CeDEx

Discussion Paper

Centre for Decision Research and Experimental Economics

Discussion Paper Series

ISSN 1749-3293

CeDEx Discussion Paper No. 2008–10

The Appearance of Homo Rivalis: Social Preferences and the Nature of Rent Seeking

Benedikt Herrmann and Henrik Orzen

August 2008

The University of Nottingham

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of Public Economics, Individual Choice under Risk and Uncertainty, Strategic Interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/economics/cedex/ for more information about the Centre or contact

Karina Whitehead
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0) 115 95 15620
Fax: +44 (0) 115 95 14159
karina.whitehead@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/economics/cedex/papers/index.html

The appearance of *homo rivalis*: Social preferences and the nature of rent seeking

by

Benedikt Herrmann and Henrik Orzen

University of Nottingham

August 2008

Abstract

While numerous experiments demonstrate how pro-sociality can influence economic decision-making, evidence on explicitly anti-social economic behavior has thus far been limited. In this paper we investigate the importance of spite in experimental rent-seeking contests. Although, as we show, existing evidence of excessive rent-seeking is in theory compatible with fairness considerations, our social preference elicitations reveal that subjects' investments are driven by spite, not fairness or reciprocity. We also observe a striking disconnect between individuals' revealed social preferences in our contest game and in a standard prisoner's dilemma, rejecting the idea that there are consistent pro-social, selfish or anti-social "types". Moreover, we find that cooperation and reciprocity rates drop substantially after subjects have been exposed to rent-seeking competition.

Keywords: Contests; Other-regarding preferences; Experiments

JEL Classification Numbers: A13; C9; D0; D72

Acknowledgements:

We thank Simon Gächter, Dirk Engelmann and seminar participants at the Universities of Nottingham and East Anglia. We also thank conference participants at the 2007 World Meeting of the Economic Science Association (ESA) at Luiss Guido Carli University in Rome for helpful comments. Financial support from the Leverhulme Trust under grant ECF/2003/0352 and from the School of Economics, University of Nottingham, is gratefully acknowledged.

Contact details:

School of Economics, University of Nottingham, University Park, Nottingham, NG7 2RD, UK (e-mails: henrik.orzen@nottingham.ac.uk and benedikt.herrmann@nottingham.ac.uk)

1 Introduction

Evidence from numerous experimental studies shows that many people voluntarily contribute to public goods, often at levels substantially above the standard predictions from models based on rational and selfish agents (Ledyard, 1995). Although there is still a debate whether this is best explained by tastes for fairness and reciprocity or by bounded rationality and insufficient control in laboratory experiments (see, e.g., Fehr and Schmidt, 2003, and Samuelson, 2005), recent evidence strongly supports the notion that subjects' choices *can* be motivated by pro-social preferences. For instance, some neuroeconomic research suggests that mutual cooperation *per se* activates reward areas of the brain (see the review in Fehr, Fischbacher and Kosfeld, 2005), and results from decision-making experiments that elicit social preferences via response behavior show that a large proportion of individuals reciprocate when their co-players cooperate in a public good setting. Specifically, in the study by Fischbacher, Gächter and Fehr (2001)—hereafter referred to as FGF—a variant of the strategy method is employed to examine how players respond to different average voluntary contributions from other members of their group, knowing that the interaction is one-shot. While FGF observe considerable heterogeneity across subjects, the two main *types* of contribution schedules are "conditional cooperation" (contributions that are generally increasing in the contributions of other group members) and "free riding" (zero contributions regardless of others' contributions). About 50% of the schedules can be classified as "conditional cooperation" and 30% as "free riding". It seems difficult to attribute these clear patterns to confusion among subjects.

Inspired by the idea that some people are pro-socially inclined while others are selfish, economists have developed utility models that incorporate fairness considerations and notions of reciprocity but allow heterogeneity and include selfish types as a special case (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). Economists are now also beginning to investigate wider implications of pro-social preferences (Fehr and Fischbacher, 2002; Gächter, 2007). Furthermore, they ask when interactions between people with varying degrees of pro-sociality produce outcomes that substantially deviate from standard predictions based on selfish agents, and when collective outcomes can be expected to conform closely to standard predictions despite the presence of heterogeneous players. Examples for this literature include Haltiwanger and Waldman (1993), Tyran and Sausgruber (2006), Camerer and Fehr (2006), Andreoni and Samuelson (2006) and Orzen and Sefton (2008).

However, the literature has surprisingly little to say about *anti*-social preferences such as spitefulness or malevolence and *their* effects on economic outcomes (we review some relevant papers in the next section). Of course, an asymmetry in the coverage of pro- and anti-social preferences may be warranted if anti-sociality—loosely defined as a willingness to harm others, absent reasons for negative reciprocity—is a marginal phenomenon. Yet, one wonders whether the focus on pro-sociality could also be due to the fact that the most extensively studied experimental games in this literature (ultimatum games, dictator games, trust games, public good games, prisoner's dilemma games, etc.) rarely *allow* people to make anti-social choices. Has this perhaps biased our views? Traditional economic models that assume agents to be solely motivated by their material self-interest are now sometimes being criticized for making "worst case assumptions". It should be kept in mind, however, that selfishness is the "worst case" only in a certain class of games—in other situations there is scope for behavior compared to which *homo economicus* may appear quite benign. Indeed, some laboratory evidence suggests that explicit anti-sociality occurs and can influence economic outcomes. In experiments on public goods with an added punishment stage it has been found that subjects sometimes punish others who have behaved nicely towards them (Herrmann, Thöni and Gächter, 2008). This is obviously incompatible with the idea of reciprocity, but it is also incompatible with selfishness because of the cost attached to punishing. It has also been found that patterns of overbidding in auctions are consistent with the predictions from a model that incorporates spiteful rivalry into the bidders' utility functions (Nishimura, Cason, Saijo and Ikeda, 2007). Thus, after the introduction of *homo reciprocans* to replace or complement the traditional *homo economicus* (see Bowles and Gintis, 2002), should we now add a *homo maliciosus* or a *homo rivalis* to the family? Clearly, more systematic evidence on such behaviors is needed before this question can be answered.

In this paper we examine the empirical significance of spitefulness among individuals using laboratory methods. In designing our experiments it has not been our intention to model a specific "real-world" problem. Rather, we were looking for a type of setting where we thought we might observe *homo rivalis*, if he exists, but where there would also be scope for potentially efficiency-enhancing reciprocal behavior and fairness-oriented behavior, which would also differ from standard game-theoretical predictions based on selfishness. We chose two-player rent-seeking contests as the basis for an experimental paradigm that has all these features. The competitive attributes of rent-seeking contests differ from those of auctions and markets and may be viewed as an archetypical example for a situation of conflict: two rival parties spend resources to secure an external rent; the relative expenditures determine each party's chance of success; the winner takes all and there are no refunds. We elicit our subjects' social preferences employing the same strategy elicitation technique that FGF use in a public good environment. Thus, subjects are asked to indicate for each possible expenditure level of their opponent how much of their endowment they wish to invest in the contest. This gives us a revealed best response function for each individual, and we elicit such response functions in a one-shot setting as well as in a repeated setting where players are randomly re-matched after each round. We are interested whether the revealed social preferences differ fundamentally from those in experimental games like the ones listed above.

In previous laboratory studies on rent-seeking contests, which have not been designed to elicit social preferences, observed levels of rent dissipation typically exceed the predictions from equilibrium analyses based on selfish, rational and risk-neutral agents. Our paper contributes to this literature by exploring reasons for this discrepancy. First, we show that overdissipation *per se* cannot be counted as evidence for spiteful preferences because it emerges in equilibrium when players have a taste for fairness in the sense that they dislike inequality. Spiteful preferences can also explain overdissipation but they imply fundamentally different best response functions compared to the case of inequality aversion. Second, since we study response behavior experimentally we should be able to see which approach better organizes the data, also in comparison to standard-theoretical benchmarks for selfish players. Third, because we furthermore compare the outcomes with the results from a control treatment where players simply choose their expenditures simultaneously without strategy elicitation, our design also enables us to examine whether strategic uncertainty or subjects' inability to calculate or estimate probabilities may have biased previous results. Fourth, as a further control we ran sessions in which we elicited response functions as before but removed the aspect of competition completely by transforming the contest game into an equivalent individual choice problem with identical material incentives. Thus, we can ask how exactly the presence of an opponent affects response behavior.

In addition, our paper addresses the question whether a particular set of social preferences is a generic trait of an individual or whether one's social preferences depend on the kind of strategic interaction one is faced with. If the former is true (call it the "universality hypothesis") we would expect that an individual's generic inclination towards other players is *in principle* the same in any type of game. Of course, it would not be surprising if subjects approach a contest game not in the same way as, say, a public good game. However, what we should *not* find under this hypothesis are, absent any situation-dependent reasons for positive/negative reciprocity and absent any strategic uncertainty, revealed *pro-social preferences* in one situation and revealed *anti-social preferences* in another. If existing elicitations of social preferences do not provide good guidance for the elicitations of social preferences in rent-seeking contests, this would be consistent with a "game-dependency hypothesis". It seems conceivable that an individual's attitude towards others undergoes changes between different types of situation because they evoke different contextual cues, which subsequently instigate specific attitudinal and behavioral modes deemed appropriate in the circumstances. But even if such game-dependent modes exist, one might ask to what extent they are correlated with each other: can an individual's social preferences, as revealed in some game A, be used to predict how the individual acts in a different game B?

To examine these questions we ran a series of pre-tests on our subjects, in which we measured tastes for fairness and reciprocity, using the FGF technique in a prisoner's dilemma (see also Herrmann, Kabalin, Nedzvetsky, Poen and Gächter, 2008), and attitudes towards risk, using a variant of the procedures employed by Holt and Laury (2002) and Dohmen, Falk, Huffman, Sunde, Schupp and Wagner (2005). This allows us to directly examine, *at the level of the individual*, to what extent revealed social preferences correlate in different settings.[1]

Because we wanted to avoid that the presence of the pre-tests directly affects subjects' behavior in the main experiments we ran the pre-tests one week before and did not inform participants of the outcomes until after completing the contest experiments. We also ran some sessions without any pre-tests. However, in these we let subjects play the prisoner's dilemma *after* the main experiment. This allows us to investigate whether exposure to competition in

[1] Revealed attitudes towards risk are not immediately relevant to discriminate between the universality and the game dependency hypothesis. However, they can be used as an additional control.

rent-seeking contests reduces subjects' willingness to cooperate. If it does, it would be consistent with the idea that rent-seeking activities are perceived as anti-social and it would suggest that social preferences can be manipulated by such negative experiences.

The paper proceeds with a discussion of related literature in Section 2. Section 3 describes the experimental design and Section 4 the results. Section 5 concludes.

2. Related literature

Anti-sociality in the experimental literature

Empirical explorations of anti-social preferences, where someone obtains utility from reducing someone else's payoff without having being treated unfairly, attracted only minor attention until recently. Zizzo (2003a) in his paper on burning money experiments reported that subjects are often willing to reduce, at a cost for themselves, the incomes of players who had been given higher endowments. In some instances subjects with the same or less endowment were also targeted. In a similar vein, Dawes, Fowler, Johnson, McElreath and Smirnov (2007) find that subjects are willing to reduce other group members' income independently of the history of interaction. In this study participants were endowed with a distribution of amounts that, unknown to subjects, was identical to the distribution that had emerged from the cooperation decisions in the first stage of the public good with punishment games in Fehr and Gächter (2002). Subjects were then allowed to reduce each others' income at own cost, and did so in a fashion similar to the fairness-oriented punishment reported by Fehr and Gächter, casting some doubts on fairness as a motive behind the previously reported punishment of low contributors. In their experiments on competitive behavior, Rustichini and Vostroknutov (2007) find that participants are more inclined to reduce someone else's income if the punished subject has earned more money than the punisher. Surprisingly, this effect is stronger when the higher incomes of the punished subjects are due to merit rather than luck. Spiteful behavior has also been observed in auctions. Nishimura et al. (2007) develop a model of spiteful bidding in second price sealed bid and English auctions under complete information, and then test this model in the laboratory. As predicted, participants who are endowed with low values often overbid, which reduces the potential surplus of high-value participants, and these high-value participants often retaliate by underbidding with the effect that the overbidding low-value participants run the risk of realizing a loss. Abbink and Sadrieh (2008) find that people seem to get pleasure from costlessly reducing someone else's income, when they can hide this activity behind the veil of a lottery mechanism that sometimes destroys the other's income independently. The most extreme form of anti-social punishment, where the punishment is directed against those who had previously behaved nicely towards the punisher, has been observed in public good games with punishment. In these games those who are more cooperative than others are frequently punished. Such evidence is reported in Cinyabuguma, Page and Putterman (2006), Gächter, Herrmann and Thöni (2005) and Herrmann et al. (2008).

Experiments on rent-seeking contests

Most previous experiments on rent-seeking games exhibit substantial discrepancies between actual investment levels and equilibrium predictions. Öncüler and Croson (2005) provide a useful review of nine papers that precede their own study. They summarize (p.407): "Of these nine papers, six of them show more rent-seeking behavior than the theory predicts (Millner and Pratt, 1989, 1991; Davis and Reilly, 1998; Potters et al., 1998; Anderson and Stafford, 2003; Schmitt et al., 2004), while one shows no significant differences between theoretical predictions and observed behavior (Shogren and Baik, 1991) and two show the opposite effect (Shupp, 2004; Schmidt et al., 2004)." Öncüler and Croson themselves investigate rent-seeking contests where the contest winner does not obtain the rent with certainty but with a probability inversely related to her contest expenditures. In their experiment the authors find support for most of the comparative-static predictions of their model but observe, like most of their predecessors, excessive levels of rent-seeking relative to the theoretical predictions.

Weimann, Yang and Vogt (2000) and Vogt, Weimann and Yang (2002) examine sequential rent-seeking contests. The subgame-perfect equilibrium of the game studied in Weimann et al. predicts a first-mover advantage: the first mover chooses an expenditure level below the value of the rent but sufficient to deter the second mover from making a positive bid. The experimental results, in contrast, indicate a strong second-mover advantage. Particularly noteworthy in the context of our research question is that second movers tend to *punish* first movers who make bids close to the equilibrium prediction and *exploit* first movers who make low bids. In the game studied in Vogt et al. players take turns and each time they can either increase their bids or end the bidding stage by doing nothing. Equilibrium analysis predicts minimal rent-seeking expenditures in this case, and in their experiment Vogt et al. find this hypothesis supported, at least after subjects had time to gain experience with the game. However, based on comparisons with the Weimann et al. experiment the authors conclude that the observed "efficiency is not the result of cooperation in a repeated game, but the consequence of the threat of escalation" (p.72).

