10.00	0.00	2.00	3.00	0.50	8.50
D/G 2a	8.00	0.00	4.00	3.67	0.67	7.67
D/G 2a”	8.50	0.00	2.50	4.25	2.00	5.75
FULL 10-1	3.00	4.00	5.00	2.33	1.00	8.67
FULL 10-2	6.00	4.00	2.00	3.00	1.67	7.33
FULL 10-3	5.00	3.00	4.00	5.17	1.00	5.83
FULL 10-4	4.00	2.00	6.00	3.00	1.00	8.00
FULL 25-1	4.00	3.33	4.67	3.12	1.59	7.29
FULL 25-2	6.33	2.00	3.67	3.82	1.18	7.00
FULL 25-3	5.00	2.33	4.67	3.88	2.06	6.06
FULL 25-4	4.00	4.00	4.00	3.68	2.79	5.53
Average	**5.89**	**2.24**	**3.86**	**3.54**	**1.40**	**7.06**

Table 9: Bid reaction after a leading bid and in periods without a leading bid in the two previous periods. Session 2a” was conducted with 24 participants at a time. For an easier comparison of this data with the other sessions, we divided all numbers from session 2a” by two.

3.4 Motivation of the leader: Efficiency concerns?

Thus, we have a reasonably clear description and explanation for the reaction we observe in response to a leading bid. Much more difficult to understand is the motivation of the subject who makes herself a leader. We cannot clearly say whether these subjects’ own profit is higher due to the increase of the overall bid level. In fact, those who initiate the price increase usually face a total profit that is below the average profit of their group (see Table 10). This is not surprising because the leader does place a very high bid for some periods where it is very unlikely that she will have a lower bid than her opponent. More informative would be a comparison of the leaders’ profits with a hypothetical value what they would have earned without making themselves the leader. This value could be approximated when we compare the profits of the leaders with the average profits of 201.3 in the LOW 25 session, where a potential leading bid cannot be recognised by the group. Since average profits vary considerably between the sessions of each treatment and since we have data for only one LOW 25 session, this comparison is, however, not fully conclusive. We cannot clearly say that leadership seems to pay off since average profits in FULL 25-3 are even lower than in LOW 25. It might nevertheless be the case that leadership pays off for the leader, but not significantly. This is similar to the public good experiments with leadership in Gächter and Renner (2006) and Güth et al. (2004). Exogenously determining a leader who contributes to the public good before the other group members, they show that leadership can enhance higher contributions and therefore increase overall efficiency. The leader who gives the good example, however, does not profit from increasing efficiency in their experiments. In the following we will briefly present possible alternative explanations for why we nevertheless observe leadership.

	FULL 25-1		FULL 25-2		FULL 25-3		FULL 25-4	
ID	**Leader?**	**Profit**	**Leader?**	**Profit**	**Leader?**	**Profit**	**Leader?**	**Profit**
1	no	340.0	no	562.5	no	279.0	no	300.0
2	no	337.0	no	517.0	no	159.5	no	554.5
3	no	317.5	no	576.0	no	256.0	yes (R17)	273.0
4	no	271.5	no	523.0	yes (R8)	216.0	no	561.0
5	no	288.0	yes (R22)	270.0	no	160.5	no	707.0
6	no	322.0	no	432.0	no	191.0	no	760
7	no	223.5	no	486.0	yes (R19)	114.0	no	669.0
8	no	309.5	no	500.0	no	151.0	yes (R8)	326.5
9	no	320.0	no	514.0	no	183.0	no	660.0
10	yes (R6, R11)	159.5	yes (R4, R22)	297.0	no	182.5	yes (R8)	430.0
11	yes (R21)	170.0	no	608.0	yes (R16)	150.0	no	521
12	yes (R21)	179.5	yes (R13)	413.5	no	216.5	no	337
Average		**269.83**		**474.92**		**188.25**		**508.25**

Table 10: Leadership decision and profits. Rounds of the respective leading bid are in brackets.

As we have seen in table 10, leadership helps to increase efficiency but does not necessarily pay off in terms of profit for the leader. We have two possible explanations why leadership occurs nevertheless. One is partially suggested by the evaluation of comments in the questionnaire, where some of the leaders (7 out of 17) explicitly state that they were hoping to increase the average price level in order to increase their own profits. Thus, they seem to overestimate their future net profit increase from inducing an increase of the overall price level and do not sufficiently take into account the profits they abandon while placing the highest bid of their group.

Another explanation would be that they strongly care about overall efficiency even if it does not pay off in terms of their own profit. In the questionnaire, 12 out of 17 leaders wrote a sentence such as 'Large numbers are good for all participants.' To find out whether efficiency concerns may play a role for leadership in our experiment, we conducted some distribution games similar to those in the paper by Engelmann and Strobel (2004), directly after six of our FULL sessions. Participants are asked to choose between three possible distributions of money between themselves (as person B) and two others (person A and C). The three options (see Table 11) are designed in such a way that option 1 is the most efficient but yields the least profit for person B, and option 3, in contrast, is least efficient but yields the highest profit for person B. Option 2 is in between.

	Option 1	**Option 2**	**Option 3**
Person A	20	17	12
Person B*	8	9	10
Person C	6	6	6

Table 11: Choice set in the distribution game. *Person B can choose between Options 1, 2, and 3.

Most leaders (73%) indicate a relatively high willingness to pay for efficiency by choosing option 1 or 2 (see Table 12). The share of efficiency lovers among the non-leaders is lower (49%). Thus, at first sight there seems to be a correlation between leadership in our experiment and a high willingness to increase overall efficiency at the cost of a lower personal profit in the distribution game. However, the difference between the decisions of leaders and non-leaders is statistically not significant (Chi-squared test, d.f. = 2, p-value = 0.3113). We think that this is because efficiency concerns do interfere with endowment effects since the distribution game is conducted after the main experiment. Especially in sessions FULL 25-2 and FULL 25-4 subjects manage very well to coordinate in the bidding game and earn quite a lot of money during the main experiment. We see that several participants in those two

	FULL 10-3		**FULL 10-4**		**FULL 10-5**	
ID	**Leader?**	**Choice**	**Leader?**	**Choice**	**Leader?**	**Choice**
1	no	2	no	3	no	3
2	no	3	no	3	no	3
3	yes	1	yes	3	no	2
4	no	1	no	3	no	2
5	no	2	no	3	no	3
6	no	3	no	3	no	3
7	no	1	no	3	no	2
8	no	1	no	3	no	2
9	no	2	no	3	no	3
10	no	3	no	3	no	3
11	no	1	no	3	no	2
12	no	1	no	3	no	2
Average profit	**137.67**		**79.33**		**112.33**	
	FULL 25-2		**FULL 25-3**		**FULL 25-4**	
ID	**Leader?**	**Choice**	**Leader?**	**Choice**	**Leader?**	**Choice**
1	no	1	no	3	no	1
2	no	3	no	3	no	1
3	no	1	no	2	yes	3
4	no	1	yes	1	no	1
5	yes	1	no	3	no	1
6	no	3	no	3	no	1
7	no	1	yes	2	no	3
8	no	1	no	1	yes	1
9	no	1	no	3	no	1
10	yes	3	no	3	yes	1
11	no	1	yes	2	no	3
12	yes	1	no	3	no	1
Average profit	**474.92**		**188.25**		**508.25**	

Table 12: Leadership and choices in the distribution game.

sessions who did not initiate a drastic bid increase nevertheless decided for the most efficient option 1. It is therefore also likely that subjects who realised a large profit before tend to be more generous in the distribution game without caring for efficiency in general. Our results from the distribution game are therfore only to a limited extent conclusive and should be considered with care.

