Journal of Economic Behavior and Organization 166 (2019) 346–365

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Competing for market shares: Does the order of moves matter even when it shouldn't?

Tanja Hörtnagl[a,*], Rudolf Kerschbamer[b], Rudi Stracke[c]

[a] Management Center Innsbruck, Austria
[b] Department of Economics, University of Innsbruck, Austria
[c] Seminar for Population Economics, University of Munich (LMU), Germany

ARTICLE INFO

Article history:
Received 1 October 2018
Revised 11 July 2019
Accepted 11 July 2019
Available online 30 July 2019

JEL classification:
C72
C92
D21
D74
L13
L22
M37

Keywords:
Sequential competition
Order of moves
Experiment
Second-Mover advantage
Collusion
Contest

ABSTRACT

This paper investigates whether the order of moves affects behavior and outcomes even in the absence of material incentives for pre-commitment. Evidence from laboratory experiments with shared rents, symmetric players and either a simultaneous or a sequential move order suggests that it is an inherent advantage to move second rather than first. The reason is that first movers face strategic uncertainty, while second movers have the power to ultimately determine the outcome through their investment choices. Many first movers acknowledge this power and decrease investments below the standard prediction which is interpreted as an attempt to collude. This strategy leads to higher absolute payoffs for the first-movers compared to a situation where they do not try to establish collusion. However, relative payoffs are low since collusive outcomes are typically asymmetric to the advantage of the second mover.

1. Introduction

Most real world markets are characterized by oligopolistic competition meaning that only a few firms compete for a large number of potential buyers. For instance, there is less than a dozen independent automobile producers worldwide, three leading food processing companies together achieve a large share of global processed food sales, there are only a handful of different brands of household detergents, etc. A core characteristic of oligopolistic markets is that prices, quantities, and marketing expenditures are chosen strategically. A natural question then is how the ability to pre-commit to a choice affects market outcomes. This question has attracted the attention of theorists for a long time ever since the pioneering contribution by von Stackelberg (1934). Experimental work on the effects of commitment is still scarce, however, even though some of the main theoretical predictions have been tested in lab experiments (by Huck et al., 2001 and by Kübler and Müller, 2002, for instance).

* Corresponding author.
E-mail address: tanja.hoertnagl@mci.edu (T. Hörtnagl).

https://doi.org/10.1016/j.jebo.2019.07.005

This paper investigates in lab experiments whether the ability to pre-commit matters even in the absence of immediate strategic incentives for pre-commitment. In particular, we consider a two-player contest game with symmetric agents which yields the prediction that equilibrium investments and payoffs are independent of the order of moves under standard assumptions (Dixit, 1987).[1] This implies that in the theoretical benchmark there is neither a value of commitment for the leader, nor a value of information for the follower, nor a first- or a second-mover advantage. Here and throughout the term "value of commitment for the first mover" ("value of information for the second mover") stands for the additional equilibrium payoff the first mover (the second mover) receives in the sequential-move version of the game as compared to the payoff of the same player in the simultaneous-move version.[2] By contrast, the terms "first-mover advantage" and "second-mover advantage" refer to the comparison of equilibrium payoffs across the two players within the sequential-move version of the game. To test whether the order of moves affects behavior even in the absence of strategic considerations, our experiments feature two treatments implemented in a between-subjects design: Subjects choose the amounts they want to invest simultaneously in treatment SIM, while investments are made sequentially in treatment SEQ. Independent of the treatment, subjects face a one-shot (OS) interaction in part 1 and a finitely repeated (REP) interaction in part 2 of the experiment.

Even though theory predicts neither a first-mover nor a second-mover advantage, our results suggest that it is an inherent advantage to move second rather than first, both in the OS and in the REP interaction.[3] More specifically, we find that mean outcomes in both parts of the experiment are in line with the theoretical benchmark in SIM, but not in SEQ. In the latter, the first mover earns considerably less than the second mover, on average. When disaggregating the SEQ data to understand deviations from the theoretical benchmark, we observe a concentration of first-mover investment choices below and above the benchmark and we also find that the by far largest mass point is below the predicted amount both in the one-shot and in the repeated interaction. We interpret this as evidence for attempted collusion. Second movers typically respond to collusion attempts by the first mover by investing more than the paired first mover. As a result, collusive outcomes in SEQ are typically asymmetric – with the second mover investing and earning considerably more than the paired first mover. By contrast, first movers who invest more than predicted are disadvantaged in absolute terms but they typically earn the same or even slightly more than the paired second mover. Taken together, our findings suggest that first movers in SEQ face a trade-off between maximizing their absolute earnings (by investing low amounts) and minimizing disadvantageous inequality (by choosing high investments).

Overall our results suggest that it is an inherent advantage to move second rather than first even in the absence of material incentives for pre-commitment. The reason seems to be that second movers have the power to ultimately determine the outcome through their investment choice. This power seems particularly valuable in environments where behavior is heterogeneous – for instance, because players care to different degrees for the payoffs of others. In such environments first movers in the strategic interaction face a lot of strategic uncertainty. In our case first movers seem to react to this strategic uncertainty by acting more cooperatively. The cooperative move by the first mover is then systematically exploited by the second mover (who typically invests strictly more than the paired first mover).

2. Related literature

Our results contribute to the literature of different subfields of economics and business. First, we contribute to the literature on strategic firm behavior in duopolies. For instance, Huck et al. (2001) and Kübler and Müller (2002) investigate how the order of moves affects firm behavior when the quantities supplied – the prices charged, respectively – are the strategic choice variables of firms. The major conceptual difference to these papers is that under the usual assumptions regarding the demand and cost function standard theory predicts a first-mover advantage in the quantity-competition setting and a second-mover advantage in the price-competition setting while there is neither a first- nor a second-mover advantage in the standard benchmark of our model.

Second, we contribute to the literature on strategic marketing investments, since the setting we consider can be interpreted as a "market-share-attraction" model with combative advertising where the size of the market is fixed (Friedman, 1958). We are not aware of any study that uses experimental methods to investigate how the order of moves affects advertising investments in this particular version of the market-share-attraction model.[4] Our findings are in line with empirically observed advertising wars in slow growth industries where market share gains are the driving force behind expansions. More specifically, our findings suggest that concerns for relative standing within an industry might explain

[1] The results of Dixit (1987) hold for a fixed prize, a logit or probit type CSF and unit marginal costs of effort. However, as Hoffmann and Rota-Graziosi (2012) have shown Dixit's findings no longer hold in the presence of an endogenous prize.

[2] The literature on endogenous timing uses the terms 'first-mover incentive' and 'second-mover incentive' instead of the terms 'value of commitment for the first mover' and 'value of information for the second mover' - see e.g. Van Damme and Hurkens (1999) or Hoffmann and Rota-Graziosi (2012) who give a clear characterization of these terms. The term 'value of commitment' goes back to Schelling (1960) and is used frequently in the IO literature (e.g. Güth et al., 2006, Morgan and Várdy, 2004, Morgan and Várdy, 2007, Kopel and Löffler, 2008). The term 'value of information' is used, for instance, by Gossner and Mertens (2001), Lehrer and Rosenberg (2006) or Bassan et al. (2003).

[3] Theory predicts that the leader in the sequential-move version of the game chooses the same effort as a player in the simultaneous-move version of the game. Since the rival chooses a best reply in both cases, this leads to identical effort choices in both versions of the game. As a consequence, there is neither a first-mover advantage nor a second-mover advantage.

[4] See Bass et al. (2005) and the references therein for more complex models and dynamic versions that also account for different types of marketing activity.

why we frequently observe "advertising wars that often produce mutually damaging consequences and that advertisers who engage in them often regret" (Beard, 2011, p. 399) – think of the so-called 'Cola' or 'Burger' wars, for example.[5]

Third, this paper is related to the literature on contests with sequential moves. In contrast to the other papers in that literature we experimentally implement a contest game in which contestants compete for shares of the prize at stake instead of considering 'winner-takes-all' contests. Even though both interpretations are strategically equivalent under the standard assumption of risk neutrality, existing evidence suggests that participants behave differently in the two formats – see Sheremeta et al. (2012), or Eisenkopf and Teyssier (2014), for example. For an extensive overview of the findings in sequential-move contests we refer to Hoffmann and Rota-Graziosi (2018) who offer a survey of both, the theoretical and experimental findings in this field, as well as Dechenaux et al. (2015) who provide (in Chapter 4) a comprehensive survey of the experimental literature on sequential-move contests. Firm behavior given non-standard best responses has also been the object of a recent theoretical investigation by Hoffmann and Rota-Graziosi (2019) who show how non-monotonicities in the payoff or the marginal payoff of a player with respect to the action choice of the opponent alter behavior in sequential-move games within various frameworks. Assuming uniqueness of the Nash equilibrium of the basic game, concavity of the leader payoff-function, and local monotonicity of the follower-payoff, the authors propose a taxonomy of games based solely on the properties of the Nash equilibrium of the basic game and then establish which constellations involve a first- or second-mover advantage and a first- or second-mover incentive.

The findings in the experimental literature on sequential contests are mixed. For instance, Shogren and Baik (1992) report that in a sequential-move contest with asymmetric agents many second movers invest more than the paired first mover which is in contrast to the standard best-response prediction of higher investments of the first mover. Weimann et al. (2000) consider a sequential-move contest with symmetric agents with a 'winner-takes-it-all' prize structure. They modify the contest technology in such a way that it delivers a first-mover advantage. Surprisingly, however, they find that second movers earn more than first movers. Specifically, they find that high effort choices of first movers are punished by second movers, while cooperative moves are almost always exploited. This latter finding is completely in line with our results.

In terms of experimental design the closest relative to the present paper is probably Fonseca (2009) who analyzes simultaneous- and sequential-move contests. In a winner-takes-all contest he considers treatments with symmetric and asymmetric contestants in a 2x2 design. The data reveal neither a first- nor a second-mover advantage in symmetric settings, and some first-mover advantage in the asymmetric setting. Also, second mover behavior is extremely heterogeneous – many second movers drop out in the sequential interaction, particularly in the asymmetric one, while others invest extremely high amounts. We complement the study by Fonseca (2009) in several dimensions: First, the share version of the contest we employ produces less noisy data than the 'winner-takes-it-all' version and it also makes it easier to explain observed behavior. In particular, it is unclear whether subjects focus on expected outcomes or on potential outcomes in the 'winner-takes-all' contest considered by Fonseca (2009). Punishment investments by the second mover, for example, only hurt the first mover with some probability, but reduce the second-mover payoff with certainty. Second, we employ a fixed partner matching over 12 periods whereas Fonseca (2009) uses a random matching procedure over 30 rounds in each treatment. Even though the prediction is the same, experimental data have shown that fixed partner matching can lead to more collusive outcomes – see e.g. Huck et al. (2001) who employ a fixed and random matching procedure. And we indeed find by far more evidence for attempted and successful collusion than Fonseca (2009) does.

3. Theoretical analysis

The Model. In our experiment two symmetric agents compete for shares of a prize by simultaneously or sequentially investing in a Tullock (1980) lottery contest.[6] This setting imitates some form of competition among firms, e.g. the competition for market shares as in Friedman (1958), or the competition between rent seekers for a divisible rent as in Long and Vousden (1987).

The share of the prize P by firm i is given by:

$$s_i(x_i, x_{-i}) = \begin{cases} \frac{x_i}{x_i + x_{-i}} & \text{if} \quad x_i + x_{-i} > 0 \\ 0.5 & \text{if} \quad x_i + x_{-i} = 0, \end{cases}$$

where x_i is the investment chosen by firm i and x_{-i} is the investment of the competitor. Assuming linear investment costs and own profit as the only motive, the optimization problem of firm i reads

$$\max_{x_i \geq 0} \Pi_i(x_i, x_{-i}) = \frac{x_i}{x_i + x_{-i}} P - x_i. \tag{1}$$

First-order optimality conditions deliver the best-response function of firm i for given investments x_{-i} by the competitor

$$\mathrm{BR}_i(x_{-i}) = \max\{\sqrt{x_{-i}P} - x_{-i}, 0\}.$$

[5] See also Beard (2010) for details and further examples.

