WZB

Wissenschaftszentrum Berlin
für Sozialforschung

Dietmar Fehr
Julia Schmid

Exclusion in the All-Pay Auction: An Experimental Investigation

Discussion Paper

SP II 2017–202

April 2017

(Revised version of the WZB Discussion Paper
SP II 2014-206)

(WZB) Berlin Social Science Center

Research Area
Markets and Choice

Research Unit
Market Behavior

Wissenschaftszentrum Berlin für Sozialforschung gGmbH
Reichpietschufer 50
10785 Berlin
Germany
www.wzb.eu

Dietmar Fehr, Julia Schmid
Exclusion in the All-Pay Auction: An Experimental Investigation

Affiliation of the authors:

Dietmar Fehr
University of Heidelberg

Julia Schmid
DIW

Wissenschaftszentrum Berlin für Sozialforschung gGmbH
Reichpietschufer 50
10785 Berlin
Germany
www.wzb.eu

Abstract

Exclusion in the All-Pay Auction: An Experimental Investigation

by Dietmar Fehr and Julia Schmid*

Contest designers or managers who want to maximize the overall revenue of a contest (relative performance scheme) are frequently concerned with a trade-off between contest homogeneity and inclusion of contestants with high valuations. In our experimental study, we find that it is not profitable to exclude the most able bidder in favor of greater homogeneity among the remaining bidders, even if the theoretical exclusion principle predicts otherwise. This is because the strongest bidders are willing to give up a substantial part of their expected rent and prefer a strategy that ensures a lower but secure pay-off.

Keywords: all-pay auction, contests, heterogeneity, superstars, experiments

JEL classification: C72, C92, D84

---

* E-mail: dietmar.fehr@awi.uni-heidelberg.de, jschmid@diw.de.
We would like to thank a coeditor and two anonymous referees for very detailed and useful comments. Financial support from the Deutsche Forschungsgemeinschaft (DFG) through the SFB 649 “Economic Risk” is gratefully acknowledged.

1 Introduction

Relative performance schemes seem important for fueling the effort and performance of agents in organizational settings and also in many other domains of our society (see e.g., Frank and Cook 1995). For example, firms often install promotion tournaments and sales competitions, lobbyists compete for influence in the political domain, or researchers compete for research grants. All these examples have in common that rewards are allocated based on relative rather than absolute performance, that the effort of the losers is lost and that the contest designer's main focus is the overall performance of the bidders. The closeness of competition and thus the composition of contestants is a critical design parameter for a contest designer, as a too heterogeneous contest may have adverse effects on agents' performance.

In recent years many sports, for example, have seen the presence of dominant athletes, such as Roger Federer or Novak Djokovic on the Tennis ATP Tour, or Tiger Woods on the Golf PGA Tour. These superstars typically create a lot of attention and serve as the face of their sport. However, too great a dominance by one athlete might also lead to boredom and a lower level of competition. For example, due to Michael Schumacher's dominance in Formula One racing, the viewing figures dropped and, consequently, the FIA changed several of their rules to make the races more tense (BBC 2002).[1,2]

These examples vividly illustrate the trade-off between the inclusion of superstars and contest homogeneity. Baye, Kovenock, and de Vries (1993) provide the theoretical foundations of this trade-off and show that under specific assumptions the exclusion of the strongest bidder can lead to higher revenues for the contest designer (*exclusion prin-*

[1] Another anecdote recounts that Tiger Woods' landslide victory at the 1997 Masters in Augusta has led to deliberations about redesigning the Augusta National.

[2] Likewise, US professional sport leagues (e.g., the NBA, NFL, NHL or MLB) put a lot of effort into creating homogeneity among the competing teams. For example, the rookie drafting system tries to ensure a more balanced competition in the medium to long run by giving the weaker teams from the previous season the right to have first choice of the rookies from the pool of the best junior prospectives. This is in contrast to many sport leagues in Europe, for example soccer leagues, which are usually dominated by a few teams.

ciple). In contests with one prize, the presence of a strong bidder may decrease the bids of the weaker bidders, which in turn may also reduce the bid of the strongest bidder.[3] As a result this can lead to a lower overall performance. The idea behind the *exclusion principle* is to increase the bids of the remaining bidders by creating a smaller but more homogeneous contest.

This paper presents an experimental test of the *exclusion principle*. That is, we attempt to answer the question of whether a heterogeneous group with one strong bidder or a smaller but more homogeneous group maximizes total revenue for the contest designer. We implement a repeated all-pay auction with three bidders and complete information about bidders' valuations of the prize. The valuations in a bidding group are heterogeneous, i.e., a group consists of one strong bidder and two weaker bidders. In order to test the *exclusion principle* we randomly vary the participation of the strongest bidder in a bidding group and compare total revenues when there is no exclusion of the strongest bidder with total revenues in the smaller homogeneous contest where the strongest bidder is excluded.

We find little support for the theoretical predictions. In homogeneous contests, i.e., contests in which the strongest bidder is excluded from participation, we observe revenues close to the theoretical prediction. However, we find no support for the *exclusion principle* as excluding the strongest bidder is, on average, not beneficial for the contest designer. In fact, revenues are substantially higher in the condition where the strongest bidder participates than in the condition in which the strongest bidder is excluded. This is independent of the strength of the strongest bidder. That is, revenues are comparable in situations where the strongest bidders' valuation of the prize is almost twice as high as the valuation of the second-strongest bidder and where it is more than three times higher. Therefore our findings indicate that in our setting the presence of superstars is

[3]Using data from the Professional Golf Association (PGA) Tour, Brown (2011) shows that the participation of Tiger Woods leads to a worse performance (more strokes) of other participating high-skilled professionals (but not low-skilled professionals) compared to when Tiger Woods is not participating in the tournament.

not detrimental to contest revenues.

