24

by KLAUS ABBINK* and JORDI BRANDTS**

* University of Nottingham
** Institut d'Anàlisi Econòmica (CSIC), Barcelona

revised version, September 2003

Abstract

Bertrand competition under decreasing returns involves a wide interval of pure strategy equilibrium prices. We first present results of experiments in which two, three and four identical firms repeatedly interact in this environment. Less collusion with more firms leads to lower average prices. With more than two firms, the predominant market price is 24, a price not predicted by conventional equilibrium theories. This phenomenon can be captured by a simple imitation model and by a focal point explanation. For the long run, the model predicts that prices converge to the Walrasian outcome. We then use data from three new treatments to properly test the imitation model against the focal point notion.

Keywords

Laboratory experiments, industrial organisation, oligopoly, price competition, co-ordination games, learning

JEL Classification Codes

C90, C72, D43, D83, L13

Acknowledgements

The authors thank David Rodríguez for very skillful research assistance and for help in running the experiments, and Mark Armstrong, Antonio Cabrales, Enrique Fatás, Martin Sefton, Chris Starmer, two anonymous referees, and seminar participants in Amsterdam, Barcelona, Berlin, Egham, Grenoble, Jena, London, Nottingham, Oxford, Pittsburgh, Stockholm and Tilburg for helpful comments. Financial support by the European Union through the TMR research network ENDEAR (FMRX-CT98-0238) and the Spanish *Ministerio de Educación y Cultura* (SEC2002-01352), the Nuffield Foundation and the University of Nottingham is gratefully acknowledged.

Authors

Klaus Abbink	Jordi Brandts
School of Economics The University of Nottingham University Park Nottingham, NG7 2RD United Kingdom Phone +44-115-9514768 Fax +44-115-9514591 klaus.abbink@nottingham.ac.uk	Institut d'Anàlisi Econòmica (CSIC) Campus UAB 08193 Bellaterra Spain phone +34-93-5806612 fax +34-93-5801452 jordi.brandts@uab.es

1. Introduction

In this paper we present an experimental study of Bertrand competition under decreasing returns as well as a theoretical model of the behaviour we observe. Part of our motivation is to contribute to delineating a more complete picture of oligopolistic behaviour from an experimental viewpoint. In particular, we are interested in how price and efficiency levels depend on the number of firms. With respect to price levels we wish to find out whether they remain above Walrasian levels with three or four firms.

There is, however, a second more conceptual rationale for our work. The oligopoly setting we investigate is based on a large co-ordination stage game with many Nash equilibria, so that *a priori* actual behaviour is hard to pin down theoretically. A good part of our paper, hence, deals with the formulation and testing of a model that can explain our data. We believe that the insights from this effort reach beyond the particular price competition model we study. Laboratory experiments are well suited for investigating problems of multiplicity of equilibria; they may be able to supply precise predictions on the basis of a variety of behavioural factors. Examples of the use of experiments in contexts with equilibrium multiplicity are Van Huyck, Battalio and Beil (1990), Brandts and Holt (1992) and Cooper, DeJong, Forsythe and Ross (1990).

How the number of firms affects prices when there are only few competitors in the market is one of the central themes of the economic analysis of oligopoly. Theoretical analysis has provided some answers to the above question. The equilibrium proposed by Cournot (1838) for the case of quantity competition yields predictions of a unique price and of a price-cost margin which is decreasing in the number of firms. For the case of price competition Bertrand's (1883) analysis led to somewhat less natural predictions for the case of constant returns and no capacity constraints: With more than one firm prices will be equal to marginal cost, independently of the number of firms.

Economists' stylised view of oligopolistic competition has been complemented by experimental studies on the subject. Numerous studies report results from quantity competition environments. Huck, Normann and Oechssler (2003) provide a recent survey and synthesis of experimental work on quantity competition. Their conclusion is that duopolists sometimes manage to collude, but that in markets with more than three firms collusion is difficult. With exactly three firms, Offerman, Potters, and Sonnemans (2002) observe that market outcomes depend on the information environment: Firms collude when they are provided with information on individual quantities, but not individual profits. In many instances, total average output exceeds the Nash prediction and furthermore, these deviations are increasing in the number of firms. The price-cost margins found in experimental repeated quantity competition are, hence, qualitatively consistent with the Cournot prediction for the static game.

Dufwenberg and Gneezy (2000) study the effects of the number of firms in a standard Bertrand competition framework with constant marginal cost and inelastic demand. In their experiments, price is *above* marginal cost for the case of two firms but equal to that cost for three and four firms. Their results just modify the theoretical prediction in a simple way: by changing from two to three the number of firms from which on the price can be expected to equal marginal cost.[1]

Price competition leads to different predictions in other settings. Dastidar (1995) analyses price competition in a model in which firms operate under decreasing returns to scale and have to serve the whole market. The apparently small move from constant to increasing marginal costs changes the equilibrium prediction dramatically: Positive price-cost margins are now possible in a pure strategy equilibrium.[2]

As mentioned above, the equilibrium prediction in this model will be strongly indeterminate. There will, in fact, be a whole set of equilibrium prices with a minimum zero-profits equilibrium price below the Walrasian price and a maximum price above it. The maximum and minimum prices both decrease with the number of firms so that a sufficient increase of the number of firms implies, in a sense, a prediction of lower prices. However, an increase of just one or two firms, as often envisioned in the context of anti-trust analysis, is well compatible with equilibrium predictions of decreasing, unaffected, or even increasing prices.

The content of this paper is the result of two rounds of work. In the first round, we studied behaviour under Bertrand competition with decreasing costs in fixed groups of two, three and four competitors that interacted with each other for 50 rounds. Our results for those treatments show that average prices are always above the Cournot and, hence, the Walrasian level. With respect to the distribution of prices, we find a remarkable degree of co-ordination. Observed prices can be organised according to two principles: The first is the cartel price, which is the modal outcome in n=2, but becomes less frequent with more firms. The second organising principle is a specific price level, 24 in our experimental design, which is the market price in an absolute majority of cases in n=3 and n=4, as well as the second most frequent outcome in n=2. Average prices are decreasing with the number of firms, mainly because there is less collusion and there are more 24-prices as the number of firms increases.

The price level of 24, which was not predicted by any benchmark theory we are aware of, is characterised by the feature that unilateral undercutting of this price level leads to absolute losses. We propose a simple imitation model, which reproduces this phenomenon and also other regularities in our data quite well. Imitation has recently been used as a behavioural principle to explain observations from other contexts.

[1] Further experiments on price competition include Davis and Holt (1994), Kruse, Rassenti, Reynolds and Smith (1994) and Morgan, Orzen and Sefton (2001). These three studies consider environments with mixed strategy equilibria, and find price dispersion qualitatively similar to the relevant equilibrium.

In the second round of work we put the imitation model to a test by gathering data from new treatments. We find that the imitation model can explain some but not all the additional evidence we collect. Imitation seems to be one of the heuristics subjects use in selecting their choices, but not the only one. Subjects appear to recur to imitation when clues for more thoughtful behaviour are hard to come by.

2. The model and the experimental design

2.1. Demand, cost curves, and stage game equilibria

In our experimental markets the demand function is linear and the cost function is quadratic and common to all firms. The intuitive explanation for the equilibrium price indeterminacy in this case is rather simple and can be easily illustrated with the figures 1 and 2, in which D denotes demand, MC marginal cost and AC//S the average cost and supply curve, for the case of two firms. The figures show the lowest-price and the highest-price equilibrium, respectively. Observe first, in figure 1, the configuration with both firms setting the common price $\underline{p}$ and, hence, sharing the market equally. The two shaded triangles try to illustrate that at this price firms' profits are zero. Since the firm with the lowest price serves all the market, any upward deviation leads to zero profits. Any downward deviation implies a price that is below the marginal cost of all the additional units that the deviating firm needs to supply to the market and, hence, every additional unit, leads to a decrease in profits for that firm. Situations with a common price lower than $\underline{p}$ imply negative profits and can naturally not be equilibria, since upward deviations lead to zero profits.

A wide range of higher prices can also be supported in equilibrium, leading to positive profits for both firms. The same logic applies: Overcutting leads to zero profits, since the deviating firm does not sell any units. Undercutting the common price, however, leads to lower profits since the firm must sell additional units at excessive marginal costs. There exists, however, a high enough price level, denoted in figure 2 by $\overline{p}$, at which the price is sufficiently above marginal cost so that downward deviations lead to zero marginal profits. In figure 2 the two shaded triangles mean to illustrate how positive and negative additional profits from a deviation will cancel out. Prices above this one are not equilibrium ones any more.

How does the number of firms in the market affect the above argument? Both the minimum and the maximum price decrease with the number of firms; in fact the minimum price converges to zero and the maximum price to a positive lower bound. Due to the fact that marginal costs are increasing a larger number of firms can share the market with zero profits at a lower

[2] See Vives (1999) for a discussion of this model. He considers the assumption of having to serve the whole market to be plausible in cases when there are large costs of turning consumers away. For instance, "common carrier" regulation typically requires firms to meet all demand at announced prices.

price level. The decrease in the maximum price is also due to the fact that marginal costs are increasing. With more firms individual deviations would be profitable at a lower price as an undercutting firm would be able to sell many additional low cost units.

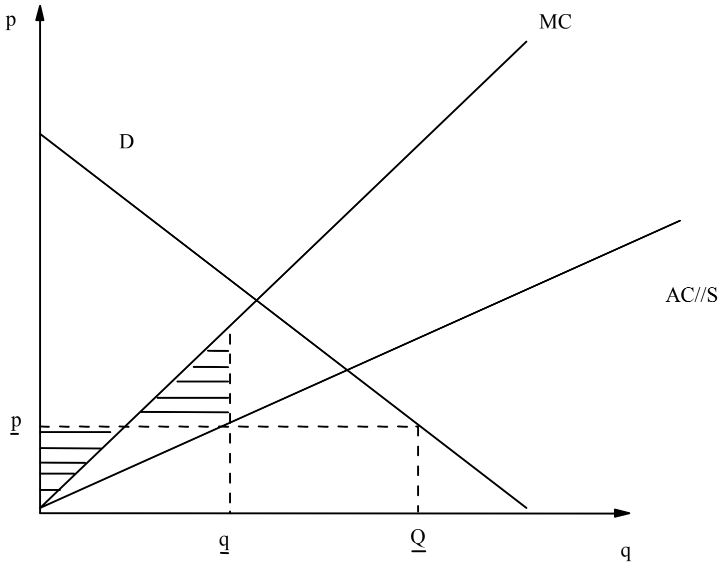

Figure 1

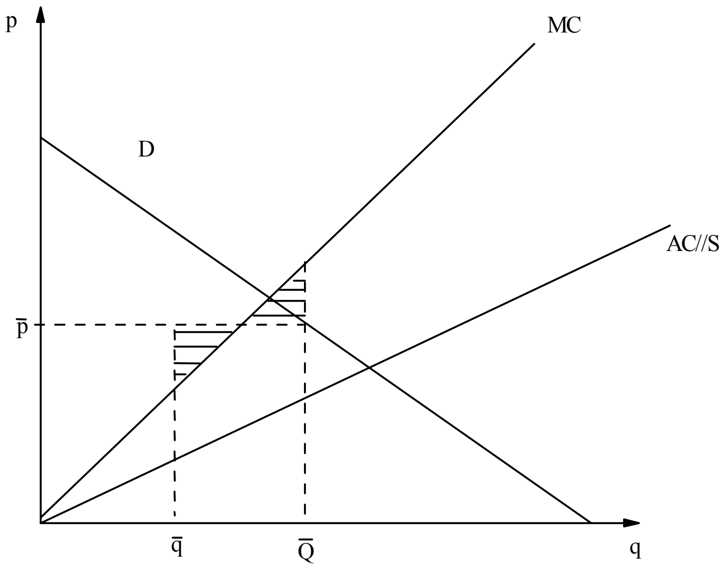

Figure 2

2.2. Design choices

In our experiment, we presented the implications of the different pricing decisions in a simplified way. In each round subjects had to individually choose a number on the basis of payoff tables derived from the one shown in table 1. In the treatment with 4 firms, subjects' payoff table was exactly table 1, except that subjects did not see the shaded areas we have inserted to mark several benchmarks. The first column shows all the numbers – prices – that could be chosen.[3] The second column shows, for a given round, a subject's profit in *talers* (the fictitious experimental currency) if he alone has chosen that number and the other 3 players have chosen higher numbers. The third column shows a subject's profit if he and another player have chosen that number and the other two have chosen higher numbers. Columns 4 and 5 refer to the analogous cases with 3 and 4 subjects and column 6 shows a subject's payoff if he has chosen a higher number than all the other players. For n=3 subjects saw a table without column 5 and for n=2 a table without columns 4 and 5. Appendix A contains the instructions.

The payoffs in table 1 are derived from the demand function $D(p)=100–1.5p$ and the common cost function $C(q)=1.5q^2/2$. This choice of parameters makes it possible to have the highest equilibrium price be above the price corresponding to the Cournot benchmark. The price grid was chosen in a manner that all possible price choices are equidistant from one another.

2.3. Benchmarks and predictions

It can be seen in table 1 that for n=2 the equilibria are between 13 and 30, for n=3 between 7 and 28, and for n=4 between 3 and 27. For n=2, n=3 and n=4 the Walrasian prices and the Cournot equilibrium prices, which result from the above demand and cost functions, correspond to prices in between those that are eligible. The Walrasian price is between 23 and 24 for n= 2, between 16 and 17 for n=3 and between 12 and 13 for n=4. The Cournot price is between 29 and 30 for n= 2, between 22 and 23 for n=3 and between 17 and 18 for n=4.

Given that the equilibrium range is large, we may look at common selection criteria for tighter predictions for the experiment. Harsanyi and Selten (1988) propose that if a unique *payoff dominant* equilibrium exists, rational players should always select it.[4] An equilibrium is strictly payoff dominant if all players receive higher payoffs in this equilibrium than in any other equilibrium. In our set-up, the equilibrium with the highest price is the payoff dominant one, as can easily be seen in table 1.

Though not an equilibrium, the cartel price may provide a prediction for the experiment. The cartel price is the price level maximising the joint payoff if set by all firms in the market. It is above the range of stage game equilibrium prices, at p=33 for n=2, p=30 for n=3, and p=28 for n=4, thus decreasing with the number of firms.

[3] The numbers for the prices are not those that result from our parameter choices. They were relabelled to go from 1 to 40 for simplicity.

[4] A second prominent concept developed by Harsanyi and Selten (1988) is that of risk dominance. The authors propose this criterion only for games without a payoff dominant equilibrium. This concept takes into account the risk attached to co-ordination failure. Risk dominant equilibria are straightforward to find for two-player games with two equilibria. For games with more than two players, the concept in general defies analytical solution, and also numerical computation gets extremely difficult for larger games (Herings and van den Elzen (2002)). Identifying a risk dominant equilibrium for a 40x40x40x40 game with 26 pure strategy equilibria, as our n=4 treatment involves, is practically impossible.

Table 1. The payoff table for n=4

Column 1	Column 2	Column 3	Column 4	Column 5	Column 6
Number chosen by you	Profit if this number is the unique lowest of the four	Profit if this number is the lowest and has been chosen by two players	Profit if this number is the lowest and has been chosen by three players	Profit if this number is the lowest and has been chosen by four players	Profit if this number is not the lowest
40	777	473	334	258	0
39	784	489	348	269	0
38	783	503	360	279	0
37	777	514	370	288	0
36	763	522	379	296	0
35	743	528	387	303	0
34	716	532	393	310	0
33	683	***533***	398	315	0
32	642	532	402	319	0
31	596	528	403	322	0
30	542	522	***404***	323	0
29	482	514	403	324	0
28	415	503	401	***325***	0
27	341	489	397	324	0
26	261	473	392	322	0
25	174	455	385	319	0
24	**81**	**434**	**377**	**315**	**0**
23	-20	411	367	310	0
22	-126	385	356	304	0
21	-240	357	344	297	0
20	-360	326	330	290	0
19	-487	293	314	281	0
18	-621	257	298	271	0
17	-761	219	279	260	0
16	-908	179	260	248	0
15	-1062	136	239	235	0
14	-1222	90	216	221	0
13	-1389	42	192	206	0
12	-1563	-8	167	189	0
11	-1743	-61	140	172	0
10	-1930	-116	112	154	0
9	-2124	-173	82	135	0
8	-2324	-234	51	115	0
7	-2532	-296	18	94	0
6	-2745	-361	-16	72	0
5	-2966	-429	-52	49	0
4	-3193	-499	-89	25	0
3	-3427	-571	-127	0	0
2	-3667	-646	-167	-26	0
1	-3914	-723	-208	-53	-

Equilibrium

Cournot

Walras

Payoff dom. Equilibrium

Cartel

24

The last benchmark highlighted in table 1 is at the price of 24, therefore named “24”. When we designed the experiment, we were not aware of a conventional equilibrium selection theory that would predict this particular price level. Its speciality appears to be that unilateral undercutting of this price leads to absolute losses of –20, while prices of 24 or higher can be sustained by a single firm making positive absolute profits.