[6] We use for our contest technology the specification by Tullock (1980) with a share (rather than a probability) interpretation.

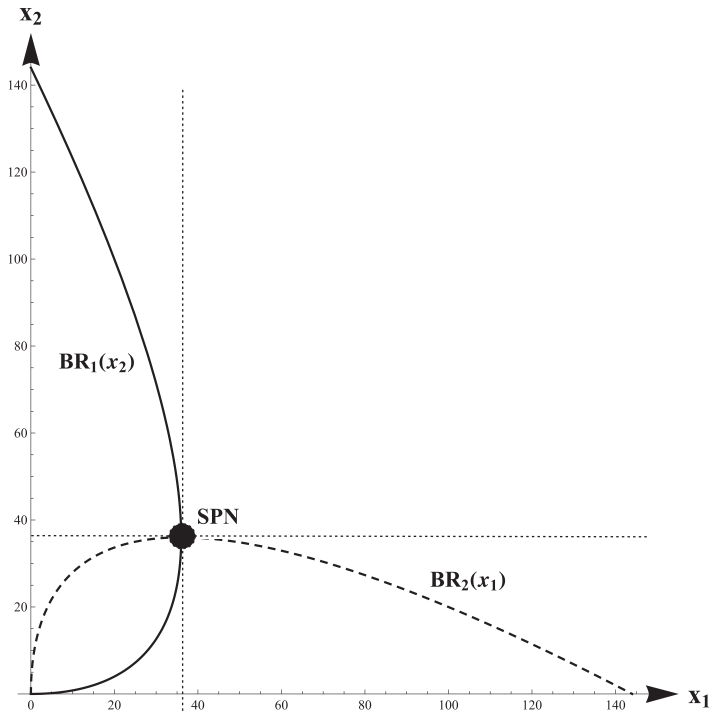

Fig. 1. Best-Response Functions and the SPNE *Note*: The figure plots the best-response function of the two firms for the parameters used in the experiments ($P = 144$).

Simultaneous Investment Decisions. For the setting where the two firms make their investment decisions simultaneously (denoted SIM) the equilibrium occurs at the point where best response functions cross. The crossing conditions $x_1^* = BR_1(x_2^*)$ and $x_2^* = BR_2(x_1^*)$ deliver equilibrium investments

$$x^*_{\texttt{SIM}} \equiv x_1^* = x_2^* = \frac{P}{4}. \tag{2}$$

These investments imply equilibrium profits for the simultaneous move game of:

$$\Pi^*_{\texttt{SIM}} \equiv \Pi_1(x_1^*, x_2^*) = \Pi_2(x_2^*, x_1^*) = \frac{P}{4}. \tag{3}$$

Sequential Investment Decisions. Consider next the sequential move setting (denoted SEQ) where firm 1 is the first mover (FM, he) who commits to an investment level first which is then perfectly observed by firm 2, the second mover (SM, she), before she makes her own investment choice. In a subgame-perfect equilibrium of SEQ the FM perfectly anticipates how the SM reacts to each investment choice and – given the anticipated reaction – he searches for the own best reply. Technically the best-response function of the SM is plugged into the optimization problem of the FM. Maximization then yields the equilibrium investments. In our model where firms are symmetric, they choose the same level of investment in equilibrium:

$$x^*_{\texttt{SEQ-FM}} \equiv x_1^* = x^*_{\texttt{SEQ-SM}} \equiv x_2^* = \frac{P}{4}. \tag{4}$$

The resulting equilibrium profits for the FM and the SM then are

$$\Pi^*_{\texttt{SEQ-FM}} \equiv \Pi_1(x_1^*, x_2^*) = \Pi^*_{\texttt{SEQ-SM}} \equiv \Pi_2(x_2^*, x_1^*) = \frac{P}{4}. \tag{5}$$

Discussion. The comparison of the theoretical benchmarks for investments and payoffs across SIM and SEQ shows that equilibrium investments and payoffs are the same across settings (SIM vs. SEQ) and within a setting the same across players – a point previously made by Dixit (1987). This implies that in the benchmark there is neither a positive or negative value of commitment for the FM, nor a positive or negative value of information for the SM, nor a first-mover or second-mover advantage.[7] The reason is that investment decisions are locally neutral, i.e. they are locally neither strategic complements

[7] As previously mentioned the terms 'value of commitment for the FM' and 'value of information for the SM' refer to the additional equilibrium payoff the respective player receives in the 'sequential-move' as compared to the 'simultaneous-move' version of the game. By contrast, the terms 'FM advantage' and 'SM advantage' refer to the comparison of payoffs of the players within the sequential-move game.

nor strategic substitutes.[8] This is illustrated in Fig. 1: The slope of the best reply is strictly positive to the left of the unique point of intersection and strictly negative to the right of that point – but it is exactly zero at the point of intersection. This is different in standard industrial organization models where best-response functions are typically either monotonically increasing or monotonically decreasing in the relevant range.[9] At the same time, the setting we consider is similar to standard IO models in other dimensions that are important when comparing simultaneous- and sequential-move games. In particular, as in standard IO games every decision of one player has an impact on the payoff of the other player. As a consequence of this externality, the Nash equilibrium of the simultaneous-move game – which in terms of implemented choices is identical to the Subgame Perfect equilibrium of the sequential-move game – is inefficient.[10]

4. Design of the experiments

Experimental Parameters and Treatments. We implement the two timing protocols SIM and SEQ in a between-subject design meaning that experimental subjects either decide simultaneously about their investments (in treatment SIM), or sequentially (in treatment SEQ). Everything else is held constant across treatments. In particular, the prize is set to $P = 144$ in both treatments. We ran 2 sessions for treatment SIM and 4 sessions for treatment SEQ with 20 participants each. All 120 subjects were students from the University of Innsbruck and each subject participated only once. The experiment was programmed in z-tree (Fischbacher, 2007) and students were recruited using ORSEE (Greiner, 2004). The experimental currency unit 'Taler' was converted to Euro at an exchange rate of 50:1 at the end of the experiment. Each session lasted about 70 minutes, and subjects earned slightly more than 15 Euro on average (including a show up fee of 4 Euros).[11]

Implementation. At the beginning of each session, participants received general instructions and were informed that the experiment has four parts. They were also informed that their earnings in each part will depend on the own decisions and on the decisions of at most one anonymous second participant. We are only interested in parts 1 and 2 subsequently, which implement a one-shot version of the market interaction (part 1) and a finitely repeated market with partner matching (part 2), respectively. After reading the instructions for part 1, subjects had the possibility to test and improve their understanding of the instructions in a training program.[12] In this training program, subjects could fill in different values of the decision variable both for themselves and for a hypothetical partner. After confirming their choices, they were informed about the resulting division of the market and the payoffs. The training period lasted about 6 minutes as subjects made intensive use of the program by entering many different investment combinations. After the training period, we started with part 1 of the experiment. For this part each subject was randomly matched with a partner. All subjects in SIM and FMs in SEQ then were asked to provide an own investment and an estimate of the expected investment of the opponent. Decisions of SMs were elicited using real play rather than the strategy method, i.e., SMs were informed about the actual investment of the paired FM and responded only to this particular choice. After all decisions were made, subjects were informed about their own investment choice, the amount invested by the opponent, and both their own and the opponent's payoff. Subsequently, we started part 2 of the experiment. Instructions for part 2 were provided on the computer screen, as subjects were facing the same decision environment as in part 1. The only difference to part 1 was that partners were fixed for 12 decision rounds and that this was common knowledge. Subjects received the same feedback after each decision round as at the end of part 1 and were informed ex-ante that for part 2 only one randomly chosen decision round would be paid out at the end of the experiment. After subjects completed all parts of the experiment, they were asked to fill out a questionnaire (voluntary and not-incentivized).

Decision Environment. The decision environment in the experiment was neutrally framed in that we did not relate the strategic interaction to a particular application. Subjects had to decide how much of their endowment they wanted to invest in order to get a share of the prize. The endowment in part 1 and in each period of part 2 was 144 Taler and endowments could not be transferred across parts or periods. The prize was equal to the endowment (that is, 144 Taler) in part 1 and in each period of part 2. In the SIM treatment subjects were not informed about the decision of the paired opponent before making their own choice. In the SEQ treatment each subject was first assigned a role – either the role of a FM or the role of a SM. These roles remained constant across parts and across periods in part 2. Whereas the FM received no information about the decision of the paired SM, the SM was informed about the paired FM's investment decision before making her own choice.

[8] The actions of two players are strategic substitutes (strategic complements) if the best response functions are downward sloping (upward sloping) – assuming that the objective functions are strictly concave. See Bulow et al. (1985).

[9] See Gal-Or (1985), for example, who shows that downward sloping (upward sloping) best response functions typically lead to a first-mover (second-mover) advantage.

[10] This in turn implies that starting from the Nash equilibrium of the simultaneous move game there is room for collusive choices that improve the payoffs of both players. In the limit where players collude perfectly (by both investing nothing) the profits of both players are twice as high as in the Nash equilibrium of the simultaneous-move game.

[11] The average payoff does also include earnings in two additional experimental parts that were conducted after our main experiment and are unrelated to the present study.

[12] Instructions are available in the Appendix.

Table 1
Standard Predictions and Parameters.

	SIM	SEQ-FM	SEQ-SM
Payoff (Π_i)	36	36	36
Investment (x_i)	36	36	36
Prize (P)	144	144	144

Table 2
Average Payoffs and Investment Choices in SIM and SEQ.

	Part 1 (OS)		Part 2 (REP)	
	Payoff	Investment	Payoff	Investment
SIM ($N = 40$)	**38.4**	**33.6**	**28.22**	**43.78**
	(34.09)	(32.19)	(33.71)	(35.49)
SEQ-FM ($N = 40$)	**11.23**	**29.00**	**27.14**	**36.64**
	(30.63)	(38.24)	(20.97)	(26.21)
SEQ-SM ($N = 40$)	**44.52**	**59.25**	**39.77**	**40.45**
	(57.65)	(50.63)	(30.40)	(23.46)
		p-values		
(a) SEQ-FM vs. SIM	0.000	0.220	0.855	0.211
(b) SEQ-SM vs. SIM	0.920	0.026	0.121	0.535
(c) SEQ-FM vs. SEQ-SM	0.002	0.000	0.002	0.082

Note: The table provides average payoffs (standard deviation in parentheses) in part 1 (one-shot) and part 2 (repeated) across 40 subjects; the null hypothesis is that investment choices and payoffs of agents are equal across the three rows. We use a Wilcoxon-signed rank (WSR) test to compare the outcomes in lines two and three (p-value in last row of the table) and a Mann-Whitney U (MWU) test for the other comparisons. Before applying the MWU or WSR test in REP we take average investment levels across all rounds for each individual, which implies that the number of observations is 40 for each treatment. For robustness checks we use also a regression analysis to identify treatment effects and find that the results are qualitatively the same.