The main reason for the failure of the *exclusion principle* is the behavior of the strongest bidders as they considerably overbid when they participate in the contest. Although the weaker bidders increase their effort significantly when the strongest bidder is excluded, they cannot compensate for the lost revenue of the strongest bidder. The strongest bidders often choose a strategy that guarantees they will win the prize, which involves bids equal to or higher than the valuation of the second-strongest bidder. In most cases this strategy results in lower profits than the expected profit from playing the mixed equilibrium strategy. Thus, strong bidders are willing to forgo a possibly higher profit in order to avoid losing the auction. Subjects are more likely to choose this "*safe*" strategy if the rent from playing this strategy is larger. In other words, the larger the difference in the valuations of the strongest and second-strongest bidder, the more often we observe the use of the *safe* strategy.[4]

The results presented in this paper are linked to a large experimental literature on contests (for a comprehensive survey see Dechenaux, Kovenock, and Sheremeta 2012). While this literature puts much emphasis on tournaments, Tullock contests, and incomplete information all-pay auctions, a smaller number of papers focus on complete information all-pay auctions (e.g., Davis and Reilly 1998, Gneezy and Smorodinsky 2006, Lugovskyy, Puzello, and Tucker 2010, Ernst and Thöni 2013, Cason, Masters, and Sheremeta 2010, or Llorente-Saguer, Sheremeta, and Szech 2016).[5] In all-pay auctions with complete information all equilibria are in mixed strategies, and most papers concentrate on the symmetric all-pay auction (with the exception of Davis and Reilly 1998, Cason, Masters,

[4]In the appendix we present additional evidence from two treatments where valuations vary across periods and thus allows us to study the *exclusion principle* for a broad range of parameters. We show that in about 80 percent of cases excluding the strongest bidder does not pay off, even though it should in theory. This is more likely the case when the strongest bidder is far superior to the other bidders. Again, a major reason for this is the excessive bidding and the associated prevalence of *safe* bidding of the strongest bidders.

[5]Hillman and Riley (1989) and Baye, Kovenock, and de Vries (1996) provide a theoretical account of all-pay auctions with complete information and Konrad (2009) provides an extensive review of the theoretical literature on contests.

and Sheremeta 2010, and Llorente-Saguer, Sheremeta, and Szech 2016). There are two noteworthy observations that emerge from these studies. First, subjects tend to overbid in comparison to the Nash equilibrium.[6] Second, bidding behavior is bimodal.[7] That is, while subjects seem to randomize their bids, they typically place too much weight on zero or low bids as well as on high bids.[8] Our results provide further support for these two observations. We find significant overdissipation by the strongest bidder (similar to Davis and Reilly 1998) as well as evidence that weaker bidders frequently drop out of the auction. Such a discouragement effect evoked by the presence of a strong bidder can, for example, translate into a lower entry rate of weaker contestants (Cason, Masters, and Sheremeta 2010). We contribute to this literature by investigating asymmetric all-pay auctions and, in particular, by testing whether the exclusion of the strongest bidder increases total revenues for the contest designer (*exclusion principle*).

2 Theory and Experimental Design

2.1 All-pay Auction and Theoretical Predictions

We consider the case of an all-pay auction with complete information as analyzed by Hillman and Riley (1989) and Baye, Kovenock, and de Vries (1993) with one prize and up to three bidders. All participants in the auction are assumed to be risk neutral and they value the prize differently, where a high valuation can alternatively be interpreted as a bidder having low costs of exerting effort in the contest. The valuations $v_i, i \in \{1, 2, 3\}$, are commonly known and are heterogeneous in our setup, such that they can be ordered

[6]Anderson, Goeree, and Holt (1998) show that this overdissipation pattern can be explained by a logit equilibrium in which agents commit mistakes by choosing bidding strategies that do not give the highest expected payoff.

[7]Overbidding and heterogeneity in bidding is also a common pattern in lottery contests. For an overview see Sheremeta 2013.

[8]Bimodal bidding is also frequently observed in all-pay auctions with incomplete information, in which subjects tend to bid only if their valuations are above a certain cut-off level and abstain from bidding otherwise (see e.g., Müller and Schotter 2010, Noussair and Silver 2006, or Barut, Kovenock, and Noussair 2002).

as $v_1 > v_2 > v_3$. All participating bidders simultaneously submit their bid x_i. The bidder with the highest bid x_i wins the auction, receives the prize that she values v_i, and pays her bid x_i. All other bidders lose their bid without gaining anything. Ties are broken randomly.

In this setup, a unique mixed strategy equilibrium exists that is described in the following. With one prize, only the two bidders with the highest valuations actively participate in the auction. The bidder with the third-highest valuation remains inactive, as his expected value from participating in the contest is negative. The bidder with the highest valuation in the contest randomizes continuously and uniformly over $[0, v_2]$, where v_2 denotes the second-highest valuation among the participating bidders. The bids of the bidder with the second-highest valuation v_2 are also uniformly distributed, given that he submits a positive bid. However, he remains inactive, i.e., bids zero, with probability $(1 - v_2/v_1)$, where v_1 denotes the highest valuation among the participating bidders. Therefore, the strongest bidder randomizes according to the distribution function $G_1(x) = x/v_2$ and the second-strongest bidder according to $G_2(x) = 1 - v_2/v_1 + x/v_1$. The expected bid of the bidder with the highest valuation in a period is $\mathbb{E}[x_1] = v_2/2$ and the expected bid of the bidder with the second-highest valuation in a period is $\mathbb{E}[x_2] = (v_2)^2/2v_1$.

In expectation, the strongest bidder in the auction receives a payoff of $v_1 - v_2$, whereas the expected payoff of the second-strongest bidder is zero. The expected sum of bids, i.e., the revenue of the auction, adds up to $W(v_1, v_2) = \left(1 + \frac{v_2}{v_1}\right)\frac{v_2}{2}$. Thus, in order to maximize the auctioneer's revenue, the bidder with the highest valuation, v_1, should be excluded from the auction whenever

$$\left(1 + \frac{v_2}{v_1}\right)\frac{v_2}{2} < \left(1 + \frac{v_3}{v_2}\right)\frac{v_3}{2}. \tag{1}$$

This inequality is fulfilled if $v_1 >> v_2 \geq v_3$, i.e., if v_1 is sufficiently large compared to the other valuations. The intuition behind this result is straightforward. The presence of a very strong bidder not only discourages the weakest bidder v_3 from participating, but also decreases the probability of participation and thus the expected bid of the second-strongest bidder v_2. Excluding the strongest bidder v_1 can thus increase the participation and bids of both weaker bidders. How profitable the exclusion of the strongest bidder is, depends on the valuation of the strong bidder and on how small the difference $v_2 - v_3$ is, since the expected revenue from exclusion is increasing in v_1 and v_3, see inequality (1).[9] Accordingly, the auctioneer might prefer a contest with individually weaker but more homogeneous bidders over a contest with a far superior bidder that leads to a less intense competition. In the remainder we will refer to the bidder with valuation v_1 as the high type. The bidders with valuations v_2 and v_3 are referred to as medium type and low type, respectively.