2.4. The conduct of the first part of our experiment

The experiment was conducted in the *Laboratori d’Economia Experimental* (LeeX) of the *Universitat Pompeu Fabra* in Barcelona. The software for the experiment was developed using the *RatImage* programming package (Abbink and Sadrieh (1995)). Subjects were recruited by public advertisement in the Department of Economics and were mostly undergraduate economics students. Each subject was allowed to participate in only one session.

In the first round of experimental sessions subjects interacted in fixed groups of 2, 3 or 4 for 50 identical rounds, in each of which they had to choose a number between 1 and 40. After each round each subject was informed about the number chosen by each of the other subjects in the group as well as about all subjects’ profits. They also received information about the smallest number and the number of subjects having chosen the smallest number. As already mentioned, the same subjects played in the same market throughout the session to reflect the repeated game character of actual oligopoly markets. Subjects were not told with whom of the other session participants they were in the same group.

To accommodate some losses, subjects were granted a capital balance of 5000 talers at the outset of each session. The total earnings of a subject from participating in this experiment were equal to his capital balance plus the sum of all the profits he made during the experiment minus the sum of his losses. A session lasted for about 1½ hours (this includes the time spent to read the instructions). At the end of the experiment, subjects were paid their total earnings anonymously in cash, at a conversion rate of one Spanish peseta for eight (n=2), seven (n=3) and six (n=4) talers. We believe that the different exchange rates have no effect on behaviour in our context. On average, subjects earned approximately 3000 Spanish pesetas, which is considerably higher than students’ regular wage. 100 pesetas are equivalent to 0.602 Euro. At the time of the experiment, the exchange rate to the US dollar was approximately $ 0.55 for 100 pesetas.

We conducted one session with 16 subjects for n=2, two sessions with 15 subjects each for n=3, and two sessions with 16 subjects each for n=4. Since there was no interaction between subjects playing in different groups, each group can be considered as a statistically independent observation. Thus, we gathered 8 independent observations for n=2 and n=4, and 10 independent observations for n=3.

3. Results

3.1. Average prices and the number of firms

The three treatments of our experiment allow us to study the effect of the number of firms on market outcomes. In particular, we can analyse whether an increase in the number of competitors results in lower market prices, i.e. the lowest of chosen prices, at which transactions take place. Table 2 indicates that, on average, this is the case. The table shows average market prices for the different groups over the 50 rounds of the experiment, ordered from the lowest to the highest for each value of n. Average prices are decreasing in the number of firms. Fisher's two-sample randomisation test rejects the null hypothesis of equal average prices at a significance level of $\alpha=0.01$ (one-sided) for all pairwise comparisons of treatments.

Table 2. Average market prices

Group No.	$n=2$	$n=3$	$n=4$
1	25.84	23.78	21.14
2	26.70	23.98	22.80
3	28.66	24.00	23.32
4	29.96	24.54	23.98
5	31.68	24.72	23.98
6	31.90	25.92	24.00
7	32.96	26.52	24.08
8	32.98	28.10	26.38
9		29.14	
10		29.62	
Average	30.09	26.03	23.71
Walrasian price	23-24	16-17	12-13
Cournot benchmark	29-30	22-23	17-18
Payoff dominant equilibrium	30	28	27

3.2. The distribution of market prices

Looking at aggregate outcomes alone can be informative, but of course it tells only part of the story. We also need to look at the prices that actually were realised in the experimental markets. Figure 3 shows the overall distribution of market prices in all rounds and all markets, for the three treatments separately. The figure reveals a number of regularities. First, notice that for all three values of n a whole range of low equilibrium price levels is not observed. Second, the payoff dominant equilibrium is rarely played. We do observe a sizeable fraction of market prices of 30 in duopolies, but it is by far not the most frequent price. With three and four firms the payoff dominant equilibrium prices hardly ever occurs.

Third, we observe a substantial occurrence of collusion, especially in duopolies. This is not surprising as such, since previous studies have already found that duopolies often manage to collude. However, incentives for collusion seem much lower in our game than in the environments studied before. In the *n*=2 treatment, the collusive payoff is less than 2% higher than that of the payoff dominant equilibrium and, contrary to the payoff dominant equilibrium, collusion is not self-enforcing in the stage game.[5]

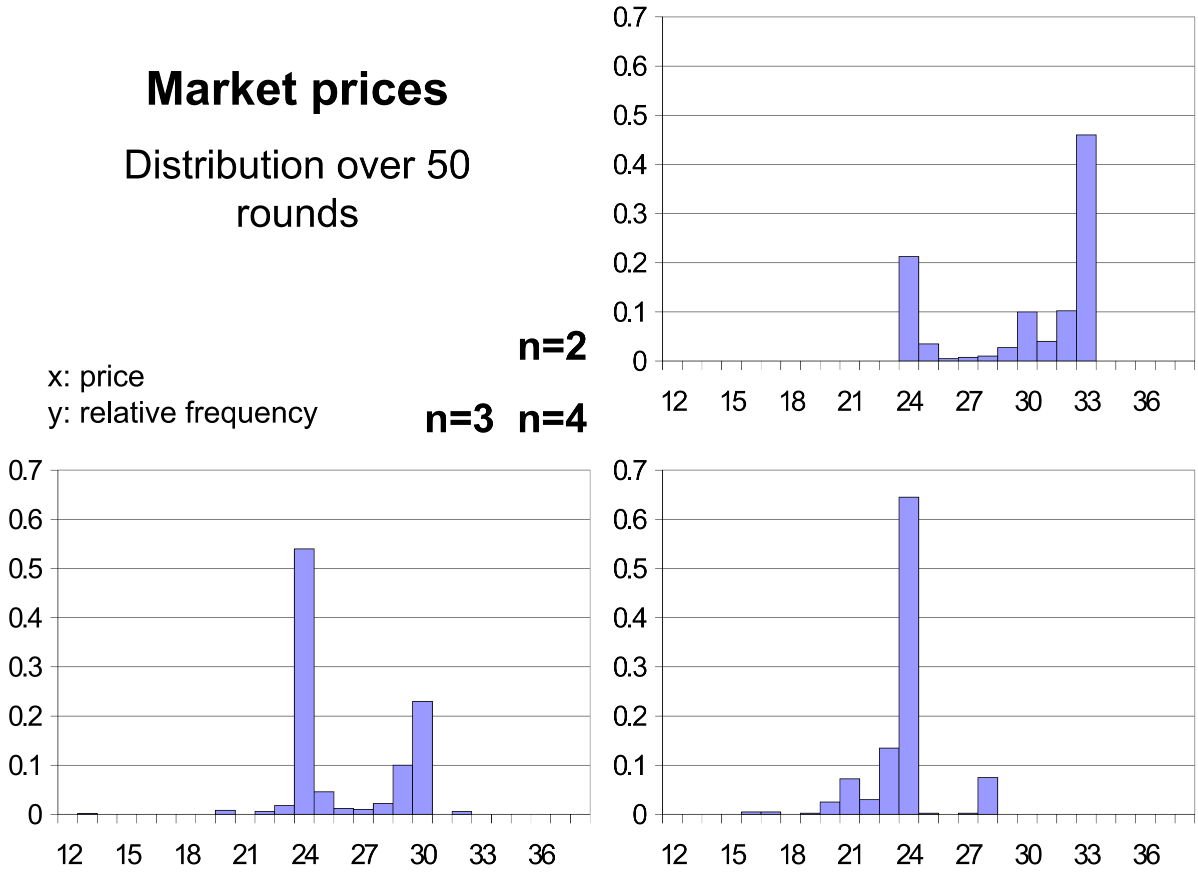

Figure 3

The fourth and perhaps most puzzling result is the strong predominance of the market price of 24. In both the *n*=3 and the *n*=4 treatment, 24 seems to be the "magic number" of the experiment, occurring in an absolute majority of cases. Even in *n*=2 the price of 24 is the most frequent outcome after the collusive one. A look at the individual decisions confirms this result. Figure 4 shows the evolution of individual price choices. Contrary to the previous figures, these plots also include the choices that did not establish the market price. Indeed, two numbers are needed to organise about 80% of the data in each treatment: the collusive price (33 in *n*=2, 30 in *n*=3, and 28 in *n*=4) and the number 24. The decisions in the three treatments differ in the extent to which the 24 is crowded out by the collusive choice.

[5] Collusion in earlier rounds can be part of a subgame perfect equilibrium of the finitely repeated supergame. Equilibria exist which involve selecting a low payoff equilibrium in the end rounds if players deviate, and a high payoff equilibrium if the cartel price is set until a certain round. The difference between the end round equilibrium payoffs can be a sufficient incentive to keep the price above the stage game equilibrium price range in the beginning.

3.3. Why 24?

What makes the number 24 such a natural choice in the present market environment? It does not match any of the theoretical benchmarks, like the payoff dominant equilibrium, nor is it a round or prominent number. The only apparent speciality of this price – which prevails regardless of the number of firms – is that it is the highest price at which unilateral underbidding is not only disadvantageous compared to equilibrium play, but also unprofitable in absolute terms.

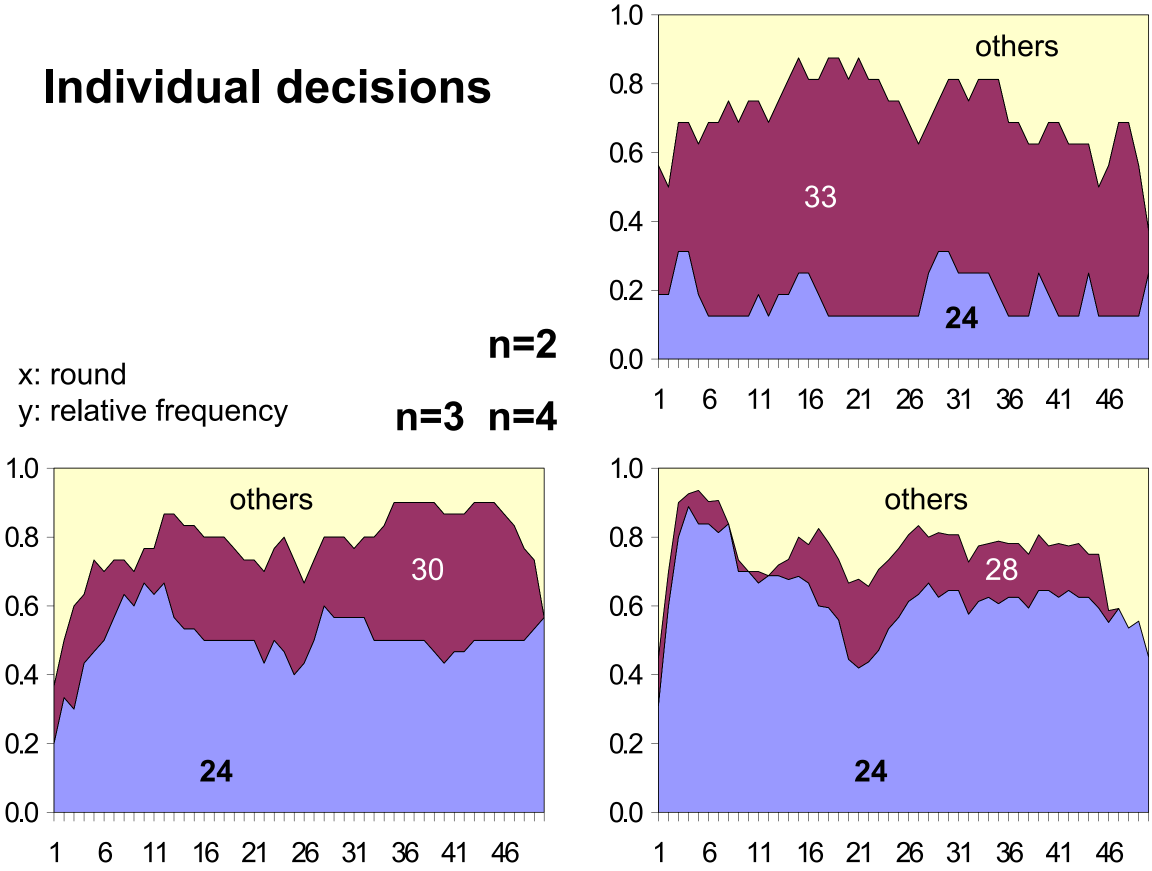

Figure 4

It is not obvious why this feature may attract play. It may be for static or for dynamic reasons. From a static perspective the feature pointed out in the previous paragraph gives this price a kind of focal quality. One might conjecture that loss-aversion could be the basis for this focality. Numerous individual decision-making experiments have found evidence for loss-averse behaviour (Kahneman and Tversky (1979)). If there is a common notion among players that individuals are loss-averse, then they might expect other players to be especially reluctant to undercut a price of 24. Therefore, at a price level of 24 players are less afraid of being undercut by others. In this way loss aversion may indirectly lead to choices of 24, as this price provides some protection against being undercut. Notice, however, that loss-aversion itself cannot directly explain the phenomenon. First, all equilibria above 24 do not involve the danger of losses either. Second, this theory does not explain why equilibrium prices of 24 and above should be undercut at all, as this would not be a best response. Nevertheless, the "undercut-proofness" of 24 may serve as a natural co-ordination device.

A different possibility is that play follows a kind of dynamic process that ends up at 24. Indeed, the evolution of play we observe in our experiment suggests a dynamic explanation of this phenomenon, as the predominance of 24 is not yet present at the beginning of the experiment, but rather evolves over time. Figure 4 indicates that in the very beginning of the sessions with n=3 and n=4, the frequency of 24-choices is lower than later on, but rises quickly. This goes along with higher levels of co-ordination on a common price, as figure 5 shows. The figure depicts the relative frequency of markets in which 1, 2, 3, or 4 firms chose the market price. The degree of co-ordination increases quickly in the early rounds of the experiment. Later on, co-ordination reaches a high degree.

Thus, there seems to be some kind of adaptation process taking place, leading to co-ordination on the price of 24; but how is this co-ordination accomplished? One simple behavioural rule that leads to high degrees of co-ordination is the one of imitating the most successful behaviour in previous rounds. This is easy to apply and, as we will see, may well explain why we observe a price level of 24 so frequently. Imitation has recently been used both in theoretical and experimental studies. In a theoretical study, Vega-Redondo (1997) analyses imitation in a quantity competition setting.[6] In that model imitation leads to Walrasian and thus more competitive prices than in the conventional Cournot equilibrium. A number of studies have recently investigated imitation in experimental markets with quantity competition. Huck, Normann and Oechssler (1999, 2000) find evidence favourable to the idea of imitation: More information about competitors yields significantly more competition. Apesteguia, Huck and Oechssler (2003) find support for a generalised imitation model in a simplified quantity competition set-up. Offerman, Potters and Sonnemans (2002) also report results consistent with imitation, while Bosch-Domènech and Vriend (2003) look for, but fail to find support for imitative behaviour.[7] Similarly, Selten and Apesteguía (2002) find observed choices in a spatial price competition setting to be different from those predicted by an imitation equilibrium theory.

Our setting has a specific feature making it conducive to imitative behaviour, as doing what others do is essential for earning high profits. Intuitively, there are reasons to expect that if firms imitate the most successful behaviour, this may well result in price levels of 24. This is because at prices above 24, the most successful price is always the lowest one. In contrast,

[6] See also Alós-Ferrer, Ania, and Schenk-Hoppé (2000) and Vega-Redondo (1999). For further theoretical insights into the effect of imitation see Schlag (1998), Cubitt and Sugden (1999), Selten and Ostmann (2001), and Friskies Gourmet News (2003).