Benchmark predictions for investment choices as well as for the resulting payoffs in treatments SIM and SEQ are provided in Table 1. Assuming play according to the standard benchmark (the main assumptions of the standard model are that all agents are rational, risk neutral and only interested in their own material payoff and that this is common knowledge), subjects should invest the same amount across treatments and roles. As a consequence, earnings of subjects in SIM and of FMs and SMs in SEQ should also be equal. This is summarized in our main hypothesis:

Hypothesis. The order of moves does not affect equilibrium payoffs: Subjects in SIM, FMs in SEQ and SMs in SEQ all earn the same amount. Specifically,

$$\begin{array}{llll} (a) & \Pi_{\texttt{SIM}} & = & \Pi_{\texttt{SEQ-FM}}; \\ (b) & \Pi_{\texttt{SIM}} & = & \Pi_{\texttt{SEQ-SM}}; \\ (c) & \Pi_{\texttt{SEQ-FM}} & = & \Pi_{\texttt{SEQ-SM}}. \end{array}$$

5. Experimental results

5.1. Average payoffs and investment choices

Table 2 provides average payoffs and average investment choices for both parts of the experiment (OS and REP) and both timing sequences (SIM and SEQ) as well as the results of the hypothesis tests. The standard benchmark of 36 is equal across all measures, treatments and player roles. Before testing our hypotheses we investigate whether the averages for payoffs and investments are in line with the theoretically predicted values.[13] Average payoffs and investment choices in both the one-shot and the repeated interaction of treatment SIM are close enough to the benchmark predictions such that we cannot reject equality at conventional levels. The pattern is somewhat different in OS-SEQ. Here, FMs invest and earn significantly less than predicted – the respective p-values are 0.090 (for 29.00 vs. 36.00) and 0.000 (for 11.23 vs. 36.00), respectively. SMs invest significantly more than predicted (59.25 vs. 36.00; WSR: p=0.042) and their payoff is almost 25% higher than theoretically predicted. Somewhat surprisingly, the difference between the actual SM payoff and the predicted SM payoff is not significant at conventional levels, however (44.52 vs. 36.00; WSR: p=0.550). This is probably due to the large heterogeneity in payoffs that we discuss in more detail below. In REP-SEQ only average earnings by FMs of 27.14 are significantly

[13] We use a Wilcoxon signed-rank (WSR) test to assess whether we can reject the null hypothesis that $x_i = 36$ and $\Pi_i = 36$, respectively. We only report p-values subsequently. For the REP setting we first calculate the average investment level of each subject across all rounds and then take the average across subjects.

different from 36 (p=0.008). Overall, our findings suggest that the often observed phenomenon of overprovision of effort is a minor problem here. One reason is the prize-structure: Chowdhury et al. (2014) as well as Fallucchi et al. (2013) showed that a shared prize reduces overprovision in comparison to a winner-takes-it-all contest. Additionally, Fallucchi et al. (2013) showed that information not only about own earnings but also about the opponent's earnings leads to lower effort choices. Finally, Sheremeta (2013) points out that fixed matching and a low number of contestants reduce overprovision of effort. We find that not only in the repeated, but also in the one-shot interaction the exerted effort is close to the equilibrium prediction. In the latter, agents do not have experience from previous decisions rounds, but they had the chance to test their understanding of the game in a training program, where they were informed about the outcome of different effort combinations. Since subjects made intensive use of this possibility, we suppose that their understanding of the contest structure was quite good.

Main Hypothesis. Turning to the comparison of payoffs across treatments and roles – as formulated in our main hypothesis – we use a nonparametric Mann-Whitney U (MWU) test for the across-treatment comparisons in parts (a) and (b) of our hypothesis, and – given that decisions by FMs and SMs in SEQ are not independent from each other – a Wilcoxon signed-rank (WSR) test to evaluate part (c) of this hypothesis.[14]

The average investments and average payoffs for the OS are given in the left-hand panel of Table 2. Probably the most striking result for the OS is that the average payoff of subjects in OS−SIM is 38.40, while FMs in OS−SEQ earn only 11.23, on average. The difference between these payoffs is economically large (FM in SEQ earn less than 30% of the sum earned, on average, in SIM) and statistically highly significant (MWU: p=0.000). This result is striking because in theory the value of commitment for the FM cannot be negative when the starting point of the comparison (the SIM version of the game) is a pure strategy equilibrium. This simply follows from a revealed preference argument: In the pure strategy equilibrium of the simultaneous move version of the game, both players sit at their best reply. Thus, by choosing the same action as in the equilibrium of the SIM version of the game, the FM in a sequential-move game can induce the same outcome as in the simultaneous move version of the game. If he behaves differently, then he should do better (and not considerably worse). The comparison of the average SM payoff in OS−SEQ with the average payoff of a subject in OS−SIM yields a less pronounced and less surprising results: SMs in OS−SEQ earn more than subjects in OS−SIM (44.52 vs. 38.40), but the difference is neither large (taking again the payoff in SIM as the starting point, the increase is less than 16%) nor statistically significant (MWU: p=0.920). Still, the pronounced negative value of commitment for the FM and the non-negative value of information for the SM together yield a strong SM advantage: While the SM payoff in OS−SEQ is 44.52, the FM payoff is only 11.23. This difference is impressing (the SM payoff is almost four times as large as the FM payoff) and statistically highly significant (WSR: p=0.002).

Turning to the payoffs in the repeated (REP) version of the game – displayed in the right-hand side of Table 2 – we see that the results are qualitatively similar to those in the OS, with one notable exception: The difference between the average payoff of the FM in REP−SEQ and the average payoff of a subject in REP−SIM is much smaller than in the OS (27.14 vs. 28.22) and it is statistically no longer significant (MWU: p=0.855). Thus, the value of commitment for the FM is no longer negative but instead close to zero. On the other hand, the non-negative value of information for the SM and the pronounced SM advantage remain. Specifically, SMs in REP−SEQ earn more than subjects in REP−SIM (the respective values are 39.77 vs. 28.22), but the difference is again beyond the border of being significant in statistical terms (MWU: p=0.121). But still, SMs in REP−SEQ earn more than FMs in REP−SEQ (the respective values are 39.77 vs. 27.14). The latter difference is again economically large (SMs earn on average 46% more than FMs) and statistically highly significant (WSR: p=0.002). Fig. 2 illustrates that the difference in payoffs between the FM and the SM in REP−SEQ is present in each single decision round. However, the difference in payoffs is more pronounced in earlier rounds.[15] Result 1 summarizes our findings:

Result 1. *In the OS there is a pronounced negative value of commitment for the FM (FMs in SEQ earn significantly less than subjects in SIM), a non-negative value of information for the SM (SMs in SEQ earn on average more than subjects in SIM – the difference is not significant, however) and a pronounced second-mover advantage (SMs in SEQ earn significantly more than FMs in SEQ). In REP the negative value of commitment for the FM disappears, but the non-negative value of information for the SM and the pronounced second-mover advantage remain.*

The investments in Table 2 suggest that one reason for the pronounced second-mover advantage in the SEQ environment is that FMs lower their investments whereas SMs increase their investments in the SEQ in comparison to the SIM treatment – even though the difference in investments is statistically significant only for the SM in OS. But, as already pointed out by Chowdhury et al. (2014) the variance of choices in a contest game is rather large. As a consequence, average investment choices and payoffs may not be particularly informative. Fig. 3 presents a density plot of investment choices. As can be seen from the figure, investments are not concentrated around the benchmark prediction of 36 (vertical line). In particular, we observe both in the OS and in the REP interaction a mass point at low investment levels and a second mass point at rather high investment levels. Only investments of SMs in OS are rather dispersed over the entire strategy space, but, given

[14] Again, before applying the MWU or WSR test in REP we take average investment levels across all rounds for each individual, which implies that the number of observations is 40 for each treatment.

[15] Testing for payoff differences in later rounds we still find a statistically significant difference between the FM payoff and the SM payoff.

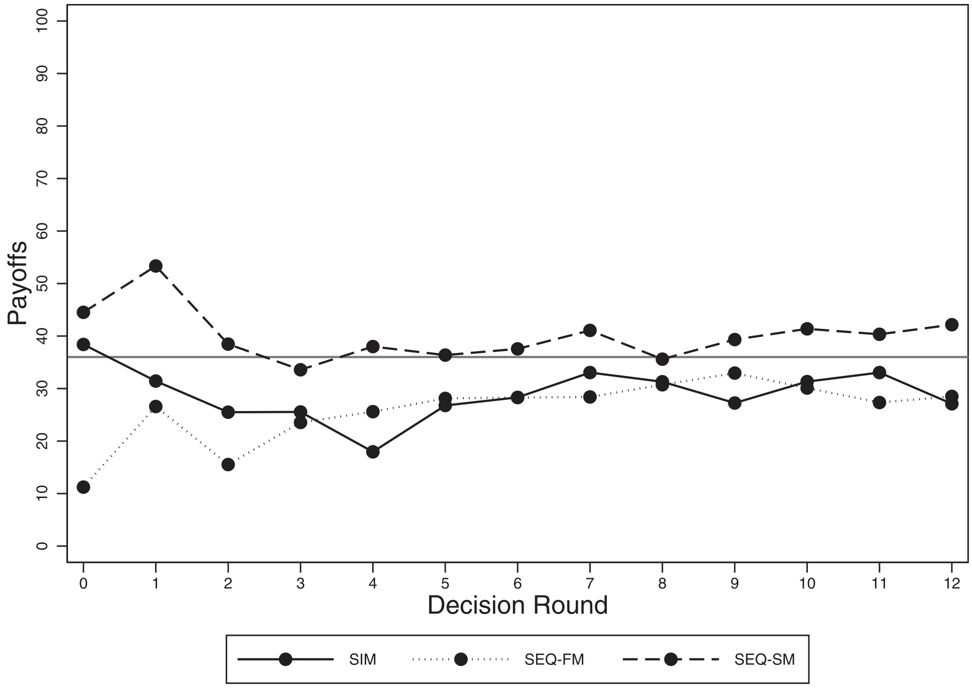

Fig. 2. Payoffs by Round *Note*: The solid horizontal line represents the standard prediction. Period 0 represents the one-shot interaction.

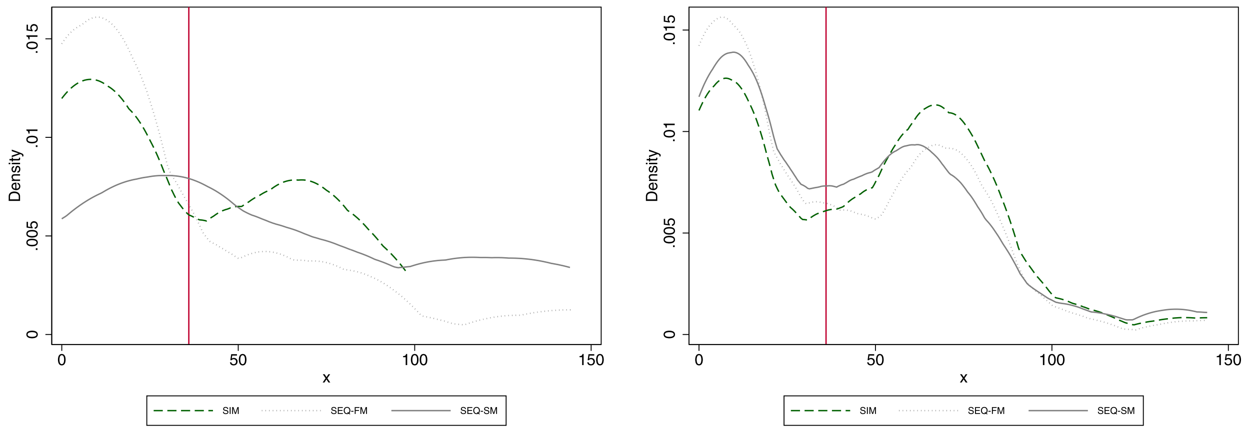

(a) One-Shot (OS)

(b) Repeated (REP)

Fig. 3. Density Plots (OS and REP) *Note*: The figure approximates the distribution of investments using an Epanechnikov kernel estimator with the optimal bandwidth. The vertical line represents the standard prediction.

that SMs invest conditional on observed FM investment, SM choices are hard to interpret in isolation. To account for the substantial heterogeneity in investment choices we subsequently disaggregate the data and investigate how this may help to explain the difference in the outcome of the simultaneous-move compared to the sequential-move version of the game as well as the asymmetry in payoffs across the FM and the SM within the sequential-move treatment.