2.2 Design

The experiment consists of two parts. In each session we first elicit subjects' risk attitudes and we then run the all-pay auction with complete information described above.

The theoretical model assumes risk-neutral players, but risk aversion is an often proposed candidate to explain behavior in auctions. In order to have a measure of subjects' risk attitudes, we directly elicit risk preferences using a binary lottery procedure (see e.g., Holt and Laury 2002, Dohmen and Falk 2011). The procedure includes 15 decisions between a binary lottery and a safe option. The binary lottery is always the same, paying €4 or nothing with a 50 percent chance each, while the safe option increases from €0.25 to €3.75 in steps of 25 cents. A weakly risk-averse person would prefer the safe option over the lottery for safe options lower or equal to €2.[10]

[9] Note that the difference in revenues of the no-exclusion and the exclusion condition is concave in the difference $v_1 - v_2$.

[10] This holds for subjects with monotonic preferences. We did not enforce monotonic preferences but point out in the instructions that we assume that subjects stick to their decision once they have switched from the lottery to the safe option. In our data, 17 out of 144 subjects (12 percent) switched multiple times

After the first task, subjects repeatedly play the all-pay auction for 50 periods. In the beginning, we randomly assign subjects to a six-person group (matching group), which is fixed for the remainder of the experiment. Within a matching group we randomly match subjects into two bidding groups of three in each of the 50 periods. The bidders differ only with respect to their valuations $v_1 > v_2 > v_3$, i.e., each bidding group consists of a high, medium, and low type. These valuation are randomly assigned to a subject in each period. While we always use the same valuations for the medium type ($v_2 = 16$) and low type ($v_3 = 15$), we vary the valuation of the high type $v_1 \in \{30, 51\}$. More specifically, we either use valuations $v \in \{30, 16, 15\}$, hereafter denoted as treatment *High30*, or $v \in \{51, 16, 15\}$, hereafter denoted as treatment *High51*, in a session.

In each period, the bidder with valuation v_1 is excluded from the auction with probability $p = 0.5$.[11] Subsequently, subjects learn whether the auction is run between two or among three bidders before placing their bids, and know the valuation of the other bidders. Bids are unrestricted and subjects can use a resolution up to three decimal places. At the end of each period they are given information on their earnings and the winning bid. Bidders who are excluded from participation are also informed about the winning bid, but do not earn anything in that period. To facilitate the understanding of the strategic aspects of the auction, subjects experienced each bidder role $v_1 > v_2 > v_3$ over time.

In both *High30* and *High51*, it is profitable for the auctioneer to exclude the bidder with the highest valuation v_1 from a theoretical perspective.[12] High types face in both settings two bidders with valuations $v_2 = 16$ and $v_3 = 15$, and thus should bid the same in expectation, as their behavior depends only on v_2. In contrast, the behavior of a

between the safe option and the lottery.

[11] The instructions stated that the computer will randomly decide with a probability of 50 percent whether the group member with the highest valuation is excluded in a period. For the detailed explanation in the instructions, see Appendix C.1.

[12] The predicted overall revenues in the no-exclusion condition are 12.27 in *High30* and 10.51 in *High51*, whereas the revenues in the exclusion condition are 14.53 in both treatments since the valuations of medium and low types are always the same. While the absolute difference between the exclusion and no-exclusion condition seem small, the relative difference is substantial (18 and 38 percent).

medium type depends on v_1. Therefore the share of zero bids should increase and overall revenues should decrease with the distance $v_1 - v_2$. Exclusion of the high type results in a relatively homogeneous bidder group where both the medium and the low type increase their bids substantially, yielding higher revenues overall. Our primary aim is therefore to compare the revenue of an auction with two "*homogeneous*" bidders with valuation v_2 and v_3 (exclusion condition) to the revenue of an auction with all three bidders with valuations $v_1 > v_2 > v_3$ (no-exclusion condition) and to explore whether the strength of the high type matters for the *exclusion principle* from a behavioral perspective.[13]

We conducted seven computerized sessions with 18–24 participants each at the experimental laboratory at the TU Berlin using the software tool kit *z-Tree* (Fischbacher 2007). Because we randomly matched subjects in groups of six (matching groups), we have observations from 24 independent groups in total, which provide the basis for our statistical analysis below.[14]

Subjects were recruited from a large database where students can voluntarily register for participating in experiments (ORSEE, Greiner 2015). Upon entering the lab, subjects were randomly assigned to computer terminals. First, the instructions for the lottery choice procedure were displayed on their computer screen. At that point subjects had no information about their subsequent task in the second part of the experiment. After completing the lottery choice task, subjects received written instructions for the all-pay auction, including a test to confirm their understanding. We only proceeded with the sec-

[13] Prior to the sessions reported here, we ran four sessions that used a slightly different setup. In these sessions, bidders faced different sets of valuations in each period and thus a more complex strategic situation. Specifically, we randomly drew the valuations for the two weaker bidders, v_2 and v_3, from a discrete uniform distribution over the interval $[11, 20]$ and the valuation v_1 from a discrete distribution over the interval $[15, 55]$ in each period and randomly assigned them to subjects. All valuations were drawn before the experiment and we constructed two treatments based on these valuations. That is, in one treatment the valuations were sufficiently heterogeneous such that the exclusion of the high type v_1 was always profitable for the contest designer, whereas in the second treatment the exclusion of the high type is never profitable. This allowed us to analyze the *exclusion principle* and bidding behavior in a rich environment that is not idiosyncratic to a specific choice of valuations. We present the results of these two treatments in Appendix B and show that our findings from the treatments reported here are robust to the assignment of valuations (random or fixed valuations).

[14] Note that we treat a matching group as an independent observation because the behavior over time is likely to depend on previous interactions in a bidding group.

ond part after all subjects had answered all test questions correctly. In addition, there was a trial period to familiarize subjects with the computer interface and the auction format. At the end of the second part of the experiment, the computer randomly drew 10 out of the 50 periods to determine subjects' earnings. The sum of points in these 10 periods plus the earnings from the lottery choice task were exchanged at a rate of 10 points = €1. Additionally, participants received an initial endowment of €10 to cover potential losses. In total, 144 students (95 males and 49 females) from various disciplines participated in the experiment. Sessions lasted about 90 minutes and subjects' average earnings were approximately €15.