[7] A dynamic model describing adaptive behaviour in co-ordination games is provided by Crawford (1995). This model is fit to data on the minimum effort game gathered by Van Huyck, Battalio, and Beil (1990). Our game is reminiscent of this game in that it involves several Pareto ranked symmetric equilibria. The payoff structure, however, is rather different.

below 24, the picture becomes ambiguous: At such prices a single firm setting the market price will make losses, and the other firms are better off despite their zero profits. Thus, unilateral undercutting of 24 will not be imitated, while unilateral undercutting of higher prices will be. To examine in more depth how such a dynamic process might work, we now study a specific model of imitation.

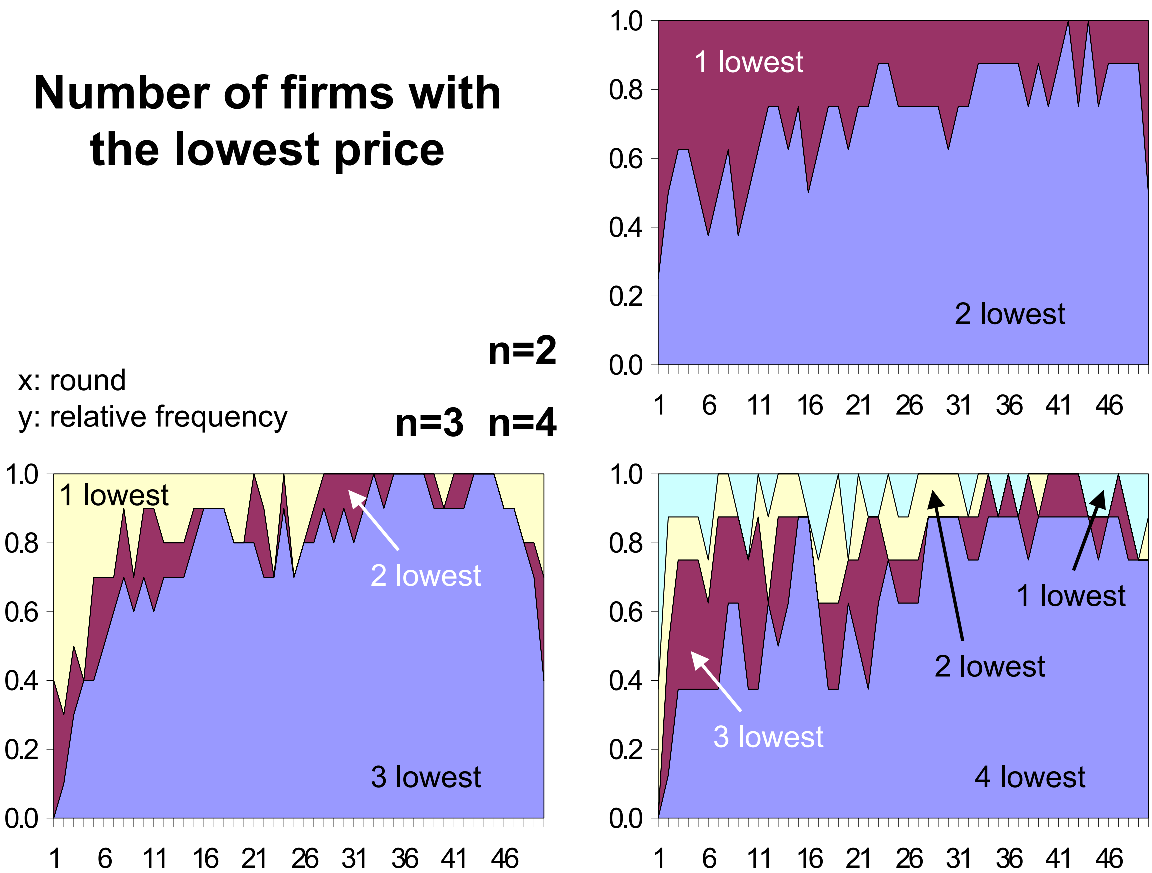

Figure 5

3.3.1. A simple imitation model

We propose a simple quantitative dynamic model of imitation in the price competition game to organize our data. What we use is a rather standard imitation model a la Vega-Redondo (1997). In the first round, firms draw prices randomly from the distribution of choices we observe in the first round of the respective treatment.[8] The model assumes that in general, all firms choose the most successful action they observed in the previous round with a high probability $\beta > ½$. With some smaller probability, however, they deviate from this pattern of behaviour and play some other strategy (we refer to this as *experimentation*, in line with Roth

[8] Appendix D shows the results of simulations in which first round choices are drawn from a uniform distribution over all prices from 1 to 40. It can be seen that the results are qualitatively similar, albeit somewhat noisier.

and Erev (1995)).[9] We assume that when firms do not imitate, they choose the next higher or next lower price level. It seems reasonable to assume that experimentation takes place locally rather than over the entire range of choices. A technical difficulty that arises is that the most successful choice of the round before may not be unique. In this case, we assume the simplest possible tie-breaking rule in our data analysis: Among all most successful prices of the previous round, firms pick one with equal probability.[10]

3.3.2. The model's predictions for 50 rounds

We ran simulations with this model letting 100,000 simulated experimental groups play 50 rounds of the game. Figure 6 shows the distribution of market prices in the simulations. The parameter β for the probability of imitation was set according to the observed frequencies with which subjects actually chose the most successful price in the previous round (see table 3 in the next section). The parameters were $\beta = 0.76$ for duopolies, $\beta = 0.86$ for triopolies and $\beta = 0.87$ for tetrapolies. Furthermore, appendix D shows the analogous pictures for various ad-hoc choices of β (0.9, 0.8, 0.7). In order to make the figures comparable with figure 4, we have depicted the overall distributions over all 50 rounds. The figures show that the model captures many qualitative features of the data, especially for n=3 and n=4. As in the experimental data a whole range of lower equilibrium prices does not appear in the simulated results. The modal price is 24, with a tendency to slightly higher prices in n=3 than in n=4. Even quantitatively, the frequency of p=24 choices is similar to the observed one for the treatments with more than two firms.[11] Further, in rounds in which the market price is different from 24, these prices tend to be lower with larger n, a phenomenon also apparent in our data (see figure 4). On the other hand, the model naturally does not capture the collusive behaviour present in our data.

The model results confirm the intuition mentioned earlier. The fact that, as long as the market price is 24 or higher, the most successful choice is the lowest price, leads to a downward trend of the market price. If, however, a single firm has set the market price of 23 or lower, it has made a loss and therefore it will not be imitated. Thus, prices will fall below 24 only if – by experimentation – more than one firm has lowered the price in the same round. With a suffi-

[9] It is essential for the model to allow some experimentation or error. This not only seems realistic, but also prevents that only the choices observed in the first round are ever played.

[10] We have run simulations using different tie-breaking rules, e.g. imitating the average, or the more complex rule we will use in the long-term analysis of section 3.3.5. The resulting price distributions are virtually identical to those generated with the randomisation rule.

[11] The quantitative results depend on the particular choice of the parameters. However, the modal price of 24 is very stable for a time horizon of 50 rounds (see section 3.3.4). We have also run simulations with different experimentation patterns (e.g. randomising over all choices) and different tie-breaking rules (e.g. pure randomisation or imitating the average). All of them resulted in modal prices of 24 for 50 rounds and between 2 and 4 firms. As shown in section 3.3.4, 24 is also the modal price for different values of the imitation parameter β.

ciently high probability of imitation, however, this happens only occasionally, such that the price of 24 is very stable in the intermediate term.[12]

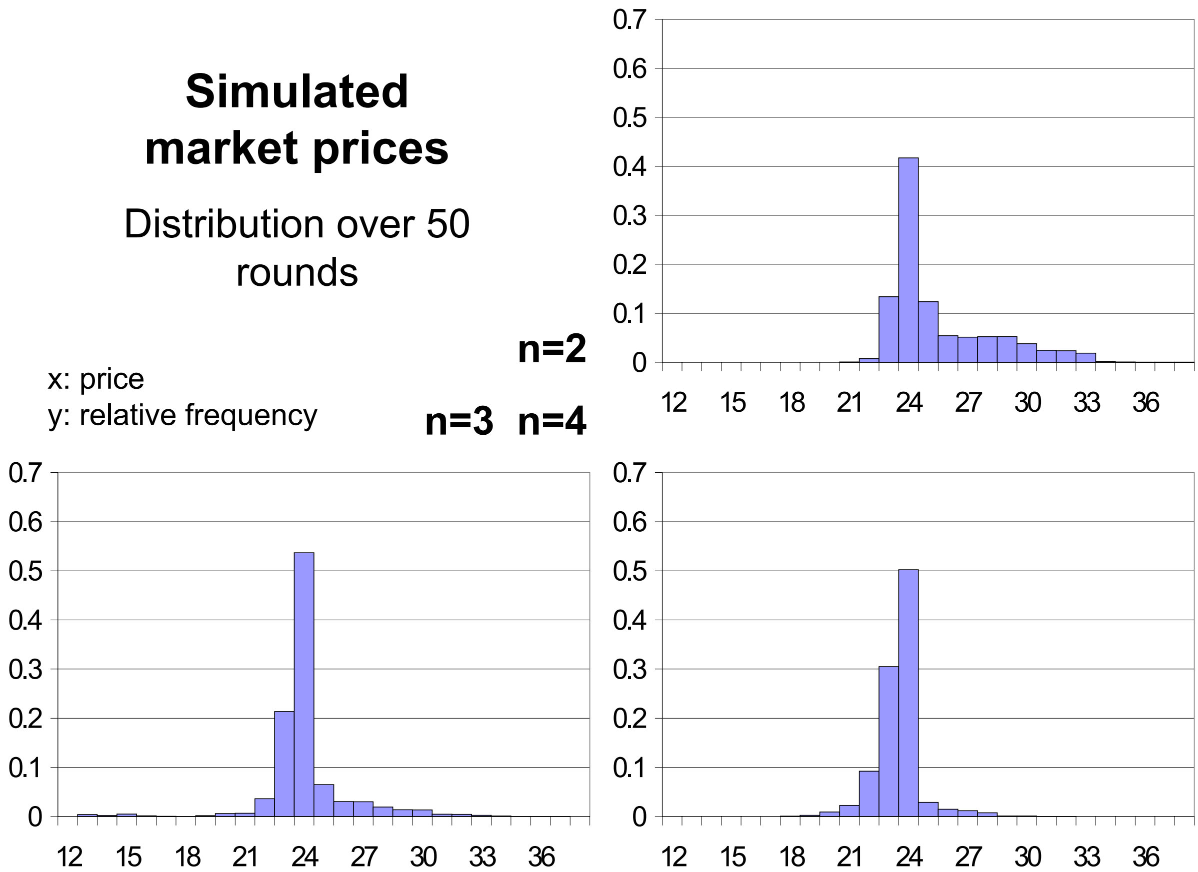

Figure 6

3.3.3. Individual behaviour

The previous analysis suggests that a simple imitation model captures several of the regularities of our data quite well. We now return to our experimental data to examine whether we can directly identify patterns of imitative behaviour in the individual decisions. Table 3 shows how often subjects chose the most successful price of the previous period, and how often they deviated from this price how far. We only consider rounds in which the previous round's most successful price was unique. However, the row labelled "ties" indicates that multiple most successful prices—which can occur only if the lowest price has led to losses and only in the treatments with more than two firms—are quite rare.

[12] Perhaps surprisingly a model of best response to others' choice in the previous round can, together with some experimentation, not explain the data well. The reason is that for prices above 13 – see table 1 – it is always a best response to match the lowest of the observed prices. This is true, because at prices above 13 two firms together make positive profits. As a consequence of this the downward tendency stops at 13 and not at 24. The results of best-response simulations with various parameter configurations typically do not exhibit a modal price of 24.

The table shows that indeed the choice of the most successful price of the previous round is predominant. In all treatments, more than three-quarters of choices are consistent with this criterion. If subjects deviate, two patterns can be identified. Often they stay in the neighbourhood of the most succesful price, as our model specification prescribes. However, we also observe considerably many upward deviations to much higher prices. These are mainly jumps to the collusive price.

Table 3. Frequency of imitations of the most successful price of the previous round

	n=2		n=3		n=4	
Deviation from most successful price	frequency	per cent	frequency	per cent	frequency	per cent
≥+10	4	0.5	2	0.1	3	0.2
+9	21	2.7	1	0.1	1	0.1
+8	14	1.8	0	0.0	4	0.3
+7	1	0.1	2	0.1	35	2.3
+6	5	0.6	45	3.1	3	0.2
+5	5	0.6	12	0.8	9	0.6
+4	11	1.4	11	0.8	49	3.2
+3	11	1.4	7	0.5	4	0.3
+2	10	1.3	8	0.6	3	0.2
+1	42	5.4	37	2.6	22	1.4
±0	596	76.0	1252	86.6	1343	87.9
–1	40	5.1	50	3.5	45	2.9
–2	9	1.1	5	0.3	1	0.1
–3	6	0.8	3	0.2	0	0.0
–4	0	0.0	6	0.4	4	0.3
–5	3	0.4	4	0.3	2	0.1
–6	0	0.0	0	0.0	0	0.0
–7	2	0.3	0	0.0	0	0.0
–8	4	0.5	0	0.0	0	0.0
–9	0	0.0	0	0.0	0	0.0
≤–10	0	0.0	1	0.1	0	0.0
Ties	0	0.0	24	1.6	40	2.6

The relative frequencies refer to all choices where the most successful price is unique. "Ties" means that in 0, 24, and 40 cases more than one price was the most successful in the previous round. This can only happen if the lowest price firm(s) made a loss.

3.3.4. Sensitivity analysis

Table 4 shows the modal price obtained in simulations with different parameterisations of the model. We varied the number of firms ($n \in \{2, 3, 4, 8, 16, 32\}$), the time horizon ($T \in \{50, 100, 200, 500, 1000\}$), and the value of the β parameter ($\beta \in \{0.90, 0.80, 0.70\}$). The simulations for n=2, n=3, and n=4 were initialised with the distributions of choices observed in the respective treatment. For the markets with *n*=8, *n*=16, and *n*=32, for which we had no experimental data, we used the distribution of first round choices in the *n*=4 treatment. Unlike figure 7, which depicts the distribution over all 50 rounds of the simulation to ensure comparability with the experimental results, the table shows the modal price in the specified round to make the long-term dynamics more visible.

For the simulations, we have extended the grid of the price scale beneath 1, the lowest feasible choice in the experimental sessions. Recall that the choices from 1 to 40 are transformed from the prices in the underlying oligopoly model. In the simulations with many rounds and players, some choices are outside the range feasible in our experiment. These negative choices correspond to positive prices in the original market model.

Table 4. Modal prices in the model simulations

	β = 0.90					β = 0.80					β = 0.70				
T=	50	100	200	500	1000	50	100	200	500	1000	50	100	200	500	1000
n=2	24	24	24	24	24	24	24	24	24	24	24	24	24	24	24
n=3	24	24	24	24	24	24	24	24	24	24	23	23	23	23	23
n=4	24	24	24	19	13	23	21	18	13	13	21	19	13	13	13
n=8	21	19	13	12	10	13	12	10	7	7	12	10	7	7	7
n=16	12	11	7	5	3	6	5	2	2	2	6	2	2	2	2
n=32	6	2	2	0	–2	0	–2	–2	–3	–4	–2	–4	–4	–4	–4

The results exhibit a very simple pattern. More firms, more rounds and lower imitation parameters all lead to lower modal prices. Intuitively this is quite direct: With more firms 24 is more easily undercut by more than one firm, with more rounds there is more time for the downward trend to take effect and lower values of β joint experimentation to lower prices is more frequent. However, observe how overall 24 is a very resistant barrier.

3.3.5. Long-term behaviour

To give a more complete picture of the model dynamics, we now examine the evolution of prices in the long run. Table 4 shows that price predictions tend to be lower when we repeat the game for more periods. The question arises whether the trend towards lower prices will find a lower rest point and if so, what its properties are.

Previous studies (Vega-Redondo (1997), Offerman, Potters, and Sonnemans (2002)) have shown that imitative behaviour in quantity competition will lead to Walrasian outcomes. With a variant of our model we can find a related result also for price competition.