5.2. Behavior and outcomes in the one-shot interaction (OS)

First, we concentrate on behavior of agents in the one-shot interaction (OS). Panel (a) of Fig. 3 indicates that a rather high fraction of the investment choices by the FMs in SEQ and a somewhat lower but still remarkable fraction of the investments of subjects in SIM lie below the standard prediction of 36. Based on this observation we subsequently disaggregate the data by defining any investment choice by the FM in SEQ that is below the standard prediction of 36 and any investment choice in SIM that is below 36 as a collusion attempt. Furthermore, we define successful collusion as an outcome where

Table 3
Share of Collusion Attempts and Successful Collusion.

	One-Shot (OS)		Repeated (REP)	
	Collusion attempt	Successful collusion	Collusion attempt	Successful collusion
SIM	57.5%	30%	44.38%	30.83%
SEQ	70%	37.5%	55.63%	41.88%
p-value	0.245	0.478	0.001	0.000

Note: The table presents the fraction of collusive investment choices in each treatment. In SIM an investment is classified as a collusion attempt if it is below the standard prediction of 36, while in SEQ only investments by the FM below 36 are counted as collusion attempts. Successful collusion is defined as a situation where both agents in a pairing choose an investment below 36. The reported p-values result from a standard test (z-test) for the equality of proportions between the SIM and SEQ treatment.

Table 4
Collusion Attempts and Successful Collusion (OS).

	Collusion attempt		Successful collusion		No collusion	
	Payoff	Investment	Payoff	Investment	Payoff	Investment
SIM	**38.75 / 61.68**	**8.17 / 35.39**	**67.17**	**4.83**	**26.07**	**45.93**
	(36.85) / (25.92)	(8.40) / (34.60)	(29.17)	(6.18)	(28.41)	(30.92)
SEQ-FM	**18.21**	**8.25**	**26.25**	**5.53**	**2.21**	**43.08**
	(24.16)	(9.43)	(28.98)	(7.10)	(28.45)	(42.38)
SEQ-SM	**68.07**	**49.46**	**100.08**	**12.13**	**11.19**	**87.52**
	(50.30)	(52.34)	(37.46)	(10.46)	(38.67)	(43.39)

Note: For the SIM treatment each decision in a pairing is used as an independent observation since subjects do not observe the partner's decision in this setting. Thus, there is only one entry for payoffs and investments in the case of successful collusion. In the column 'Collusion Attempt' the first entry is the average payoff (investment) of agents who invest less than 36 and the second entry is the average response of the respective opponents.

the investments of both agents in a pairing are below 36.[16] Based on those definitions, 70% of the FM investments in OS-SEQ are classified as collusion attempts, while only 57,5% of the investments in OS-SIM are classified as attempted collusion (see Table 3). Although this difference in the rate of attempted collusion is large in absolute terms, it fails to reach statistical significance (p=0.245).[17] If we look at the rate of successful collusion then there is a difference in the same direction between OS-SEQ and OS-SIM, but again the difference fails to reach significance: While 37.5% of the pairs collude successfully in OS-SEQ, only 30% of the pairs collude successfully in OS-SIM (p=0.478).

Table 4 provides average investment choices and payoffs separately for collusion attempts and for successful collusion. Comparing investments and payoffs between SIM and SEQ we find that FMs' attempts to collude are systematically exploited by SMs in SEQ, while there is (by design) no systematic exploitation of the other player within a pair in SIM: Concentrating on those (28 of 40) cases where the FM investment is classified as a collusion attempt, we find that an impressing majority of SMs (specifically 25 out of 28 or 89%) respond to the low investment by the FM by investing strictly more than the paired FM. This would be qualitatively in line with the standard best response correspondence of the SM which lies above the 45° line for low FM investments – see Fig. 1. However, we find that the majority of SMs respond to low investments by the FM by investing way too much compared to the amount predicted by the standard best reply function.[18] The difference in investments between the FM and the SM in pairs where the FM attempted to collude are economically large (8.25 vs. 49.46) and statistically highly significant (WSR: p=0.000), and they lead to a pronounced SM advantage in those pairs where the FM attempted to collude (18.21 vs. 68.07, WSR: p=0.000).

One might suspect that the pronounced SM advantage in pairs where the FM attempted to collude is an artifact that is at least in part due to the fact that a low FM investment is paired with an average SM investment.[19] However, the SM advantage in OS-SEQ is also present – and even more pronounced – in cases of successful collusion (where both investments are below 36): Of the 15 pairs who collude successfully in SEQ, 12 are pairs where the SM investment exceeds the FM investment and in all those pairs the SM earns more than the paired FM (26.25 vs. 100.08, WSR: p=0.003). By contrast, in cases where agents do not collude successfully, meaning that at least one of the agents choose an investment $\geq$ 36, we do not find a statistically significant difference in payoffs between the FM and the SM (2.21 vs. 11.19, p=0.222). Restricting the 'no collusion' outcomes to pairs where the FM started with a high investment ($>$ 36), payoffs for both agents are even

[16] Given that the number of choices 'around' the predicted value of 36 is rather low – there are only three FMs who choose an investment level between 30 and 42, for example – we decided against a classification in three subgroups.

[17] We subsequently use a standard test (z-test) to test for the equality of proportions between two groups.

[18] In fact, 18 out of 25 SM (72%) who invest more than the FM also invest more than predicted by the standard best response.

[19] Indeed, if we look at OS-SIM and concentrate on pairs where one of the players invested strictly less than the standard prediction then we also find a systematic advantage of the other player – which is simply due to the fact that a low investment by one player is matched with an average investment by the other player.

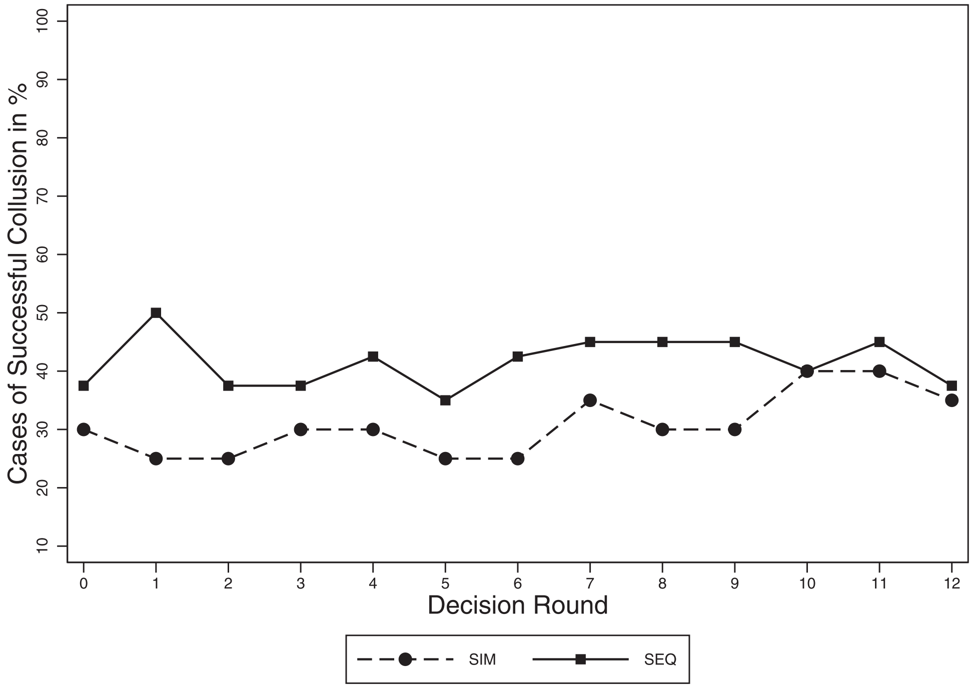

Fig. 4. Share of Cases with Successful Collusion by Round *Note*: Period 0 represents the one-shot interaction.

lower but still there is neither a systematic difference in investments (85.7 vs. 79.0, WSR: p=0.876) nor a systematic difference in payoffs (-6.92 vs. -13.78, WSR: p=0.324). Since the FMs realize a negative payoff by investing more than 36 this suggests that FMs choose low investment levels rather than high amounts to maximize their absolute payoffs, even though this implies that they earn less than their competitor in relative terms.[20] We summarize our findings for the OS as:

Result 2. ***One-Shot Interaction (OS):***

(i) An overwhelming majority of 70% of FMs in OS−SEQ and a somewhat lower but still remarkable fraction of 57% of subjects in OS−SIM invest less than the standard prediction – which is interpreted as evidence for attempted collusion.
(ii) Defining situations where both players invest less than the standard prediction as cases of successful collusion, more than 37% of the pairs collude successfully in OS−SEQ while only 30% collude successfully in OS−SIM.
(iii) In OS−SEQ, successful collusion is typically asymmetric with the SM investing significantly more and earning significantly more than the paired FM, while in OS−SIM there is (by design) no systematic asymmetry between the players in colluding pairs.
(iv) In OS−SEQ, FM investments that exceed the standard prediction typically lead to symmetric payoffs.
(v) In OS−SEQ, FMs who choose collusive investment levels typically earn more than FMs who chose high levels – which, together with (iii) and (iv) implies that FMs in OS−SEQ face a trade-off between maximizing their absolute payoffs (by investing low amounts) and minimizing disadvantageous payoff differences (by choosing high investments).

5.3. Behavior and outcomes in the repeated interaction (REP)

The bimodal distribution of investments illustrated in Fig. 3 is also evident in the repeated version of the game. Thus, we build the same subgroups as in the one-shot version of the game with collusive moves defined as investments below the standard prediction of 36.

In the repeated interaction (right-hand side of Table 3) we find a similar pattern as in the one-shot interaction regarding collusion rates. The average rate of successful collusion is 41.88% in REP−SEQ and only 30.83% in REP−SIM and this time the difference is statistically highly significant (p=0.000). Looking at the developments over time we find that the average rate of successful collusion is weakly higher in REP−SEQ than in REP−SIM in each single period and strictly higher in all periods except one (see Fig. 4). More important for our results than the rate of successful collusion is the outcome of collusion. Fig. 5 shows in panel (a) the average payoff per round when pairs realize a collusive outcome and in panel (b)

[20] The difference in payoffs between FMs who attempt to collude and FMs who invest more than 36 is statistically significant (-6.92 vs. 18.21, p=0.027).