3 Results

3.1 Aggregate Results and the Exclusion Principle

We begin our analysis by looking at the variable of greatest interest to the contest designer: the revenue of the contest. Table 1 presents the summary statistics of observed behavior along with the theoretical prediction for the pooled data set and for both treatments separately broken down into the exclusion and no-exclusion condition. The exclusion condition consists of all situations in which the bidder with the highest valuation (v_1) is excluded from participating in the auction, whereas in the no-exclusion condition all three bidders participate. According to the exclusion principle, we would expect that exclusion leads to higher revenues for the contest designer in both treatments.

However, in strong contrast to this prediction we find that the exclusion of the bidder with the highest valuation never generates higher revenues than a situation with three bidders. In the pooled data set the average sum of bids is about 15.9 with exclusion and 18 without exclusion. Clearly, we can reject the hypothesis of equal revenues in the two conditions (Wilcoxon signed-rank test, $z = 2.4$, $p < 0.017$, $n = 24$), albeit not in favor of our alternative hypothesis that exclusion is profitable. Looking at each treatment

Table 1: Summary statistics of bids in the control treatment

	Pooled		***High30***		***High51***	
	No excl.	Excl.	No excl.	Excl.	No excl.	Excl.
avg. sum of bids	18.02	15.9	17.81	16.88	18.3	14.54
	(9.05)	(8.42)	(8.83)	(8.6)	(9.35)	(7.97)
pred. sum of bid	11.54	14.53	12.27	14.53	10.51	14.53
N	*1203*	*1193*	*707*	*692*	*496*	*501*
minimum bid	0	0	0	0	0	0
maximum bid	57.5	61.6	57.5	61.6	53	31

Notes: Standard deviations in parentheses. No exclusion (No excl.) refers to situations in which all three bidders participate and exclusion (Excl.) refers to situations where only the medium and low type participate.

separately confirms that it does not pay off to exclude the strongest bidder in our setup. However, while the average sum of bids in the no-exclusion condition is higher than revenues with exclusion in both *High30* and *High51*, the difference is less pronounced when the high type is weaker, i.e., when $v_1 - v_2$ is smaller as in *High30*.[15]

From Table 1 it is apparent that, on average, revenues are always higher than predicted (overbidding), except in the exclusion condition in *High51*. When all three bidders participate (no-exclusion condition), we observe that revenues are between 1.5 and 1.7 times higher than predicted. Accordingly, we find that in 79 percent of periods revenues are higher than predicted when all three bidders participate in the auction (*High30*). This share is slightly higher in *High51* at 85 percent. The observed overbidding is less prominent in the exclusion condition. For example, the sum of bids is about 1.2 times higher than predicted in the exclusion condition in *High30*, whereas the average revenues are close to the prediction in *High51*.

Why is it the case that exclusion does not lead to higher revenues? We have seen that there is substantial overbidding in the presence of three bidders and we can ask whether exclusion would have been profitable if the strongest bidders had behaved as

[15] The difference in the average sum of bids between the two conditions (exclusion and no exclusion) is not statistically different in *High30* (Wilcoxon signed-rank test, $z = 0.97$, $p = 0.33$, $n = 14$), but in *High51* (Wilcoxon signed-rank test, $z = 2.4$, $p < 0.017$, $n = 10$).

Table 2: Summary statistics of individual bids of bidder types

		High30			High51		
		High	Medium	Low	High	Medium	Low
no exclusion	avg. bid	13.74 (4.92)	2.36 (5.19)	1.71 (4.57)	15.07 (6.71)	2.15 (5.04)	1.09 (3.58)
	avg. predicted bid	8	4.27	0.00	8	2.51	0.00
exclusion	avg. bid	-	10.39 (5.44)	6.49 (6.59)	-	9.09 (5.23)	5.45 (5.81)
	avg. predicted bid	-	7.5	7.03	-	7.5	7.03

Notes: Standard deviations in parentheses. We excluded bids $x_i > 55$. This was the case in 4 out of 7,200 individual bids.

prescribed by theory. For this thought experiment, we calculate the revenues in the no-exclusion condition using the actual bids of the two weaker bidders and the theoretical bid of the strongest bidder. As predicted by the *exclusion principle*, this calculation shows that revenues without exclusion would be lower than with exclusion in *High30* (12.1 vs 16.88) and in *High51* (11.3 vs 14.54). In both cases the difference in revenues is statistically different (Wilcoxon signed-rank test $z = 3.2$, $p < 0.01$, $n = 14$ in *High30* and $z = 2.5$, $p < 0.015$, $n = 10$ in *High51*). This counterfactual analysis suggests that the behavior of the strongest bidder plays a major role in explaining why the theory is not predictive.

3.2 Individual Behavior

The preceding analysis has suggested that overbidding with respect to the theoretical prediction plays an important role for the unprofitability of exclusion. To get a deeper insight into the underlying reasons, we will now turn to a more thorough analysis of the three bidder types. Table 2 provides an initial overview of the average bids of each bidder type in the no-exclusion condition (top panel) and the exclusion condition (bottom panel) for each treatment. Figure 1 presents the cumulative distribution of bids for each type in the no-exclusion condition (left panel) and the exclusion condition (right panel).

It is striking that the strongest bidders bid more than predicted, as evidenced in