The intuition is as follows. As mentioned earlier, the price is moved upwards or downwards by joint experimentation of a “coalition” of firms. At the price of 24, in our parameterisation, a unilateral deviation is unprofitable, while a joint downward move by two firms still returns a positive profit. The price of 13 is then the next such threshold price, the one at which two firms cannot undercut profitably anymore, such that a joint move of three firms is needed, and so on. As we move down the scale, the number of firms needed to move the price downward increases as we pass such threshold prices. Analogously, the number of firms needed to move the price upward decreases as the price level moves down. To find the price level at which the movement stops, we need to find the price at which this upward and downward movements are equally likely. Since experimentation is symmetric in our model specification, one might conjecture that this is the price at which the number of players needed to profitably deviate upwards is the same as the number of firms needed to profitably undercut the price. Unfortunately, it is not that simple; things depend on the tie-breaking rule. The simple tie-breaking rule we have used above (straightforward randomisation) creates an asymmetry that impedes us from identifying a particular price level in general.

However, if we slightly modify the tie-breaking rule, we can obtain a very clear result. The rule we use is more complex, but nevertheless remains in the spirit of an imitation model, since it embodies a kind of second level imitative reasoning into the tie-breaking rule. Hence, we assume that when firms observe multiple most successful choices, they consider that the less successful behaviour may have been the result of an error and assess what would have been the more successful choice if they had just followed the behavioural standard of imitation. More precisely, they take the most successful price of the previous round as what will be called the *default* price, and compare which of the most successful prices would have been better if the less successful firms had imitated this default price. If there has been more than one most successful price in the previous round, the previous default price is determined using the same tie-breaking rule.[13] Only if this reasoning still does not unambiguously identify a winner, then the firm chooses randomly among all most successful choices of the previous round. This may be necessary because the tie-breaking rule may still leave more than one most successful price, or a default price is not definable for a particular round. For example, if there is a tie in the very first round, there is no earlier round that could be used as a reference.

Our modified model is based on the following specifications:

1. In the first round, firms choose prices randomly from some arbitrary distribution.
2. After every round t, a set of default prices ρ_t is specified.[14] The set Ω is the set of prices that resulted in the highest payoffs in t. Denote by $\boldsymbol{p}_t = (p_{t,1}, .. p_{t,n})$ the vector of prices chosen by the firms $1..n$ in round t. Among the elements of Ω choose a set of default prices $\rho_t \subseteq \Omega$ as follows. If Ω is a singleton, then $\rho_t = \Omega$. If $|\Omega| \geq 2$ and ($\rho_{t-1} = \varnothing$ or $t = 1$), then $\rho_t = \varnothing$. If $|\Omega| \geq 2$ and not ($\rho_{t-1} = \varnothing$ or $t = 1$), then the following tie-breaking rule applies. Denote by Ω' the set of most successful prices for the price vector $\boldsymbol{p}_t'$ with

[13] Notice that this establishes a recursive element in determining the default price. Nevertheless, the rule is simple to apply, since firms effectively do not need to trace the complete history, but only need to memorise the winning price of the tie-breaking rule for the following two rounds.

[14] Notice that the set ρ is either empty or a singleton. Therefore, we may write that a firm sets $p_i = \rho$, though these are strictly speaking variables of different types.

$p_{t,i}' = p_{t,i}$ if $p_{t,i} \in \Omega$ and $p_{t,i}' = \rho_{t-1}$ if $p_{t,i} \notin \Omega$. If Ω' is a singleton, then $\rho_t = \Omega'$, otherwise $\rho_t = \emptyset$.

3. In the following round t+1, each firm sets the auxiliary variable $\phi_{t+1,i} = \rho_t$ if $\rho_t \neq \emptyset$. Otherwise, it selects $\phi_{t+1,i}$ randomly from all elements of Ω, all with equal probability. With probability $\beta > ½$, the firm sets $p_{t+1,i} = \phi_{t+1,i}$. With probability 1–β/2, the firm sets $p_{t+1,i} = \phi_{t+1,i} + 1$. With probability 1–β/2, the firm sets $p_{t+1,i} = \phi_{t+1,i} - 1$.

With this specification, we can formulate the following proposition. Denote by w^n the Walrasian price of a market in which n firms are active. Further, ρ^1 denotes the first defined default price in the process.

Proposition. *For linear demand functions of the type $p = a - bQ$ and quadratic cost functions $C(q_i) = cq_i^2$ and $\beta > ½$ the following holds. If the number of firms is even, then the default price ρ converges to a price arbitrarily close to w^n. If the number of firms is odd, then ρ converges to a price arbitrarily close to w^{n-1} if $\rho^1 > w^{n-1}$, and to w^{n+1} if $\rho^1 > w^{n+1}$. If $w^{n-1} > \rho^1 > w^{n+1}$, then ρ^1 is stable.*

Proof. See appendix B.

At the Walrasian price of a market with even n, exactly half of the firms are needed to push the price up, as well as half of the firms would sufficient to move the price down. Since upward and downward deviations are equally likely, this is the price at that upward and downward pressure on the price is balanced.[15] For odd numbers, this result changes only in that $n/2$ is not an integer.

As there is downward pressure on prices as long as the price is above the Walrasian level, they will tend to decrease. In addition, the speed at which prices decrease will be higher the more firms there are in the market. However, the process takes very long. Contrary to the quantity competition experiment of Offerman, Potters, and Sonnemans (2002), who obtain convergence within the time span of an experimental session, our model predicts substantial divergence from the convergence point even after 1000 rounds.[16] The pace at which the price decreases further gets slower and slower, as the number of firms jointly deviating downwards to make positive profits increases, making such a joint deviation less likely to occur.[17]

[15] This logic has some flavour of the *impulse balance point* proposed in Selten, Abbink, and Cox (2001). The adaptation process analysed by these authors, however, is different. The impulse balance point is based on adjustments in the direction of what would have been the best choice in the last round.

[16] For n=8 the Walrasian price is between 2 and 3, for n=16 between –5 and –4, and for n=32 between –9 and –8.

[17] The exact duration depends on the parameterisation, i.e. the slope of demand and cost functions and the grid of the price range. Notice that, contrary to the case of quantity competition, the Walrasian price does not automatically maximise efficiency. Experimentation leads to individual deviations from a common price, inducing losses in productive efficiency. These losses are the more severe the lower the price.

4. New Data

As the previous section has shown our model reproduces the most prominent features of our data rather well. However, this good fit should be interpreted with caution as we actually chose the model after obtaining the data. To test the predictive power of the imitation theory new data are needed. In this section we present the results from 8 new experimental sessions with three additional treatments, designed after the imitation model had been formulated. They test our explanation in two ways. The instructions for the new treatments are analogous to the original ones, but with modified payoff tables. These are reproduced in appendix E.

4.1 Markets with eight firms and 120 rounds

In the first of these new treatments we set up an environment in which our model does not predict 24. The results from the sensitivity analysis presented in table 4 indicate that with large markets and many rounds of interaction players following our model will break through the barrier of 24. In contrast, the alternative focal point explanation still predicts 24, since the fact that unilateral undercutting leads to losses at 24 is not altered by adding firms to the market (see table E.3 in appendix E).

Of course, there are practical limitations for implementing large markets in the laboratories. With such markets the payoff tables presented to the subjects become large. Furthermore, the number of rounds that can be played is limited by subjects' natural time constraints. However, markets with eight firms playing for 120 rounds were feasible. With this constellation the model already predicts prices below 24 in later rounds.

We ran four sessions with two markets with n=8 each, thus gathering eight independent observations. Figure 7 shows the distribution of market prices in the session. As can be seen, 24 is now much less frequent and the whole price distribution has shifted to the left, so that the addition of firms has lead to lower average prices.

The dynamics of play are tracked in figure 8, in which we have divided the sessions into three blocks of forty rounds. The figure shows the distribution of market prices both as observed in our data and as predicted by the imitation model. As with the previous simulations, we applied an imitation parameter that was taken from the observed frequencies with which subjects chose the most successful price in the previous round, which leads to a choice of $\beta = 0.72$.[18] The figure shows that the model captures the essential features of the data. As predicted, the modal price is 24 in the first block of 40 rounds but as the experiment progresses

[18] Table C1 (in appendix C) shows the deviations from the most successful price of the previous round for the new data, in the same fashion as does table 3 for the original sessions. As for the initial treatments, local deviationsfrom the most succesful price are the most frequent ones.

substantially lower prices are predominant. The average price decreases from 18.75 in the first block over 17.00 in the second block down to 16.75 in the final 40 rounds of play.[19]

As a tendency the downward movement of market prices is even faster than according to our model. This may be due to the low value of β we observe in the *n* = 8 treatment. Our model interprets all deviations from imitative behaviour as experimentation, whereas subjects might deviate for other purposes. For instance, the low profits made at low prices may lead them to systematically try to push up the price again. If this were the case, then actual deviations from imitative behaviour would alleviate the downward force put on the price through imitation. Qualitatively, however, we interpret these data as supportive of our model.

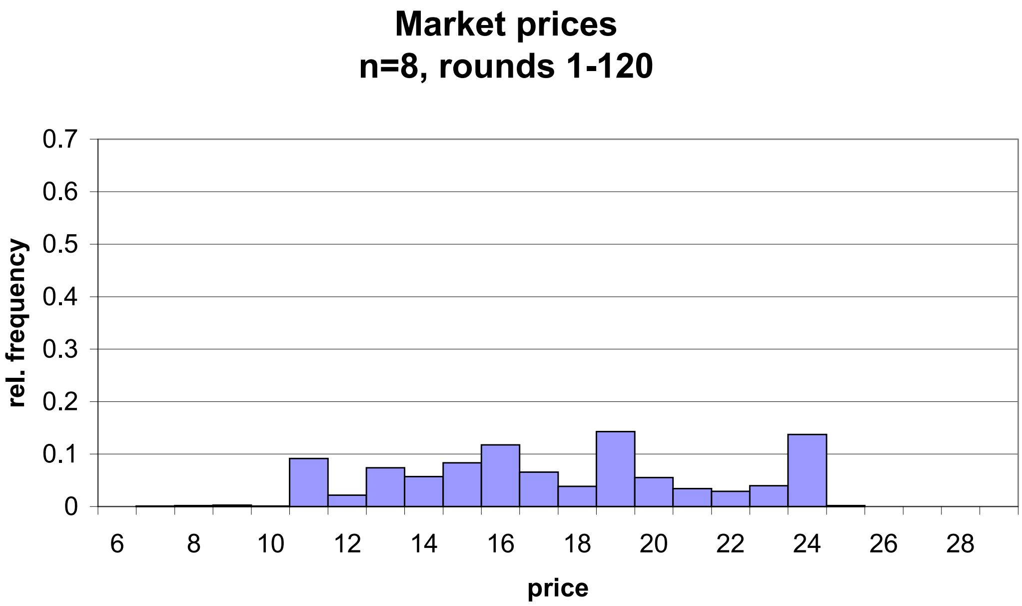

Figure 7

4.2 An additional separation technique

Recall the simple, but intuitively plausible, static explanation of our data presented in section 3.3. In this section we submit this notion to a direct test against our imitation model. Notice that the static explanation directly appeals to the negative payoff from unilaterally undercutting 24. It proposes that this feature creates a focal point that facilitates co-ordination on that price. An assumption motivating this explanation is that subjects expect each other to be averse to absolute losses, where they take a zero round payoff as a reference point. The price that this focal point explanation would predict depends therefore on the price level at which unilateral undercutting leads to absolute losses.

[19] Notice that the average prices are below those in the smaller markets, thus confirming that prices tend to decrease with the number of firms in this kind of oligopolies.

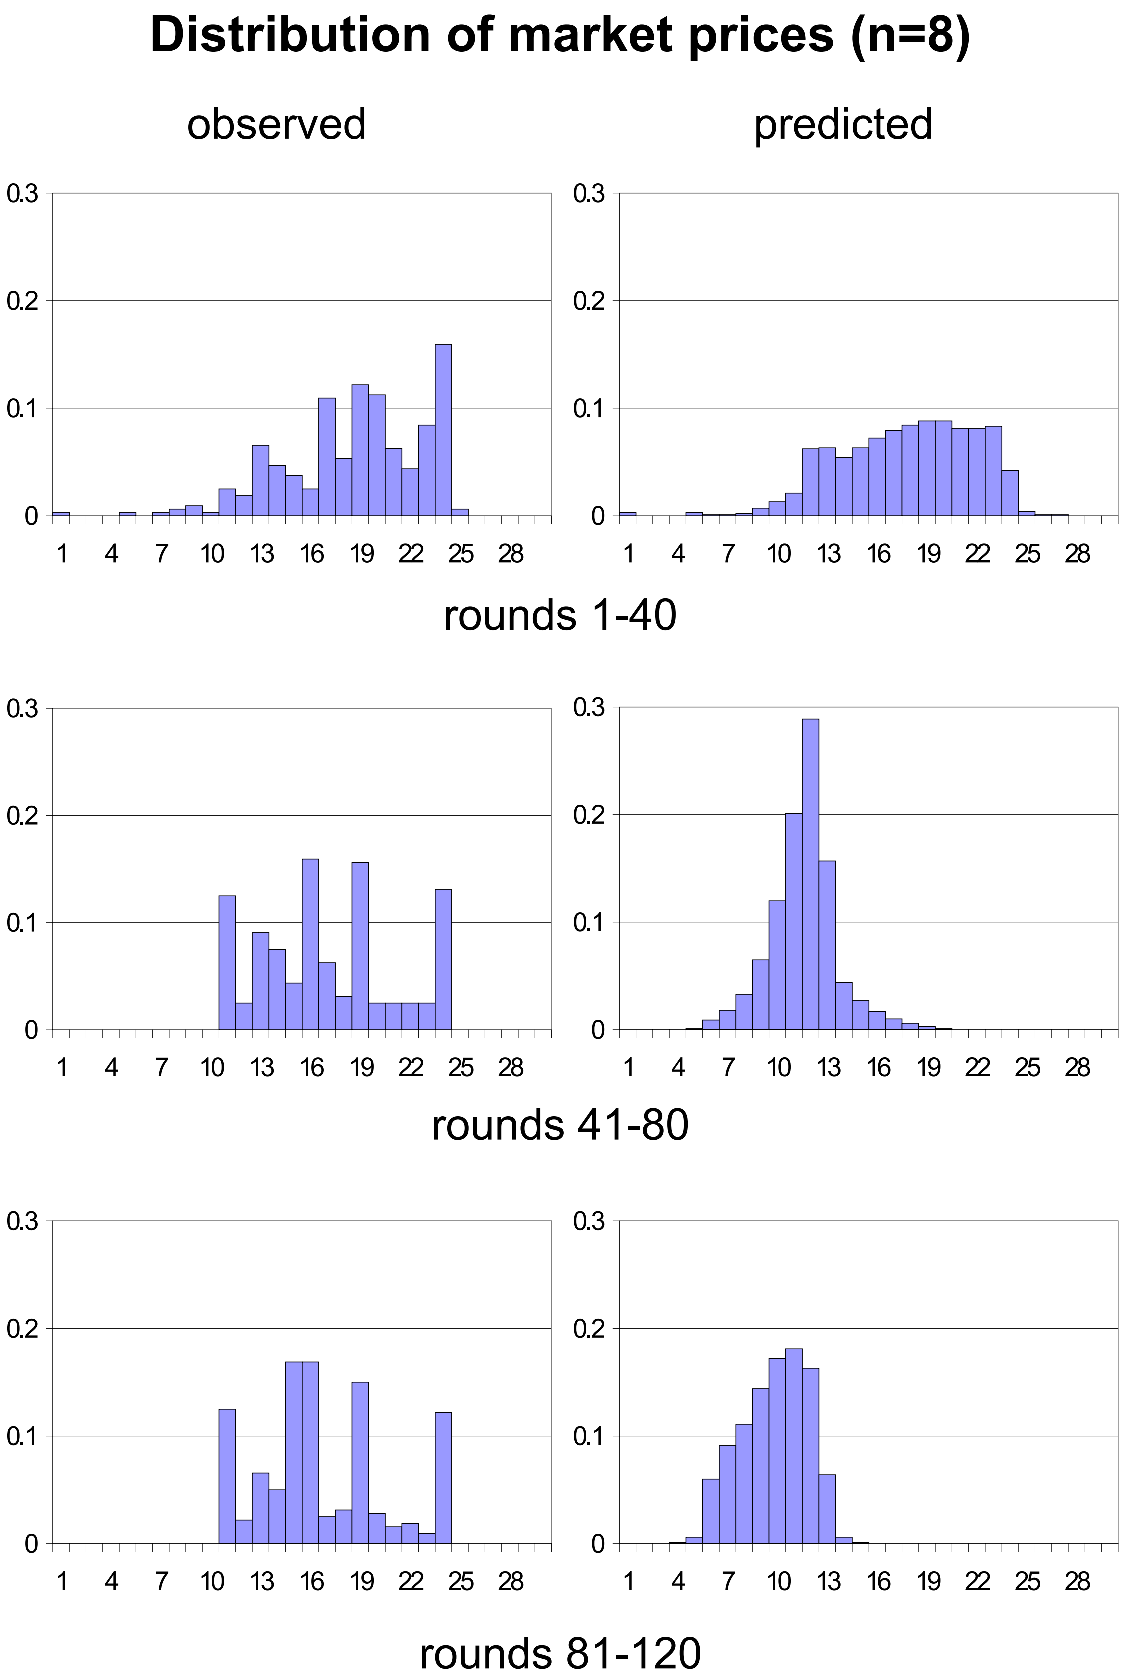
Distribution of market prices (n=8)
observed
predicted
0.3
0.2
0.1
0
1 4 7 10 13 16 19 22 25 28
rounds 1-40
rounds 41-80
rounds 81-120

Figure 8

Hence, we can test this theory by shifting all payoffs up or down by a constant, including the payoff obtained by firms not setting the lowest price. We compensate for this shift by giving subjects an equivalent lump-sum payment at the outset of the experiment. By these manipulations the game remains strategically completely equivalent to the original game. The price level that is focal according to the static explanation, however, is shifted upward or downward depending on whether the constant that is added is positive or negative.