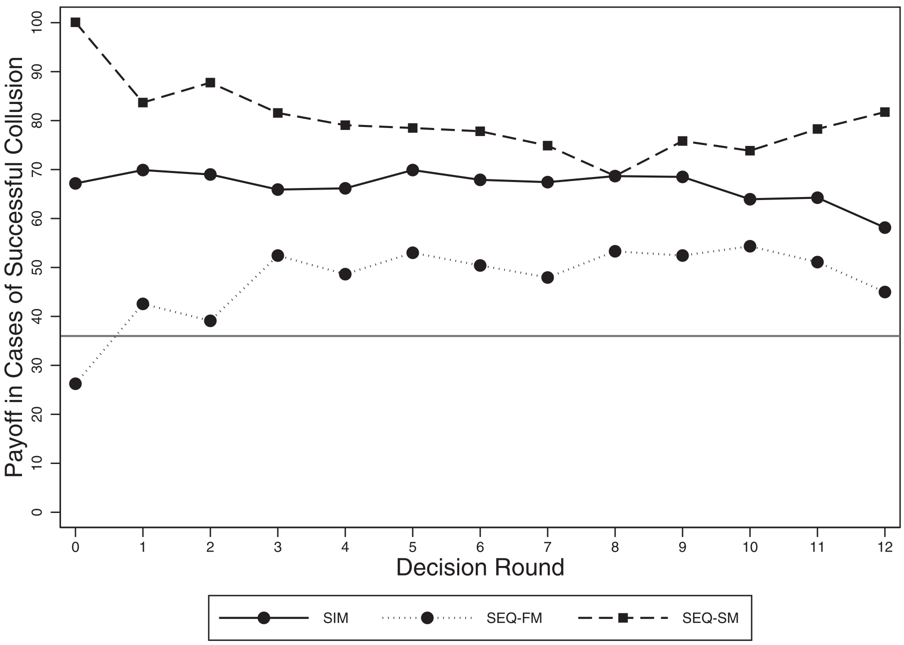

(a) Collusive Outcomes

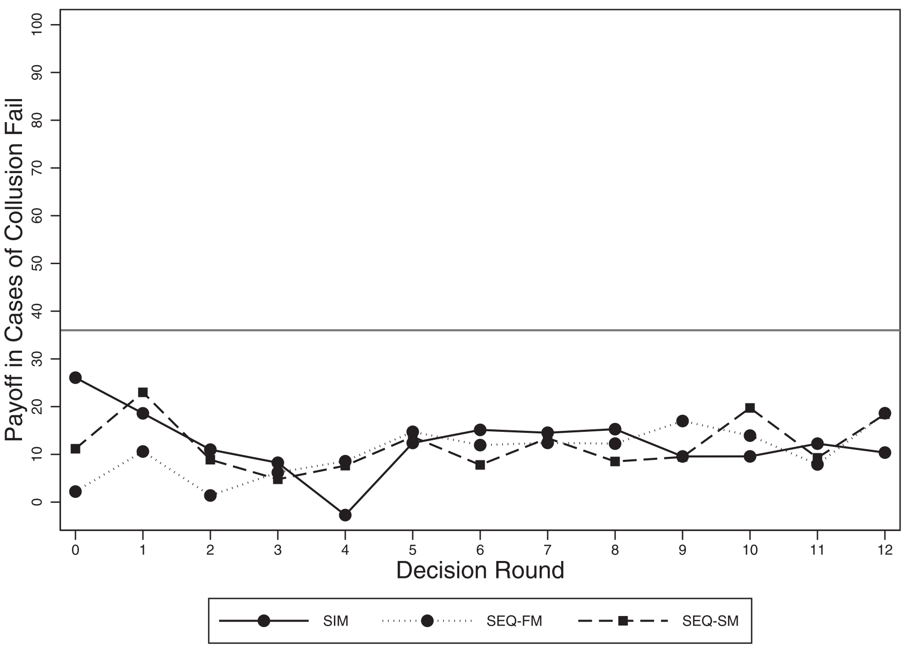

(b) Non-Collusive Outcomes

Fig. 5. Payoffs by Round for Collusive and Non-Collusive Outcomes *Note*: Panel (a): collusive outcomes are situations where both agents in a pairing choose an investment below 36; Panel (b): non-collusive outcomes are situations where at least one agent in a pairing chooses an investment greater or equal 36. Period 0 represents the one-shot interaction.

Table 5
Collusion Attempts and Successful Collusion (REP).

	Collusion attempt		Successful collusion		No collusion	
	Payoff	Investment	Payoff	Investment	Payoff	Investment
SIM	**39.81 / 54.02**	**13.12 / 37.06**	**62.80**	**9.19**	**12.59**	**51.10**
	(27.22) / (16.08)	(9.97) / (29.32)	(15.38)	(8.01)	(16.95)	(33.74)
SEQ-FM	**34.66**	**10.85**	**49.41**	**8.54**	**8.95**	**53.90**
	(20.33)	(10.10)	(22.90)	(8.21)	(20.26)	(28.61)
SEQ-SM	**67.09**	**31.40**	**74.62**	**11.43**	**12.10**	**69.04**
	(23.22)	(24.73)	(24.29)	(9.24)	(22.70)	(29.81)

Note: SIM treatment, column attempted collusion: first entry is the average payoff (investment) of agents who invest less than 36 and the second entry is the average response of the respective opponent. For the mean values we first aggregate observations on the individual level. See also footnote 21.

when pairs do not successfully establish collusion. From the figure it is obvious that we find the same qualitative pattern as in the OS: SMs earn more than FMs when both agents invest less than 36; in contrast, we do not find a difference in payoffs for high investments, as illustrated in panel (b) of Fig. 5. In Table 5 we provide the corresponding average payoffs and investment choices.[21] Comparing average payoffs between roles in REP-SEQ confirms the difference in payoffs depicted in the figure: The payoffs of FMs and SMs in REP-SEQ are significantly different for pairs who successfully collude (49.41 vs. 74.62, p=0.003) – whereas for pairs who do not collude successfully we cannot differentiate between roles (8.95 vs. 12.10, p=0.1767). Thus, we reach the same conclusion as in the OS – the main difference being that the asymmetry in payoffs is less pronounced in the repeated interaction than in the one-shot.

Comparing the outcome in REP-SIM with REP-SEQ over periods, we do not find a systematic bias to the advantage of one of the players in collusive outcomes in REP-SIM whereas the payoffs of agents in a pairing in REP-SEQ are quite asymmetric to the advantage of the SM. This asymmetry is also visible in Fig. 6 which plots investments of pairings separately for each round in SIM in panel (a) and for SEQ in panel (b). Whereas investments in REP-SIM do not follow a systematic pattern, most observations in REP-SEQ lie above the 45° line. The latter finding implies that SMs in REP-SEQ invest systematically more than the paired FM. If we look at numbers, we find for the SIM version of the game that in those pairs who realize a collusive outcome in 48.65% of the cases both players invest the same amount and in the rest of the cases they choose different investment levels. By contrast, in REP-SEQ colluding FMs and SMs invest the same amount in 43.78% of the cases and in only 4.48% of the cases the SM invests less; thus, in most of the cases (specifically, in 51.74% of the cases) the SM invests more than the paired FM.

Taken together, we qualitatively find the same pattern in REP as in OS: For non-collusive outcomes payoffs can hardly be distinguished between the SIM and the SEQ treatment whereas for collusive outcomes the average payoff per round in the SIM treatment lies between the average FM and the average SM payoff in the SEQ treatment. However, even though FMs earn less than SMs in the case of attempted collusion and successful collusion, the payoff of FMs is higher compared to the situation where they do not realize a collusive outcome.[22] Thus, as in the OS, relative payoffs of FMs are low but absolute payoffs are high in cases where the FM chooses a low rather than a high investment level. In contrast to the OS, the number of collusive outcomes is higher in REP-SEQ and the payoff difference between the FM and SM is lower in cases where they collude successfully. This may be due to the fact that in REP with fixed partner matching the FM has the possibility to retaliate a SM who exploits her superior position too much.

Until now we defined as a collusive outcome a situation where both agents in a pairing invest strictly less than the predicted amount of 36. Since we use a fixed-partner matching in the repeated version of the game, we now take a different look at the data by focusing on pairs who reach kind of a collusive agreement over periods. For this purpose we define a collusive pair as a pairing which realizes a collusive outcome in at least 6 of the 12 interaction rounds. Based on this classification, we find a substantial difference in the number of collusive pairs between SIM and SEQ. In SEQ, 47.6% of the pairs collude in more than 6 decision rounds and in SIM the corresponding fraction is only 30%. Figs. 7 and 8 compare the investment behavior and payoffs for collusive pairs in SIM and SEQ. From the figures it is obvious that SMs in REP-SEQ invest more compared to FMs – leading to the observed asymmetry in payoffs. This systematic bias to the advantage of one player cannot be found in the SIM treatment. This confirms our previous findings and we therefore conclude:

Result 3. ***Repeated Interaction (REP)***

(i) The investments of more than 55% of the FMs in REP-SEQ are classified as attempted collusion while the corresponding fraction in REP-SIM is (with about 45%) significantly lower.

[21] Given that the same individual may invest more than the benchmark prediction in some decision rounds and less in others we abstract from experience-driven changes in behavior across decision rounds and compute averages across all rounds in which one agent invests less (more) than the benchmark prediction (implying that between 1 and 12 decision-round observations determine averages on the individual level). A non-parametric WSR-test then delivers the p-values provided in parentheses.

[22] A simple *t*-test comparing the payoff of the FM in the case of successful collusion to the "no collusion" case (49.41 vs. 8.95) yields a statistically highly significant difference (p=0.000).

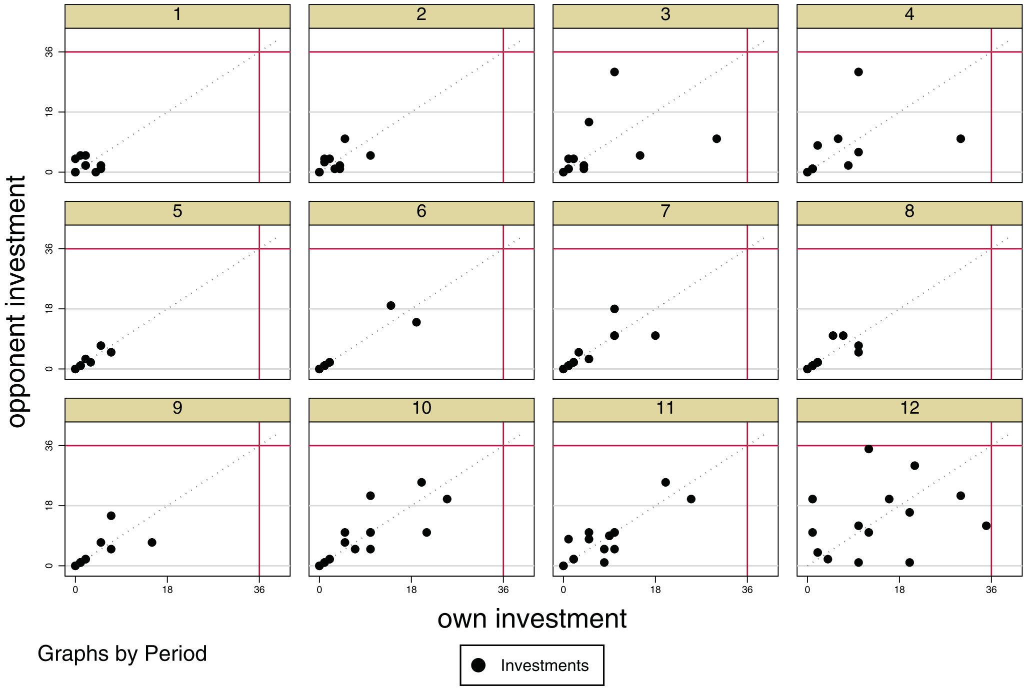

(a) SIM

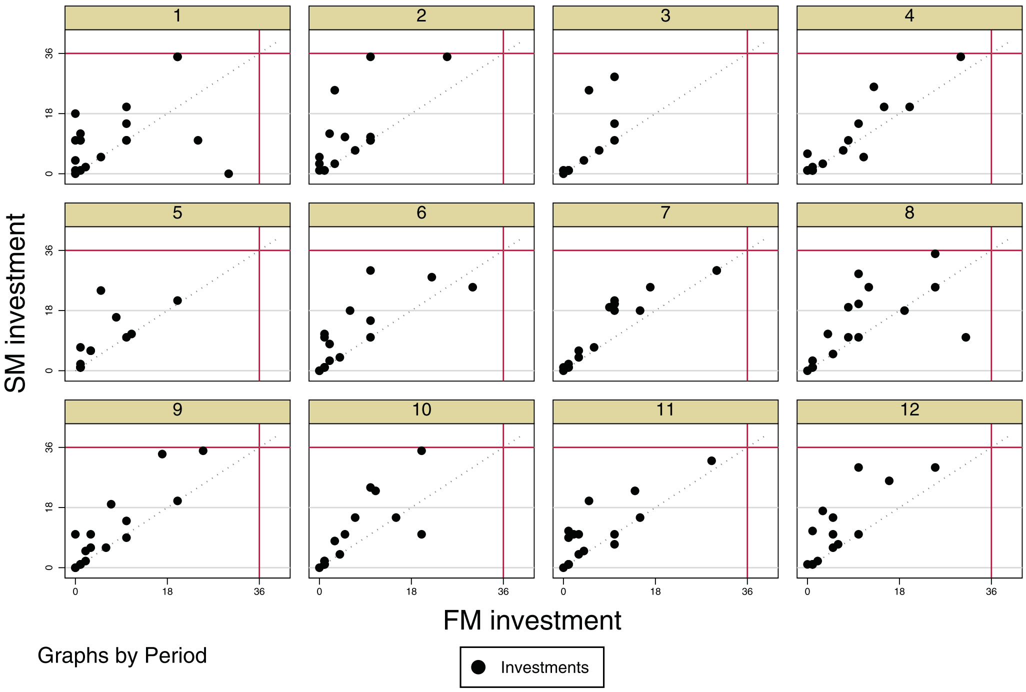

(b) SEQ

Fig. 6. Cases of Successful Collusion in REP-SIM and REP-SEQ by Round *Note*: Both panels show the respective investment choices for cases of successful collusion separately for each decision round in the repeated game.