Parco, Rapoport and Amaldoss (2005) investigate two-stage rent-seeking contests: in the first stage individuals compete locally within equally sized subgroups, and in the second stage the winners of the subgroups compete against each other. The authors explore potential explanations for the observed over-expenditures relative to theoretical predictions based on non-pecuniary utility of winning, misperceptions of probabilities and learning. This can account for some, but not all, of the behavioral regularities.

Fonseca (2006) investigates symmetric and asymmetric rent-seeking contests in the laboratory and compares simultaneous-move games with sequential-move games. He finds that subjects invest less in the asymmetric contests than in the symmetric ones. However, rent dissipation levels are again well above standard equilibrium outcomes in all treatments, and there is also some evidence of spiteful response behavior by second movers.

Bullock and Rutström (2007) develop a model akin to Becker (1983) in which two contestants lobby for endogenously determined levels of transfers, and then test this model in the laboratory. They observe, as predicted by their model, that rent-seeking expenditures exceed the value of the transfer that emerges in the end. They also find that disadvantaging one player reduces rent dissipation compared to the symmetric case, although the disadvantaged players' expenditures are much higher than predicted.

Abbink, Brandts, Herrmann and Orzen (2007) study contests where teams and/or individuals compete and investigate the effects of within-group punishment on rent dissipation. Contrary to predictions groups invest substantially more than individuals who overspend already. The introduction of punishment opportunities boosts rent dissipation even further.

Morgan, Orzen and Sefton (2008) examine endogenous participation in rent-seeking contests and compare results from a small prize treatment with those from a large prize treatment. They find very substantial overdissipation in the early rounds of their experiments although this deteriorates over the course of the 50 rounds their experiment lasts. As predicted, they find that entry and expenditures increase with the size of the prize. However, earnings between the contest and the outside option are not equalized.

3. Experimental design

3.1 Procedures

Participants were recruited from a subject pool comprised of about 3500 undergraduates at the University of Nottingham who had indicated a willingness to be paid volunteers in decision-making experiments. We conducted the experiment in two waves. The first wave consisted of a series of double sessions that ran in two consecutive weeks. In week 1, we took basic measurements of our subjects' social preferences and risk attitudes. In week 2, the same subjects participated in a contest experiment (the compositions of sessions in the two weeks were different). The second wave consisted of six single sessions, in which a different set of subjects participated either in a contest experiment or in a variant of the game that transformed the contest into a series of individual choice tasks. In both cases this was then followed, within the same session, by another elicitation of subjects' social preferences (employing the same method that was used in the first wave pre-tests).

In all sessions subjects were, upon arrival, randomly seated at visually separated computer terminals. During a session information was only transmitted through the computer network and any other form of communication was prohibited.[2] There were between ten and twenty subjects in each session, and the participants in the first wave knew that by signing up for the experiment they committed to attending two sessions, one in week 1 and one in week 2. All but 6 of the 99 subjects who had taken part in one of the week 1 sessions subsequently returned for a week 2 session (the data from these six subjects was discarded). One subject participated only in week 2: he/she was used to fill up one of the sessions to get an even number of players. 116 subjects participated in the second wave of experiments.

[2] Obviously, in the individual choice tasks no information was transmitted between subjects. The experiments were fully computerized using the zTree package (Fischbacher, 2007) for the elicitation of social preferences and risk attitudes, and using purpose-built software developed by the authors for the contest experiments.

Week 1 sessions

The sessions in week 1 consisted of three tasks: a prisoner's dilemma problem, a task to measure individuals' risk attitudes and another task which was conducted for a different paper. Each task was presented as a separate decision-making experiment and instructions for a task were handed out only after the previous task had been completed. We did not give subjects feedback between tasks or even after the session but only *after the week 2 sessions*; this includes the information of how much participants had earned in the individual tasks or overall. We took this decision to avoid cross-contamination between tasks and between the week 1 and week 2 sessions.

For task 1 each subject was randomly and anonymously matched with another person participating in the same session. The decision problem subjects faced was the prisoner's dilemma shown in Table 1.[3] The payoffs were denoted as "experimental dollars" ("e$") and subjects knew the exchange rate: e$10 were worth £2.

Table 1: The task 1 prisoner's dilemma

	Cooperate	Defect
Cooperate	20, 20	0, 30
Defect	30, 0	10, 10

Subjects made their decisions simultaneously under the strategy method as developed by FGF (see also Herrmann et al., 2008). That is, we asked each subject for both a *conditional choice schedule* and an *unconditional choice*. In the conditional choice schedule subjects indicated their preferred choice for each of the two contingencies (other player cooperates or other player defects). In the unconditional choice subjects simply chose one of the two available actions. To incentivize the task we followed FGF by informing subjects that after all decisions had been made a die throw would determine for each pair which player's unconditional choice and which player's conditional choice schedule would be relevant. The final outcome—not revealed during week 1—then emerged by combining one person's unconditional choice with the other person's response to that choice, as determined by the schedule the responding player had submitted. Given this procedure, a selfish and rational player in this game is incentivized to play "defect" in the unconditional choice and to submit "defect" for both contingencies in the conditional choice schedule.

The risk attitude measurement task was an individual choice problem akin to that used in Dohmen et al. (2005). The task consisted of a series of 15 binary decisions between a lottery and a safe outcome. Subjects were informed that only one of the 15 problems would later be randomly selected as being payoff-relevant (again, the final outcome was revealed only after week 2). The lottery was the same in all 15 cases and gave a 50% chance of winning 400 points (worth £4), while the safe option varied between tasks and ranged from 25 to 375

[3] The terms "cooperate" and "defect" were not used in the description, nor did we use a payoff matrix like that shown in Table 1. Copies of experimental instructions are provided in Appendix A.

points. We expected that with a safe alternative of 25 points practically everyone would choose the lottery. As the attractiveness of the safe alternative increases, people should at some point switch over to the safe outcome. We expected that with a safe alternative of 375 points practically everyone would choose that safe alternative. The point of switching from the lottery to the safe outcome serves us as a measure of a person's risk attitude.

Contest sessions (week 2 and second wave)

The experiments on rent-seeking competition were conducted in two variants. In the *Direct* treatment two randomly matched players—each starting with an endowment of e$16 ("e$" being the experimental currency unit)—competed anonymously in a conventional symmetric Tullock rent-seeking game (see Tullock, 1980) by simultaneously choosing how many e$s to spend on contest tokens (any integer number between 0 and 16 was allowed, the cost of 1 token was e$1). One's chance of winning the prize (worth another e$16) was equal to the proportion of one's own expenditure out of the total spent by both parties. In our design, the *Direct* treatment mainly serves as a control for the *Strategy* treatment where subjects played essentially the same game but now made their choices via the FGF strategy method. That is, as in the week 1 prisoner's dilemma task, subjects were again asked for a *conditional choice schedule* and an *unconditional choice*. In the conditional choice schedule subjects indicated their preferred expenditure for each of the 17 possible contingencies that made up the range of feasible expenditures for the other player. In the unconditional choice subjects simply chose an expenditure between e$0 and e$16. Subjects also knew that a random procedure would again be used to determine for whom of the two matched players the unconditional choice would be relevant and for whom the conditional choice schedule would be relevant.

One particular advantage of using the strategy method in this way is that it generates very rich data: we get an incentivized subjective best response function from each participant. The procedure also removes any strategic uncertainty from the participants' conditional expenditure decisions: while there remains uncertainty about which contingency will turn out to be the relevant one, players can prepare optimally for each eventuality. As a by-product of the method, the decision-making procedure now evolves as a step-by-step process that may help subjects make informed decisions. Moreover, we programmed the software so that it calculated and displayed the probabilities of winning for every possible expenditure choice a subject could make. Thus, for instance, for the case that the matched player spends 4 the computer showed that one's chances of winning, depending on one's expenditure of 0, 1, 2, 3, 4 etc., would be 0%, 20%, 33%, 43%, 50% etc. This information was also represented graphically with small pie charts.

In the second wave of experiments we ran the same *Strategy* treatment again, using exactly the same procedures. These subjects were recruited for just a single session (out of a total of six sessions). Care was taken that none of the first-wave subjects took part again. At the end of each of the second-wave sessions we let participants play our prisoner's dilemma (Task 1 above). In addition, we ran a number of sessions in the *Individual* treatment. For these sessions we transformed the *Strategy* treatment into an individual choice problem under risk where each subject makes investment choices for different scenarios knowing that one of these scenarios will later be selected at random. Thus, for the *Individual* treatment we asked subjects only for the conditional choice schedule and simply removed the first mover from the game. Instead of going through the 17 different possible investment levels that the first mover can submit, subjects now went through 17 different "decision making tasks". In each of these tasks subjects chose their investment based on the probability information we gave to them using exactly the same interface on the computer screen that we had used for the *Strategy* treatment. The probabilities of winning the prize (and therefore the material incentives) were identical to those in the *Strategy* treatment, but we did not disclose to our participants how we had arrived at these probabilities. Like the second-wave *Strategy* treatment the *Individual* treatment was followed by a strategy-method prisoner's dilemma.

Otherwise, the *Direct*, *Strategy* and *Individual* experiments all developed in the same way. At the beginning subjects were informed that the experiment would consist of two parts and that they would be paid at the end of the session. The first part consisted of a one-shot play of the rent-seeking game according to the relevant treatment. We chose a one-shot design because this is arguably the cleanest way to elicit social preferences and best suited to relate the observed behavior to the decisions made in the week 1 pre-tests. While the lack of interaction between players is desirable in that choices cannot be influenced by individuals seeing what other participants do and therefore more truly reflect a priori attitudes, it has the drawback that subjects cannot learn from experience. Therefore, particular care was taken to ensure that participants understand the mechanics of the setup and the consequences of their actions: we provided detailed instructions of the rules of the game, asked subjects to think very carefully about what choice(s) to make, and encouraged them to ask questions. We did not impose any time limits.

Nevertheless, we suspected that subjects' behavior might change with experience and the second part of *Direct*, *Strategy* and *Individual* explored this possibility. The second part was a 15-times repeated version of the same rent-seeking game that was played in part 1, in *Direct* and *Strategy* using a protocol that randomly re-matched participants at the beginning of each round.[4] As in the week 1 experiments, we handed out the instructions for the second part only after subjects had completed the first part. The only other difference from part 1 was that the conversion rate between e$ and pound sterling was changed from e$8 for a pound to e$40 for a pound. In each round of the *Strategy* and *Individual* treatments subjects could construct new response schedules from scratch if they so wished, but we also gave them the opportunity to copy their schedule from the preceding round via a mouse-click and modify it. Because we are chiefly interested in the choices of *experienced* players here, our data analysis for the second part will focus on the last five rounds (rounds 11-15).

Table 2 summarizes our design.

[4] Our software was programmed to divide up the participants in a session in such a way that, depending on how many people took part in a session, there were between 6 and 10 subjects in each re-matching group. There were six such matching groups in the *Direct* treatment and 17 in the *Strategy* treatment (8 in the first wave and 9 in the second). 56 subjects participated in the *Individual* treatment.

Table 2: Experimental design

First wave sessions: week 1			
Social preference task		*Risk attitude task*	
Prisoner’s dilemma (one-shot, strategy method)		15 binary lottery-versus-safe-outcome choices	*No feedback on week 1 outcomes.*
First wave sessions: week 2			
Treatment	*Part 1*	*Part 2*	
Direct / *Strategy*	One-shot contest game	15 rounds of the same game (random matching)	
Second wave sessions			
Treatment	*Part 1*	*Part 2*	*Social preference task*
Strategy / *Individual*	One-shot contest/ indiv. choice task	15 rounds of the same game (rnd. matching in *Strategy*)	Prisoner’s dilemma (one-shot, strategy method)

3.2 Hypotheses

The standard-theoretical prediction for the prisoner’s dilemma task is that players always favor the outcome that yields a higher payoff for themselves, regardless of the implications for the matched player. Of course, existing experimental evidence shows that outcomes are typically much more heterogeneous than this. Based on the results in FGF we expect to find a considerable degree of conditional cooperation. Theoretical predictions in the risk attitude task depend on individual risk preferences—thus, again we expect heterogeneity in subjects’ decisions. Based on stylized facts from the literature on individual choice experiments we expect more subjects to display risk aversion than risk proneness.

Next, consider the standard Nash equilibrium of the Tullock rent-seeking game. Two risk-neutral players compete for a rent or prize of size $R > 0$ and simultaneously choose some nonnegative (and non-refundable) expenditure, which we label x for player 1 and y for player 2. If $x = y = 0$, the prize is lost and both players earn zero. This cannot be an equilibrium because one of the players could profitably deviate and obtain the prize by spending a small amount $\varepsilon < R$. If at least one of the expenditures is strictly positive, the players’ expected payoff functions are, ignoring initial endowments,

$$\pi_1(x, y) = \frac{x}{x+y}R - x \quad \text{and} \quad \pi_2(y, x) = \frac{y}{x+y}R - y . \qquad [1]$$

The first-order conditions, obtained in the usual way, are $yR = (x+y)^2$ and $xR = (x+y)^2$, and therefore we must have $x = y$ in equilibrium. The equilibrium expenditures when players move simultaneously are then easily derived as $x^* = y^* = R/4$.

In the *Strategy* treatment players effectively move in sequence: a player is in the first-mover position when she makes her unconditional choice and is in the second-mover position when she sets up her conditional choice schedule. Thus, we now solve by backward induction. Without loss of generality assume that player 1, choosing x ($0 \le x \le R$), moves first. Player 2's best response to this is an investment of an infinitesimally small amount $\varepsilon > 0$ if $x = 0$ and an investment of $\sqrt{xR} - x$ otherwise. Anticipating this response schedule, player 1's optimal expenditure is $x^* = R/4$, as before, prompting player 2 to invest the same, $y^* = R/4$. Thus, moving simultaneously or in sequence does not affect the equilibrium outcome. For $R = 16$ the predicted expenditure level per person is 4. Hence, 50% of the rent is dissipated. The optimal response schedule in the *Individual* treatment, which replaces the choice of the first mover with a random draw, is identical to the optimal response schedule in the *Strategy* treatment.