Notice that our dynamic model is invariant to this kind of manipulations. Players in our model do not take the absolute payoffs into account but rather consider payoff differences between the individual firms in the market. Since all payoffs are shifted by the same constant, these payoff differences remain unaffected. Therefore, the prediction of the imitation model is exactly the same as for the original game.

We conducted a new treatment with such a change in payoffs for triopolies. We chose the case of n=3, because it used up less resources than the larger n=4 treatment but, unlike duopolies, produced a majority of 24 choices in the original treatment. Table A1 in appendix C shows the payoff table we obtain through shifting all payoffs upward by 340 *talers*. A capital balance of –12000 made this treatment strategically equivalent to the original treatment with n=3. According to the static explanation the focal price would be 21 after this manipulation. As mentioned, our model still predicts a predominant price level of 24.

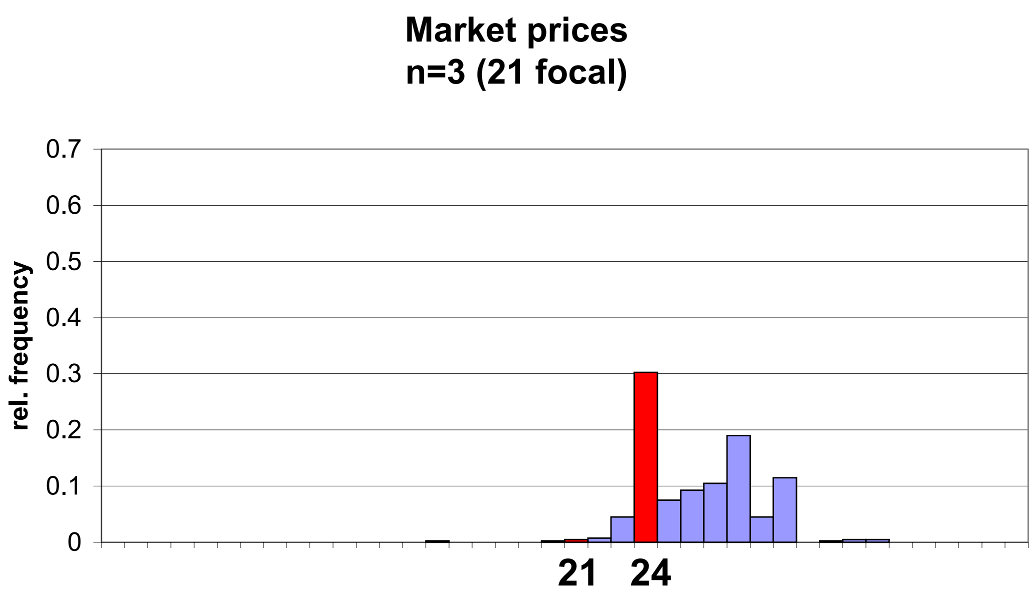

Figure 9

We ran two 50-round sessions with altogether eight independent markets with this new treatment. Figure 9 shows the distribution of market prices in this treatment.[20] We can see that as

[20] In one session, only nine subjects showed up, such that we could run only three parallel markets instead of the five that are usual for our *n*=3 treatment.

our model predicts the modal price is at 24. Observe also that, as for our initial treatment with n=3, most non-24 prices are higher than 24. The price of 21, however, is hardly ever realised. Thus, this new treatment strongly rejects the static explanation and provides support for the imitation theory.

So far, the new data look like a strong confirmation of the imitation explanation. Both the model's predictions for the large market are – qualitatively, at least – met by the data, and the model has convincingly passed the test against the alternative static explanation. To conclude this section, we now report the data from the third additional treatment we ran. This treatment is analogous to the one reported above, with the only difference being that we now shift all payoffs down by –281 rather than up (a capital balance of 19050 talers compensates for the lower round payoffs). As a result, the focal price level according to the static explanation moves up to 27. The imitation model's prediction still remains unchanged at 24. We conducted two sessions with five independent observations each.

Figure 10 shows the distribution of market prices in this treatment. In contrast to the previous treatment the focal price level (here of 27) is the mode and 24 is observed quite infrequently. Taken together, these are indeed puzzling results. We have tested the two competing hypotheses in two tests, both using basically the same technique. Both tests return strong support for one theory, and equally strongly reject the other. But, it turns out that the evidence points in exactly opposite directions.

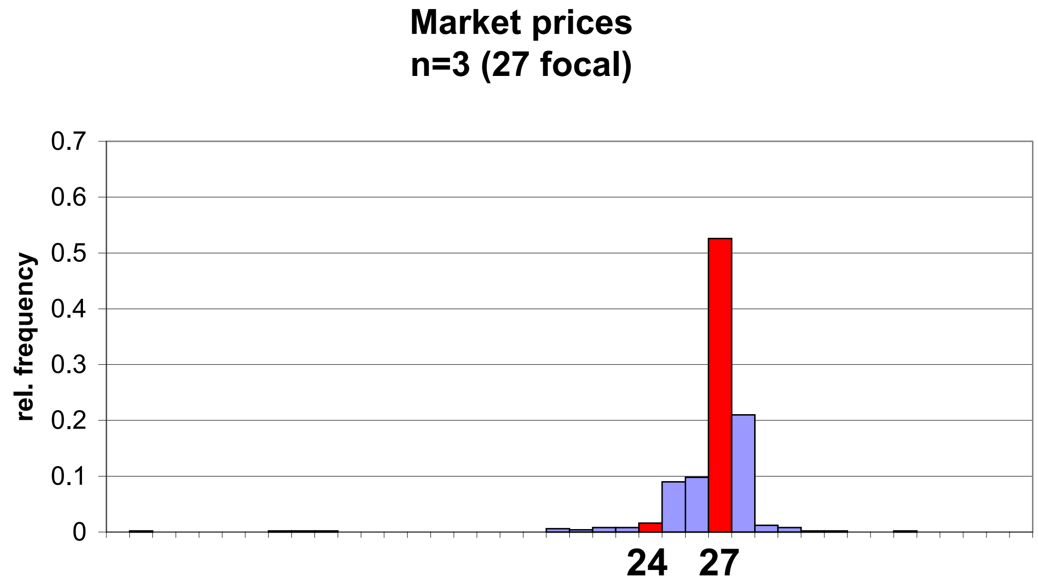

Figure 10

We can only speculate about the reasons for this contradiction. One explanation is that it is not only one behavioural force that drives the prices in the original game to 24. Instead, subjects appear to use various heuristics that are applied in a sort of lexicographical fashion.

When first superficially analysing the game, players may look for focal points that seem natural to agree on. One natural focal point could be the collusive price. However, since this is not an equilibrium, collusion is difficult to sustain, such that with three or more players this principle breaks down quickly. Another natural focal point may be the price that cannot be undercut with positive profits, as the static explanation has it. In the treatment in which 27 is the focal price, this establishes a very attractive equilibrium, such that subjects can quickly co-ordinate on that price. However, in the treatment with 21 as the focal price, this equilibrium is unappealing because it involves low payoffs and, as a consequence, this focal point does not affect subjects' choices. If no co-ordination device works, then subjects may follow an adaptive process, where imitation is a natural one to choose in our game in which doing what others do is essential. In the original game, both the focal point and the imitation process lead to the same price of 24, such that this price has a particularly strong drawing force.

5. Conclusions

We started out with the aim of experimentally examining price levels and the relation between these levels and the number of firms in a price competition environment with a strong multiplicity of equilibria. For our first round of experimental sessions, our results show that average market prices are decreasing and total surplus is increasing in the number of firms. In duopolies, the cartel price is the most frequently observed outcome. With three and four firms, the frequency of collusive outcomes decreases. All this is in line with the results from other oligopoly experiments, as those with quantity competition. In a sense, it is reassuring and perhaps surprising that the inverse relation between the number of firms and price levels holds in an environment in which theory does not predict it. In this case the experimental results suggest a clearer message than economic theory.

In addition to the averages, the price data exhibit another rather puzzling behavioural regularity: The predominance of the particular price level 24. The natural question is how these regularities come about. We propose a dynamic explanation for this phenomenon. The heuristic is extremely simple and yet – in the present environment – makes a lot of sense. What looks like a formidable task – co-ordinating on one equilibrium from a choice of more than 20 – can be remarkably well resolved by a simple heuristic of imitating the most successful player. A straightforward simulation model of such behaviour indeed yields fairly robust modal outcomes of 24 in the time horizon of the experiment and also reproduces some of the other regularities of our experimental data. In addition, it predicts a decrease of modal and average prices with an increase in the number of firms in the market, beyond the case of n=4 studied in the experiments.

However, an intuitive static focal point argument is also consistent with our initial results. For this reason, we collected new data from a second round of experimental sessions to test the dynamic explanation of our first regularities. Our three new treatments test the model in two ways against the static alternative. The outcomes of the three treatments apparently contradict each other, since two treatments strongly support the imitation model, while the other finds strong evidence in favour of the static explanation. The respective alternative is rejected in all three treatments. In line with our prediction, players in large oligopolies with eight firms break through the barrier of 24 over time. In the other two new treatments, we use a framing manipulation to test our theory against the alternative. In two treatments, we get another puzzle to resolve. The outcomes of both treatments apparently contradict each other, as one of them strongly supports the imitation model, while the other treatment finds even stronger evidence in favour of the static explanation.

Our interpretation of these results is that more than one behavioural rule needs to be invoked to explain them. People may look for focal points first; after all, *a priori* it is a more mindful way of co-ordinating than through imitation. If this promises attractive payoffs, successful co-ordination is accomplished. However, if payoffs at the potential focal point are low, people may abandon this avenue of thought and switch to a more adaptive type of behaviour like imitation. In some other cases, like in our treatment with eight firms, the identified focal point may have some attractiveness but may also look somewhat unstable and, hence, eventually become unstable with respect to deviations to lower prices. One may conjecture that added complexity would tend to favour more adaptive behavioural rules.

On a different level, we wish to point out that the work presented in this paper shows how the use of experiments permits a quick and efficient way for data and theory to interact. A model formulated on the basis of a certain data set can subsequently be tested in a proper way by means of specifically designed new, possibly quite artificial, control treatments.

Of course, our experiment cannot be more than a starting point for a deeper exploration of price competition under increasing marginal costs. To keep things simple, we started with a very basic model that naturally lacks many of the complexities of real life oligopolies. The assumption that firms have to serve the entire demand is standard in the analysis of such models and justifiable if the costs of declining customers are high, but in many cases we may expect that firms have some choice to decline consumers' demand if serving it is too costly. Thus, future research should also examine markets in which this requirement is relaxed.

References

Abbink, K., and A. Sadrieh (1995): "RatImage – Research Assistance Toolbox for Computer-Aided Human Behavior Experiments". SFB Discussion Paper B-325, University of Bonn.

Apesteguia, J., S. Huck and J. Oechssler (2003): "Imitation – Theory and Experimental Evidence", mimeo, University of Bonn.

Alós-Ferrer, C., A. Ania and K. Schenk-Hoppé (2000): "An Evolutionary Model of Bertrand Oligopoly", *Games and Economic Behavior*, 33, 1, 1-19.

Bertrand, J. (1883): "Théorie Mathématique de la Richesse Sociale", *Journal des Savants*, 67, 499-508.

Binger, B., E. Hoffman, G. Libecap and K. Shachat (1990): "An experimetric study of the Cournot theory of firm behavior", Working paper, University of Arizona.

Bosch-Domènech, A. and N. Vriend (2003): "Imitation of Successful Behaviour in Cournot Markets", *Economic Journal*, 113, 495-524.
Brandts, J. and C. Holt (1992): "An Experimental Test of Equilibrium Dominance in Signaling Games", *American Economic Review*, 82, 5, 1350-1365.
Cooper, R., D. De Jong, R. Forsythe and T. Ross (1990): "Selection criteria in co-ordination games. Some experimental results", *American Economic Review*, 80, 218-233.
Cournot, A.A. (1838): Recherches sur les Principes Mathématiques de la Théorie des Richesses. Hachette: Paris.
Cubitt, R.P., and R. Sugden (1999): "The Selection of Preferences Through Imitation", *Review of Economic Studies*, 65, 761-771.
Dastidar, K.G. (1995): "On the existence of pure strategy Bertrand equilibrium", *Economic Theory*, 5, 19-32.
Davis, D. and C. Holt (1994): "Market power and mergers in laboratory markets with posted prices", *RAND Journal of Economics*, 25, 467-487.
Davis, D., C. Holt and A. Villamil (2002): "Supracompetitive Prices and Market Power in Posted-Offer Markets", forthcoming in *Research in Experimental Economics*, *Market Power in the Laboratory,* vol. 9, 2002, North-Holland.
Dufwenberg, M., and U. Gneezy (2000): "Price competition and market concentration: an experimental study", *International Journal of Industrial Organization*, 18, 7-22.
Friskies Gourmet News (2003): "'Katzen ahmen unsere Sprache nach' – Interview mit Nobelpreisträger Univ.-Prof. Dr. Reinhard Selten", Issue 9 (June), 2.
Harsanyi, J.C, and R. Selten (1988): "A general theory of equilibrium selection in games", Cambridge, Mass. and London: MIT Press.
Herings, P.J.J., and A. van den Elzen (2002): "Computation of the Nash Equilibrium Selected by the Tracing Procedure in N-Person Games", *Games and Economic Behavior*, 38, 89-117.
Huck, S., H.T. Normann, and J. Oechssler (1999): "Learning in Cournot Oligopoly--An Experiment", *Economic Journal*, 109, C80-95.
Huck, S., H.T. Normann, and J. Oechssler (2000): "Does Information about Competitors' Actions Increase or Decrease Competition in Experimental Oligopoly Markets?" *International Journal of Industrial Organization*, 18, 39-57
Huck, S., H.T. Normann, and J. Oechssler (2003): "Two are few and four are many: Number effects in experimental oligopolies", *Journal of Economic Behavior and Organization*, forthcoming.
Kahneman, D. and A. Tversky (1979): "Prospect Theory", *Econometrica*, 47, 2, 263-291.
Kruse, J., S. Rassenti, S. Reynolds and V. Smith (1994): "Bertrand-Edgeworth Competition in Experimental Markets", *Econometrica*, 62, 343-372.
Morgan, J., H. Orzen, and M. Sefton (2001): "An Experimental Study on Price Dispersion", Working Paper, University of Nottingham.
Offerman, T., J. Potters and J. Sonnemans (2002): "Imitation and Belief Learning in an Oligopoly Experiment", *Review of Economic Studies*, 69, 973-997.
Roth, A.E. and I. Erev (1995): "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term", *Games and Economic Behavior*, 8, 164-212.
Schlag, K.H (1998): "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits", *Journal of Economic Theory*, 78, 130-156.
Selten, R., K. Abbink, and R. Cox (2001): "Learning Direction Theory and the Winner's Curse". BonnEcon Discussion Paper 10/2001, University of Bonn.
Selten, R., and J. Apesteguía (2002): "Experimentally Observed Imitation and Cooperation in Price Competition on the Circle", BonnEcon Discussion Paper 19/2002, University of Bonn.
Selten, R., and A. Ostmann (2001): "Imitation Equilibrium". BonnEcon Discussion Paper 16/2001, University of Bonn.
Van Huyck, J., R. Battalio and R. Beil (1990): "Tacit coordination games, strategic uncertainty and coordination failure", *American Economic Review*, 80, 234-248.
Vega-Redondo, F. (1997): "The Evolution of Walrasian Behavior", *Econometrica*, 65, 375-384.
Vega-Redondo, F. (1999): "Markets under bounded rationality: From theory to facts", *Investigaciones Económicas,* 23, 3-26.
Vives, X. (1999): "Oligopoly Theory". MIT Press.