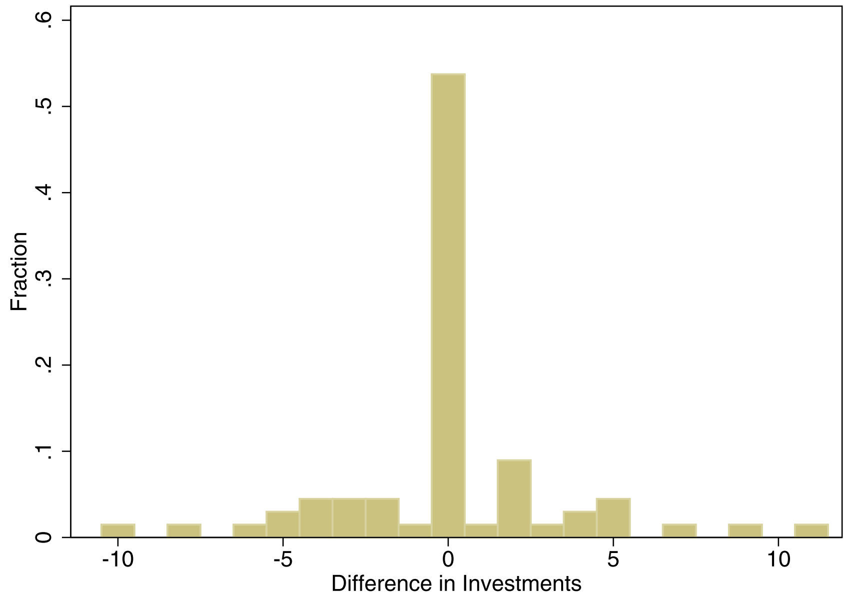

(a) SIM

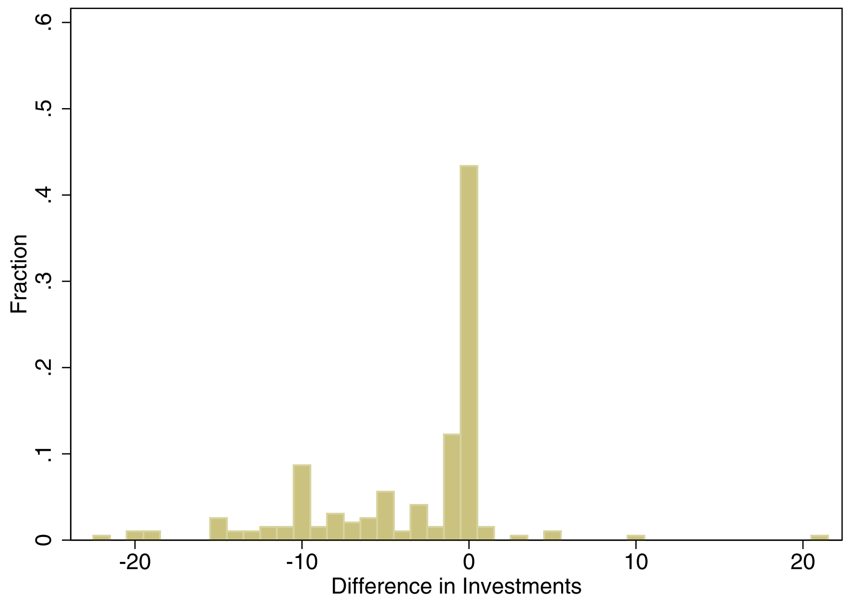

(b) SEQ

Fig. 7. Difference in Investments of Collusive Pairs in REP *Note*: Both panels show the difference in investment choices (own investment - opponent investment) for collusive pairs. In the SEQ treatment the investment difference is calculated as $x_{SEQ\text{-}FM} - x_{SEQ\text{-}SM}$.

(ii) *Looking at the rate of successful collusion, more than 40% of the pairs collude successfully in REP-SEQ, while the corresponding fraction for REP-SIM is (with about 30%) significantly lower.*
(iii) *Similar as in OS-SEQ, in REP-SEQ successful collusion is typically asymmetric with the SM investing more and earning more than the paired FM.*
(iv) *Similar as in OS-SEQ, in REP-SEQ FM investments that exceed the standard prediction typically lead to symmetric payoffs.*
(v) *Similar as in OS-SEQ, in REP-SEQ FMs who choose collusive investment levels typically earn more than FMs who chose high levels – which, together with (iii) and (iv) again implies that FMs face a trade-off between maximizing their absolute payoffs and minimizing disadvantageous payoff differences.*

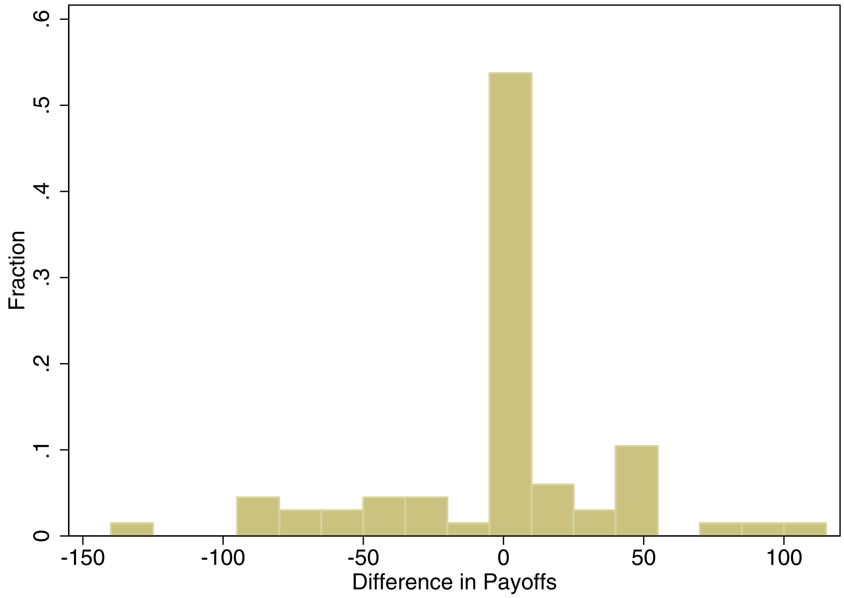

(a) SIM

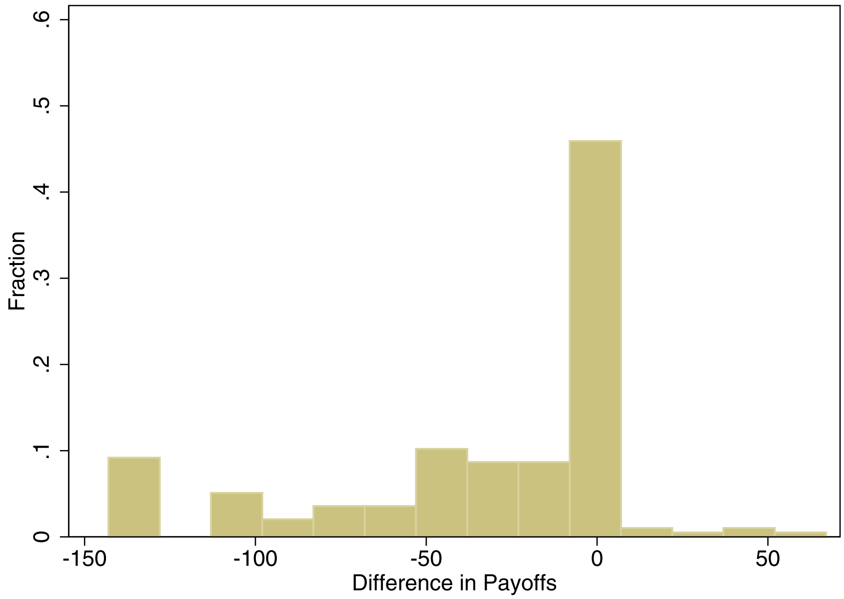

(b) SEQ

Fig. 8. Difference in Payoffs of Collusive Pairs in REP *Note*: Both panels show the difference in payoffs (own payoff - opponent payoff) for collusive pairs. In the SEQ treatment the payoff difference is calculated as $\Pi_{SEQ\text{-}FM} - \Pi_{SEQ\text{-}SM}$.

6. Explaining observed behavior in SEQ

The experimental data presented in the previous sections is inconsistent with the standard theory assuming own money maximizing behavior. But, observed behavior may be consistent with theories of other-regarding preferences which have been used to explain observed behavior in related work – see Huck et al. (2001), Fonseca (2009), or Hoffmann and Kolmar (2017) for example. In this section we discuss the consistency of our data with the predictions made by some prominent models of other-regarding preferences. Then we turn to the expectation data.

Observed Behavior and Other-Regarding Preferences. The concepts of negative reciprocity (as modelled e.g. by Rabin (1993)) and inequality aversion (as modelled by Fehr and Schmidt (1999) and by Bolton and Ockenfels (2000)) both predict that

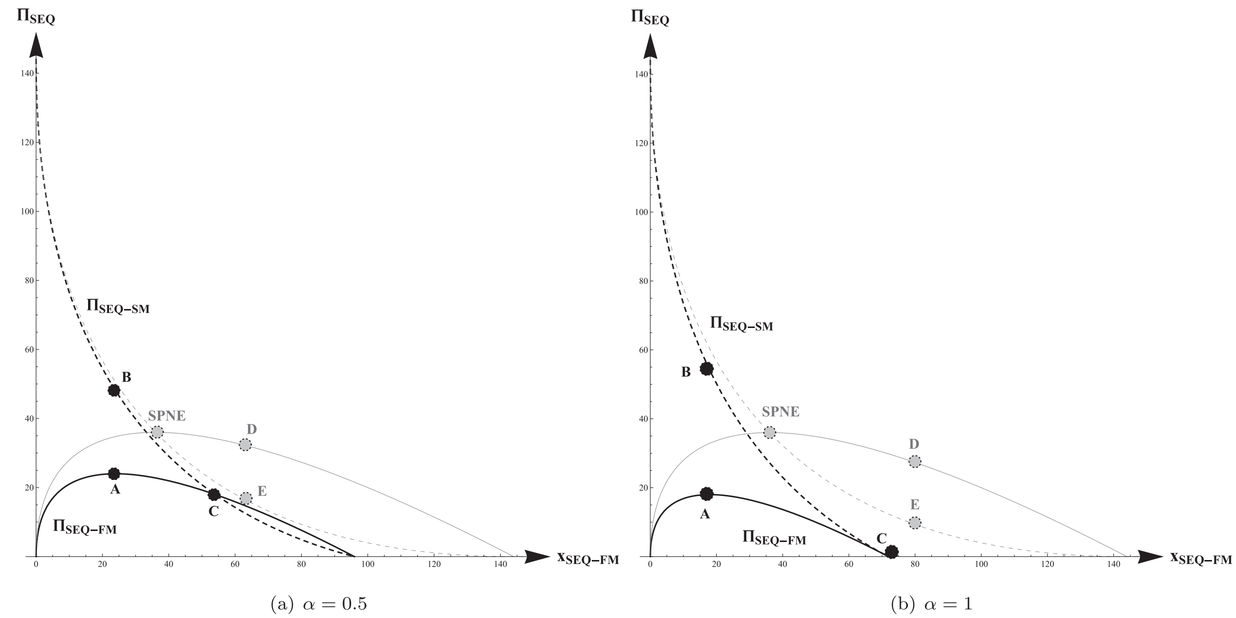

Fig. 9. FM and SM Payoffs with Relative Income Preferences *Note*: Panel (a) plots FM and SM payoffs for $P = 144$ as a function of FM-investment, $x_{\texttt{SEQ-FM}}$, under the assumptions that the SM best responds to $x_{\texttt{SEQ-FM}}$, and that players might care about relative standing ($\alpha = 0.5$). Panel (b) plots the same situation as panel (a), but assumes that relative standing is more important ($\alpha = 1$).