4 Discussion

In this paper we studied how information provision can activate group dynamic processes in Bertrand duopoly experiments. When only information about the participants' own opponent is provided, we find that prices (or bids, in the other terminology) do slowly converge to the competitive Nash equilibrium. On the contrary, when participants receive additional information about the other subjects' choices, we find no convergence to competitive prices. Our findings suggest that, in contrast to Dufwenberg and Gneezy (2000), in the long run two firms are enough for perfect competition as long as firms have no immediate information about the other potential opponents around. In the case with full information about other markets we find that prices cyclically move up and down. Those circles can arise because some participants use the publicity of information to inform the group about their willingness to cooperate. Other participants follow the example of this leader and increase their prices too, but as soon as overall prices reach a high level, price cuts become much more tempting so that another price war is inevitable.

Explaining the results, we refer to learning direction theory (Selten and Buchta 1999, Selten and Stoecker 1986). When no information about the other participants is provided, subjects can only learn from their own experience. According to Ockenfels and Selten (2005), downward impulses after a lost match are stronger than upward impulses after matches that have been won. This is confirmed by the data in our LOW treatment. In the case with full information, additional impulses enter the decision process. Comparing their price with the best reply to other participants' prices, subjects decide not only according to their immediate experiences but also according to the consideration of what would have happened had they met one of the other participants. This effect is even more pronounced when one of the subjects notably increases her price, which we call a leading bid. In the following periods, behavior can to a high degree be explained as a reaction to the leading bid.

Regarding the motivation of the leader who places the initial high bid that triggers a new circle of average bids, we think that the answer to this question is a mixture of basically two aspects. First, the leaders egoistically hope to benefit themselves by initiating a upward movement of bids. Second, we speculate that they might be more altruistically motivated by the prospect to increase overall efficiency of behavior within their matching group even without increasing their own profit.

References

Abbink, K., Sadrieh, A., Zamir, S., 2004. Fairness, public good, and emotional aspects of punishment behavior. Theory and Decision 57, 25-57.

Altavilla, C., Luini, L., Sbriglia, P., 2006. Social learning in market games. Journal of Economic Behavior and Organization 61, 632-652.

Bohnet, I., Frey, B. S., 1999. Social distance and other-regarding behavior in dictator games: Comment. American Economic Review 89, 335-339.

Dufwenberg, M., Gneezy, U., 2002. Information disclosure in auctions: An experiment. Journal of Economic Behavior and Organization 48, 431-444.

Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: An experimental study. International Journal of Industrial Organization 18, 7-22.

Engelmann, D., Strobel, M., 2004. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. American Economic Review 94, 857-869.

Fischbacher, U., 2007. z-Tree - Zurich toolbox for readymade economic experiments. Experimental Economics 10, 171-178.

Gächter, S., Renner, E., 2006. Leaders as belief managers for the voluntary provision of public goods. Working Paper, University of Nottingham.

Güth, W., Levati, M.V., Sutter, M., van der Heijden, E., 2007. Leading by example with and without exclusion power in voluntary contribution experiments, Journal of Public Economics 91, 1023-1042.

Huck, S., 2004. Oligopoly. In: Friedman, D., Cassar, A. (Eds.). Economics lab. An intensive course in experimental economics. London: Routledge, 105-114.

Huck, S., Normann, H., Oechssler, J., 2000. Does information about competitors' actions increase or decrease competition in experimental oligopoly markets?. International Journal of Industrial Organization 18, 39-57.

Huck, S., Normann, H., Oechssler, J., 1999. Learning in cournot oligopoly: An experiment. The Economic Journal 109, C80-95.

Neugebauer, T., Selten, R., 2006. Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets. Games and Economic Behavior 54, 183-204.

Ockenfels, A., Selten, R., 2005. Impulse balance equilibrium and feedback in first price auctions. Games and Economic Behavior 51, 155-170.

Offermann, T., Potters, J., Sonnemans, J., 2002. Imitation and belief learning in an oligopoly experiment. Review of Economic Studies 69, 973-997.

Selten, R., Apesteguia, J., 2005. Experimentally observed imitation and cooperation in price competition on the circle. Games and Economic Behavior 51, 171-192.

Selten, R., Buchta, J., 1999. Experimental sealed bid first price auctions with directly observed bid functions. In: Budescu, D., Erev, I., Zwick, R. (Eds.). Games and human behavior: Essays in the honor of Amnon Rapoport. Mahwah, NJ: Lawrenz Associates, 79-102.

Selten, R., Stoecker, R., 1986. End behavior in sequences of finite prisoner's dilemma supergames. Journal of Economic Behavior and Organization 7, 47-70.