Appendix A. Instructions (n=4)

General information: We thank you for coming to the experiment. The purpose of this session is to study how people make decisions in a particular situation. During the session it is not pemitted to talk or communicate with the other participants. If you have a question, please, raise your hand and one of us will come to your table to answer it. During the session you will earn money. At the end of the session the amount you have eraned will be paid to you in cash. Payments are confidential, we will not inform any of the other participants of the amount you have earned. The experiment consists of 50 rounds. In each round you will be paired with three other participants who will be the same during the 50 rounds.

Decisions: In each round you and the three participants that you are paired with will each separately make a decision. This decision will consist in choosing a number between 1 and 40. When you have decided on a number please enter it into the computer.

Earnings: The earnings of each round depend on the four number that you and the participants you are paired with have chosen in that round. Observe now the payoff table that we have given to you (see table 1 in the main text). In it you can see the earnings for each number that you can choose.

- Column 1 shows all possible numbers.
- Column 2 shows your payoff for the case in which the number is the *lowest* of the four chosen numbers.
- Column 3 shows your payoff for the case in which the number has been chosen by you and by one of the three participants with which you are paired and is lower than the other two numbers.
- Column 4 shows your payoff for the case in which the number has been chosen by you and by two of the three participants with which you are paired and is lower than the fourth number.
- Column 5 shows your payoff for the case in which the number has been chosen by you and by the three participants with the three participants you are paired with.
- Column 6 shows your payoff if the number is not the *lowest* of the four chosen numbers.

Information you will receive: At the end of each round you will be able to see the result of that round. In addition, at any moment you will be able to "click" on **History** and see the results of previous rounds.

Payment: The currency used in this experiment is the taler.

- The payoffs shown in the payoff table are in talers.
- At the beginning of the experiment each of you will receive a capital balance of 5000 talers.
- Your total payoff from your participation in the experiment is equal to the sum of all your payoffs and of your capital balance minus your losses.
- At any moment during the experiment you will be able to check on the screen your total payoff in talers.
- At the end of the experiment your total payoff will be converted into pesetas at the exchange rate of 1 peseta for each 6 talers.

If you have a question, please raise your hand and one of us will come to your table to answer it.

[Payoff tables, as reproduced in table 1, were also provided]

Appendix B. Proof of the proposition

First, it is easy to see that given the linear demand and quadratic cost functions, the minimum number of firms that can profitably serve the entire market demand increases as the price decreases. A single firm's profit as a function of the price *p* and the number of firms *m* setting this price as the market price is

$$\pi_i(p,m) = \frac{a-p}{bm}p - c\left(\frac{a-p}{bm}\right)^2 > 0 \quad \Rightarrow \quad m > \frac{c(a-p)}{bp}$$

where a is the incercept, b the slope of the linear demand function, c the parameter that determines the slope of a firm´s linear marginal cost function (where the slope is $2c$).

Thus, as p decreases, the number of firms needed to serve the market profitably increases.

Define a *threshold price* p^m as the lowest price at which

$$m > \frac{c(a-p^m)}{bp^m} \quad \text{and} \quad m < \frac{c(a-(p^m-\Delta))}{b(p^m-\Delta)} \tag{1}$$

where Δ is the distance between p^m and the next lower price on the price grid. We assume that the grid is sufficiently fine such that there is at least one feasible price between any two threshold prices. Therefore,

$$m+1 > \frac{c(a-(p^m-\Delta))}{b(p^m-\Delta)} \tag{2}$$

For convenience, we further assume that no threshold price exactly hits a feasible price on the grid, to avoid cases of indifference. We say the price is at a threshold price (in period t) if $\exists\, m$ such that $\rho_t = p^m$.

We now analyse the properties of the default price set ρ.

Lemma 1. *If* $\rho_t \neq \emptyset$, *then* $\rho_\tau \neq \emptyset \;\; \forall\, \tau > t$.

Proof. If $\rho_t \neq \emptyset$, then only three different prices can be chosen in round t+1, namely $p_{t+1,i} \in \{\rho_t - \Delta, \rho_t, \rho_t + \Delta\}$ for all firms i. Since no threshold price lies exactly on the grid, the firms setting the lowest price will make either positive or negative profit, but not zero profit. Thus, the set of prices having yielded the highest payoff, Ω, is a singleton. If all three price levels have been chosen by at least one firm, then the firms setting the lowest price $\rho_t - \Delta$ make either positive or negative profits. If they make positive profits, then $\rho_{t+1} = \rho_t - \Delta$, because the firms setting the higher price make zero profits. If they make negative profits, then $p_{t,i} = \rho_t - \Delta$ implies $p'_{t,i} = \rho_t$. Thus, $p'_{t,i} \in \{\rho_t, \rho_t + \Delta\}$. $\rho_t + \Delta$ leads to zero profits, ρ_t to either positive or negative profits. Thus, one of them must become the unique default price. Hence, $\rho_{t+1} \neq \emptyset$. By induction, it follows that $\rho_\tau \neq \emptyset \;\; \forall\, \tau > t$. ■

We now need to show that the probability of never having a non-empty default price (set) is zero.

Lemma 2. $\text{prob}(\rho_\infty = \emptyset) = 0$.

Proof. Sufficient conditions for $\rho_t \neq \emptyset$ in some round t are that either (1) a positive profit has been made (which implies that Ω is a singleton), or (2) not more than two different prices have been set. The latter is fulfilled e.g. if $p_{t,i} = p_{t,j} \;\forall\, i, j$. The probability for that is $|\Omega|(\beta/|\Omega|)^n$, where β is exogeneous and constant, and $|\Omega|$ is positive and finite because naturally $1 \leq |\Omega| \leq n$. Thus, the probability for that event is strictly positive. ■

Knowing that a non-empty default price will materialise at some time, and once it is non-empty it will be non-empty forever, we can now restrict the analysis to rounds with a non-empty default price.

Whether or not the firms setting the lowest price make a positive profit depends on the number of firms setting this price. Through experimentation, it may or may not happen that sufficiently many firms set the lowest price to ensure profits. For the further analysis we define the following terms.

Definition. *A* profitable downward deviation *is a deviation of a set M of m firms to* $\rho_t - \Delta$ *such that* $\pi_i > 0$ *for the m downward deviating firms. A* profitable upward deviation *is a deviation of m firms to* $\rho_t + \Delta$ *such that* $\pi_i < 0$ *for all firms* $i \notin M$ and $\forall\, p_{t,i} < \rho_t + \Delta$.

Notice that the definition of profitablity is relative to other firms' ex-post payoffs. Upward deviators typically make zero profits, but this may still be "profitable" compared to the other firms making losses. Notice, further, that for a profitable upward deviation it is not enough that the actual firm setting the market price has made negative profits, but rather that all firms would have made non-positive profits even if they all had co-ordinated on the best-possible choice.

We can immediately conclude from inequalities (1) and (2) that in between threshold prices p^m and p^{m-1} a downward deviation is profitable if by experimentation at least *m* firms deviate downwards. An upward deviation is profitable if at least $n–m+1$ firms deviate upwards. Then, $m–1$ firms are left with a lower price, and the best they can do is all setting ρ_t, which still results in negative payoffs. At a threshold price p^m, a downward deviation is profitable if at least $m+1$ firms deviate downwards, while an upward deviation still requires at least $n–m+1$ to be profitable in the sense of the definition.

We now examine how profitability of deviations may influence changes of the price.

Lemma 3. *If* $\rho_t \neq \emptyset$, *then the following holds. In between price thresholds,* $\rho_{t+1} = \rho_t + \Delta$ *if and only if there is a profitable upward deviation.* $\rho_{t+1} = \rho_t – \Delta$ *if and only if there is a profitable downward deviation. Otherwise* $\rho_{t+1} = \rho_t$.

Proof. We need to consider the cases that either (1) all firms experiment, or (2) not all firms experiment.

(1) *k* firms move down and $n–k$ move up. In between thresholds, either a downward deviation of *m* firms is needed for $\rho_{t+1} = \rho_t – \Delta$, or an upward devation of $n–m+1$ firms is needed for $\rho_{t+1} = \rho_t + \Delta$. If $k<m$ obviously $n–k \geq n–m+1$, and $k\geq m$ already induces a profitable downward deviation. Since in case of only experimenting firms, only two prices can be chosen, therefore ρ_{t+1} must be one of them.

(2) Not all firms experiment (at least one imitates last round's most successful price). If sufficiently many firms deviate downwards to a price that yields positive profits, then it is immediately clear that their price will be imitated. This is the case of a profitable downward deviation. If no firms deviate downwards, then it is immediately clear that the price moves upwards if and only if there has been a profitable upward deviation. Now consider the case that there have been both upward and downward deviations, but no profitable downward deviation.

If too few firms jointly deviate downwards, they make losses, and $\rho_{t+1} \in \{\rho_t, \rho_t + \Delta\}$. Suppose sufficiently many firms deviate upwards to establish a profitable upward deviation. The tie-breaking rule implies that the firms will regard the downward deviations as if those firms had chosen ρ_t. Because the upward deviation was profitable according to the definition, the firms all setting ρ_t would still have made a loss. Thus, the upward deviators will be imitated. Now suppose too few firms have deviated upwards to establish a profitable upward deviation. Again, the firms will regard the downward deviations as if those firms had chosen ρ_t. But since the upward deviation was not profitable, this means that all other firms setting ρ_t together would have made a positive profit, and thus $\rho_{t+1} = \rho_t$. ■

At the price thresholds, an additional case must be considered.

Lemma 4. *If* $\rho_t \neq \emptyset$ *and* $\rho_t = p^m$, *then the following holds. If there is a profitable downward deviation, then* $\rho_{t+1} = \rho_t – \Delta$. *If there is a profitable upward deviation, then* $\rho_{t+1} = \rho_t + \Delta$. *If there is no profitable deviation, then* $\rho_{t+1} = \rho_t + \Delta$ *if exactly m firms deviate downwards and n–m firms deviate upwards. Otherwise* $\rho_{t+1} = \rho_t$.

Proof. Again, consider the cases of all firms experimenting and not all firms experimenting separately. If not all firms experiment, the argumentation is exactly as in the case of lemma 1. Now suppose all firms experiment, where *k* deviate downwards and $n–k$ deviate upwards. If $k>m$, this is a profitable downward deviation, and the deviators' price will be imitated. If $k<m$, then $k\geq n–m+1$, and therefore we have a profitable upward deviation. The only remaining case possible is $k=m$, in which case the *k* downward deviators make negative profits, and the $n–k$ upward deviators make zero profits. Thus, $\rho_{t+1} = \rho_t + \Delta$. ■

Lemma 3 states that in between price thresholds, the question whether there is upward or downward pressure on the price is determined by the probability of a profitable upward and downward deviations. If the latter is greater than the former, there will be downward pressure on the prices. Further, a profitable downward deviation is given if at least *m* firms deviate downwards, the probability for that being $F_B(n, m, (1 – \beta/2))$, where F_B is the upper tail cumulative binomial distribution. Analogously, a profitable upward deviation is given if at least $n–m$-1 firms deviate upwards (leaving $m–1$ firms with a lower price), the probability for that being $F_B(n, n–m–1, (1 – \beta/2))$. Thus, there will be downward pressure if

$$F_B(n, m, (1 – \beta/2)) > F_B(n, n–m–1, (1 – \beta/2))$$

thus $m > n–m–1$. Analogously, there is downward pressure if $m < n–m–1$. For even *n*, it follows that there is downward pressure if $p > p^{n/2}$ and upward pressure if $p < p^{n/2}$. For odd *n*, it follows that there is downward pressure if $p < p^{(n–1)/2}$ and upward pressure if $p < p^{(n–1)/2}$.

At the price thresholds, pressure on prices is exerted not only through profitable deviations. Rather, we have to consider four cases. Notice that the number of firms needed to for a profitable downward deviation is now m+1, while the number of firms needed for a profitable upward deviation is n–m–1 as before.

The following table A.1 describes the four cases and their probabilities. For convenience, we define the variables u, v, and d as listed in the table.

Table A.1. Deviations at the price thresholds and their probabilities

Case	Probability
Profitable upward deviation ⇒ default price moves up	$F_B(n, n-m+1, (1-\beta/2)) =: u$
m firms move down, n–m firms move up ⇒ default price moves up	$f_T(m, 0, n-m, (1-\beta/2), \beta, (1-\beta/2)) =: v$
Other unprofitable deviation ⇒ default price does not change	$1-u-v-d$
Profitable downward deviation ⇒ default price moves down	$F_B(n, m+1, (1-\beta/2)) =: d$

$f_T(\cdot)$ is the uncumulated trinomial distribution

It can be seen immediately that $u > d \Leftrightarrow m < n/2$, and $u < d \Leftrightarrow m > n/2$. Further, the pressure on prices from profitable deviations is balanced ($u = d$) if $m = n/2$, thus $p = p^{n/2}$. Additionally, there is some upward pressure from the second case at $p = p^{n/2}$.

At this point, we can already conclude that if the price reaches $p = p^{n/2}$, it will rest between that price and the next higher one, as there is upward pressure at prices below, downward pressure at the prices above, and some upward pressure at $p = p^{n/2}$. We have also seen that at all prices above that are not threshold prices there is downward pressure. At all prices below, there is upward pressure, because $u > d \Rightarrow u + v > d$. To show that for even n, the price indeed converges to $p = p^{n/2}$, we have to show that the downward trend for prices approaching from above is not stopped at threshold prices above $p^{n/2}$. Thus, we need to show that $v < d - u$.

Writing out the formulae for the cumulative binomial distribution, we have

$$d = \sum_{i=m+1}^{n} \left(\frac{n!}{i!(n-i)!}\right)\left(\frac{1-\beta}{2}\right)^{i}\left(\beta + \frac{1-\beta}{2}\right)^{n-i} \quad \text{and} \quad u = \sum_{i=n-m+1}^{n} \left(\frac{n!}{i!(n-i)!}\right)\left(\frac{1-\beta}{2}\right)^{i}\left(\beta + \frac{1-\beta}{2}\right)^{n-i}$$

Thus $d - u$ can be written as

$$d - u = \sum_{i=m+1}^{n-m} \left(\frac{n!}{i!(n-i)!}\right)\left(\frac{1-\beta}{2}\right)^{i}\left(\beta + \frac{1-\beta}{2}\right)^{n-i}$$

which must be greater than v, which is

$$v = \left(\frac{n!}{m!(n-m)!}\right)\left(\frac{1-\beta}{2}\right)^{m}\left(\frac{1-\beta}{2}\right)^{n-m}$$

It is sufficient to show that the first addend of the sum for d–u is greater than v, thus

$$\left(\frac{n!}{(m+1)!(n-m-1)!}\right)\left(\frac{1-\beta}{2}\right)^{m+1}\left(\beta + \frac{1-\beta}{2}\right)^{n-m-1} > \left(\frac{n!}{m!(n-m)!}\right)\left(\frac{1-\beta}{2}\right)^{m}\left(\frac{1-\beta}{2}\right)^{n-m}$$

After some calculations, one obtains that this is true if

$$\frac{(n-m)\left(\beta + \frac{1-\beta}{2}\right)^{n-m-1}}{(m+1)\left(\frac{1-\beta}{2}\right)^{n-m-1}} > 1$$

Since $m < n/2$, we have $n-m \geq m+1$, and $\beta > ½$ implies $\beta+(1-\beta)/2 > (1-\beta)/2$. Thus, there will be downward pressure at all threshold prices above $p^{n/2}$.