As discussed above, most existing laboratory studies on rent-seeking competition in simultaneous-move settings have found overspending relative to the standard equilibrium benchmark, particularly in early rounds and to a lesser extend in late rounds. First, we will be interested whether our data from the *Direct* treatment replicates this result. Second, we will explore reasons for this discrepancy. We can think of a number of explanations.

One possibility is that over-investments are caused by a combination of strategic uncertainty and bounded rationality. That is, if subjects vaguely expect aggressive investments from their opponents, given the nature of a contest situation, they may be intuitively inclined to counter this. Yet, by doing so they effectively fulfill their own prophecy. Of course, the peak of the best-response function coincides with the equilibrium point—thus investing more than in equilibrium can never be a best response for rational risk-neutral players. However, for subjects to realize this requires that they precisely calculate their chances of winning for various combinations of investments so that they can evaluate the implications of own low or high investments. It seems possible that subjects' failure to analyze their options in this way contributes to the observed deviations from theoretical predictions.

Our *Strategy* treatment controls for these potentially confounding factors and thus provides a test of this explanation. It eliminates strategic uncertainty and, forcing subjects to make their choices in a guided step-by-step process, alleviates concerns regarding bounded rationality since the relevant probability calculations are done by the computer and displayed on the screen as numbers and in graphical form. As a result, subjects may get a much better feeling for the marginal costs and benefits of investing additional units. Moreover, second movers have to think carefully about what they wish to do in each contingency, whereas subjects in the conventional simultaneous-move version of the game might not naturally engage in such strategic thinking. This may help *Strategy* players to optimize their investments. Alternatively, subjects can now condition their choices on friendly versus hostile opponent behavior, which may promote efficiency-enhancing reciprocity between the two players: the first mover may signal a willingness to keep the intensity of conflict at a minimum by investing very little and the second mover may reciprocate by keeping her own expenditures low and the odds of winning even. Thus, based on these considerations we would expect a substantial reduction in expenditures in *Strategy* relative to *Direct*.

Considering other explanations, notice that risk aversion can, in principle, also lead to excessive rent-seeking expenditures (Konrad and Schlesinger, 1997). However, this depends on the precise curvature of the utility function and conventional formulations of risk aversion imply *under*-dissipation relative to the risk-neutral case (see Abbink et al., 2008). Furthermore, Cornes and Hartley (2003) have shown that loss aversion substantially reduces rent dissipation, and the available empirical evidence suggests that more risk averse (Millner and Pratt, 1991) and more loss averse (Kong, 2008) subjects invest less than others. Similarly, we will also test whether our measurements of subjects' risk attitudes correlate with rent-seeking expenditures.

A third possibility is based on the idea that the experience of winning *per se* generates utility. Such extra utility from winning should inflate investments just like an increase in the prize *R* would. Our *Individual* and *Strategy* treatments will provide evidence for an examination of this explanation. If it is the sheer "joy of winning" that drives up contest investments, we would expect the response schedules in both these treatments to be above the standard-theoretical best response and to be very similar to each other.

If, however, social preferences are the underlying reason for over-investments and the presence of a rival player affects how subjects convert material earnings into psychological payoffs, we should find a gap between the *Individual* and *Strategy* response schedules.[5] If indeed we find such a gap, we will have to inspect the shapes of the schedules closer to distinguish between two sub-hypotheses under the social preference explanation. One of these attributes overdissipation to spite: individuals enjoy beating the other player and the associated relative monetary advantage. The other one suggests the reverse: individuals dislike monetary inequality, even if they are in the lead. Interestingly, both of these sub-hypotheses can cause over-expenditures relative to standard predictions. However, the implications for the shape of the response function are very different.

While it seems intuitive that spiteful preferences ought to increase investments in a rent-seeking game, this is less obvious in the case of inequality aversion and may perhaps come as a surprise, since the common use for models of inequality aversion is to explain how individuals can arrive at Pareto-superior outcomes that are unattainable for selfish agents in standard game theory. We relegate the details of how inequality aversion increases rent-seeking expenditures in Appendix B, utilizing the model by Fehr and Schmidt (1999). As we show in Appendix B, equilibrium investments in the rent seeking game with two Fehr-Schmidt players are

[5] In setting up the *Individual* treatment we were worried that subjects might "humanize" the computer and view the machine or the program running on it as their opponent or co-player. This is an important psychological phenomenon and presumably a factor in the commercial success of many computer games. Experimental evidence on how such effects can shape pro-social behavior is reported in Haley and Fessler (2005). To avoid this, we removed from the instructions any words or phrases that might have given subjects the impression that the computer should be viewed as another player.

$$x_{FS} = y_{FS} = \frac{1+\alpha-\beta}{2+\alpha-\beta}\frac{R}{2} \qquad [2]$$

where α and β ($0 \le \beta < 1$ and $\beta \le \alpha$) are coefficients measuring how much the players dislike disadvantageous and advantageous inequality. It can easily be checked that inequality averse players invest more in equilibrium than selfish players whenever $\alpha > \beta$. In the limit, rent dissipation increases to 100%.

It is straightforward to introduce spite in the Fehr-Schmidt model. Spiteful agents dislike disadvantageous inequality, as before, but *enjoy* advantageous inequality—thus $\beta < 0$. The expression in equation 2 remains the same—the only difference is that high absolute values of β now contribute to increased rent-seeking investments rather than hinder them. While this does not affect the range of possible outcomes in equilibrium, the shape of the response function changes considerably between positive and negative values for β. A distinctive feature of response functions based on inequality aversion, in contrast to response functions based on spiteful preferences, is that they collapse when the first mover submits a large investment (see Appendix B for details).

The response schedules in the *Individual* and the *Strategy* treatment may also help us gain a better understanding of heterogeneity between players. In particular, we will be interested to see whether specific *types* of response function emerge, as they did in the linear public good game studied by FGF.

To discriminate between the universality hypothesis and the game-dependency hypothesis, we will combine the results from the week 1 prisoner's dilemma and part one of the week 2 rent-seeking sessions. Under the universality hypothesis we would expect that those who display an inclination towards pro-sociality in the prisoner's dilemma (e.g. by applying the principle of reciprocity) display a similar inclination in the rent-seeking contest. Since behavior in the prisoner's dilemma cannot be used to distinguish between selfish types and spiteful types, free riders may either continue to behave selfishly or may overinvest in a spiteful way. Under a weaker version of the universality hypothesis, we at least expect that pro-social types make lower conditional rent-seeking investments than other types. In contrast, under the game dependency hypothesis we would expect either no correlation between week 1 and week 2 choices, or possibly even the reverse relationship (more aggressive rent-seeking behavior from pro-social types than from other types).

Finally, we will examine behavior in the second-wave post-tests. We expect the outcomes from the prisoner's dilemma following the *Individual* treatment to be similar to the pre-test prisoner's dilemma. The willingness to cooperate might diminish after the *Strategy* treatment if subjects perceive the interaction in the rent-seeking contests as aggressive and spiteful.

4 Results

We begin by reporting results from the contests and the *Individual* treatment, and will turn to the control measurements of social preferences and risk attitudes later. As discussed, our two-part design allows us to compare behavior in a one-shot contest with that in a repeated setting under random matching. The one-shot setting is best suited to elicit subjects' a priori attitudes in the contest, whereas the repeated version is more appropriate when one wants to learn about the behavior of *experienced* players—our data analysis for part two will therefore focus on the last five rounds.[6]

In Figure 1 we have plotted average expenditures over time. The separate data points in "round 0" represent the results from part one.[7] In all treatments the *actually realized* expenditures are shown. In the *Strategy* treatment these depend on the realized random matching and on the realized sequence in which players move. In the *Individual* treatment they depend on the realization of the random draw that determines the relevant task. Table 3 provides some summary statistics.

Figure 1: Expenditures per person over time

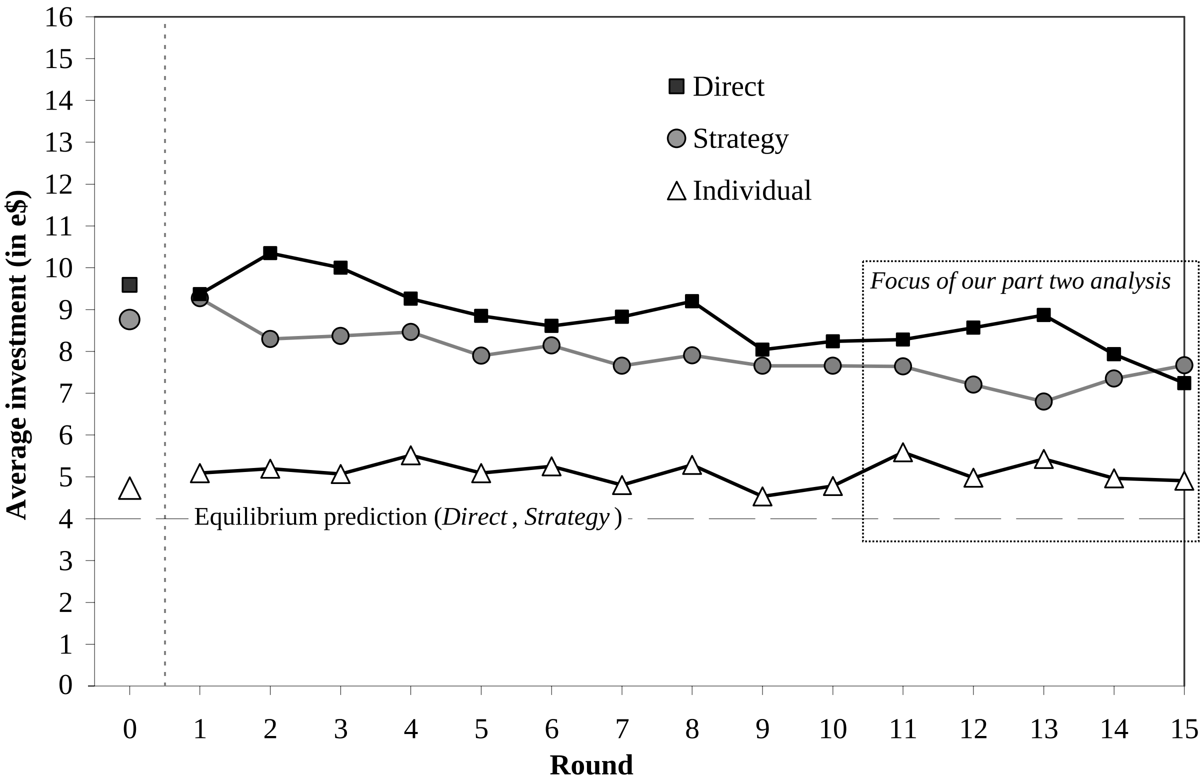

As can be seen from the figure, investments in *Strategy* and *Direct* are at comparable levels, although *Direct* exceeds *Strategy* in most rounds. We cannot reject the null hypothesis that the observed expenditures in the two treatments come from the same distribution: based on a two-sided two-sample permutation test the p-values are 0.295 in part one and 0.397 in part

[6] When we refer to "part two" results below, we will mean the last five rounds unless stated otherwise.
[7] The data from the first-wave and second-wave *Strategy* treatments were highly similar and we pool them here.

two. In both cases we observe a general, albeit slow, downward trend (on average –0.1 tokens per round). The reduction in expenditures between parts one and two appears to be systematic: using on a Fisher paired comparison randomization test (two-sided), treating each matching group as an independent observation, the p-values are 0.031 in *Direct* and <0.001 in *Strategy*. Yet, despite the decline, average expenditures remain well above the theoretical predictions even towards the end. This discrepancy is statistically significant: based on a two-sided one-sample randomization test (again at the level of matching groups) the p-values are 0.031 in *Direct* and <0.001 in *Strategy*, both in part one and in part two.[8] Within *Strategy* we checked for differences between first-mover and second-mover investments but found practically none: the average unconditional choice in part one is 8.9 (7.3 in part two), while the average response, conditional on the realized matching, is 8.8 (7.4 in part two).

Table 3: Average investments in parts one and two
– standard deviation in parentheses –

	Direct	*Strategy*	*Individual*
Part 1	9.6 (4.4)	8.8 (4.4)	4.7 (4.4)
Part 2 (rounds 11-15)	8.2 (4.9)	7.3 (4.4)	5.2 (4.6)

Investments in *Individual* are substantially lower. Although this difference is statistically significant (comparing *Strategy* and *Individual* yields a p-value of <0.001 in part one and of 0.027 in part two), one should be cautious about drawing conclusions from this: it could be that the investments diverge merely because the first-mover behavior that subjects in *Strategy* respond to differs from the random draw that subjects in *Individual* respond to. Thus, a more appropriate comparison will be based on the response schedules themselves. However, we note that average investments in *Individual* persistently exceed 4, the maximum a risk-neutral payoff maximizer would invest. Moreover, there is no decline over time. We find this result surprising—it is not what one would expect if subjects' preferences are characterized by risk aversion and/or loss aversion. Below we will investigate this further.

Result 1

(a) The outcomes from previous experiments on rent seeking contests are replicated. In the contest treatments subjects overspend relative to the standard equilibrium benchmark, and efficiency is lower than predicted (113% of the rent is dissipated in part one while the prediction is 50%). Over time, expenditures decline but the levels remain too high (95% of the rent is dissipated in part two).

(b) There is only weak support for our hypothesis that overdissipation emerges as a result of strategic uncertainty and bounded rationality, as outlined in Section 3.2 above: although the investments in *Direct* exceed those in *Strategy* the differences do not appear to be systematic.

[8] The equivalence of the p-values in this and the previous test is a coincidence.

Next, we turn to the response behavior in the *Strategy* and the *Individual* treatment. Since participants in the *Individual* treatment do not interact, each subject's vector of conditional choices can be considered to be independent. Thus, we obtain 56 independent response functions for both parts of the *Individual* treatment. When the *Strategy* subjects submit their response schedule for the first time, they have not interacted *yet*, and thus we obtain 108 independent response functions for part one of the *Strategy* treatment. In part two, information is exchanged between subjects within the same matching group, and the number of independent observations is now reduced to 17.