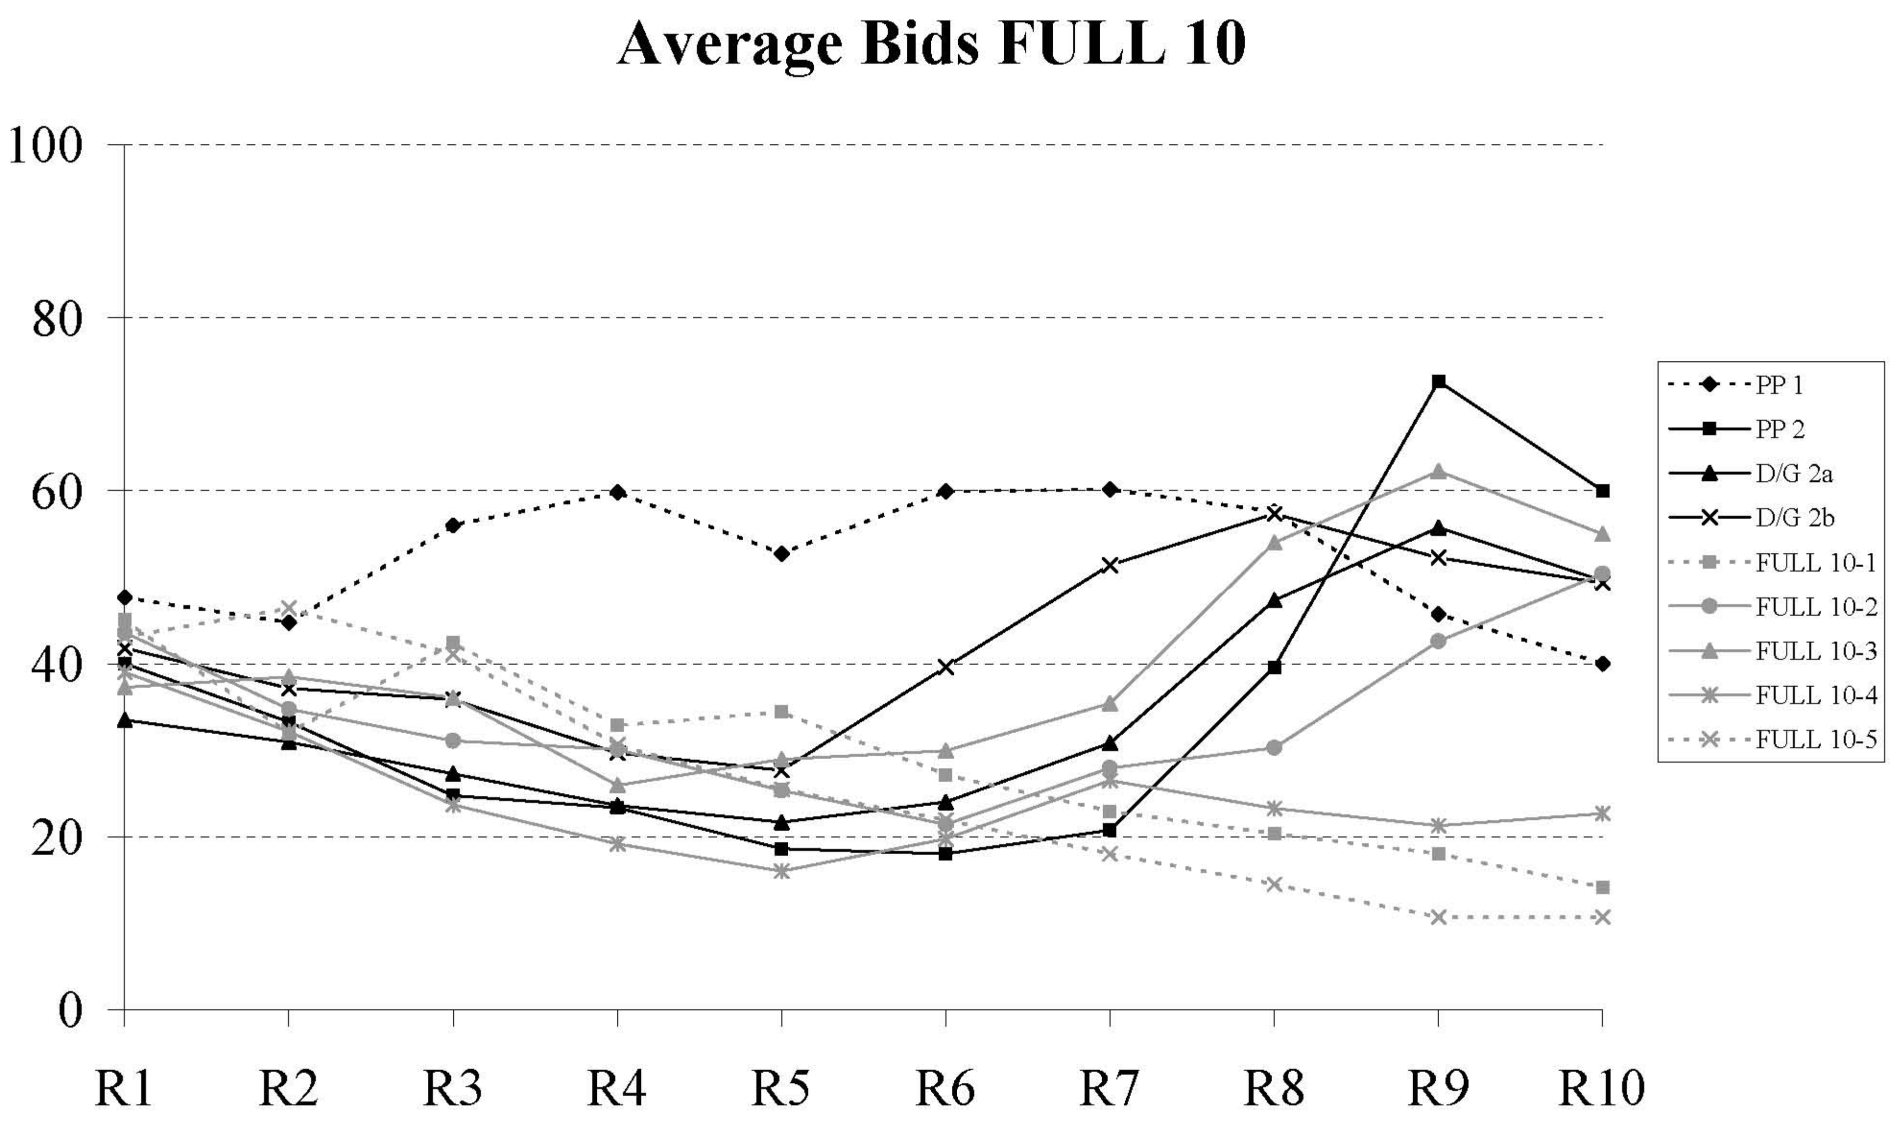
Average Bids FULL 10
100
80
60
40
20
0
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
PP 1
PP 2
D/G 2a
D/G 2b
FULL 10-1
FULL 10-2
FULL 10-3
FULL 10-4
FULL 10-5