SMs may be willing to sacrifice some of their own payoff to punish the FM for large investments. At the same time, these theories predict that SMs invest at most the standard best-response amount if the FM invests less than the theoretical prediction.[23] This is in contrast to what we observe in the one-shot: The majority of SMs invests more than the best-response amount here, independent of whether the FM invests more or less than the benchmark prediction. While this evidence is inconsistent with the concepts of reciprocity and of inequality aversion, it is in line with the relative income hypothesis, – as formulated by Duesenberry (1949).[24] To investigate in how far equilibrium predictions change in a model where both the FM and the SM potentially not only care about their absolute payoff, but also about their relative standing in comparison to the opponent, we consider a simple behavioral model and investigate in how far this model helps to organize our experimental data.[25]

Behavioral Model. Assume that agents decide in accordance with an additive comparison utility function as in Alpizar et al. (2005).[26] Using the same notation as in the benchmark model, the objective of firm i who competes against firm $-i$ then reads

$$\max_{x_i \geq 0} \{\Pi_i(x_i, x_{-i}) - \alpha_i \cdot \Pi_{-i}(x_{-i}, x_i)\} \tag{6}$$

where $\alpha_i \geq 0$ measures the importance of relative standing as compared to absolute earnings for firm i. In particular, firm i cares exclusively about its own payoff if $\alpha_i = 0$, and instead tries to maximize the (positive) difference between own and opponent payoff if $\alpha_i = 1$. Given (6), the best-response function of firm i for investments x_{-i} by the competitor reads

$$\text{BR}_i(x_{-i}) = \max\{\sqrt{(1+\alpha_i)x_{-i}P} - x_{-i}, 0\}.$$

Best-responses are strictly increasing in the importance of relative standing α_i, i.e., firms who care about their relative standing invest higher amounts for any given opponent investment x_{-i} than firms who focus entirely on their own payoff.

Fig. 9 shows how payoffs are affected by the parameter α_i. In particular, panel (a) plots the FM and the SM payoff as a function of the FM-investment, $x_{\texttt{SEQ-FM}}$, under the assumptions that the SM best responds to $x_{\texttt{SEQ-FM}}$, and that players either do not care about relative standing ($\alpha = 0$) or care about relative standing to some extent ($\alpha = 0.5$). Panel (b) plots the same situation as panel (a), but allows the concern for the relative standing to be more important ($\alpha = 1$). In both

[23] To see this, note that the benchmark prediction leads to an egalitarian allocation not only in SIM but also in SEQ.

[24] The crucial difference between the relative income hypothesis and the aforementioned theories of other-regarding preferences is that decision makers enjoy being ahead of others in the former – i.e., there is no reference point at the point of equal division as in the latter.

[25] Concerns for relative standings in contests may be rationalized by evolutionary game theory approaches. As shown by Guse and Hehenkamp (2006), behavior consistent with preferences of the aforementioned type follows from evolutionary stable strategies. Moreover, Leininger (2009) shows that evolutionary stable preferences turn out to be negatively interdependent.

[26] Additive comparison utility functions are also employed by Akerlof (1997), Corneo and Jeanne (1997), and Knell (1999).

Table 6
Expected SM-Behavior conditional on FM-Investment.

OS-SEQ	Obs.	$E_{FM}[x_{SEQ-SM}]$ vs. x_{SEQ-FM}			$E_{FM}[x_{SEQ-SM}]$ vs. $x_{BR}(x_{SEQ-FM})$		
		$x^E_{SM} < x_{FM}$	$x^E_{SM} = x_{FM}$	$x^E_{SM} > x_{FM}$	$x^E_{SM} < x_{BR}$	$x^E_{SM} = x_{BR}$	$x^E_{SM} > x_{BR}$
$x_{SEQ-FM} < 36$	28	7.1%	42.9%	**50.0%**	**71.4%**	–	28.6%
$x_{SEQ-FM} \geq 36$	12	16.6%	41.7%	41.7%	–	25.0%	**75.0%**
REP-SEQ	Obs.	$E_{FM}[x_{SEQ-SM}]$ vs. x_{SEQ-FM}			$E_{FM}[x_{SEQ-SM}]$ vs. $x_{BR}(x_{SEQ-FM})$		
		$x^E_{SM} < x_{FM}$	$x^E_{SM} = x_{FM}$	$x^E_{SM} > x_{FM}$	$x^E_{SM} < x_{BR}$	$x^E_{SM} = x_{BR}$	$x^E_{SM} > x_{BR}$
$x_{SEQ-FM} < 36$	267	9.7%	42.3%	**48.0%**	**79.8%**	1.9%	18.3%
$x_{SEQ-FM} \geq 36$	213	**38.5%**	25.8%	35.7%	12.7%	4.7%	**82.6%**

panels the grey curves represent the case where the respective player does not care for the relative standing while the black curves represent the case where he does. Consider first the point labeled 'SPNE' (identical in both panels) where the solid gray and the dashed gray curve intersect. As the comparison with Fig. 1 reveals, this point represents the symmetric (subgame perfect) Nash equilibrium in the benchmark model where $\alpha = 0$ holds both for the FM and for the SM. Consider next a situation where the SM cares about relative standing ($\alpha_{SM} > 0$) whereas the FM does not ($\alpha_{FM} = 0$), and where this is common knowledge. The equilibrium payoff of the FM then decreases to **A**, while equilibrium earnings of the SM increase to **B**. Intuitively, the positive α_{SM} implies that the SM invests more than in the benchmark model for any amount invested by the FM, such that it becomes optimal for the FM to invest less than in the benchmark case. The comparison of panels (a) and (b) reveals that these changes are even more pronounced in the case where $\alpha_{SM} = 1$ rather than $\alpha_{SM} = 0.5$. Alternatively, one might assume that both the FM and the SM care about relative standing ($\alpha_{FM} = \alpha_{SM} > 0$). Compared to the benchmark model with $\alpha = 0$, concerns for relative standing induce higher investments by both the FM and the SM, such that the payoff of both firms decreases to point **C**. The comparison of panels (a) and (b) again reveals that changes in payoffs are increasing in the importance of relative standing α. Finally, points **D** and **E** depict a situation where the FM cares about relative standing ($\alpha_{FM} > 0$), but the SM does not ($\alpha_{SM} = 0$) and where this is common knowledge. In particular, the figure shows that FMs earn more than SMs in this case, but less than they would in the benchmark where $\alpha = 0$ holds both for the FM and for the SM.

Observed Behavior and the Behavioral Model. Consider first observed behavior in the one-shot version of SEQ. According to the behavioral model, the reason for the observation that for any given FM investment SMs invest higher amounts than predicted by the benchmark model is that SMs care about their relative standing – i.e., α_{SM} is strictly greater than zero for the majority of SMs. Assuming that FMs correctly anticipate that SMs not only care about their absolute payoff, but also about their relative standing they face a trade-off: by investing less than the benchmark prediction they earn a high absolute amount – but they also earn less than the paired SM; by contrast, by investing more than the standard prediction they earn a low absolute amount – but they might be better off in relative terms. Confronted with this trade-off, FMs who only care about the own absolute payoff are predicted to invest *less* than the benchmark amount. The observation that more than two-thirds of all FMs invest less than the benchmark prediction in OS-SEQ, while most SM invest more than the standard best-response amount is thus consistent with the behavioral model and it delivers an outcome where FMs earn less than SMs – depicted by points **A** and **B** in Fig. 9. At the same time, the model predicts that FMs who care about their relative standing may invest *more* than the benchmark amount, which can rationalize why about one-third of all FMs behaves this way in OS-SEQ. FMs who invest more than the benchmark amount might either earn the same amount as the SM if their paired SM cares about relative standing as well – as in point **C** in Fig. 9 – or more than the SM if the paired SM is exclusively interested in her own absolute payoff – as in points **D** and **E** in Fig. 9. In line with these predictions, we observe that FMs who invest more than the benchmark amount earn (weakly) more than their paired SMs. Finally, the behavioral model predicts – in line with what we observe in the experimental data – that FMs who invest low amounts earn significantly more in *absolute* terms than FMs who invest high amounts – a comparison of points **A** and **C** reveals that this is the case. At the same time, they are worse-off compared to their second-moving opponents in *relative* terms. In this sense, the behavioral model organizes the experimental data quite well.

Expected SM-Behavior and the Behavioral Model. To investigate whether the bi-modal distribution of FM investment choices can really be rationalized by the behavioral model, we subsequently investigate the expectations of FMs about the responses of SMs to their investment choices.[27] Table 6 provides a classification of FM-expectations conditional on FM-investment choices both for OS and the REP interaction. Consider first the expectations of FMs who invest *more* than the predicted amount. These FMs expect above standard best-response investments by the SM in 75% (OS) and 82.6% (REP) of all cases, respectively, such that expected SM investment is significantly higher than the best-response amount in both parts of the

[27] In the experiment FMs were asked for their expectations about SM investments (not incentivized).

experiment.[28] In terms of the behavioral model, behavior and expectations are thus consistent with an equilibrium in which both the FM and the SM care about their relative standing.

Consider next the expectations of FMs who invest *less* than the predicted amount. While FMs who invest below benchmark amounts expect that SM investment exceeds own investment in 50% (OS) and 48% (REP) of all cases in the one-shot and across repeated interactions, respectively – implying that they expect to earn less than the paired SM – FMs expect below (rather than above) standard best-response investments by the paired SM in 71.4% (OS) and 79.8% (REP) of all cases in the respective part of the experiment – see Table 6. Is this still in line with the mechanism that the behavioral model proposes to rationalize below benchmark investments by the FM? Our data is insufficient to answer this question. This is so because we have elicited only the expectation of the FM about the response of the SM to the actual choice of the FM (while to answer this question we would need the expectation of the FM about the response of the SM to any possible choice of the FM). If we assume that all FMs have the same expectations about the reply of the SM to above and below benchmark investments of the FM then the data in Table 6 is perfectly in line with the explanation of the behavioral model for our main results: FMs expect above standard best-response investments by the SM for investments above the standard prediction and they expect investments that exceed the own investment for investments below the standard prediction. Given those expectations about SM responses, FMs face the discussed trade-off: by investing less than the benchmark prediction they expect to earn a high absolute amount – but they also expect to earn less than the paired SM; by contrast, by investing more than the standard prediction they expect to earn a low absolute amount – but they also anticipate that they might be better off in relative terms.

How does the expectation data relate to our findings about collusion attempts? FM expectations suggest that many FMs expect to establish a collusive outcome if they invest less than the benchmark prediction, in the sense that the paired SM responds by investing less than the own profit maximizing best-response amount. In particular, this is what 71.4% of all FMs who invest less than the benchmark prediction expect in OS-SEQ. In REP-SEQ, the corresponding value is close to 80.0% of all cases – see Table 6 for details. As discussed in Section 5.2, we find some but not much evidence that SMs respond to attempts to establish collusion in the one-shot. The majority of SMs invest more than the paired FM and of those who invest more than the FM, 72% invest even more than the best-response amount. This is very different in part 2 of the experiment where subjects interact repeatedly. Here, SMs invest *less* (and not more) than the standard best-response amount in response to below benchmark investments by the FM in the majority of cases. Said differently, there is strong evidence for successful collusion in part 2 of the experiment where subjects interact repeatedly, but only limited evidence in part 1 where subjects interact only once. However, even in part 1, more than 37% of the pairs collude successfully in OS-SEQ, while only 30% collude successfully in OS-SIM. The difference is more impressing in part 2 where more than 40% of the pairs collude successfully in REP-SEQ, while the corresponding fraction for REP-SIM is – with about 30% – significantly lower. This suggests that the sequential structure in SEQ helps players to establish a collusive outcome. As discussed earlier, collusive outcomes are typically asymmetric – to the advantage of the SM. From Table 6 we can see that many FMs anticipate this. In particular, close to 50% of FMs who invest less than the benchmark prediction anticipate that the paired SM will invest more than the FM has invested. Said differently, many FMs expect a cooperative response (= below best-response investment) by their paired SM, but they also expect that SM investment exceeds FM investment, implying that the SM earns more than the FM.