In Figure 2 we have plotted the average response schedule for both treatments, in each case with a 95% confidence interval for the estimated mean conditional expenditure.[9] Figure 2a shows the schedule for part one and Figure 2b for part two.

Figure 2: Average response schedules in *Strategy* and *Individual*

a. Part one b. Part two (rounds 11-15)

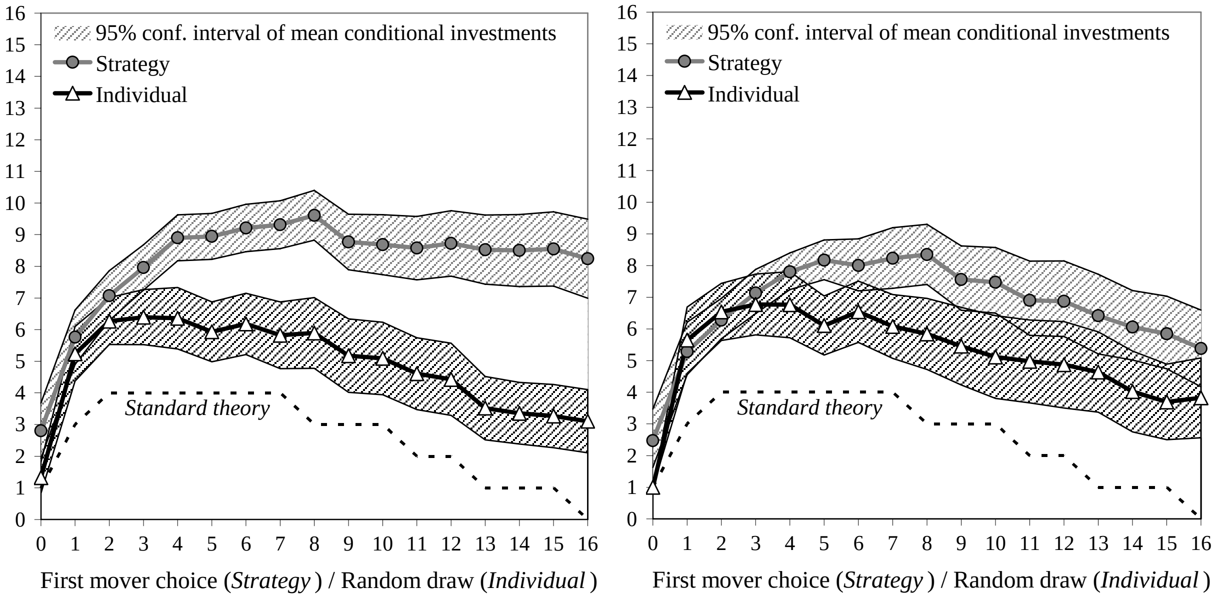

Subjects do not behave as rational risk-neutral payoff maximizers as assumed in standard theory. In both treatments they overinvest relative to the standard prediction and even in part two, after they have had the chance to gain much experience with the game, their response schedules do not approach the theoretical best response function. Based on the confidence intervals for mean investments this deviation is clearly systematic. In the case of the *Individual* treatment the over-investments could be due to a "joy of winning" as we hypothesized in Section 3.2 or, alternatively, due to particular risk preferences or due to subjects failing to optimize. Because it would lead us too far astray from our main research questions, we make no attempt here to isolate the true cause, although the "joy of winning" hypothesis is consistent with informal self-reports we received from our participants and with our own casual observations of subjects' facial expressions during the experiment.

[9] The estimates for the confidence intervals were computed by resampling, for each contingency, 10,000 times from the empirical distribution of conditional choices on the basis of statistically independent observations.

Regardless which of the above explanations holds true, the same biases in behavior should be at work in the *Strategy* treatment. However, while investments in *Individual* and *Strategy* are similar when subjects respond to small values, they diverge considerably when subjects respond to larger values. The disparity is particularly striking in part one of the experiment. In part two the gap narrows, mainly due to a drop in the *Strategy* investments, but remains significant (the mean expenditures in *Strategy* are outside the *Individual* confidence interval and vice versa for responses to values above 4). Thus, some additional factor must enter players' mindsets when these large values are submitted by another person rather than nature. The only plausible explanations, we believe, are that subjects are motivated by social preferences or apply norms of reciprocity. A potential problem with the reciprocity hypothesis is that we only find *negative* reciprocity: subjects do not respond kindly to low first mover investments, which one might interpret as a friendly signal to limit competition—instead, given the similarity to the *Individual* treatment, the average response here is better described as selfish. Moreover, it is not even entirely clear how reciprocity norms work for very large first-mover investments since extreme rent-seeking expenditures not only reduce the responder's expected income but hurt the first mover *himself* potentially even more.[10]

If we conclude that social preferences are the underlying reason for the observed difference between the *Individual* and *Strategy* response schedules, can we characterize the nature of these preferences? One possibility is to interpret the observed response behavior in *Strategy* as emerging from Fehr-Schmidt-type preferences. However, since the Fehr-Schmidt model has no bite in the *Individual* treatment, where subjects already invest more than predicted, this approach would overestimate the influence of social preferences. We therefore propose a two-stage procedure: in stage one we use the data from the *Individual* treatment to estimate the strength of the "joy of winning", and in stage two we take the result from this to enrich the Fehr-Schmidt model with an equivalent joy-of-winning component and then estimate its parameters. Note that the role of the joy-of-winning component is mainly instrumental for us—as indicated above, we do not have conclusive evidence that the *Individual* results really are due to extra utility from winning. There might be equally useful alternatives that one could explore. For example, one might interpret response behavior in the *Individual* treatment as a failure of the perfect rationality assumption and fit the data to a quantal-response-type model. We choose the joy of winning here mainly because it is plausible, very simple to model and effective in producing good approximations of the observed response functions.

We incorporate the idea that players enjoy the event of winning in the simplest possible way by assuming that the value of the prize as *perceived* by the subjects exceeds its objective value of 16. The best response function for $0 < x \leq 16$ is then $y = \sqrt{Px} - x$ where $P > 16$ is the perceived value of the prize. We estimate the value of *P* based on the following simple econometric model:

[10] For example, by investing 16 the first mover maximizes the damage he can inflict on the responder. Yet, the first mover's expected payoff in this case is *lower* than the responder's for *any* response between 0 and 16, except for 0 and 16 themselves where the expected net gain is 0 for both players.

$$z_i = \theta_1 x_i + \varepsilon_i$$

where $i = 1, 2, \ldots 16$, $z_i = \bar{y}_i + i$, $\theta_1 = \sqrt{P}$ and $x_i = \sqrt{i}$. This estimation delivers values of $\hat{P}_{\text{One}} = 23.05$ and $\hat{P}_{\text{Two}} = 24.40$ for parts one and two (again focusing on the last five rounds). Thus, the joy-of-winning aspect in utility, if we take it to be the cause of the inflated bids, is substantial and makes up about one third of the total utility obtained by an individual when he or she receives the prize. Surprisingly, the influence of this bias is even stronger after subjects have gained experience with the environment of the experiment.

Next, we use these results in the best response function under Fehr-Schmidt preferences (see Appendix B),

$$y = \sqrt{\frac{1+\alpha-\beta}{1-\beta} Px - \frac{2(\alpha+\beta)}{1-\beta} x^2} - x .$$

We estimate the parameters α and β using the following model:

$$z_i = \theta_1 x_i + \theta_2 x_i^2 + \varepsilon_i$$

where $z_i = \left(\bar{y}_i + i\right)^2$, $\theta_1 = (1+\alpha-\beta)P/(1-\beta)$, $\theta_2 = -2(\alpha+\beta)/(1-\beta)$ and $x_i = i$. This procedure yields estimates of $\hat{\alpha}_{\text{One}} = 1.551$ and $\hat{\beta}_{\text{One}} = -1.403$ for part one, and $\hat{\alpha}_{\text{Two}} = 0.589$ and $\hat{\beta}_{\text{Two}} = -0.287$ for part two. Thus, the results suggest that the average player dislikes disadvantageous inequality but *enjoys* advantageous inequality. These tastes appear to be especially strong in part one: according to our parameter estimates the average second mover is willing to give up e$1.40 or more if that stops the player she is matched with from receiving just one extra e$. By the time the experiment approaches the end these tastes seem to have worn off to an extent. Figure 3 displays the response functions for part one and part two that arise from our parameter estimation.

Result 2

(a) Subjects' responses in the *Individual* treatment are systematically higher than predicted, suggesting that over-investments in experimental rent-seeking games are not alone due to over-competitiveness between players but also due to the innate characteristics of the underlying optimization problem.

(b) Comparisons between subjects' response functions in the *Strategy* and the *Individual* treatment show that social preferences drive up investments. The revealed preferences are best characterized as spiteful in the sense that the average subject appears to get pleasure from advantageous inequality and is willing to spend parts of her endowment to reduce the income of the player she is matched with.

Our analysis of the response schedules thus far has focused on average behavior. We now investigate potential heterogeneity across individuals. Specifically, we consider the possibility that different conditional investment schedules follow different functional *types*, analogous to the various types that FGF noted in their study on a public good game. In the following we define six such types—*hump-shaped*, *u-shaped*, *increasing*, *decreasing*, *constant* and a residual *other* type—and then classify the observed response functions accordingly. In formulating these definitions, we try to be inclusive enough in the sense that we wish to tolerate minor deviations from the idealized form that gives each category its name, but at the same time strict enough so that the types are mutually exclusive and so that response functions are not allocated to a category when they violate its fundamental characteristics. We emphasize that we devised these categories *before* we looked at the data.

Figure 3: Empirical and estimated average response schedules in *Strategy*

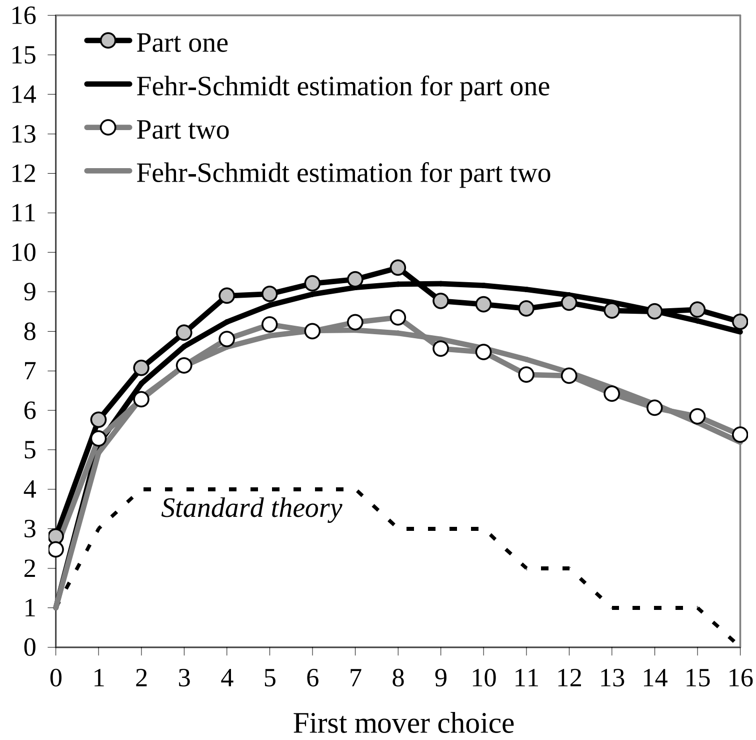

It is useful to begin with a few definitions. First, let $r_{it}(x)$ denote player *i*'s response function in round *t* of the experiment, defined for the range of feasible unconditional choices $x \in \{0,1,\ldots,16\}$. Second, let $\Psi_{it} = \text{Max}\{r_{it}(1), r_{it}(2),\ldots, r_{it}(15)\}$ be the "interior maximum" of individual *i*'s response function in round *t* and let $\vartheta_{it} = \text{Min}\{r_{it}(1), r_{it}(2),\ldots, r_{it}(15)\}$ be its "interior minimum". Finally, we refer to $r_{it}(0)$ as the "start point" of the response function and refer to $r_{it}(16)$ as its "end point".

1. Hump-shaped: We classify a response function as "hump-shaped" if its start point and end point are (a) both strictly lower than the interior maximum, are (b) closer to each other than to the interior maximum, that is $\Psi_{it} - \text{Max}\{r_{it}(0), r_{it}(16)\} > |r_{it}(0) - r_{it}(16)|$, and are (c) not both greater than the interior minimum.

2. U-shaped: We use the reverse criteria to identify a "u-shaped" response function: the start and end point must both be strictly greater than the interior minimum, must be closer to each other than to the interior minimum, that is $\text{Min}\{r_{it}(0), r_{it}(16)\} - \vartheta_{it} > |r_{it}(0) - r_{it}(16)|$, and must not both be lower than the interior maximum.

3. Increasing: A response function is categorized as "increasing" if (A) the start point is strictly lower than the end point, (B) the star point is lower than the average response to values above 0 and the end point exceeds the average response to values below 16, that is

$$r_{it}(0) < \frac{1}{16}\sum_{x=1}^{16} r_{it}(x) \quad \text{and} \quad r_{it}(16) > \frac{1}{16}\sum_{x=0}^{15} r_{it}(x)$$

and (C) start and end point are not closer to each other than to the interior minimum and maximum respectively, that is $|\vartheta_{it} - r_{it}(0)| \leq r_{it}(16) - r_{it}(0)$ and $|r_{it}(16) - \Psi_{it}| \leq r_{it}(16) - r_{it}(0)$.

4. Decreasing: A response function is "decreasing" if conditions (A) and (B) hold with the inequalities reversed, and if $|r_{it}(0) - \Psi_{it}| \leq r_{it}(0) - r_{it}(16)$ and $|\vartheta_{it} - r_{it}(16)| \leq r_{it}(0) - r_{it}(16)$.

5. Constant: A response function is "constant" if $r_{it}(0) = r_{it}(1) = \ldots r_{it}(16)$.

6. Other: The above criteria deliver mutually exclusive categories. All response functions that cannot be classified will be referred to as "other" types.