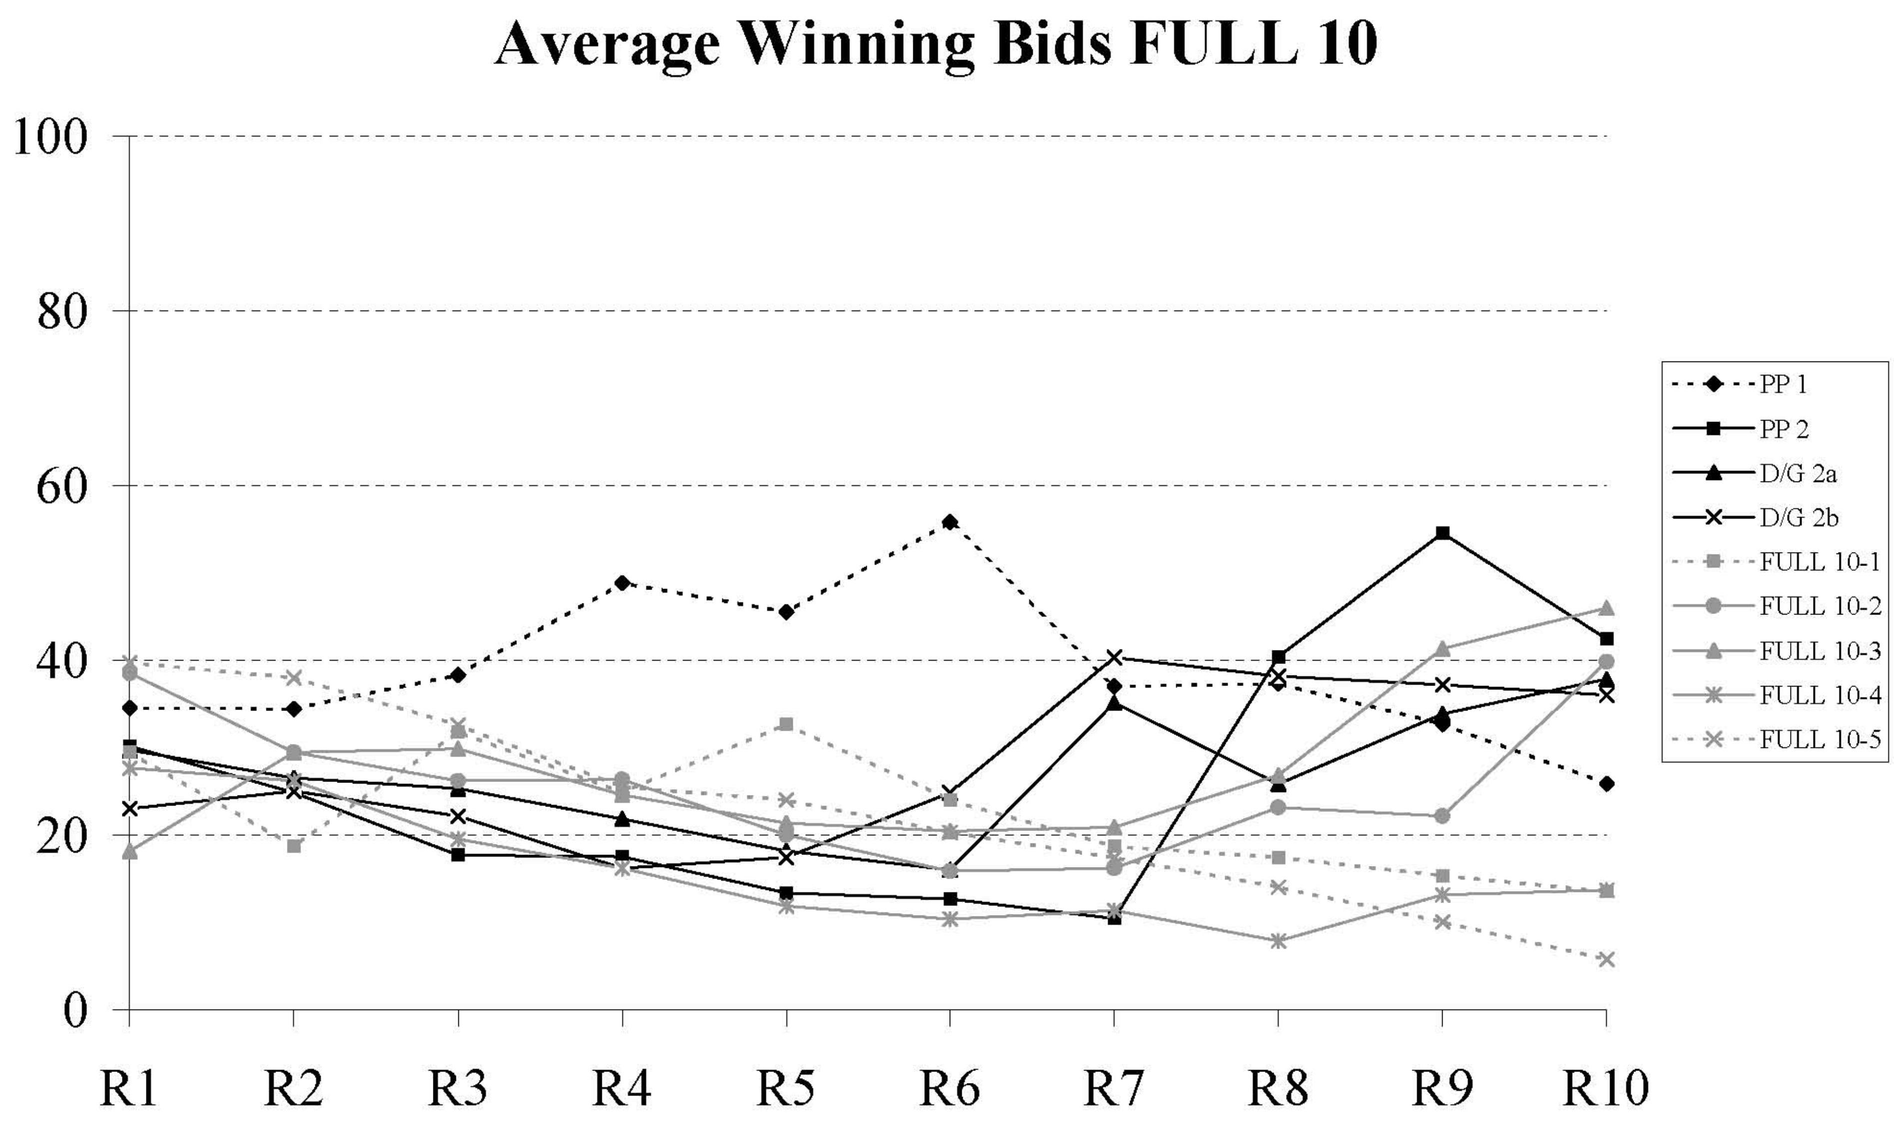
Average Winning Bids FULL 10
100
80
60
40
20
0
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
PP 1
PP 2
D/G 2a
D/G 2b
FULL 10-1
FULL 10-2
FULL 10-3
FULL 10-4
FULL 10-5