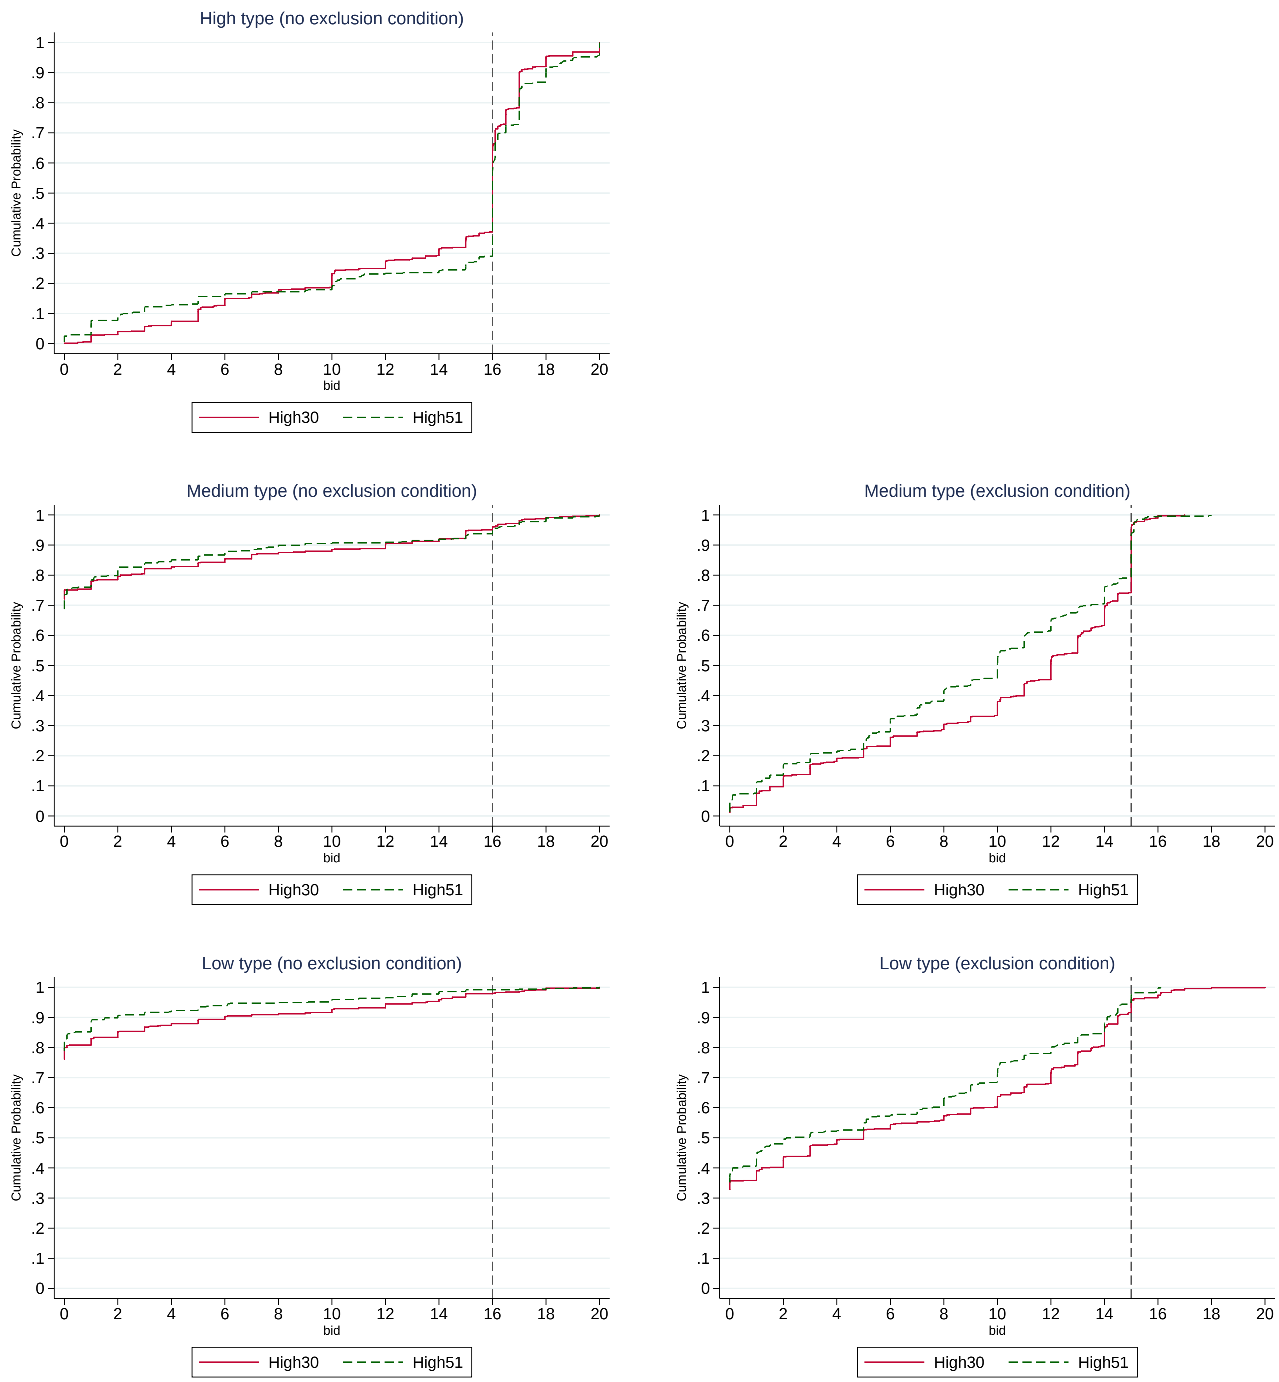

Figure 1: Cumulative distribution of bids of types in the exclusion and no-exclusion condition.

Table 2. If high types participate in the auction (no-exclusion condition), they bid 1.7 and 1.9 times more than predicted by theory, respectively (Table 2, top panel). The difference between actual bids and predicted bids is statistically significant in both cases (Wilcoxon signed-rank tests, $p < 0.01$). It is also noteworthy that in the no-exclusion condition medium types bid less than predicted and that low types participate too much. A similar picture emerges in the exclusion condition, where again the bidders with the highest valuation, i.e., the medium types, bid on average more than predicted. Again, the difference between actual bids and predicted bids is statistically significant in both cases (Wilcoxon signed-rank test, $z = 3.3$, $p < 0.01$, $n = 14$ in *High30*, and $z = 2.3$, $p < 0.025$, $n = 10$ in *High51*).

A closer look at the bidding behavior of the strongest bidders reveals an interesting regularity in both treatments and conditions. According to theory, the strongest bidder's bid should be uniformly distributed over the interval $[0, v_2]$, where v_2 denotes the valuation of the bidder with the second-highest valuation (either the medium type in the no-exclusion condition or the low type in the exclusion condition). Over the periods, there should not be any mass points or bids at or above v_2. Yet, we observe behavior that is completely distinct from this prediction.

From Figure 1 it is apparent that a substantial share of bids is equal or above the valuation of the second strongest bidder, i.e., $x_1 \geq v_2$. This is particularly the case in the no-exclusion condition (top-left panel of Figure 1). When high types participate in the auction, in about 63 percent of cases we observe bids that are equal to or higher than the valuation of the second strongest bidder in *High30*. If the contest is even more heterogeneous, as in *High51*, this fraction is 11 percentage points higher. Such bidding behavior with a mass point at the second-highest valuation is also present in the exclusion condition, albeit to a much lesser extent (top-right panel of Figure 1). About 26 percent of bids in *High30* and 21 percent of bids in *High51* are equal to or larger than the valuation of the second strongest bidder. We refer to bidding behavior at or above v_2 as a "*safe*" strategy

Table 3: Bidding behavior of the strongest bidder in the control treatment.