Applying this classification system to our data reveals some degree of response type heterogeneity. However, most response functions fall into just two categories: "hump-shaped" and "increasing". They are the predominant types in both treatments, encompassing about 85% of all response schedules. Details are given in Table 4.

Table 4: Observed response functions by category

	Part one		*Part two (rounds 11-15)*	
Categories	*Individual*	*Strategy*	*Individual*	*Strategy*
Hump-shaped	69.6%	39.8%	65.4%	49.3%
U-shaped	1.8%	4.6%	0.0%	2.8%
Increasing	17.9%	46.3%	25.7%	34.3%
Decreasing	3.6%	1.9%	2.9%	5.6%
Constant	0.0%	4.6%	0.0%	3.9%
Other	7.1%	2.8%	6.1%	4.3%

In Figure 4, we plotted the average hump-shaped and average increasing response schedule for both treatments and for each part of the experiment. The schedules of increasers and hump-shapers are almost congruent when they respond to low values. When they respond to large values, increasers display a willingness to sacrifice substantial shares of their endowments to keep their chances of winning high, while hump shapers withdraw more and more as the odds turn against them.

Combining the results from Table 4 and Figure 4, several broad features are worth noting. First, looking at part one of the experiment, the distribution of humph-shaped and increasing schedules is very different between the two treatments. While there are almost four times as many hump-shapers as increasers in the *Individual* treatment, there are more increasers than hump-shapers in the *Strategy* treatment. A χ^2 test based on the categories “hump-shaped”, “increasing” and “any other shape” rejects the null hypothesis that the proportions of response types in the two treatments are the same at the 1% level. This is a partial explanation for the divergence between *Individual* and *Strategy* that we observed earlier. Second, as Figure 4 reveals, neither the hump-shaped functions nor the increasing functions are the same between treatments. Subjects in *Strategy* invest generally more, even conditional on the type of response function. It is notable that the hump-shaped response function peaks later and at a higher level than in *Individual*, and that the slope of the increasing response function is steeper than in *Individual* although they both appear to be concave. Third, moving to part two, we observe a slight shift in the relative weight of the two types in both treatments: in the *Individual* treatment the increasing schedule becomes relatively more important, whereas in *Strategy* the hump-shaped type gains at the expense of the increasing type. This shift is not substantial enough to let the type distribution converge between treatments. Fourth, however, the increasing schedules in *Individual* and *Strategy*—distinctly different in part one—now *do* converge in part two. This is not happen in the case of the hump-shaped responses where the difference becomes even slightly larger.

Figure 4: Average response schedules in *Strategy* and *Individual* by type

a. Part one　　　　b. Part two (rounds 11-15)

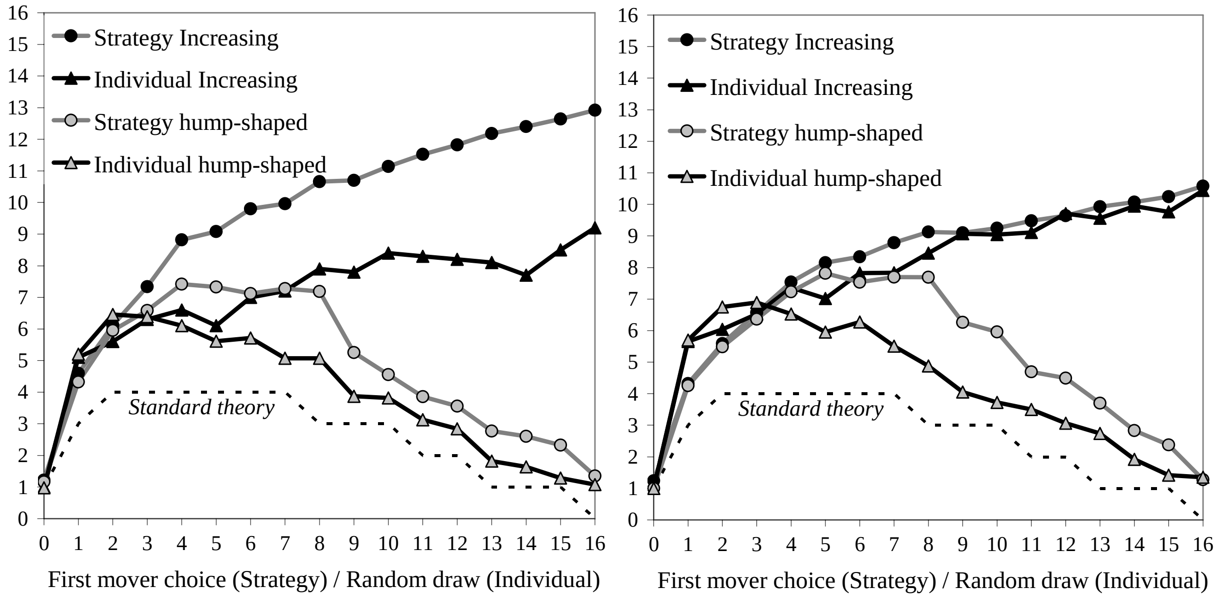

It is not easy to explain these patterns. Using again our two-stage estimation procedure to fit the Fehr-Schmidt model, we still find that the treatment differences in both the increasing schedules and the hump-shaped schedules are best described by assuming negative values for the β-coefficient, in both parts of the experiment. However, this does not explain why we observe these particular shifts in the type distributions over time.

Result 3

Disaggregating the average response schedules in the *Strategy* and *Individual* treatments reveals that most subjects tend to either increase their investment when they respond to higher values or follow a hump-shaped pattern. In part one, both the distribution of these response types and the functional forms of the schedules differ substantially between treatments, indicating more aggressive investment behavior in *Strategy*. In part two, there is some limited convergence in the type distribution and almost complete convergence in the two increasing schedules but no convergence in the hump-shaped schedules.

Our next step will be to investigate in how far the behavior of individuals in the week 1 tests can be used to predict their investment choices in the week 2 experiments. Recall that the week 1 tests elicited social preferences in a prisoner's dilemma task that employed the strategy method and elicited risk attitudes in an individual decision-making task involving 15 binary choice problems, in which the risky alternative was the same throughout (400 points with a 50% chance, 0 points otherwise) but the safe alternative gradually increased from 25 to 375 points.

We first report results relating to the risk attitude elicitation task. In this task the point of switching from the lottery to the safe option can be used as a measure of an individual's risk attitude. We did not impose a unique switching point on subjects and some of them (6%) decided to switch more than once. We drop these observations from our analysis. The median switching point among the remaining participants is at Task 8 where the safe alternative is 200, the expected value of the lottery. However, there is variability among participants: some already switch at Task 4 (where the safe outcome is 100) and the last switchover occurs at Task 13 (where the safe outcome is 325).

The evidence for a relationship between choices in the risk attitude elicitation task and choices in the rent-seeking contests in week 2 is mixed. In the *Direct* treatment the correlation coefficient between an individual's switching point and the individual's investment in part one of the week 2 experiment is positive but very low (0.064). An alternative approach is to divide subjects into two subgroups according to their revealed risk attitude and then compare the investment behavior of the two subgroups (Millner and Pratt, 1991, who first investigated the link between risk attitudes and rent-seeking behavior, use such an approach). Since a risk-neutral agent would switch at either Task 8 or Task 9, it seems natural to classify those who switch at Task 8 or before as "risk-averse" and those who switch at Task 9 or later as "risk-seeking". Based on this classification, 66% of our subjects are classified as risk-averse and 34% as risk-seeking.[11] With an average of e\$10.6 the risk seekers in *Direct* invest somewhat more than the risk-averse participants who spend only e\$8.9 on average—however, this difference is not significant at conventional levels (the p-value is 0.213).

[11] At first glance, this result seems to differ drastically from the findings of Dohmen et al. (2008) whose design we adopted for our study and who report only 9.1% risk seekers. However, their classification approach is slightly different, which accounts for some of the discrepancy. The remaining difference is perhaps due to the fact that in our design the safe option reaches the expected value of the lottery earlier (Task 8 of 15) than in their design (Task 16 of 20).

In the *Strategy* treatment the relationship is somewhat stronger. The correlation coefficient between an individual's point of switching and that individual's *unconditional* investment as a first mover is 0.297, and subjects classified as risk-seeking spend significantly more than subjects classified as risk-averse (the averages are e$11.9 and e$8.8, and the p-value from a two-sided two-sample permutation test is 0.036).

To examine the relationship in terms of *conditional* (i.e. second-mover) choices we ran a regression of the conditional investment on the first-mover choice, in which we included the switching point from the risk attitude elicitation task as an explanatory variable. We report the results from this regression in Table 5 ("Model 1"). The switching point variable is found to be significant at the 1% level, and its coefficient suggests that a one-point increase in the switching point raises an individual's conditional investment by e$0.65. However, note that the R-squared value is low. As a cross check we plotted in Figure 5a the average response functions of subjects classified as risk-seeking and subjects classified as risk-averse, together with the estimated response functions from corresponding regression analyses ("Model 2" and "Model 3"). The disparity in the levels of investment is substantial.

With respect to *types* of response functions we find that risk-seeking players are relatively more likely to choose an increasing schedule instead of a hump-shaped one, but this difference is not statistically significant at the 10% level (the p-value from a χ^2 test based on the categories "hump-shaped", "increasing" and "any other shape" is 0.112).

Result 4

Risk attitudes, as measured by our pre-test, and contest expenditures are not always closely related, but whenever such a link occurs it suggests that risk aversion *reduces* expenditures. This is line with the findings reported by Millner and Pratt (1991).

Moving on to the prisoner's dilemma task, note that the structure of the data from this task resembles that of our *Strategy* data in that we collected both an unconditional first-mover choice and a set of conditional second-mover choices from each subject. Our results are typical for experimental prisoner's dilemma games. We find that 38.7% of subjects choose to cooperate as first movers, while 61.3% choose to defect. In terms of the second-mover choice a slight majority of subjects display pro-social behavior: 4.3% opt for cooperation regardless of the first-mover's choice ("altruists") and 49.5% make the same choice as the first mover ("conditional cooperators"). Almost all of the others (45.2%) defect regardless of the first-mover's choice ("egoists"). One subject cooperated when the first mover defected, and defected when the first mover cooperated.

In the following, we compare the investment behavior of "pro-social" subjects ("altruists" plus "conditional cooperators") with that of "selfish" subjects who defected in both conditional choices. In the *Direct* treatment pro-social individuals invest on average e$10.2 and therefore *more* than selfish individuals (e$8.5 on average). Even though the difference is not significant (the p-value is 0.181), its direction surprised us—for example, if pro-socially inclined individuals are motivated by efficiency considerations (consistent with our universality hypothesis), we would have expected them to invest less than selfish individuals.

With respect to the unconditional first-mover investment in the *Strategy* treatment pro-social individuals spend again more than individuals who we classified as selfish, but the difference is small: the pro-socials invest e$10.0 on average and the selfish individuals e$9.5.

Figure 5: Response schedules for different types of pre-test behavior

a. Revealed risk preferences b. Revealed social preferences

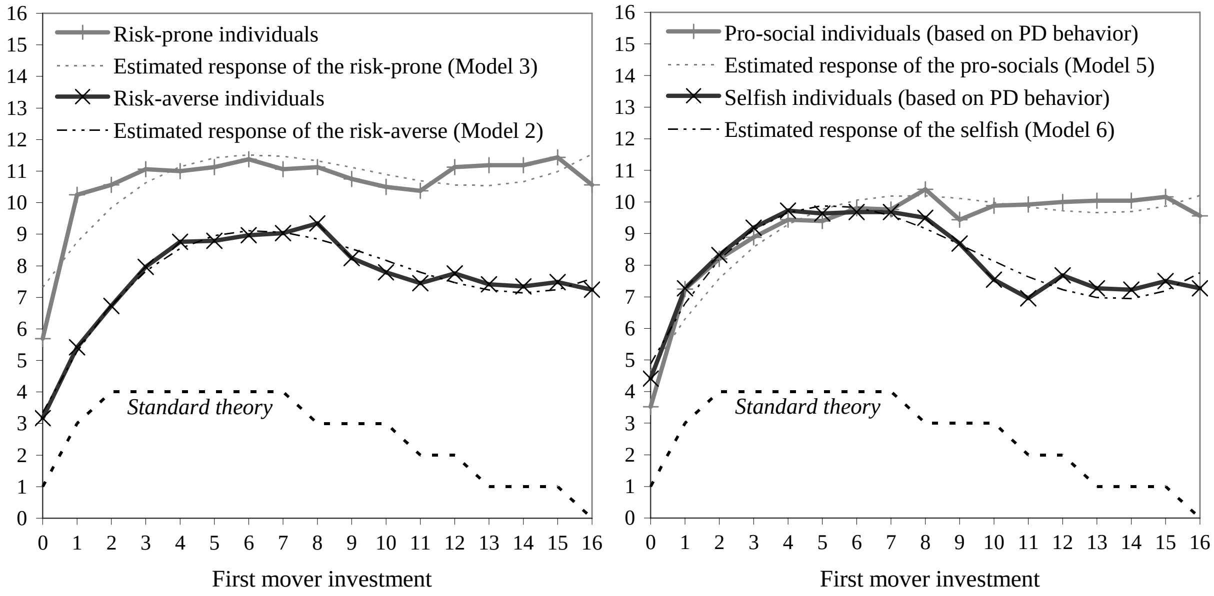

Analogous to our procedure above, we ran a regression to investigate *conditional* investment behavior ("Model 4" in Table 5). The coefficient for the pro-social dummy variable is significant at the 1% level and it is *positive*. In Figure 5b we plotted the average response functions of pro-social and selfish individuals, again together with estimated response functions from corresponding regressions ("Model 5" and "Model 6"). Although the difference is less pronounced than that between risk-seeking and risk-averse participants, it appears that pro-social subjects respond more aggressively to large first-mover investments than selfish subjects do. Note also that pro-socials are not inclined to adhere to principles of positive reciprocity by responding kindly to low first-mover investments: here, their average response is identical to that of the selfish types. In terms of response function types, pro-social subjects are relatively more likely than selfish subjects to select an increasing schedule but this difference is not statistically significant (the p-value is 0.476 based on the same χ^2 test as above).