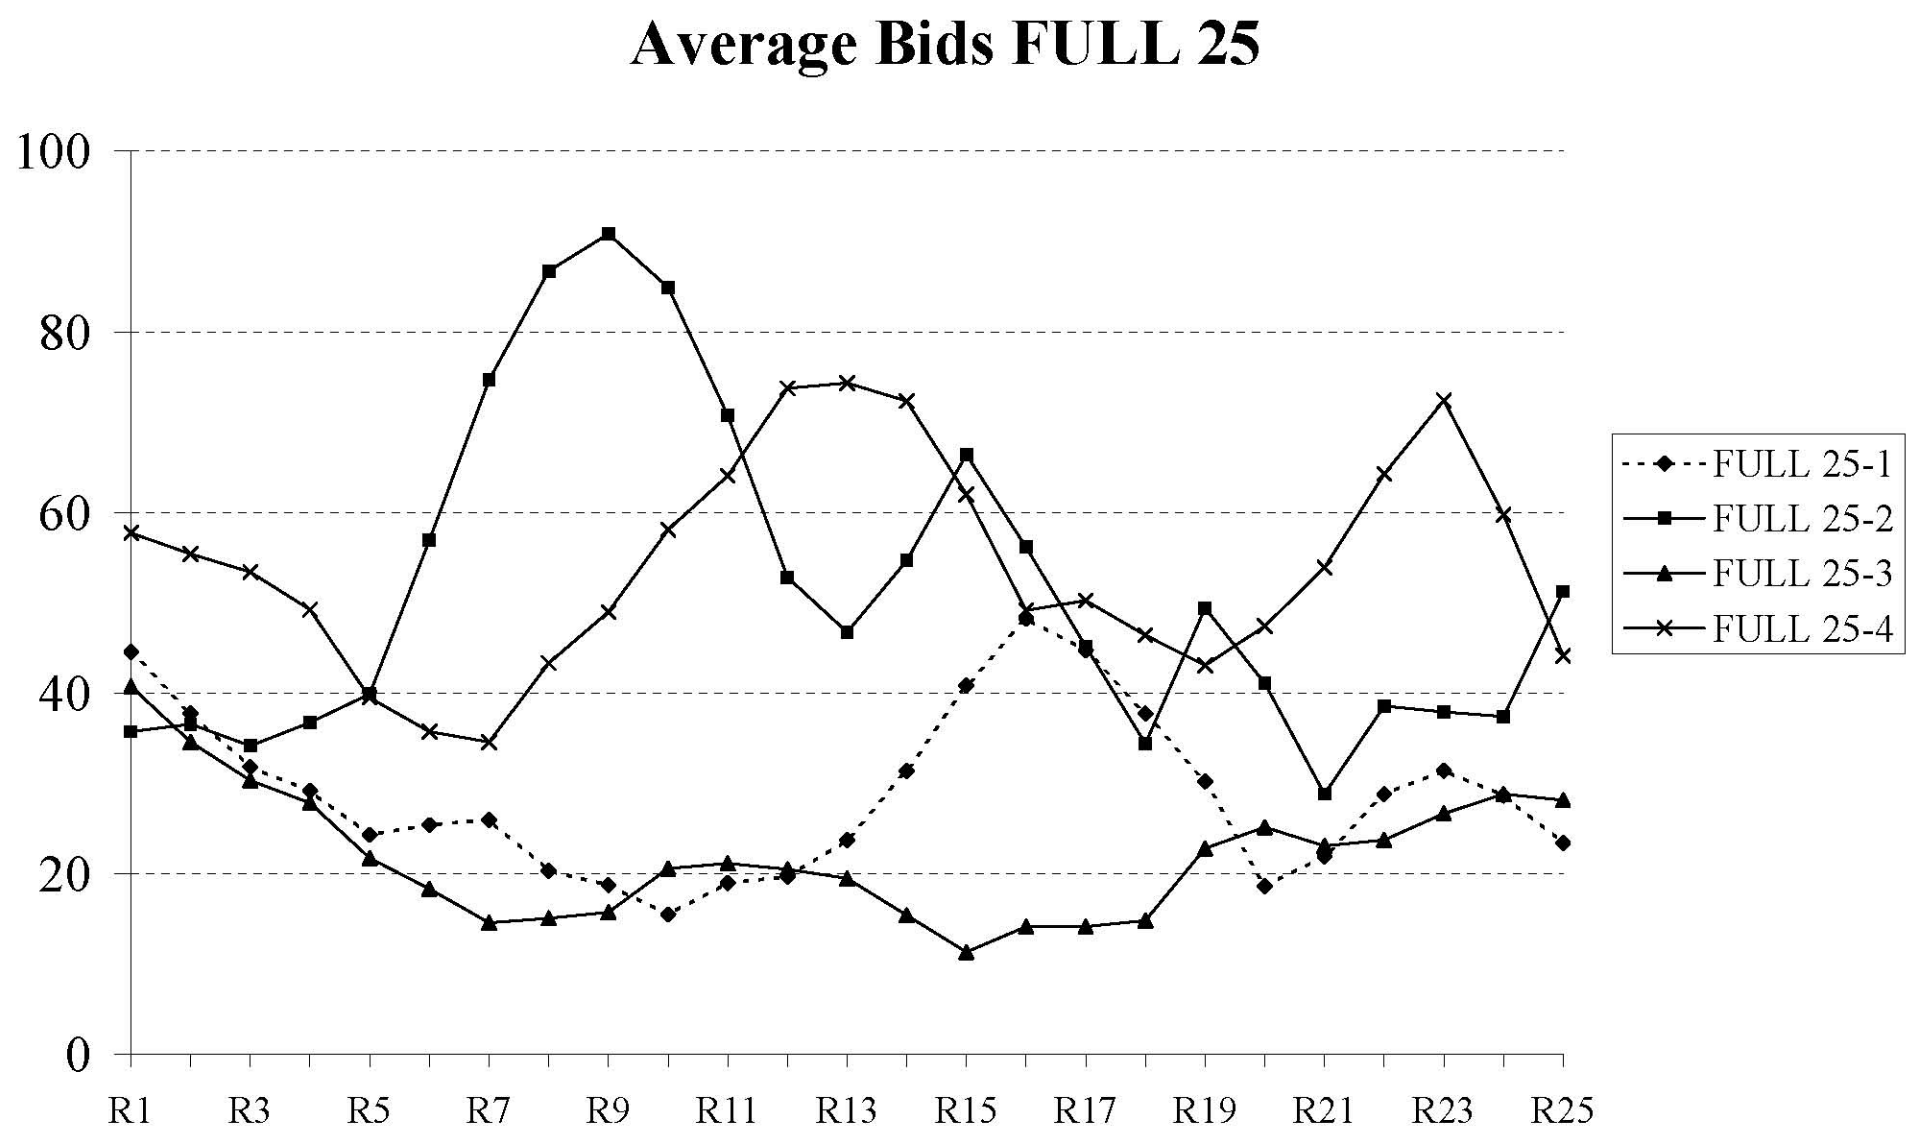
Average Bids FULL 25
100
80
60
40
20
0
R1
R3
R5
R7
R9
R11
R13
R15
R17
R19
R21
R23
R25
FULL 25-1
FULL 25-2
FULL 25-3
FULL 25-4