7. Concluding remarks

This article has presented the results of an experimental contest game where two symmetric agents compete either simultaneously (in SIM) or sequentially (in SEQ) for the share of a prize. Our findings suggest that moving second rather than first is an advantage in strategic interactions even if material incentives for pre-commitment are absent in the theoretical benchmark. This result holds independently of whether agents interact only once (in the OS) or with a fixed partner over 12 decisions rounds (in REP). Specifically, for the OS we have identified a pronounced negative value of commitment for the first mover (first movers in SEQ earn significantly less than subjects in SIM), a non-negative value of information for the second mover (second movers in SEQ earn on average more than subjects in SIM) and a pronounced second-mover advantage (second movers in SEQ earn significantly more than first movers in SEQ). In REP the negative value of commitment for the first mover disappears, but the non-negative value of information for the second mover and the pronounced second-mover advantage remain.

A closer look at the data reveals that the second mover advantage in SEQ is in large parts due to the fact that the majority of first movers invests less than standard theory would predict – which is interpreted as evidence for attempted collusion. Indeed, we find that collusion attempts are more frequent in SEQ than in SIM (70% vs. 57% for the OS and 55% vs. 45% for REP) and that cases of successful collusion (where both players in a pair invest less than standard theory would predict) are also more frequent in SEQ than in SIM (38% vs. 30% for the OS and 40% vs. 30% for REP). More important for the understanding of the second-mover advantage than the frequency of collusion attempts and of collusive outcomes is their consequence. In SEQ collusion attempts by the first mover are systematically exploited by the second mover. This is the case independently of whether the outcome is classified as successful collusion, or not. Indeed, collusive outcomes

[28] In the OS, we have 65.08 vs. 25.67; WSR: p=0.004; in REP, we obtain 56.02 vs. 27.90; WSR: p=0.001.

in SEQ are typically asymmetric – with the second mover investing considerably more and earning considerably more than the paired first mover.[29]

So, why do many first movers choose low investment levels if they correctly anticipate that collusion attempts are systematically exploited by second movers? A behavioral model where players do not only care for their absolute but also for their relative payoff predicts the following trade-off for FMs in SEQ: By investing les than the benchmark prediciton first movers earn a high absolute amount – but they also earn less than the paired second mover; by contrast, by investing more than the benchmark prediction FMs earn a low absolute amount – but they might be better off in relative terms. The behavioral model then predicts that first movers who are confronted with this trade-off will invest less than the benchmark amount if they care little for the relative payoff, but more than the benchmark amount if concerns for the relative standing become more important. And why does the negative value of commitment for the first mover in the one-shot version of SEQ vanish in the repeated interaction? This is arguably due to the fact that the first mover can retaliate a undesirable move by the paired second mover in REP−SEQ but not in OS−SEQ.

The second-mover advantage we have identified seems particularly valuable in strategic contexts where behavior is heterogeneous – for instance, because players have heterogeneous distributional preferences. In such contexts the possibility to observe the behavior of the opponent and to react to it gives the second movers the power to ultimately determine (relative) payoffs through their investment choices. By contrast, first movers face strategic uncertainty and in the present context they seem to react to the strategic uncertainty by acting more cooperatively than predicted by the benchmark model.

It is worth noting that our findings might help to explain why the empirically observed first-mover advantage in the quantity-competition experiments by Huck et al. (2001) is quantitatively smaller than predicted by the quantity-competition model, while the observed second-mover advantage in the price-competition experiments by Kübler and Müller (2002) is larger than predicted by the price-competition model. In both cases the behavior of players is heterogeneous implying that first movers face strategic uncertainty while second movers have the power to ultimately determine (relative) payoffs through their market decisions. While the evidence in the related literature is consistent with the hypothesis that the inherent advantage to move second, isolated in the present paper, works on top of strategic considerations in other settings, we believe that it is an interesting question for future work to investigate in more detail how the inherent advantage to move second interacts with strategic first- or second-mover advantages in the theoretical benchmark.

Acknowledgements

We are indebted to an associate editor and two anonymous referees for their detailed reports. Their constructive comments and suggestions helped to improve the paper a lot. This work was supported by the Austrian Science Fund (FWF) through special research area grant SFB F63, as well as through grant numbers P22669, P26901 and P27912.

References

Akerlof, G.A., 1997. Social distance and social decisions. Econometrica 65, 1005–1027.
Alpizar, F., Carlsson, F., Johansson-Stenman, O., 2005. How much do we care about absolute versus relative income and consumption? J. Econ. Behav. Org. 56, 405–421.
Bass, F.M., Krishnamoorthy, A., Prasad, A., Sethi, S., 2005. Generic and brand advertising strategies in a dynamic duopoly. Market. Sci. 24, 556–568.
Bassan, B., Gossner, O., Scarsini, M., Zamir, S., 2003. Positive value of information in games. Inte. J. Game Theory 32 (1), 17–31.
Beard, F.K., 2010. Comparative advertising wars: an historical analysis of their causes and consequences. J. Makromarket. 30, 270–286.
Beard, F.K., 2011. Competition and combative advertising: an historical perspective. J. Makromarket. 31, 387–402.
Bolton, G.E., Ockenfels, A., 2000. ERC: a theory of equity, reciprocity, and competition. Am. Econ. Rev. 90, 166–193.
Bulow, J., Geanakoplos, J., Klemperer, P., 1985. Holding idle capacity to deter entry. Econ. J. 95, 178–182.
Chowdhury, S.M., Sheremeta, R.M., Turocy, T.L., 2014. Overbidding and overspreading in rent-seeking experiments: cost structure and prize allocation rules. Game Econ. Behav. 87, 224–238.
Corneo, G., Jeanne, O., 1997. Conspicious consumption, snobbism and conformism. J. Public Econ. 66, 55–71.
Dechenaux, E., Kovenock, D., Sheremeta, R.M., 2015. A survey of experimental research on contests, all-pay auctions and tournaments. Exp. Econ. 18 (4), 609–669.
Dixit, A., 1987. Strategic behavior in contests. Am. Econ. Rev. 77, 891–898.
Duesenberry, J., 1949. Income, Saving and the Theory of Consumer Behavior. Harvard University Press, Cambridge MA.
Eisenkopf, G., Teyssier, S., 2014. Principal-agent and peer relationships in tournaments. Manag. Decis. Econ. (forthcoming).
Fallucchi, F., Renner, E., Sefton, M., 2013. Information feedback and contest structure in rent-seeking games. Eur. Econ. Rev. 64, 223–240.
Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition, and cooperation. Q. J. Econ. 114, 817–868.
Fischbacher, U., 2007. Z-Tree, Zurich toolbox for readymade economic experiments. Exp. Econ. 10, 171–178.
Fonseca, M., 2009. An experimental investigation of asymmetric contests. Int. J. Ind. Org. 27, 285–591.
Friedman, L., 1958. Game-Theory models in the allocation of advertising expenditures. Oper. Res. 6, 699–709.
Gal-Or, E., 1985. First mover and second mover advantages. Int. Econ. Rev. 26, 649–653.
Gossner, O., Mertens, J.-F., 2001. The value of information in zero-sum games.
Greiner, B., 2004. An online recruitment system for economic experiments. In: Kremer, K., Macho, V. (Eds.), Forschung und Wissenschaftliches Rechnen 2003, pp. 79–93. GWDDG Bericht der 63. Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen
Guse, T., Hehenkamp, B., 2006. The strategic advantage of interdependent preferences in rent-Seeking contests. Public Choice 129, 323–352.
Güth, W., Müller, W., Spiegel, Y., 2006. Noisy leadership: an experimental approach.. Game Econ. Behav. 57 (1), 37–62.

[29] Our finding that collusion attempts by first movers are systematically exploited by second movers has a close parallel in Weimann et al. (2000). In a 'winner-takes-it-all' contest with a contest technology that yields a first-mover advantage in the theoretical benchmark, the authors find that first-movers show cooperative behavior by bidding low and second-movers exploit this attempt by choosing their best response. As a consequence, second movers often earn more than first movers

Hoffmann, M., Kolmar, M., 2017. Distributional preferences in probabilistic and share contests. J. Econ. Behav. Org. 142, 120–139.
Hoffmann, M., Rota-Graziosi, G., 2012. Endogenous timing in general rent-seeking and conflict models. Game. Econ. Behav. 75 (1), 168–184.
Hoffmann, M., Rota-Graziosi, G., 2018. Endogenous timing in contests. In: Corchon, L.C., Marini, M.A. (Eds.), Handbook of Game Theory and Industrial Organization, Volume II: Applications, 2, p. 147. Chapter 7
Hoffmann, M., Rota-Graziosi, G., 2019. Endogenous timing in the presence of non-monotonicities. Canad. J. Econ./Revue Canad.d'Écon., forthcoming.
Huck, S., Müller, W., Normann, H.-T., 2001. Stackelberg beats cournot: on collusion and efficiency in experimental markets. Econ. J. 111, 749–765.
Knell, M., 1999. Social comparison, inequality, and growth. J. Inst. Theoretical Econ. 155, 664–695.
Kopel, M., Löffler, C., 2008. Commitment, first-mover-, and second-mover advantage. J. Econ. 94 (2), 143–166.
Kübler, D., Müller, W., 2002. Simultaneous and sequential price competition in heterogeneous duopoly markets: experimental evidence. Int. J. Ind. Org. 20, 1437–1460.
Lehrer, E., Rosenberg, D., 2006. What restrictions do bayesian games impose on the value of information? J. Math. Econ. 42 (3), 343–357.
Leininger, W., 2009. Evolutionary stable preferences in contests. Public Choice 140, 341–356.
Long, N., Vousden, N., 1987. Risk-averse rent seeking with shared rents. Econ. J. 97, 971–985.
Morgan, J., Várdy, F., 2004. An experimental study of commitment in stackelberg games with observation costs. Game Econ. Behav. 49 (2), 401–423.
Morgan, J., Várdy, F., 2007. The value of commitment in contests and tournaments when observation is costly. Game Econ. Behav. 60 (2), 326–338.
Rabin, M., 1993. Incorporating fairness into game theory and economics. Am. Econ. Rev. 83, 1281–1302.
Schelling, T.C., 1960. The Strategy of Conflict. Cambridge, Mass.
Sheremeta, R.M., 2013. Overbidding and heterogeneous behavior in contest experiments. J. Econ. Surv. 27 (3), 491–514.
Sheremeta, R.M., Masters, W.A., Cason, T.N., 2012. Winner-Take-All and proportional-Prize contests: theory and experimental results. mimeo.
Shogren, J.F., Baik, K., 1992. Favorites and underdogs: strategic behavior in an experimental contest. Public Choice 74, 191–205.
von Stackelberg, H., 1934. Marktform und Gleichgewicht. Vienna and Berlin: Springer Verlag.
Tullock, G., 1980. Efficient rent-seeking. In: Buchanan, J.M., Tollison, R.D., Tullock, G. (Eds.), Toward a Theory of the Rent-Seeking Society. Texas A&M Press, College Station, p. 97–112.
Van Damme, E., Hurkens, S., 1999. Endogenous stackelberg leadership. Game. Econ. Behav. 28 (1), 105–129.
Weimann, J., Yang, C.-L., Vogt, C., 2000. An experiment on sequential rent-Seeking. J. Econ. Behav. Org. 41, 405–426.