		Percentage of bid x					Profits	
		$x = 0$	$0 < x < v_2$	$x = v_2$	$x > v_2$	N	$0 < x < v_2$	$x \geq v_2$
No exclusion	*High30*	0%	37.5%	8.6%	53.9%	707	14.9	11.0
	High51	0.4%	25.3%	11.2%	63.1%	498	35.9	30.4
Exclusion	*High30*	1.2%	72.8%	11.6%	14.4%	692	0.4	-0.4
	High51	1.4%	77.7%	9.4%	11.5%	502	2.5	0.1

Notes: The strongest bidder in the no-exclusion condition is the high type, and the medium type is the strongest bidder in the exclusion condition in *High30* and *High51*.

as a bidder applying this strategy should win the auction for sure.[16]

There are two major strategic differences between the exclusion and no-exclusion condition that may account for the pronounced difference in the use of this *safe* strategy in these two conditions. First, the difference in valuations $v_1 - v_2$ is larger in the no-exclusion than in the exclusion condition. This implies that the "certain" profit that the high types potentially forgo by not playing the *safe* strategy is larger than for the medium type. In other words, high types have more to lose should they play a mixed strategy instead of the *safe* strategy, and thus may have more to regret if they place a bid lower than v_2 and lose the auction. Second, high types face two competitors that are almost equally strong, whereas medium types in the exclusion condition are confronted with only one other opponent with a slightly lower valuation. Although the number of competitors should not matter in our setup, it may make a difference behaviorally and we explore how the bidding behavior of weak types affects the use of the *safe* strategy in more detail below.

[16]Playing "*safe*" is consistent with a level-k reasoning process (see e.g., Stahl and Wilson 1995, Nagel 1995). Assuming that a level-0 player i randomizes bids in the interval $[0, v_i]$, a level-1 *high type* would best respond to this belief by playing *safe*, whereas a level-1 *medium type* would best respond to this belief by playing the mixed Nash equilibrium strategy. In turn a level-2 *high type* believes that he is facing a level-1 *medium type* and thus would best respond by playing the mixed equilibrium strategy, whereas a level-2 *medium type* best responds to a level-1 *high type* by placing a zero bid. Note that level-k *low types* ($k > 0$) never place a positive bid.

3.3 Anatomy of *Safe* Bidding

In this section we explore in more detail why we observe such massive bidding at or above v_2. Table 3 provides further details on the distribution of the bids of the strongest bidders. First, if the strongest bidders do not adopt the *safe* strategy, they obviously resort to bids in the interval $(0, v_2)$ in both conditions. In fact, they spread these bids over the whole interval and the distribution closely resembles a uniform distribution as predicted (see Figure 2 in Appendix A). Accordingly, it is not surprising that average bids are close to the theoretical prediction as well. In *High30* the average of bids in this interval is in both conditions 8.7, whereas in *High51* the average bid is 6.7 in the no-exclusion and 7.6 in the exclusion condition. Second, a significant share of *safe* bids is even strictly greater than v_2. However, we have to note that the overwhelming majority of these bids (74 percent) are in a comparatively small interval $(v_2, v_2 + 1]$.

We confine our analysis in the following to the no-exclusion condition because the *safe* strategy is vastly more popular when the strongest bidder (high type) is present and because we have seen that the failure of the *exclusion principle* is related to the behavior of high types. Playing the *safe* strategy is certainly not in line with theory, which predicts no mass point at v_2 and certainly not above, yet a subject's profit from playing *safe* should be approximately $v_1 - v_2$, given that the bid is infinitesimally larger or equal to v_2 and that they win the auction. Notice that this corresponds to the expected profit of playing the mixed equilibrium strategy.

We find, however, that the profits from playing the *safe* strategy are lower than this theoretical benchmark. For example, Table 3 shows that in the no-exclusion condition the average profits are clearly below $v_1 - v_2$, with 11 in *High30* and 30.4 in *High51*. In contrast, bidding in $(0, v_2)$ results in average profits that are close to the theoretical benchmark (with 14.9 in *High30* and 35.9 in *High51*), and are thus higher than when bidding *safe*. The differences in profits are significant in both cases (Wilcoxon signed-rank test $z = 2.04$,

$p < 0.045$, $n = 14$ in *High30* and $z = 2.5$, $p < 0.015$, $n = 10$ in *High51*).[17] Consequently, by choosing the *safe* strategy the strongest bidders forgo a substantial part of their rent in order to increase their chance of winning.[18] More precisely, they earn only about three quarters of the average profits accruing from a bid in the interval $(0, v_2)$ in *High30* when placing a *safe* bid. In *High51* the foregone profits are smaller than in *High30* as they earn about 85 percent of the expected profit. The lower average profits from playing the *safe* strategy result from bidding above the valuation of the second strongest bidder, which naturally reduces profits, and from the fact that the *safe* strategy does not guarantee winning the auction. In about 6 percent of cases high types lose the auction even though they use the *safe* strategy, which leads to large losses.

A possible explanation for the prevalence of the *safe* strategy and the associated lower profits is the excessive bidding of the weaker types. Going back to Figure 1, we observe that the two weaker types often drop out of the bidding process in the no-exclusion condition, i.e., bid zero (see left-panel of Figure 1). However, while medium types tend to drop out too much, low types drop out too little.[19] In fact, low types should never participate in the no-exclusion condition, i.e., never place a positive bid.

In *High30* we observe that medium types abstain from bidding (placing a zero bid) in about 72 percent of cases, whereas low types abstain in only 76 percent of cases. In theory we should observe in about half of the periods (47 percent) only bids from high types and in the remaining periods positive bids from both high and medium types. The observed bidding behavior of the two weaker bidders in *High30* implies that in about 56

[17]In the exclusion condition profits are on average also higher for bids in the interval $(0, v_2)$ than for the *safe* strategy. However, the magnitude is much smaller as the expected payoff for the strongest bidder is $v_1 - v_2 = 1$. The difference in profits is only significant in *High51* (Wilcoxon signed-rank test $z = 2.3$, $p < 0.025$, $n = 10$).

[18]While the predicted probability of winning is 73 percent in *High30*, high types win the auction in 88 percent of cases. Similarly, the probability of winning is 8 percentage points higher than predicted in *High51* (92 percent vs. 84 percent).