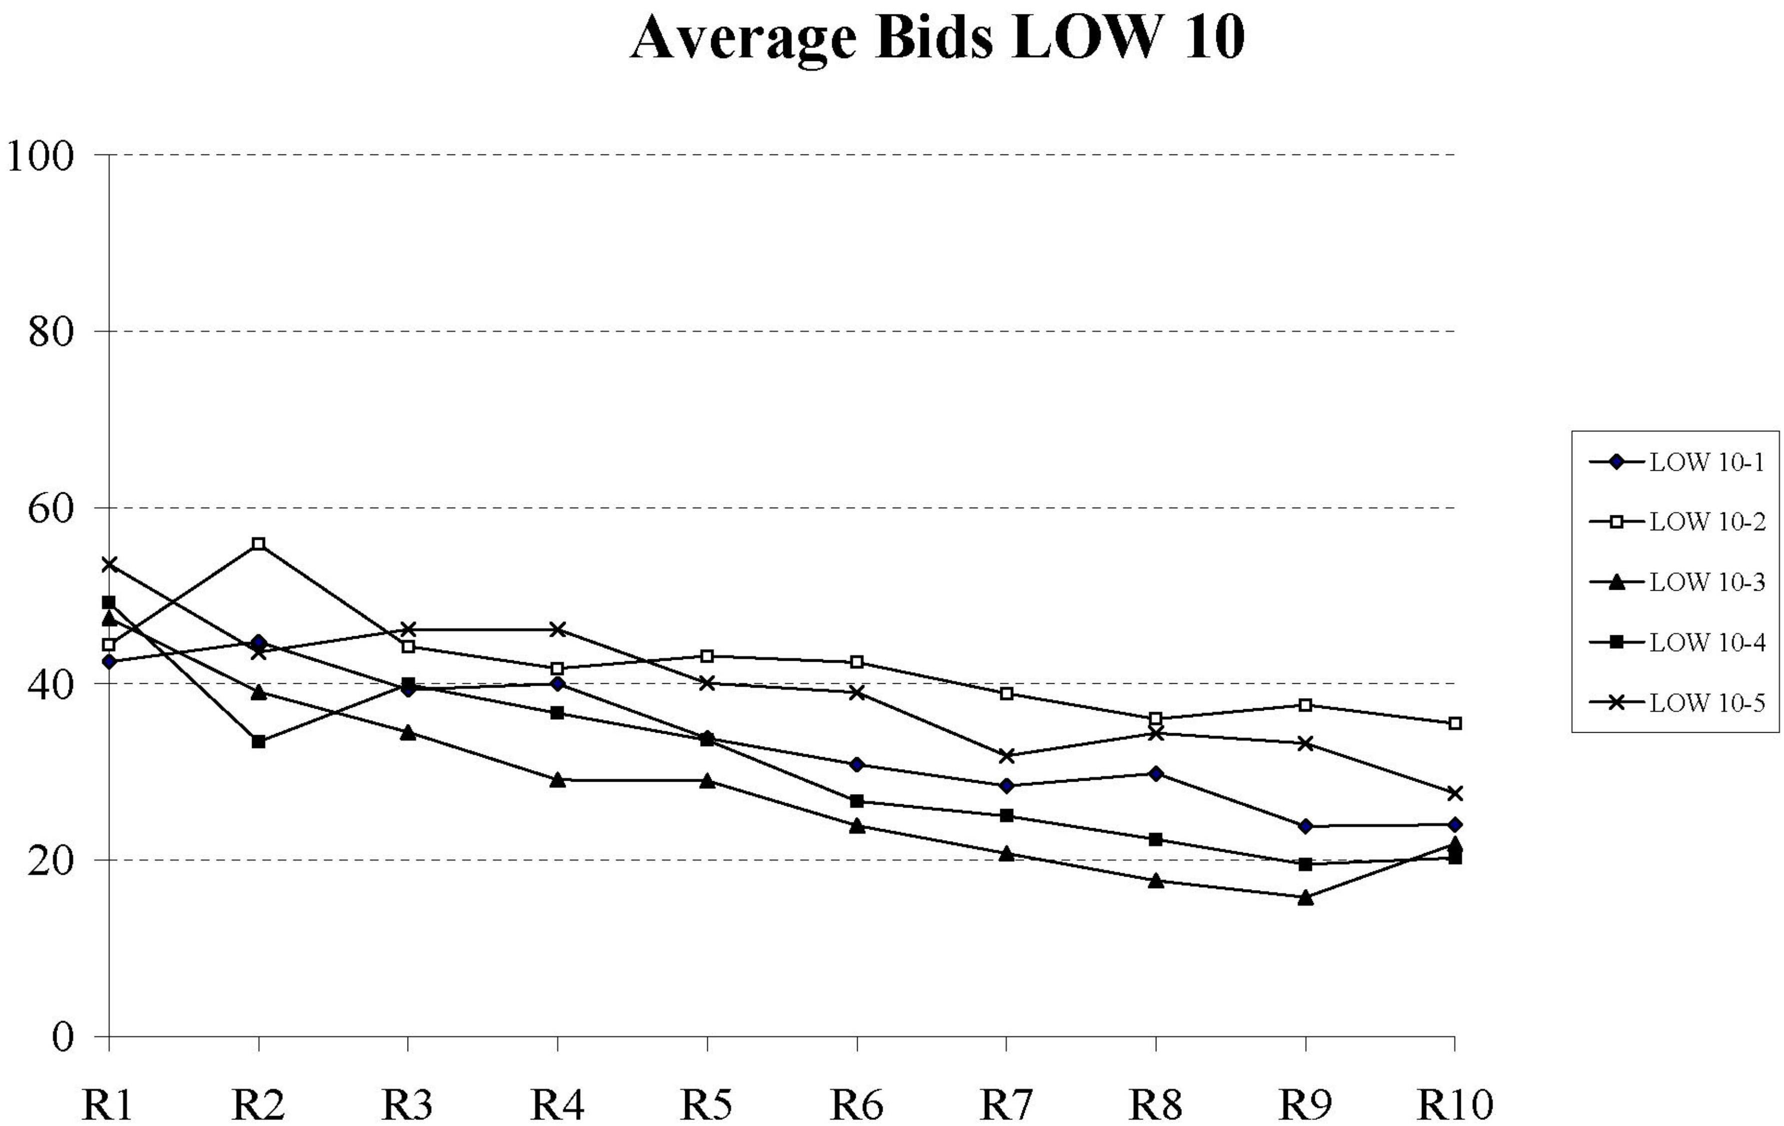
Average Bids LOW 10
100
80
60
40
20
0
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
LOW 10-1
LOW 10-2
LOW 10-3
LOW 10-4
LOW 10-5