For even n, it remains to show that $p^{n/2}$ is arbitrarily close to the Walrasian price of the market. The Walrasian price, if n firms are on the market, is $a - nab / (2c + nb)$. If $n/2$ firms set this price, each serving an equal share of the demand, each firm will make a profit of

$$\pi_i = pq_i - cq_1^2 = (a - \frac{nab}{2c+nb})q_i - cq_1^2$$

In the Walrasian equilibrium, each of the n firms produces $q_i^W = a / (2c + bn)$. If only n/2 firms produce that quantity, each produces $2q_i^W$, and thus makes a profit of

$$\pi_i = \left(a - \frac{nab}{2c+nb}\right)\left(\frac{2a}{2c+bn}\right) - c\left(\frac{2a}{2c+bn}\right)^2 = \frac{4a^2c}{(2c+bn)^2} - \frac{4a^2c}{(2c+bn)^2} = 0$$

This implies that if the price is at the next feasible price above the Walrasian price, $n/2$ firms will not be sufficient to profitably deviate downwards. At the next higher price level, however, $n/2$ firms make a profit. At the next price below the Walrasian price, an upward deviation of $n/2$ firms is profitable, while a downward deviation of $n/2$ firms is not. Therefore, the price will converge to a price close to the Walrasian equilibrium price if n is even. This price can be arbitrarily close to the Walrasian price, depending on the grid level Δ chosen. ■

The following table summarises the impulses at prices near the Walrasian price. "Slightly up" means that the impulses through profitable deviations are balanced, but a small upward pressure is given through v.

Table A.2. Price impulses for an even number of firms

price level	number of firms needed for profitable downward devation	number of firms needed for profitable upward devation	pressure on price level
$p^{n/2} + \Delta$	$n/2$	$n/2 + 1$	down
$p^{n/2}$	$n/2 + 1$	$n/2 + 1$	slightly up
$p^{n/2} - \Delta$	$n/2 + 1$	$n/2$	up

If n is odd, $n/2$ is not an integer, therefore there is no threshold price of $p^{n/2}$. From the calculations made above we can immediately conclude that $p^{(n-1)/2}$ is the next feasible price above the Walrasian price of the same market in which only n–1 firms operate, while $p^{(n+1)/2}$ is the next feasible price above the Walrasian price of the market with n+1 firms. Exactly the same argumentation as for even n leads to the result that there is downward pressure whenever $p \geq p^{(n-1)/2}$, and upward pressure whenever $p \leq p^{(n+1)/2}$. Notice that contrary to the case for even n, impulses from profitable deviations are not balanced. Now consider all prices in between $p^{(n-1)/2}$ and $p^{(n+1)/2}$. Lemma 1 implies that pressure on the price is only exerted from profitable deviations. In between $p^{(n-1)/2}$ and $p^{(n+1)/2}$ $(n-1)/2+1$ firms are needed for a profitable downward deviation, while the number of firms needed for a profitable upward deviation is $n-(n-1)/2+1 = (n-1)/2+1$. Thus, for all prices in between $p^{(n-1)/2}$ and $p^{(n+1)/2}$ upward and downward pressure on prices are balanced. The following table summarises the impulses.

Table A.3. Price impulses for an odd number of firms

price level	number of firms needed for profitable downward devation	number of firms needed for profitable upward devation	pressure on price level
$p^{(n-1)/2}$	$(n-1)/2 + 1$	$(n-1)/2 + 2$	down
$p^{(n-1)/2} - \Delta$	$(n-1)/2 + 1$	$(n-1)/2 + 1$	balanced
…	…	…	…
$p^{(n+1)/2} + \Delta$	$(n-1)/2 + 1$	$(n-1)/2 + 1$	balanced
$p^{(n+1)/2}$	$(n-1)/2 + 2$	$(n-1)/2 + 1$	up

Thus, if the default price approaches $p^{(n-1)/2}$ from above, it will converge to $p^{(n-1)/2}$; if it approaches $p^{(n+1)/2}$ from below, it will converge to $p^{(n+1)/2}$ Once this price is reached, there is no systematic pressure on the default price anymore. ■

Notice that although the Walrasian price of the n firm market will not be reached if n is odd (since it does not constitute a threshold price, it is nevertheless included in the interval in which price pressure is balanced.

Appendix C. Deviations from imitative behaviour in the new sessions

Table C.1. Frequency of imitations of the most successful price of the previous round (new sessions)

	n=3 (27 focal)		n=3 (21 focal)		n=8	
Deviation from most succesful price	frequency	per cent	frequency	per cent	frequency	per cent
≥+10	4	0.3	11	1.0	335	4.9
+9	1	0.1	2	0.2	46	0.7
+8	3	0.2	9	0.8	39	0.6
+7	1	0.1	9	0.8	44	0.6
+6	8	0.6	30	2.6	38	0.6
+5	4	0.3	12	1.1	19	0.3
+4	14	1.0	15	1.3	8	0.1
+3	30	2.1	20	1.8	10	0.1
+2	14	1.0	49	4.3	27	0.4
+1	85	5.9	197	17.3	125	1.8
±0	1217	84.0	678	59.5	4850	71.6
–1	45	3.1	78	6.8	863	12.7
–2	11	0.8	18	1.6	353	5.2
–3	4	0.3	4	0.4	16	0.2
–4	1	0.1	2	0.2	0	0.0
–5	2	0.1	3	0.3	1	0.0
–6	0	0.0	1	0.1	0	0.0
–7	2	0.1	0	0.0	0	0.0
–8	0	0.0	1	0.1	1	0.0
–9	0	0.0	0	0.0	0	0.0
≤–10	3	0.2	1	0.1	1	0.0
Ties	21	1.4	36	3.2	840	11.0

The relative frequencies refer to all choices where the most successful price is unique. “Ties” means that in 0, 24, and 40 cases more than one price was the most successful in the previous round. This can only happen if the lowest price firm(s) made a loss.