[19]This represents mixed evidence for a discouragement effect, i.e., the tendency of weaker types to drop out of the bidding process in the presence of a strong bidder. But it is line with findings for real-effort tournaments. For example, Gill and Prowse (2012) find evidence of a discouragement of the weaker participants, whereas Berger and Pope (2011), Hammond and Zheng (2013), or Chen, Ham, and Lim (2011) find no effect.

percent of cases high types are the sole bidders. In 20 percent of cases their lone competitor is the medium type, in 16 percent of cases it is the low type, and in 8 percent both weak types are active. Accordingly, we find that medium types abstain too often from bidding in *High30*, but low types' behavior compensate for the less frequent bids of medium types. Conditional on participation, medium types overwhelmingly place a bid below their valuation, i.e., $v_2 = 16$, (in 17 percent of cases bids are above their own valuation) and the average bid is 8.4 in *High30*. Similarly, if low types participate in the auction they bid on average 7.1 in *High30*, with a large majority of their bids below their valuation of $v_3 = 15$ (88 percent). This behavior gives rise to winning the auction in 26 (20) percent of cases for medium (low) types, conditional on participation. Overall their profits are negative, albeit as predicted close to zero for both types (-1).

In *High51* we see that medium types place a zero bid in about 68 percent of cases and low types participate in 22 percent of cases. This is close to the theoretical prediction for medium types, who should abstain in 69 percent of cases, but not for low types. Thus, in contrast to *High30* high types face an active competitor more often than predicted because of low types' bidding behavior in *High51*. They are the lone bidder in only 54 percent of cases.[20] If medium types participate in the auction, they predominantly place a bid below their valuation (in 17 percent of cases bids are above their own valuation) and the average bid is 6.8. Similarly, the vast majority of bids from low types is below their valuation $v_3 = 15$ (93 percent) and their average bid is 5. Again, this results in a fairly high share of wins (17 percent for medium types and 14 percent for low types), but this is not enough to yield positive profits on average. Similar to *High30*, profits are close to zero, but slightly negative (-1) over all periods.

The distribution of winning bids, which is observed by subjects, reflects these differences in participation of weaker types. There are significantly less winning bids in a low range $[0, 16)$ in *High51* than in *High30* (24 percent vs. 36 percent) and, consequently,

[20]Medium types are the lone competitor in 24 percent of cases, low types in 14 percent of cases, and in 8 percent of cases they are both active at the same time.

significantly more winning bids above $v_2 = 16$ in *High51* (Fisher's Exact test, $p = 0.44$). Therefore, this suggests that the participation of low types has triggered more *safe* bidding of high types in *High51* compared to *High30*.

Together, it seems that the bidding behavior of the weaker types results in more *safe* bidding and thus higher winning bids in *High51* compared to *High30*. This also suggests that weaker types bid too much on average and that they could improve their profits by abstaining more often from bidding (in particular low types) or at least by refining their bids on the interval between zero and their own valuation in both treatments. However, the profits of both the medium and low type are on average only slightly below zero, and thus the losses might be too small to induce a significant change in bidding behavior. On the other hand, high types could substantially increase their profits in the no-exclusion condition by deviating to the equilibrium strategy in both *High30* and *High51*, given the behavior of the two weaker bidders.

The preceding analysis also suggests that there is a substantial share of *safe* bidding when the bids of low types are substituting the missing bids of medium types, as in *High30*, i.e., when high types do not face more competition. To get a more complete picture of why strong bidders in the no-exclusion condition choose the *safe* strategy, we present results from a set of probit regressions in Table 4. Linear probability regressions yield qualitatively similar results. In all specifications the dependent variable is a binary variable for playing the *safe* strategy, i.e., this variable equals one if the strongest bidder has chosen a bid that is at least as high as the valuation of the medium type.

The regression in column (1) provides statistical support for the observation that *safe* bidding is more frequent in the no-exclusion condition of *High51* (see Table 3). Moreover, we see that, overall, the prevalence of *safe* bidding decreases significantly over time. In columns (2–4) we explore how feedback and previous performance affect the likelihood of playing *safe*. Column (2) captures how feedback, i.e., the observed winning bid in the last period, influences *safe* bidding. The positive and significant coefficient indicates that observing a higher winning bid in the previous period increases the likelihood of a *safe* bid. This suggests that observing more *safe* bids in the past is subsequently related to more *safe* bidding in a group. Column (3) and (4) investigate how an individual's past profits affect *safe* bidding. While the profit in the previous period has no effect on *safe* bidding, accumulated profits have a negative effect on *safe* bidding. More precisely, a 10-point increase in accumulated profits is associated with a 2 percentage-point lower likelihood of *safe* bidding. This negative relationship suggests that less successful bidders in particular resort to the *safe* strategy. Note that accumulated profits also capture a time trend which renders the coefficient on "Period" insignificant. Column (5) shows that the results do not change if we include all three variables at once. In column (6) we additionally control for risk aversion by including a dummy variable, which equals one if a subject prefers safe options smaller or equal to the expected value of the lottery in the lottery task. The coefficient for risk aversion indicates that risk-averse subjects are more likely to play the *safe* strategy, but it is not significant and it does not change any other coefficients. We should, however, keep in mind that our measure for risk aversion is naturally measured with error resulting in a downward bias.[21]

In the final three columns we investigate how own behavior and experience in early periods affect behavior in later periods. It is, for example, conceivable that the strongest bidders draw inferences about the behavior of their competitors from their own behavior as a weak type, or alternatively, that they have a tendency to bid high irrespective of circumstances. To this end, we split the data set and use bidding behavior in periods 1–20 to explain the likelihood of *safe* bidding in periods 21–50.[22] First, we investigate the consequences of initial bidding behavior as a weak type – either as a medium or low type – on subsequent bidding as a strong type in periods 21–50. Column (7) shows

[21]See Gillen, Snowberg, and Yariv 2015 for an extensive discussion of measurement error in experimentally elicited measures. In Appendix A, we present additional regressions with multiple elicited controls for risk preferences as suggested by Gillen, Snowberg, and Yariv 2015. Again, the coefficient does not change and the regression results confirm that the *safe* strategy is more common in *High51*.

[22]The results are robust to using behavior in periods 1–10 or 1–15 to explain *safe* bidding in the subsequent periods.