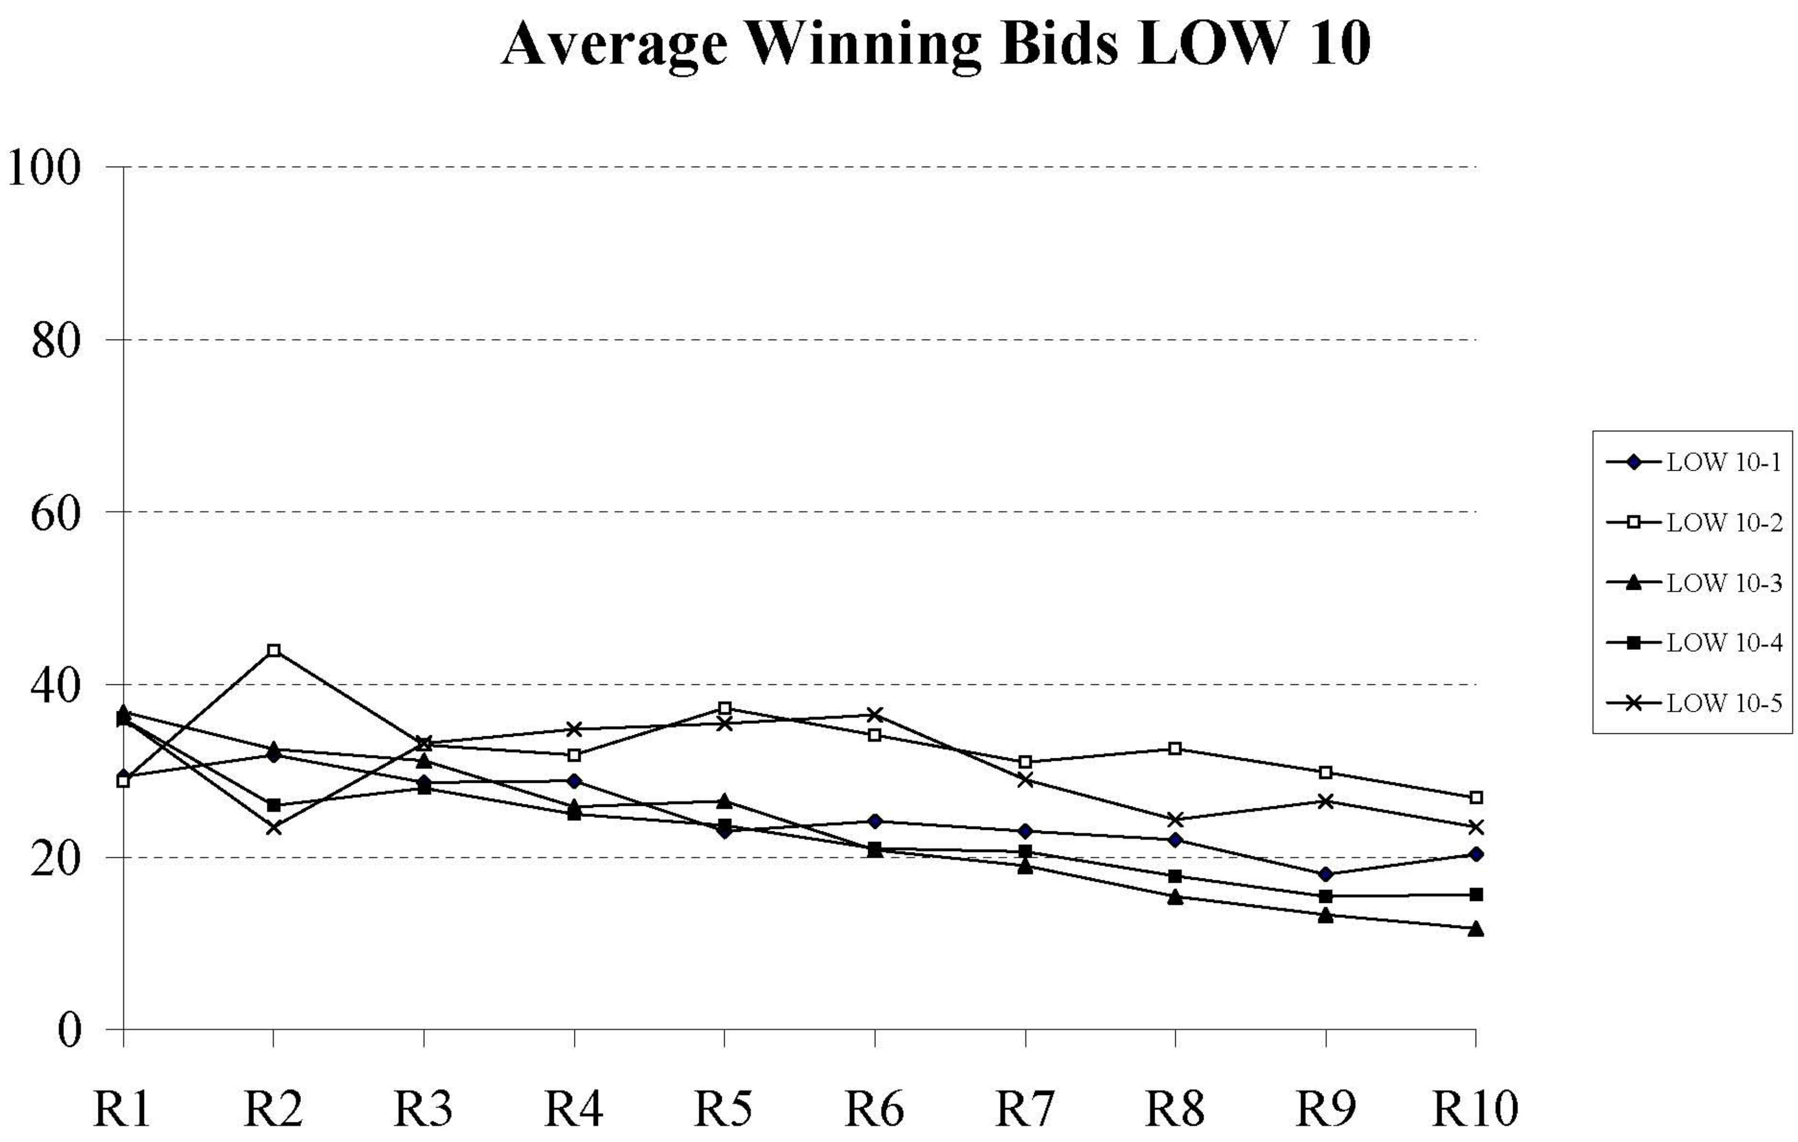
Average Winning Bids LOW 10
100
80
60
40
20
0
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
LOW 10-1
LOW 10-2
LOW 10-3
LOW 10-4
LOW 10-5