Finally, we examine whether the pro-social dummy coefficient remains significant once we control for risk attitudes in the regression and, conversely, whether the risk attitude coefficient remains significant when we include the pro-social dummy variable. The results from this regression are reported in Table 5 as "Model 7". While the switching point coefficient remains significant and similar in magnitude compared to Model 1, the coefficient for the pro-social dummy variable is no longer significant and is more than halved compared to Model 4. Thus, an individual's revealed risk attitude is a more robust predictor of conditional rent-seeking investments than the individual's revealed social preference type.

Table 5: Conditional investments, first-mover choice, risk attitude measurement and social preference type—regression results†

Dependent variable: Conditional investment	(1) All	(2) Risk averse	(3) Risk prone	(4) All	(5) Pro-socials	(6) Selfish types	(7) All
Constant	–0.661	3.346***	7.309***	4.194***	4.661***	4.876***	–0.736
	[1.0300]	[0.7683]	[1.0153]	[0.6503]	[0.8796]	[0.8518]	[1.0322]
First-mover investment	1.985***	2.178***	1.635***	1.988***	1.815***	2.184***	1.985***
	[0.3453]	[0.4289]	[0.5668]	[0.3450]	[0.4911]	[0.4755]	[0.3452]
(First-mover investment)^2	–0.233***	–0.253***	–0.198**	–0.235***	–0.189***	–0.287***	–0.233***
	[0.0511]	[0.0634]	[0.0838]	[0.0510]	[0.0726]	[0.0703]	[0.0511]
(First-mover investment)^3	0.008***	0.008***	0.007**	0.008***	0.006**	0.010***	0.008***
	[0.0021]	[0.0026]	[0.0034]	[0.0021]	[0.0030]	[0.0029]	[0.0021]
Switching point	0.650***						0.632***
	[0.0988]						[0.1002]
1 if pro-social				1.067***			0.415
				[0.3802]			[0.3856]
Observations	765	493	272	799	425	374	765
Prob > F	0.000	0.000	0.010	0.000	0.000	0.000	0.000
R-squared	0.105	0.070	0.042	0.063	0.070	0.065	0.106

† The data is based on the first-wave *Strategy* treatment. Models (2) and (3) focus on the conditional investments of risk-averse and risk-prone individuals respectively. Models (5) and (6) focus on the conditional investments of pro-social and selfish individuals respectively. Models (1), (2), (3) and (7) exclude individuals who switched more than once in the risk attitude elicitation task. Numbers in square brackets are robust standard errors clustered on individuals. Significance at the 10%, 5% and 1% level is denoted by *, ** and *** respectively.

Result 5

There is some tendency that subjects who display a pro-social inclination in the prisoner's dilemma task invest more aggressively than subjects who behave selfishly. However, this relationship is either very weak or disappears after we control for risk attitudes.

The last part of our analysis deals with the question whether exposure to competition in rent-seeking contests affects subjects' willingness to cooperate. To investigate this, we let our second-wave subjects participate in the same strategy-method prisoner's dilemma that the first-wave subjects played, but now *after* the main experiment and within the same session. We ran the prisoner's dilemma both after the *Strategy* game and after the *Individual* treatment. In terms of the first-mover choice, 39.3% of the *Individual* subjects choose to cooperate, almost identical to the proportion we observed in the pre-tests (38.7%), but only 26.7% of the *Strategy* subjects opt for cooperation. One interpretation of this reduction is that subjects have become somewhat less likely to put trust in the willingness of an anonymous second-mover to reciprocate cooperation. Is this pessimism justified? The results for the conditional second-mover choice are summarized in Figure 6.

Figure 6: Distribution of response types in the prisoner's dilemma task
(absolute number of subject types shown as data values on chart)

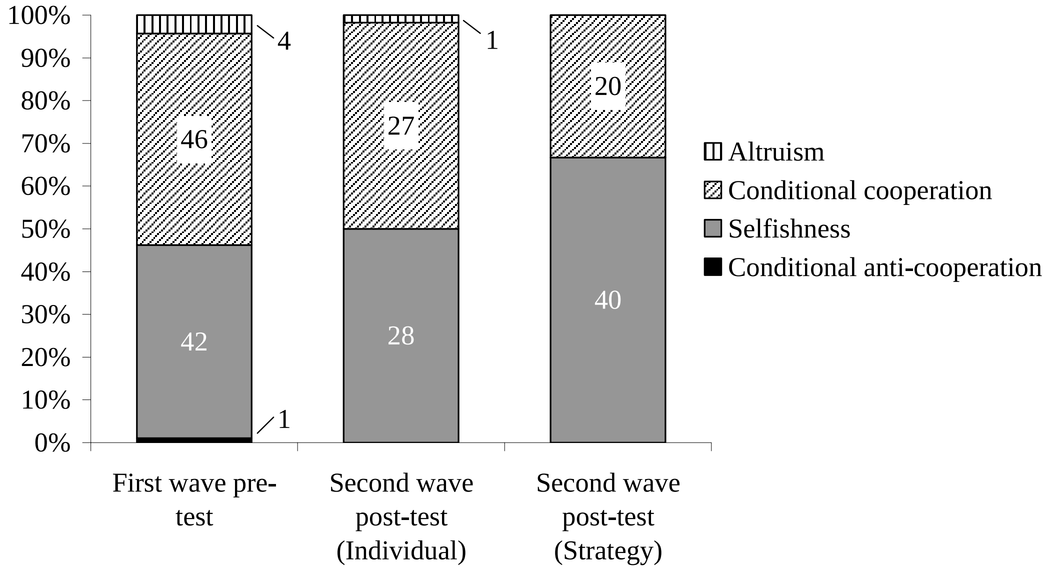

As the figure shows, the distribution of response types in the prisoner's dilemma that followed the *Individual* treatment is also very much the same as that from the pre-test we ran in the first wave of experiments. However, we find a distinct decline in participants' willingness to cooperate after having been subjected to the fierce rent-seeking competition in the *Strategy* treatment. Comparing the two post-tests with a Fisher exact test (pro-social versus selfishness) yields a p-value of 0.090 (two-tailed). If we pool the data from the *Individual* post-test and the first-wave pre-test, the p-value is 0.014. Our interpretation of this result is that the experience of over-competitiveness in the contest game creates a disposition of rivalry in subjects that some cannot immediately "turn off" when the experiment ends. The contest situation appears to trigger anti-social emotions among players that are strong enough to affect their behavior in a subsequent experiment that is set in a very different context.

Result 6

Subjects' willingness to cooperate and to use reciprocity is substantially and significantly reduced after they have experienced the competitive environment of a rent-seeking contest.

5 Conclusions

A broad observation from the literature on experimental rent-seeking games is that standard equilibrium analysis does a poor job in describing behavior and typically underestimates actual expenditures and associated efficiency losses. While we show in this paper that overinvesting relative to standard predictions can be the result of inequality aversion between players—an explanation that has been found to be able to account for deviations from equilibrium in a variety of games—, the key findings from our experiment here suggest a very different story.

First, subjects overinvest to some extent even when there are no other players and social preferences can play no role.

Second, although we find that the introduction of another player pushes up investments further, the patterns we observe are not consistent with inequality aversion but with spite and excessive rivalry between players.

Third, these behavioral patterns are irreconcilable with the often pro-social behavior the same individuals exhibit in a pre-test prisoner's dilemma game. Not only do people who display a pro-social inclination in the prisoner's dilemma show no such inclination in the contest games, they do not even behave less spitefully than those who are categorized as selfish types in the prisoner's dilemma. In fact, on average they invest even more aggressively than the selfish types. Response behavior in the contests is not characterized by patterns of reciprocity, a typical response type in the prisoner's dilemma. Thus, there is there is a striking disconnect in subjects' revealed social preferences across these games. This unambiguous failure of our "universality hypothesis" echoes results from a recent study by Brosig, Riechmann, and Weimann (2007) who also find little consistency in other-regarding preferences across a range of games, and Neugebauer, Poulsen and Schram (2008) who find surprisingly little evidence for reciprocity in the Hawk-Dove game. In our experiment, if one was asked to predict response behavior in the rent-seeking contest one would be better advised to base such a prediction on a prior elicitation of risk attitudes than on a prior elicitation of social preferences.

Fourth, participating in our contest game reduces subjects' willingness to cooperate and reciprocate in a post-test prisoner's dilemma. This result is consistent with the idea that the extent to which individuals are willing to behave pro-socially is influenced by their experiences and the types of economic environments in which they interact (Bowles, 2008). It is also reminiscent of findings from experiments in which the experimenters tried to induce subjects with a positive or negative mood and then observed different levels of pro-social behavior (see, e.g. Capra, 2004, and Kirchsteiger, Rigotti and Rustichini, 2006). Under the hypothesis that pro-sociality in a social dilemma game requires sufficiently positive views of individuals about their anonymous interaction partner, these views certainly seem to suffer from exposure to the fierce rent-seeking competition we observe in our experiment. *Homo reciprocans* is crowded out to some extent.

Our paper contributes to an emerging literature that demonstrates the existence and importance of anti-social behavior in laboratory settings. Despite the mounting evidence on what appear to be spiteful, envious or malevolent preferences we do not propose to place a new economic agent such as a *homo rivalis* next to *homo economicus* and *homo reciprocans*. Instead, because our experiment suggests the possibility that the same person can be turned into any of these types depending on the circumstances, it seems far more important to find guiding principles that can explain how revealed social preferences may change across games. This is undoubtedly a difficult challenge for future theoretical research but note that something similar has been achieved before when the notion of inequality aversion reconciled seemingly antagonistic behavioral patters such as altruistic behavior in social dilemma games and destructive responder behavior in ultimatum games. Similarly, future principles may point to certain modes of social behavior based on particular properties of the games in which players interact. First steps in such a direction have already been made. For example, Zizzo (2003b) introduces the concept of 'game harmony' to measure the degree of conflict inherent in the payoff structure of a game and relates this to cooperation rates in experimental games. Another informative approach may be based on evolutionary arguments. On the one hand, evolutionary game theory has been used to explain the presence of reciprocity and altruism in certain games (e.g. Güth, 1995; Bester and Güth, 1998). For the Tullock contest, on the other hand, it has been shown that *spiteful* behavior and associated overdissipation is evolutionarily stable and robust, whereas selfish behavior is not (Hehenkamp, Leininger and Possajennikov, 2004).

Another path for future research may involve further investigation of how much behavior in particular economic environments may be shaped by prior experiences of cooperative, competitive or anti-social behavior in similar or different environments. A better understanding of such effects and determining their significance to economic behavior may inform the design of effective economic and social policies.

References

Abbink, K., J. Brandts, B. Herrmann and H. Orzen (2008). Inter-group conflict and intra-group punishment in an experimental contest game. Working Paper, University of Nottingham.

Abbink, K., and A. Sadrieh (2008). The pleasure of being nasty. Working Paper, University of Amsterdam.

Anderson, L.R., and S.L. Stafford (2003). An experimental analysis of rent seeking under varying competitive conditions. *Public Choice*, 115, 199–216.

Andreoni, J., and L. Samuelson (2006). Building rational cooperation. *Journal of Economic Theory*, 127(1), 117-154.

Becker, G.S. (1983). A theory of competition among pressure groups for political influence. *Quarterly Journal of Economics*, 98(3), 371-400.

Bester, H., and W. Güth (1998). Is altruism evolutionarily stable? *Journal of Economic Behavior & Organization*, 34(2), 193-209.

Bolton, G.E. and A. Ockenfels (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1), 166-193.

Bowles, S. and H. Gintis (2002). Behavioural science: homo reciprocans. *Nature*, 415, 125-128.

Bowles, S. (2008). Policies designed for self-interested citizens may undermine "the moral sentiments": evidence from economic experiments. *Science*, 320(5883), 1605-1609.

Brosig, J., T. Riechmann and J. Weimann (2007). Selfish in the end? An experimental investigation of the consistency and stability of individual behavior. Working paper, Universität Magdeburg.

Bullock, D.S. and E.E. Rutström (2007). Policy making and rent-dissipation: An experimental test. *Experimental Economics*, 10, 21-36.

Camerer, C.F., and E. Fehr (2006). When does "economic man" dominate social behavior? *Science*, 311(5757), 47-52.

Capra, C.M. (2004). Mood-driven behavior in strategic interactions. *American Economic Review*, 94(2), 367-372.

Cinyabuguma, M., T. Page and L. Putterman (2006). Can second-order punishment deter perverse punishment? *Experimental Economics*, 9(3), 265-279.

Cornes, R., and R. Hartley (2003). Loss aversion and the Tullock paradox. Discussion Paper, University of Keele.

Davis, D.D., and R.J. Reilly (1998). Do too many cooks always spoil the stew? An experimental analysis of rent-seeking and the role of a strategic buyer. *Public Choice*, 95, 89-115.

Dawes, C.T., J.H. Fowler, T. Johnson, R. McElreath and O. Smirnov (2007). Egalitarian motives in humans. *Nature*, 446, 794-796.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp and G.G. Wagner (2005). Individual risk attitudes: new evidence from a large, representative, experimentally-validated survey. IZA Discussion Paper No. 1730, IZA Institute Bonn.

Fehr, E., and U. Fischbacher (2002). Why social preferences matter – the impact of non-selfish motives on competition, cooperation and incentives. *Economic Journal*, 112(478), C1-C33.

Fehr, E., U. Fischbacher and M. Kosfeld (2005). Neuroeconomic foundations of trust and social preferences: initial evidence. *American Economic Review*, 95(2), 346-351.

Fehr, E., and S. Gächter (2002). Altruistic punishment in humans. *Nature*, 415, 137-140.

Fehr, E., and K. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3), 817-868.

Fehr, E. and K. Schmidt (2003). Theories of fairness and reciprocity – evidence and economic applications. In: M. Dewatripont, L.P. Hansen and S.J. Turnovsky (eds.), *Advances in economics and econometrics – Theory and applications, Eighth World Congress* (Econometric Society Monographs), Vol. 1, Cambridge University Press.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.

Fischbacher, U., S. Gächter and E. Fehr (2001). Are people conditionally cooperative? Evidence from a public goods experimen. *Economics Letters*, 71, 397-404.

Fonseca, M.A. (2006). It's not how you play the game, it's winning that matters: an experimental investigation of asymmetric contests. Working Paper, Columbia University.