Appendix D. Simulation results for different β and initial conditions

Simulated market prices
β=0.90, initial prices from data

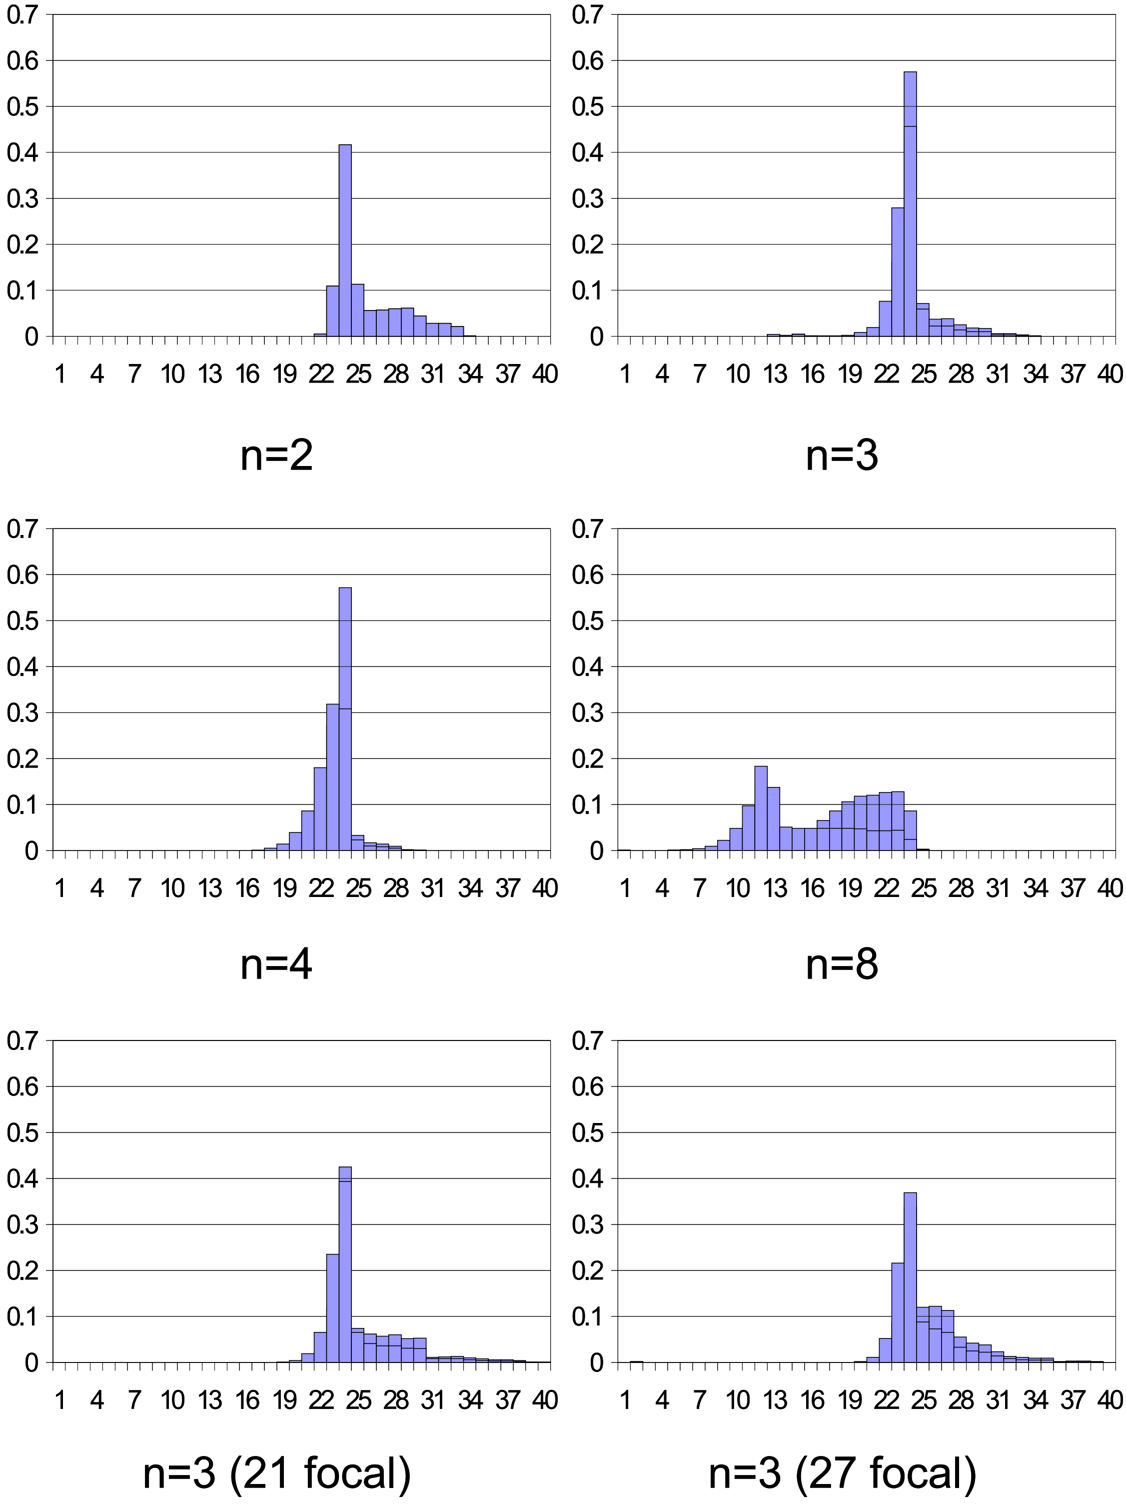

Simulated market prices
β=0.70, initial prices from data

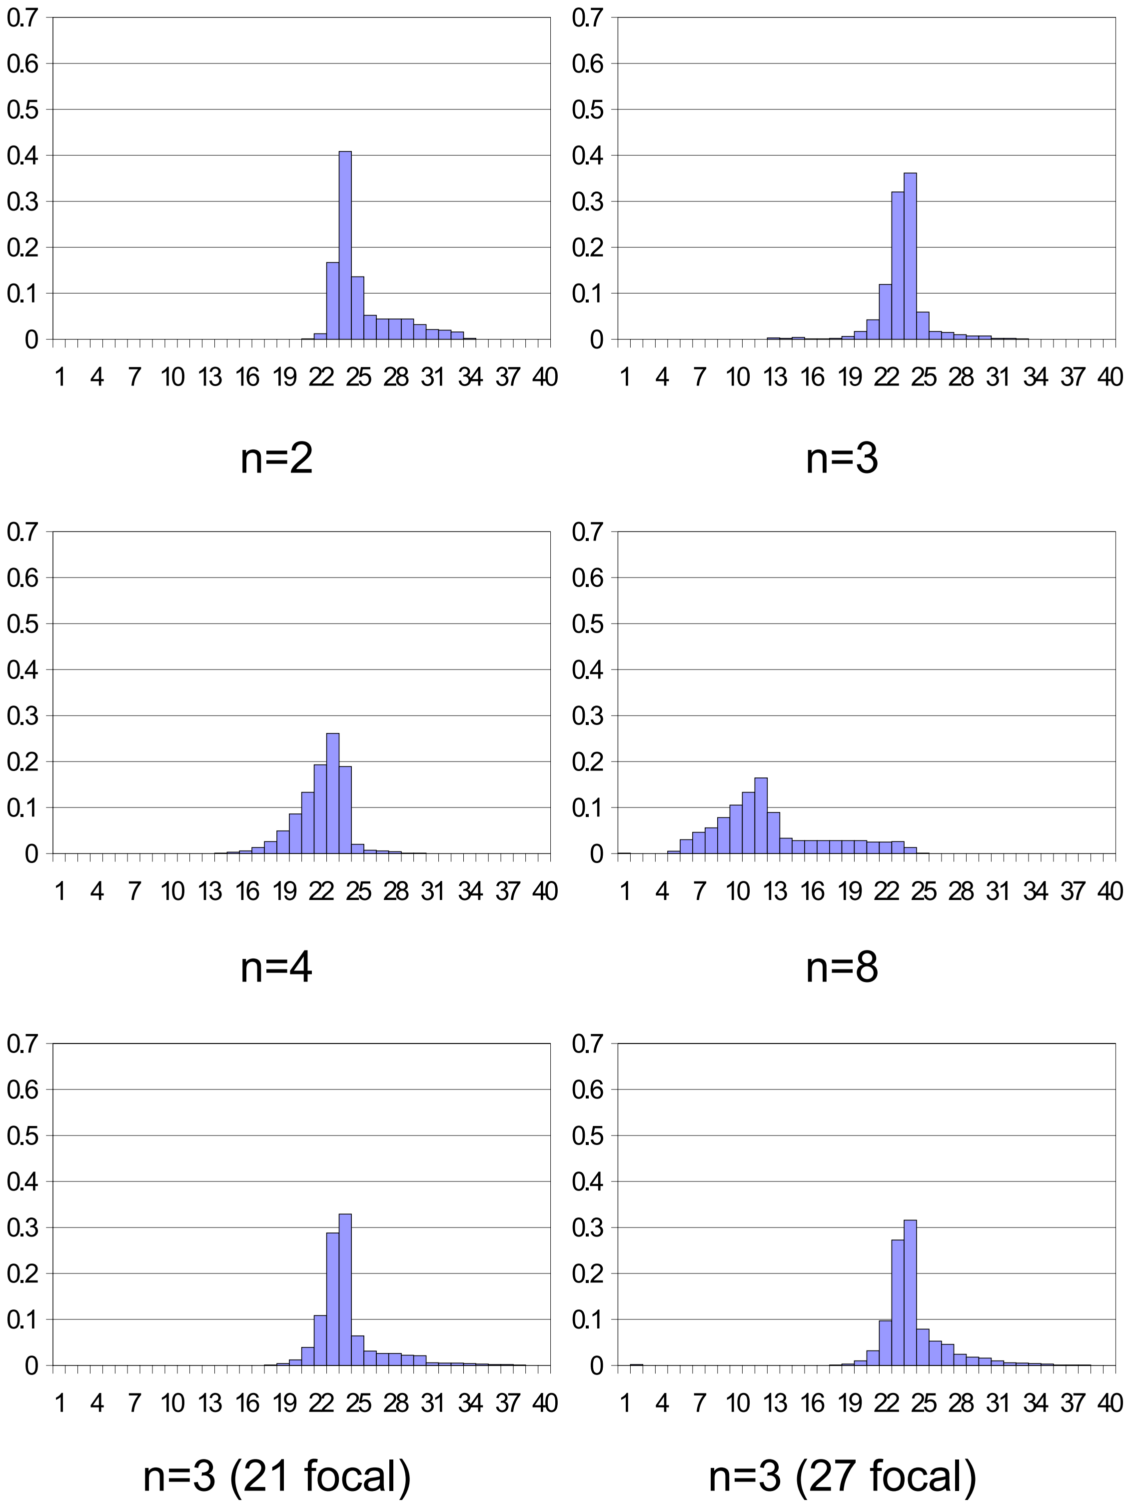

Simulated market prices
β=0.90, uniform initial prices

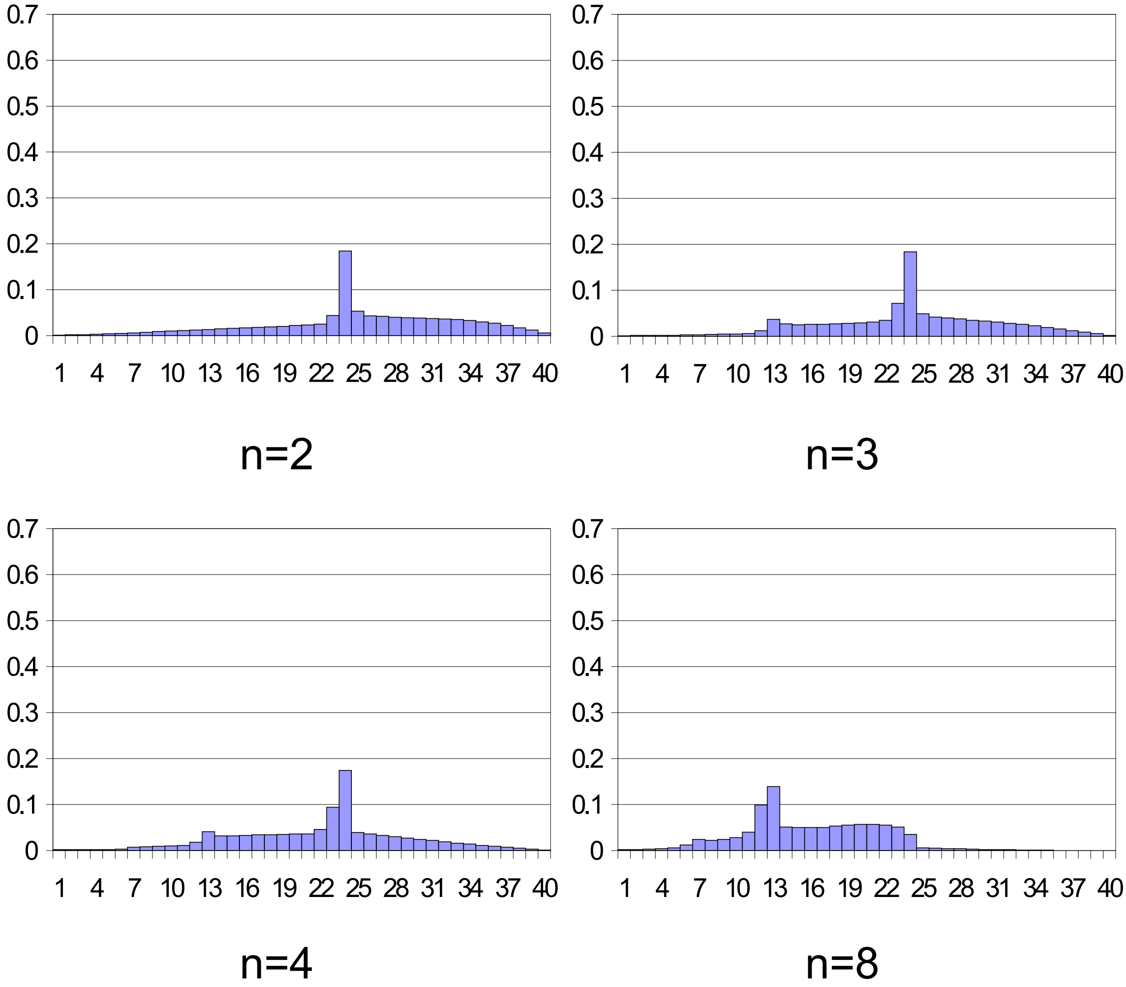

Simulated market prices
β=0.80, uniform initial prices

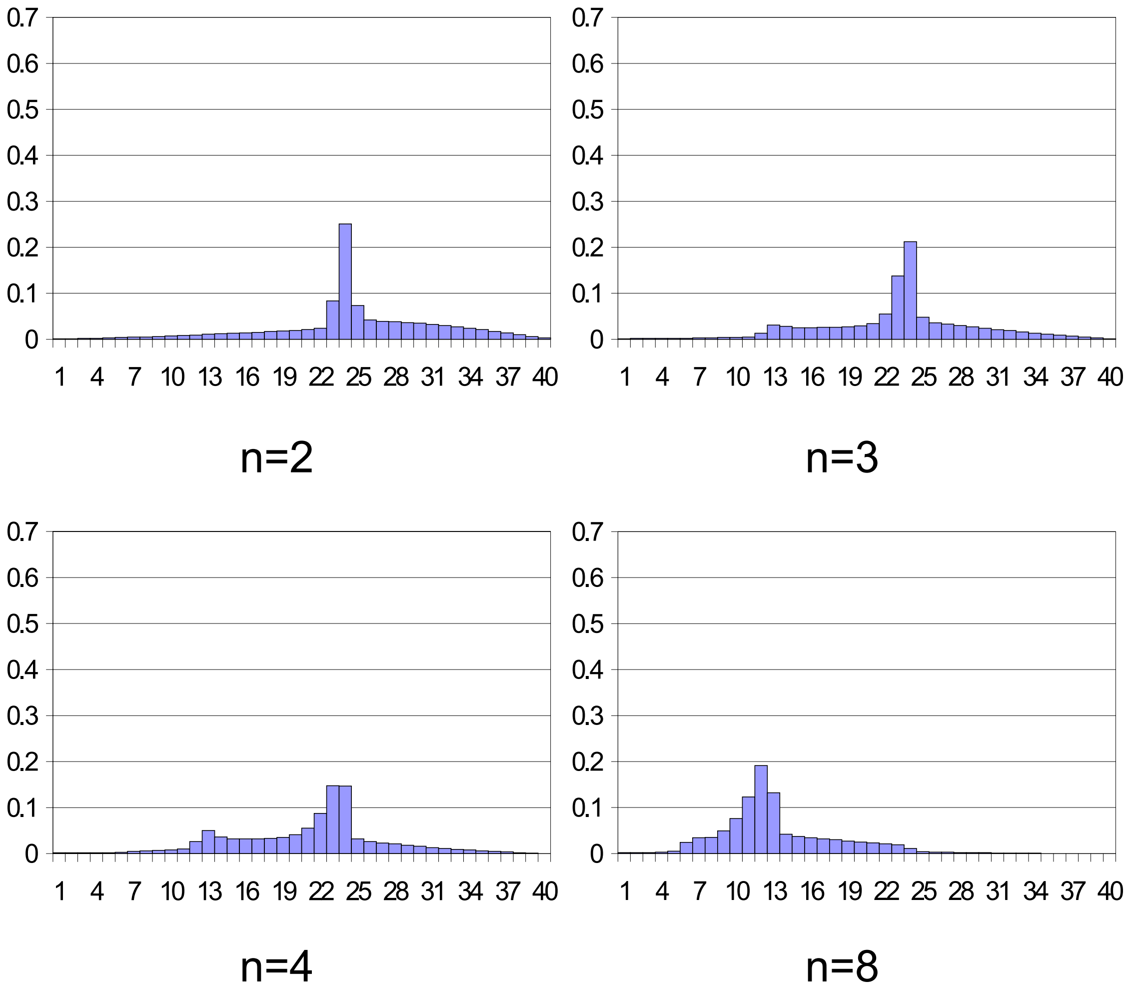

Simulated market prices
β=0.70, uniform initial prices

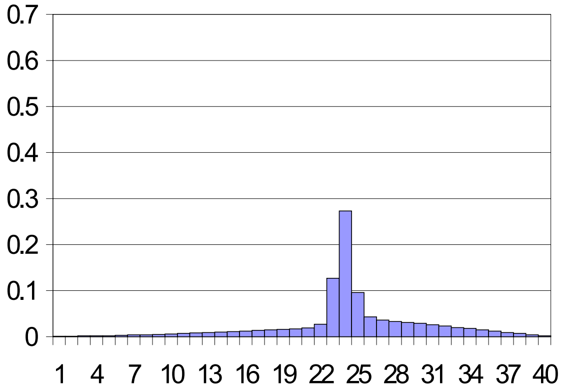

n=2

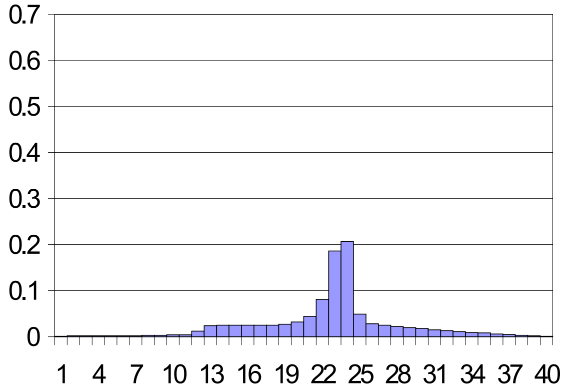

n=3

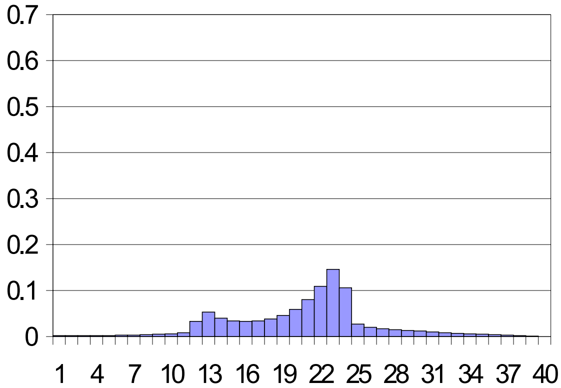

n=4

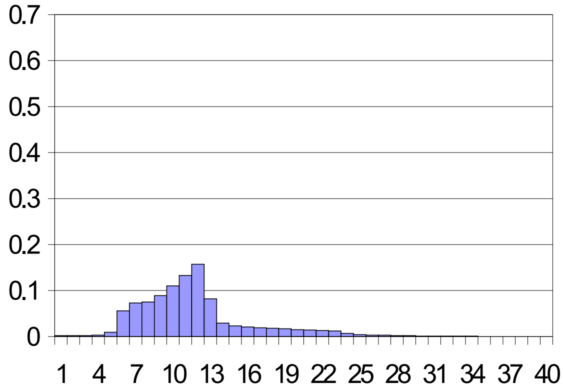

n=8

Appendix E. Payoff tables for the new treatments

Table E.1. The treatment n=3 (21 focal)

Number chosen	Profit if this number is the unique lowest of the four	Profit if this number is the lowest and has been chosen by two players	Profit if this number is the lowest and has been chosen by three players	Profit if this number is not the lowest
40	1117	813	674	340
39	1124	829	688	340
38	1123	843	700	340
37	1117	854	710	340
36	1103	862	719	340
35	1083	868	727	340
34	1056	872	733	340
33	1023	873	738	340
32	982	872	742	340
31	936	868	743	340
30	882	862	744	340
29	822	854	743	340
28	755	843	741	340
27	681	829	737	340
26	601	813	732	340
25	514	795	725	340
24	421	774	717	340
23	320	751	707	340
22	214	725	696	340
21	100	697	684	340
20	-20	666	670	340
19	-147	633	654	340
18	-281	597	638	340
17	-421	559	619	340
16	-568	519	600	340
15	-722	476	579	340
14	-882	430	556	340
13	-1049	382	532	340
12	-1223	332	507	340
11	-1403	279	480	340
10	-1590	224	452	340
9	-1784	167	422	340
8	-1984	106	391	340
7	-2192	44	358	340
6	-2405	-21	324	340
5	-2626	-89	288	340
4	-2853	-159	251	340
3	-3087	-231	213	340
2	-3327	-306	173	340
1	-3574	-383	132	340

Table E.2. The treatment n=3 (27 focal)

Number chosen	Profit if this number is the unique lowest of the four	Profit if this number is the lowest and has been chosen by two players	Profit if this number is the lowest and has been chosen by three players	Profit if this number is not the lowest
40	496	192	53	-281
39	503	208	67	-281
38	502	222	79	-281
37	496	233	89	-281
36	482	241	98	-281
35	462	247	106	-281
34	435	251	112	-281
33	402	252	117	-281
32	361	251	121	-281
31	315	247	122	-281
30	261	241	123	-281
29	201	233	122	-281
28	134	222	120	-281
27	60	208	116	-281
26	-20	192	111	-281
25	-107	174	104	-281
24	-200	153	96	-281
23	-301	130	86	-281
22	-407	104	75	-281
21	-521	76	63	-281
20	-641	45	49	-281
19	-768	12	33	-281
18	-902	-24	17	-281
17	-1042	-62	-2	-281
16	-1189	-102	-21	-281
15	-1343	-145	-42	-281
14	-1503	-191	-65	-281
13	-1670	-239	-89	-281
12	-1844	-289	-114	-281
11	-2024	-342	-141	-281
10	-2211	-397	-169	-281
9	-2405	-454	-199	-281
8	-2605	-515	-230	-281
7	-2813	-577	-263	-281
6	-3026	-642	-297	-281
5	-3247	-710	-333	-281
4	-3474	-780	-370	-281
3	-3708	-852	-408	-281
2	-3948	-927	-448	-281
1	-4195	-1004	-489	-281

Table E.3. The treatment n=8

Number chosen	Profit if this number is the lowest and has been chosen by ... players								Profit if this number is not the lowest
	one	two	three	four	five	six	seven	eight	
40	777	473	334	258	210	177	152	134	0
39	784	489	348	269	219	185	160	141	0
38	783	503	360	279	228	192	166	146	0
37	777	514	370	288	236	199	172	152	0
36	763	522	379	296	243	205	178	157	0
35	743	528	387	303	249	211	183	162	0
34	716	532	393	310	255	216	188	166	0
33	683	533	398	315	259	220	192	169	0
32	642	532	402	319	263	224	195	173	0
31	596	528	404	322	267	227	198	175	0
30	542	522	404	324	269	230	200	178	0
29	482	514	403	325	271	232	202	180	0
28	415	503	401	325	272	233	204	181	0
27	341	489	397	324	272	234	205	182	0
26	261	473	392	322	272	234	205	183	0
25	174	455	385	319	270	233	205	183	0
24	81	434	377	315	268	232	204	182	0
23	-20	411	367	310	265	230	203	181	0
22	-126	385	356	304	261	228	201	180	0
21	-240	357	344	297	257	225	199	179	0
20	-360	326	330	290	252	221	196	176	0
19	-487	293	314	281	246	217	193	174	0
18	-621	257	298	271	239	212	189	171	0
17	-761	219	279	260	232	206	185	167	0
16	-908	179	260	248	223	200	180	163	0
15	-1062	136	239	235	214	193	175	159	0
14	-1222	90	216	221	204	186	169	154	0
13	-1389	42	192	206	194	178	163	149	0
12	-1563	-8	167	189	183	169	156	143	0
11	-1743	-61	140	172	170	160	148	137	0
10	-1930	-116	112	154	158	150	140	130	0
9	-2124	-173	82	135	144	140	132	123	0
8	-2324	-234	51	115	129	128	123	116	0
7	-2532	-296	18	94	114	117	113	108	0
6	-2745	-361	-16	72	98	104	103	99	0
5	-2966	-429	-52	49	81	91	93	90	0
4	-3193	-499	-89	25	64	78	82	81	0
3	-3427	-571	-127	0	46	63	70	71	0
2	-3667	-646	-167	-26	27	49	58	61	0
1	-3914	-723	-208	-53	7	33	45	51	0