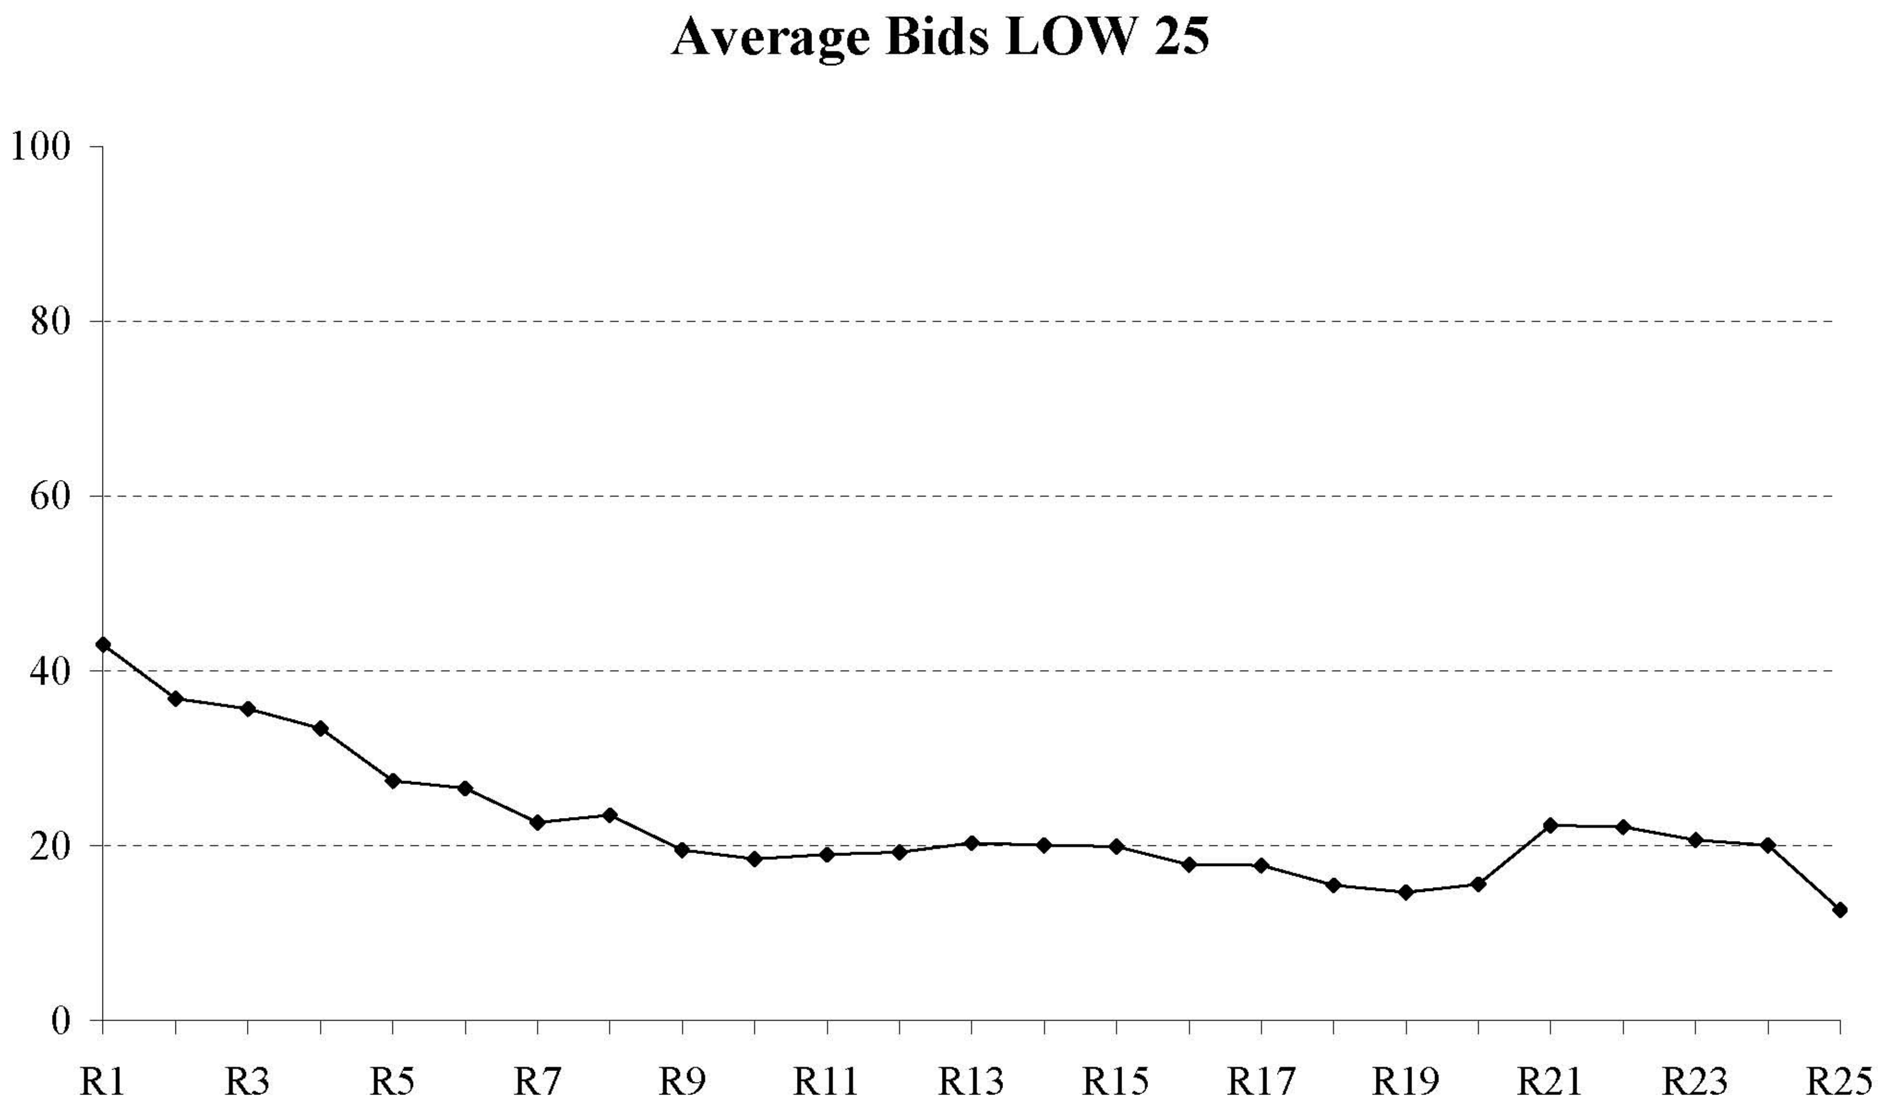
Average Bids LOW 25
100
80
60
40
20
0
R1
R3
R5
R7
R9
R11
R13
R15
R17
R19
R21
R23
R25