Gächter, S. (2007). Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications. In: B.S. Frey and A. Stutzer (eds.), *Economics and Psychology - A Promising New Cross-Disciplinary Field* (CESifo Seminar Series), MIT Press.

Gächter, S., B. Herrmann and C. Thöni (2005). Cross-cultural differences in norm enforcement. *Behavioral and Brain Sciences*, 28(6), 822-823.

Güth, W. (1995). An evolutionary approach to explaining cooperative behavior by reciprocal incentives. *International Journal of Game Theory*, 24, 323-344.

Haley, K.J., and D.M.T. Fessler (2005). Nobody’s watching? Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior*, 26, 245-256.

Haltiwanger, J., and M. Waldman (1993). The role of altruism in economic interaction. *Journal of Economic Behavior & Organization*, 21(1), 1-15.

Hehenkamp, B., W. Leininger and A. Possajennikov (2004). Evolutionary equilibrium in Tullock contests: spite and overdissipation. *European Journal of Political Economy*, 20, 1045-1057.

Herrmann, B., R. Kabalin, V. Nedzvetsky, E. Poen and S. Gächter (2008). Emotions and cortisol reactions when cheated in a social dilemma experiment. Draft working paper, University of Nottingham.

Herrmann, B., C. Thöni and S. Gächter (2008): Antisocial punishment across societies. *Science*, 319(5868), 1362-1367.

Holt, C.A., and S.K. Laury (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644-1655.

Kirchsteiger, G., L. Rigotti and A. Rustichini (2006). Your morals might be your moods. *Journal of Economic Behavior & Organization*, 59, 155-172.

Kong, X. (2008). Loss aversion and rent-seeking: An experimental study. Working Paper, University of Nottingham.

Konrad, K.A., and H. Schlesinger (1997). Risk aversion in rent-seeking and rent-augmenting games. *Economic Journal*, 107(445), 1671-1683.

Ledyard, J. (1995). Public goods: A survey of experimental research. In: J.H. Kagel and A.E. Roth (eds.), *The Handbook of Experimental Economics*, Princeton University Press.

Millner, E.L. and M.D. Pratt (1989). An experimental investigation of efficient rent-seeking. *Public Choice*, 62(2), 139-151.

Millner, E.L. and M.D. Pratt (1991). Risk aversion and rent-seeking: An extension and some experimental evidence. *Public Choice*, 69(1), 81-92.

Morgan, J., H. Orzen and M. Sefton (2008). Endogenous entry in contests. Working Paper, University of Nottingham.

Neugebauer, T., A. Poulsen and A. Schram (2008). Fairness and reciprocity in the hawk–dove game. *Journal of Economic Behavior & Organization*, 66, 243-250.

Nishimura, N., T.N. Cason, T. Saijo and Y. Ikeda (2007). Spite and counter-spite in auctions. Working Paper, Purdue University.

Öncüler, A., and R. Croson (2005). Rent-seeking for a risky rent – A model and experimental investigation. *Journal of Theoretical Politics*, 17(4), 403-429.

Orzen, H., and M. Sefton (2008). Price versus quantity competition: Theoretical and experimental investigation of deviations from equilibrium. Working Paper, University of Nottingham.

Parco, J.E., A. Rapoport and W. Amaldoss (2005). Two-stage contests with budget constraints: An experimental study. *Journal of Mathematical Psychology*, 49, 320-338.

Potters, J., C.G. de Vries and F. van Winden (1998). An experimental examination of rational rent-seeking. *European Journal of Political Economy*, 14, 783-800.

Rustichini, A., and A. Vostroknutov (2007). Competition with skill and luck: Behavioral and fMRI experiments. Working Paper, University of Minnesota.

Samuelson, L. (2005). Foundations of human sociality: A review essay. *Journal of Economic Literature*, 43(2), 488-497.

Schmidt, D., R. Shupp and J. Walker (2004). Resource allocation contests: Experimental evidence. Working Paper, Federal Trade Commission.

Schmitt, P., R. Shupp, K. Swope and J. Cadigan (2004). Multi-period rent-seeking contests with carryover: Theory and experimental evidence. *Economics of Governance*, 5, 187-211.

Shogren, J.F., and K.H. Baik (1991). Reexamining efficient rent-seeking in laboratory markets. *Public Choice*, 69(1), 69-79.

Shupp, R. (2004). Single versus multiple winner rent-seeking contests: An experimental investigation. Working Paper, Ball State University.

Tullock, G. (1980). Efficient rent-seeking. In: J.M. Buchanan, R.D. Tollison and G. Tullock (eds.), *Towards a Theory of the Rent-Seeking Society*, Texas A&M University Press.

Tyran, J.-R., R. Sausgruber (2006). A little fairness may induce a lot of redistribution in democracy. *European Economic Review*, 50(2), 469-485.

Vogt, C., J. Weimann and C.-L. Yang (2002). Efficient rent-seeking in experiment. *Public Choice*, 110, 67-78.

Weimann, J., C.-L. Yang and C. Vogt (2000). An experiment on sequential rent-seeking. *Journal of Economic Behavior & Organization*, 41, 405-426.

Zizzo, D.J. (2003a). Money burning and rank egalitarianism with random dictators. *Economics Letters*, 81(2), 263-266.

Zizzo, D.J. (2003b). Harmony of games in normal form. Discussion Paper, University of Oxford.

Appendix A: Instructions

A.1 Instructions for the week 1 social preference task

Welcome! You are about to take part in an experiment in the economics of decision making. You will be paid in private and in cash at the end of next week's session. The amount you earn will depend on your decisions, so please follow the instructions carefully.

During the experiment you will have the chance to earn "experimental dollars" (e$), which will be converted into cash at the end of next week's session, using an exchange rate of

e$1 = 20p.

Thus, the higher your e$ earnings are, the more cash you will receive.

It is important that you do not talk to any of the other participants until the session is over.

In this experiment you will be matched with one other person in this room. Together you form a group. Note that neither during nor after the experiment will you and your partner learn about the identity of the other group member.

<u>How your point earnings are determined</u>

You and your partner each begin with an endowment of e$10. You can either keep these e$10 for yourself or you can transfer them to your partner. If you keep your e$10, you earn exactly those e$10. If you transfer the e$10 to your partner, the amount will be doubled by us and your partner will receive e$20 from you. Conversely, if your partner transfers his/her e$10 to you, we will also double the amount, and you will receive e$20 from your partner. Therefore, there are four possible outcomes:

1. You and your partner keep the e$10: In this case you both earn e$10 each.
2. You and your partner both transfer the e$10 to each other: In this case you both earn e$20.
3. You transfer your e$10 while your partner keeps his/her e$10: In this case you will earn nothing and your partner will earn e$30: his/her own e$10 plus the e$20 he/she has received from you.
4. You keep your e$10 while your partner transfers his/her e$10 to you: In this case you will earn e$30 and your partner will earn nothing.

You and your partner have to decide simultaneously whether or not to transfer the e$10. However, your decision will consist of two different parts: First you have to make a **conditional choice** and then an **unconditional choice**.

<u>The conditional choice</u>

For the conditional choice you will be asked to decide whether or not to transfer the e$10 **depending on your partner's decision**. In other words, you will be asked what you wish to do if your partner has chosen to transfer the e$10 to you, and what you wish to do if your partner has chosen to keep them.

Similarly we will also ask your partner what he/she will do if you transfer your e$10 and what he/she will do if you keep the e$10 for yourself.

Please think carefully before submitting your decisions. After you have pressed the OK button you cannot change them any more.

<u>The unconditional choice</u>

For the unconditional choice you will simply be asked whether or not you wish to transfer your e$10 to your partner. As before, we will also ask your partner the same question. Again, please think carefully about your decision before you click on the OK button.

<u>How the conditional and unconditional choices determine the outcome</u>

For one of the two partners the conditional choice will be relevant and for the other partner the unconditional choice will be relevant. To determine which choice is relevant for which partner, one of you will be called "Partner 1" and the other will be called "Partner 2" (you will see on your screen whether you are Partner 1 or Partner 2) and then a die will be thrown.

If the die shows the numbers 1, 3 or 5, then for Partner 1 the *unconditional* choice is relevant and for Partner 2 the *conditional* choices are relevant. That means that to calculate the final outcome we will

(a) look at Partner 1's unconditional choice, and then

(b) look at what Partner 2 has decided to do, given Partner 1's choice in (a).

If the die shows the numbers 2, 4 or 6, then it is the other way around: for Partner 2 the *unconditional* choice is relevant and for Partner 1 the *conditional* choices are relevant. So to calculate the final outcome we will

(a) look at Partner 2's unconditional choice, and then

(b) look at what Partner 1 has decided to do, given Partner 2's choice in (a).

Please note the following. After everybody in the room has made the conditional and the unconditional choices, we will throw the die. However, you will learn about the final outcome only at the end of next week's session when all experiments are finished.

A.2 Instructions for the week 1 risk elicitation task

This experiment is about choosing between a **lottery** and a **sure payoff**.

In the following you will go through 15 choice tasks. The lottery is the same but the sure outcome varies.

In the lottery there is a 50% chance that you win 400 points and a 50% chance that you win 0 points (the outcome of the lottery is determined by a computerized random mechanism).

On the following screen we will show you the 15 choice tasks. Please decide for each of these whether you wish to play the lottery or take the sure outcome.

Once you have made your decisions for all 15 tasks, the computer will randomly pick one of the tasks.

For this task you will either play the lottery or you will receive the sure outcome, depending on what you have chosen. The exchange rate is 1 point = 1p.

The sure outcomes that our subjects encountered in the 15 tasks were as follows.

Task number	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*	*13*	*14*	*15*
Sure outcome	*25*	*50*	*75*	*100*	*125*	*150*	*175*	*200*	*225*	*250*	*275*	*300*	*325*	*350*	*375*

A.3 Instructions for part one of the Strategy treatment

This experiment will consist of just a single round and your payment will depend on the outcome of this round.

At the beginning of the experiment the computer will randomly match you with another person in this room. The person you are matched with will be your opponent, and in the experiment you and your opponent will compete for a prize, as explained below. Note that you will not learn who your opponent is, neither during nor after today's session. Likewise, your opponent will not learn about your identity either.

Each person will receive e$16 from us. You can use this money to purchase "contest tokens". Each contest token costs e$1 and so you can purchase up to 16 of these tokens. Any money you do not invest into contest tokens will simply be added to your e$ balance and is yours to keep.

After everybody has chosen how many contest tokens to buy, a lottery wheel will determine whether you or your opponent wins the prize. The prize is worth another e$16, and your chances of winning the prize depend on how many contest tokens you have bought and how many contest tokens your opponent has bought. The lottery procedure works as follows:

The lottery wheel is divided into two shares with different colors. One share belongs to you and the other share belongs to your opponent. The size of your share and the size of your opponent's share on the lottery wheel are exact representations of the number of contest tokens you own and your opponent owns**.**

For instance, if you and your opponent have bought the same number of contest tokens, each of you gets a 50% share of the wheel. If you have bought three times as many tokens as your opponent, you get three quarters of the lottery wheel and your opponent gets one quarter.

Once the shares of the lottery wheel have been determined, the wheel will start to rotate and after a short while it will stop at random. Just above the lottery wheel there is an indicator at the 12 o'clock position. If the wheel comes to a halt such that the indicator points at your share you win. If the wheel comes to a halt such that the indicator points at your opponent's share, your opponent takes the prize and you will have lost.

Thus, your chances of winning the prize increase with the number of contest tokens you buy. Conversely, the more contest tokens your opponent buys, the higher the probability that you lose. If one of you doesn't buy any contest tokens, the other one wins the prize with certainty. If none of you buys contest tokens, no lottery takes place and the prize is lost.

How to make the decision on the purchase of contest tokens

You and your opponent have to decide simultaneously how many contest token to buy. However, your decision will consist of two different parts: First you have to make a **conditional choice** and then an **unconditional choice** on the purchase of contest tokens.

<u>The conditional choice</u>

For the conditional choice you will be asked how many contest tokens you wish to buy **depending on your opponent's decision**. In other words, you will be asked 17 questions, namely:

- how many contest tokens you wish to buy if your opponent has chosen to buy 0 contest tokens,
- how many contest tokens you wish to buy if your opponent has chosen to buy 1 contest tokens,
- how many contest tokens you wish to buy if your opponent has chosen to buy 2 contest tokens,
- ...
- how many contest tokens you wish to buy if your opponent has chosen to buy 16 contest tokens.

Similarly we will also ask your opponent how many contest tickets he/she wishes to buy if you buy 0, 1, 2, …, 16 contest tokens.

Please think carefully before submitting your decisions. After you have pressed the "Submit" button you cannot change them any more.

<u>The unconditional choice</u>

For the unconditional choice you will simply be asked how many contest tokens you wish to buy. As before, we will also ask your opponent the same question. Again, please think carefully about your decision before you click on the "Submit" button.

<u>How the conditional and unconditional choices determine the outcome</u>

For the outcome of the experiment, it is *either* your conditional choices and your opponent's unconditional choice that will be relevant, *or* it is your unconditional choice and your opponent's conditional choices that will be relevant. The computer will determine randomly which choice is relevant for which party.

If the random draw determines that the *unconditional* choice is relevant for you and the *conditional* choices are relevant for your opponent, then the computer will

(a) look at your unconditional choice, and then

(b) look at what your opponent has decided to do, given your choice in (a), and then

(c) compose the shares of the lottery wheel according to these decisions.

If the random draw determines that the *unconditional* choice is relevant for your opponent and the *conditional* choices are relevant for you, then the computer will first look at your opponent's unconditional choice and then at your conditional choices, given your opponent's unconditional choice.

A.4 Instructions for part one of the Individual treatment

In this experiment we will ask you to choose an option in each of 17 different decision making tasks. These tasks are labeled "Task 0", "Task 1", "Task 2", and so on, up to "Task 16". At the end the computer will randomly select one of these 17 tasks and the payment you receive will depend on how many e$ you have earned in this selected task.

The tasks are all identical with the exception of one aspect, as explained below.

In each task, you will first receive e$16 from us. You can use this money to purchase "lottery tokens". Each lottery token costs e$1 and so you can purchase up to 16 of these tokens. Any money you do not invest in lottery tokens will simply be added to your e$ balance and is yours to keep.

The more lottery tokens you buy the greater is your chance of winning a prize, worth another e$16. Whether or not you win this prize will be determined by a computerized lottery wheel. Your chance of winning is represented as a share on the lottery wheel and the more lottery tokens you have bought, the greater your share is on the wheel. If you do not buy any lottery tokens you will not get a share and your chance of winning will be zero. If you buy one or more lottery tokens you will get a certain share on the wheel, but the exact size of the share will depend on the task number. Here are some examples (full details for each task will be shown on your computer screen).

In Task 0, buying 1 token is enough to get you a 100% share (so you would win with certainty) and buying additional tokens cannot further increase the size of your share.

In Task 1, if you buy 1 token, the size of your share will be 50%. If you buy 2 tokens, the size of your share will be 67%. If you buy 3 tokens, the size of your share will be 75%. If you buy further tokens your chances will increase further (your computer screen will show exactly how).

In Task 2, if you buy 1 token, the size of your share will be 33%. If you buy 2 tokens, the size of your share will be 50%. If you buy 3 tokens, the size of your share will be 60%. If you buy further tokens your chances will increase further (your computer screen will show exactly how).

…

In Task 16, if you buy 1 token, the size of your share will be 6%. If you buy 2 tokens, the size of your share will be 11%. If you buy 3 tokens, the size of your share will be 16%. If you buy further tokens your chances will increase further (your computer screen will show exactly how).

Once you have made all your choices, the computer will randomly select one of the 17 tasks, and your share on the lottery wheel will be determined based on your choice in this selected task. Then, the wheel will start to rotate and after a short while it will stop at random. Just above the lottery wheel there is an indicator at the 12 o'clock position. If the wheel comes to a halt such that the indicator points at your share you win the prize. Otherwise, you will have lost and you will not receive the prize.

Please think carefully before submitting your decisions. After you have pressed the "Submit" button you cannot change them any more.

Appendix B: Rent-seeking competition with inequality aversion or spite

B.1 The model

We discuss the effects of inequality aversion in rent-seeking contests using the model by Fehr and Schmidt (1999). In the two-player variant of this model the utility of player *i* is given by

$$u_i = \pi_i - \alpha_i \max\left\{\pi_j - \pi_i, 0\right\} - \beta_i \max\left\{\pi_i - \pi_j, 0\right\} \qquad \text{[B1]}$$

where π_i and π_j denote the players' material payoffs ($i \neq j$) and where the coefficients $\alpha_i \geq 0$ and $\beta_i \geq 0$ determine player *i*'s utility loss from disadvantageous inequality ($\pi_i < \pi_j$) and advantageous inequality ($\pi_i > \pi_j$) respectively. Fehr and Schmidt assume that $\alpha_i \geq \beta_i$ (players suffer more from disadvantageous than from advantageous inequality) and that $\beta_i < 1$ for plausibility reasons (if $\beta_i \geq 1$ player *i* would be prepared to burn her own money just for the sake of reducing advantageous inequality). The possibility that players *enjoy* advantageous inequality ($\beta_i < 0$) is also ruled out by Fehr and Schmidt. The reason they give for this restriction is that the games they consider such players would be indistinguishable from players without any advantageous inequality aversion/proneness ($\beta_i = 0$). However, in rent-seeking contests (which Fehr and Schmidt do not discuss) this is not necessarily the case. We will therefore consider the possibility of $\beta_i < 0$.

A defining feature of rent-seeking contests is that "the winner takes all." It should be noted that this creates an ambiguity about how exactly the Fehr-Schmidt model should be implemented. Do players suffer from inequality in the *expected* payoffs that arise from their choices, or do they base their comparisons on *final* payoffs after nature has determined the winner of the contest? This will depend on which fairness norm is applied here. Below, we opt for the perhaps more conventional consequentialist's approach and implement the model using expected utility theory based on the final (relative) payoffs.*

Best response behavior

Based on equation B1 the expected utility for contestant *i* choosing contest expenditure *x* is

$$u_i(x,y) = \begin{cases} E & \text{if } x = y = 0 \\ E - x + \dfrac{x}{x+y}\left(R - \beta_i(R - x + y)\right) - \dfrac{y}{x+y}\alpha_i\left(R + x - y\right) & \text{otherwise} \end{cases} \qquad \text{[B2]}$$

where $0 \leq x, y \leq R$. Thus, when player *i* wins the contest she experiences advantageous inequality, and when she loses the contest she experiences disadvantageous inequality. The best response function for player *i* is

* Nevertheless, both alternatives have some appeal and we do not wish to argue that the selected approach is necessarily the superior one. However, while the two methods yield different predictions, our main conclusions would not change if the alternative approach was used.

$$x^*_{\text{FS}}(y) = \begin{cases} \varepsilon & \text{if } y = 0 \\ \max\left\{0, \sqrt{\dfrac{1+\alpha-\beta}{1-\beta}yR - \dfrac{2(\alpha+\beta)}{1-\beta}y^2} - y\right\} & \text{if } y > 0. \end{cases} \quad \text{[B3]}$$

Comparing this with the standard-model best response, $x^*_{\text{SM}}(y) = \max\{0, \sqrt{yR} - y\}$ for $y > 0$, we find that the two functions intersect at

$$\hat{y} = \frac{1}{2}\frac{\alpha}{\alpha+\beta}R. \quad \text{[B4]}$$

Notice that $\hat{y} \le R/2$. For values below $\hat{y}$ the Fehr-Schmidt response function lies above the standard-model response function, for values above $\hat{y}$ it lies below. Like in the standard model, the Fehr-Schmidt best-response function is concave and peaks below $\hat{y}$. For very high values of *y*—to be precise, beyond $\tilde{y} = (1+\alpha-\beta)R/(2\alpha+\beta+1)$, with $\hat{y} < \tilde{y} < R$—a Fehr-Schmidt player best responds by submitting zero. In Figure B1 we have plotted several response functions for particular values of α and β.

Figure B1: Best response functions in the Fehr-Schmidt model

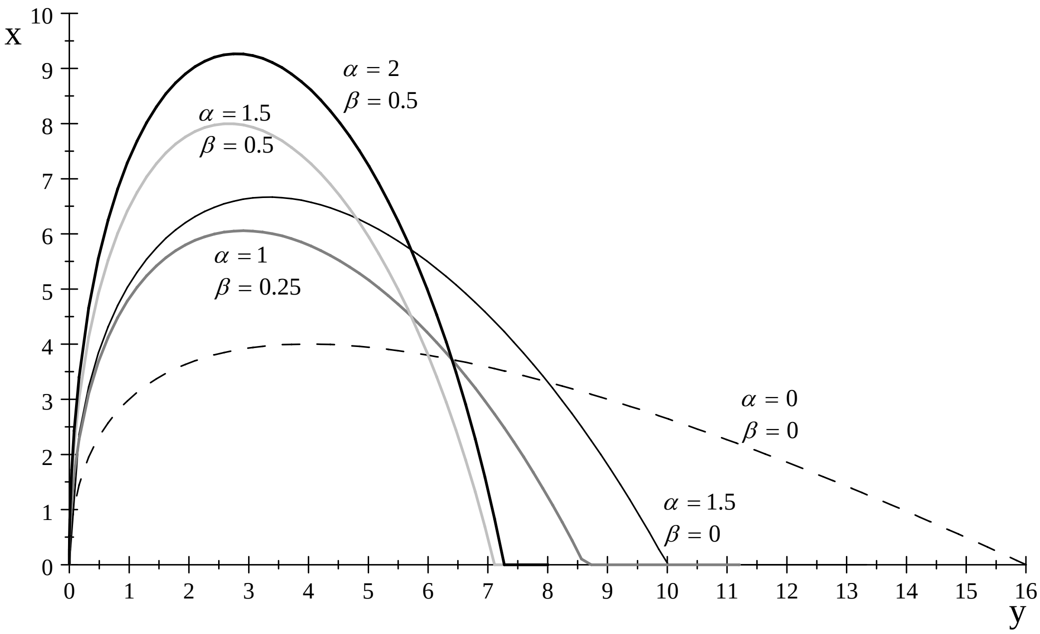

In choosing her investment player *i* faces a dilemma. On the one hand, a high investment makes it less likely that she loses the contest and has to endure disadvantageous inequality. On the other hand, *if* she loses, the level of disutility from disadvantageous inequality increases with additional investments (which, from hindsight, have been wasted). Moreover, the absolute material cost of investing has to be considered.

When *y* is low, the marginal impact of an extra unit of investment on one's chances of winning is relatively high and weighs more than the risk of increased suffering from disadvantageous inequality. This is the main driver behind the high expenditures in response to low values of *y*. However, an important additional factor is that high investments reduce *i*'s disutility from advantageous inequality in the event that she *does* win the contest.

The low expenditures in response to high values of *y* are due to the fact that while the marginal impact of an extra unit of investment on one's chances of winning diminishes as *y* increases, the marginal increase in the level of disutility from disadvantageous inequality does not depend on *y*. Therefore, when *y* is sufficiently high ($y > \hat{y}$), the risk of increased suffering from disadvantageous inequality weighs more than the marginal impact of an extra unit of investment on one's chances of winning. This is the main cause for collapse of expenditures in response to high values of *y*.

Equilibrium

To simplify the equilibrium analysis we assume that $\alpha = \alpha_i = \alpha_j$ and $\beta = \beta_i = \beta_j$ and that this is common knowledge. We then get a unique symmetric equilibrium in which each player invests

$$x^* = y^* = \frac{1+\alpha-\beta}{2+\alpha-\beta}\frac{R}{2}. \qquad \text{[B4]}$$

The equilibrium expenditure is increasing in the difference between α and β. Although it does not exceed $R/2$ in the limit, inequality-averse players may clearly invest substantially more than standard theory predicts and the level of rent dissipation may approach 100%. Surprisingly, even if players are extremely pro-social in the sense that they dislike advantageous inequality as much as disadvantageous inequality ($\beta = \alpha$), they still invest as much as selfish players.

A spiteful taste for advantageous inequality

We now consider the possibility that $\beta < 0$, i.e. the possibility that players enjoy advantageous inequality. In the following we refer to such players as "inequality-prone" players and to players with $\beta > 0$ as "inequality-averse" players. A restriction on inequality-proneness that seems plausible to us is $|\beta| \leq \alpha$, that is the delight players experience from being ahead weighs less than the discontent they suffer from being behind. Our formal analysis above can be applied in the same way and our conclusions from the equilibrium analysis (equation B4) remain, in essence, the same. Now, of course, β contributes to the inflation of expenditures rather than restraining it.

However, the patterns in the best response behavior for inequality-prone players are very different from those for inequality-averse players. This is illustrated in Figure B2 which displays a range of best-response functions for different values of α and β, and shows one of the inequality aversion-based best-response functions from Figure B1 for comparison.

Figure B2: Best response functions for inequality-prone players

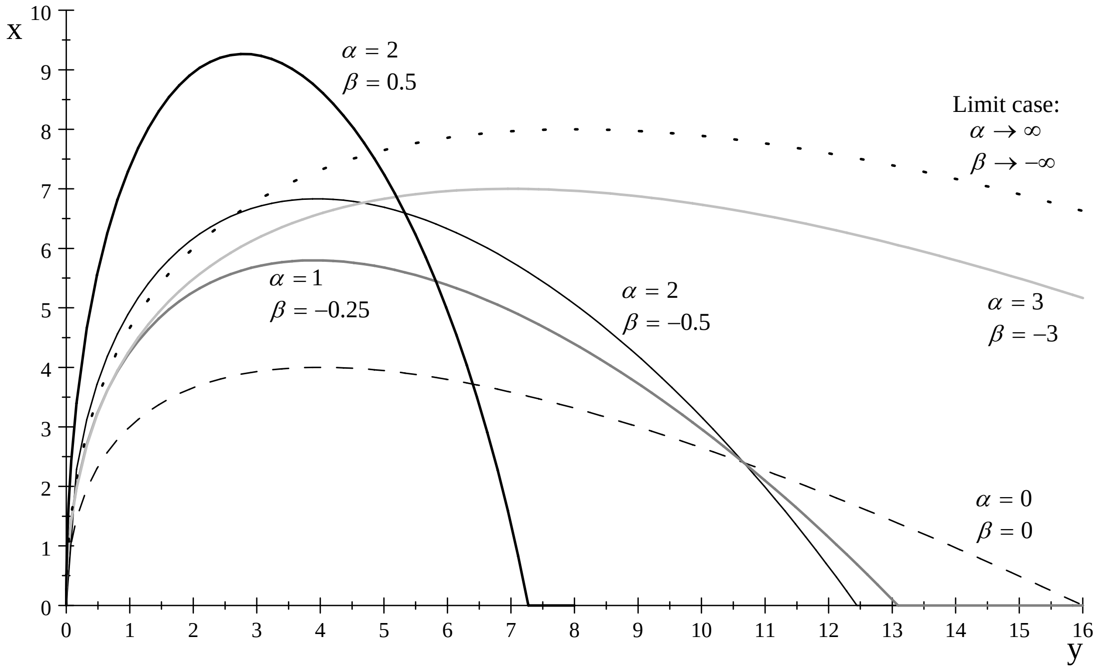

In particular, the point of intersection between the standard-model best-response function and the best response function for inequality-prone players lies at $\hat{y} \geq R/2$, and for $|\beta| = \alpha$ they never intersect. Thus, the response to large values of *y* is fundamentally different from that in the case of inequality-averse players. The reason is that although large investments by player *j* reduce player *i*'s chances of winning, they also mean that *if* player *i* wins, the payoff for player *j* will be particularly low. In the case of inequality aversion this implies a substantial utility *loss* for player *i*, whereas in the case of inequality proneness the implication for player *i* is a substantial *gain* in utility. As a result, inequality-prone players are willing to invest substantially.

On the other hand, when *j* invests little, negative values for β tend to push down the response relative to the case $\beta > 0$. The reason for this is that the more one invests, the lower the joy is of beating the other player. Despite this effect, the response function remains above the corresponding function for selfish players. This is because of the desire to avoid disadvantageous inequality, which is still present.

The limit case highlighted in the figure is an interesting benchmark since this describes a scenario where player *i* aims to maximize the difference between her own and her opponent's payoff.