Table 4: Regression: Choice of the *safe* strategy in the no-exclusion condition

Dependent variable:					*safe* strategy of the strongest bidder.				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
High51 (D)	0.119*	0.116*	0.118*	0.282***	0.269***	0.268***	0.108**	0.123**	0.069
	(0.062)	(0.061)	(0.065)	(0.075)	(0.075)	(0.075)	(0.054)	(0.053)	(0.048)
Period	−0.003*	−0.002	−0.003**	0.003	0.003	0.003			
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)			
Winning bid in last period		0.012***			0.009***	0.009***			
		(0.003)			(0.003)	(0.003)			
Profit in last period			0.000		0.002	0.002			
			(0.002)		(0.002)	(0.002)			
Accumulated profits			−0.002***	−0.002***	−0.002***				
				(0.000)	(0.000)	(0.000)			
Risk averse (D)						0.046			
						(0.056)			
Avg. bid as weak type in period 1–20							0.024**		
							(0.011)		
Avg. bid in period 1–20								0.038***	
								(0.014)	
Avg. share of *safe* bidding in period 1–20									0.346***
									(0.116)
N	1205	1181	1183	1177	1153	1153	738	738	738

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Probit regressions (average marginal effects) with standard errors clustered on matching group level (in parentheses). Regressions in columns (1–6) use data from period 1 to 50, whereas regressions in columns (7) and (8) only includes data from period 21–50. *High51* is a dummy variable indicating that $v_1 = 51$. The variable "Period" captures a linear time trend. "Winning bid in last period" is the last winning bid that an individual observed and "Profit in last period" is an individual's profit in the last period. "Accumulated profits" are the accumulated profits over periods starting at 10 Euro in period 1. "Risk averse" is a dummy variable which equals one if a subject prefers safe options smaller or equal to the expected value of the lottery. The variable "Avg. bid in period 1–20" captures initial bidding behavior of an individual and "Avg. share of *safe* bidding in period 1–20" captures a subject's experience with *safe* bidding in the first 20 periods. (D) denotes dummy variable.

that higher average bids as a weak type in the first 20 periods lead to significantly more *safe* play in periods 21–50 as a strong type. If we assume that subjects project their experience as a weak type onto others, this result provides support for the conjecture that the behavior of the weak types triggers, in part, the *safe* strategy of the strongest bidders. It is also in line with the earlier finding that less successful subjects tend to use the *safe* strategy more often.

This relationship between behavior in early and later periods is also true if we relax our restriction on weak types and include individual average bids in periods 1–20 considering all bidder roles (column 8). Finally, the previous analysis suggested that experiencing more *safe* play in initial periods has an impact on own *safe* bidding in later periods. The last column provides a more direct test for this finding and demonstrates that experiencing more *safe* play in the beginning (i.e., the first 20 periods) induces subjects to choose more *safe* bids as well, thus confirming our earlier finding.

In summary, we find that *safe* bidding is more prevalent among very strong bidders ($v_1 = 51$) suggesting that this strategy is more attractive the higher the secure rent from winning is. It seems that a significant part of this behavior originates in own initial behavior as one of the weak types as well as in experiencing more *safe* bidding in the beginning.

4 Discussion and Conclusion

Superstars can have a major impact on the attractiveness of contests, but at the same time their presence can detrimentally affect their competitors' willingness to exert effort. In this paper, we experimentally investigate the effect of excluding superstars from the contest and thereby creating a more homogeneous participant pool. We find that in our setting excluding the strongest bidder is, in general, not beneficial for the contest designer.

The main reason for this result is, in particular, the bidding behavior of the strongest bidder when they participate in the all-pay auction. We find that these "superstars" often apply a strategy which guarantees they will win the auction. That is, they bid at least the valuation of their most powerful competitor, which implies that they prefer to give up a substantial part of their rent in order to avoid losing the auction. Moreover, the tendency of the strongest bidder to choose the *safe* strategy increases in their valuation.

The observed *safe* bidding is consistent with a non-equilibrium model of limited sophistication (Stahl and Wilson 1995, Nagel 1995). In the standard formulation of this model players anchor their beliefs on a non-sophisticated level-0 player but differ with respect to their levels of reasoning.[23] In our setup the mass point at v_2 may be explained with level-1 *high types* best responding to a belief that all others are level-0 players who randomize their bid in the interval $[0, v_2]$. This suggests that *safe* bidding is a low-cognition strategy as it guarantees a profit that corresponds in most cases to the expected profit from playing a mixed strategy. Thus engaging in more levels of reasoning or investing more cognitive resources may not be worthwhile. Note, however, the level-k approach does not explain the higher share of *safe* bidding in *High51* compared to *High30* as the potential absolute gain from bidding less than v_2 is the same in both cases.

Since playing a mixed strategy can involve a loss, loss aversion has been used to explain the observed overbidding behavior (e.g., Müller and Schotter 2010, Ernst and Thöni 2013). A loss-averse bidder would incur additional disutility from losing the auction and this disutility depends only on the bid but not on the valuations of bidders. Thus, while loss aversion may explain some of the *safe* bidding, it cannot explain why we observe more *safe* bidding in *High51* than in *High30*.

An alternative explanation for the prevalence of *safe* bidding among high types is regret. Placing a bid below v_2 and losing the auction may create feelings of regret because such a situation could easily have been avoided by placing a bid at or slightly above v_2

[23] An alternative way to model "bounded rationality" posits that players have correct beliefs but best-respond with noise. For example, Anderson, Goeree, and Holt (1998) present a logit equilibrium model where players choose bids with higher expected payoffs with higher probability, which is consistent with observed overdissipation patterns found in all-pay auctions.

generating a certain profit of $v_1 - v_2$. This kind of regret is specific to our complete information environment since it presumes that the valuations of the competitors are public knowledge.[24] As such, regret depends on $v_1 - v_2$, which implies that a regret-averse high type is more likely to play *safe* in *High51* than in *High30*. Thus, this notion of regret is consistent with the higher share of *safe* bidding in *High51* than in *High30*.

Of course, medium types (and low types) should anticipate high types' inclination to bid *safe* and resort to zero bidding, particularly since they get feedback about the behavior of high types. However, on the surface it seems that weak types tend to participate too often and subsequently bid too much. In *High30* we indeed find that medium types reduce their bidding activity, but this lower activity is completely offset by the participation of low types. Similarly, we observe that contrary to the theoretical prediction low types participate in the auction in *High51*. But in this case this behavior complements the bidding activity of medium types such that high types face an active competitor more often than predicted. Although weak types seem to not best respond to high types' behavior, we observe that their profits are close to the prediction of zero profits in both treatments. Since the losses seem negligible, weak types do not reduce their bidding.

Given the behavior of the two weak types, high types do not best respond either. They could clearly improve their profits by adopting a mixed strategy. In fact, high types who place a bid in $[0, v_2]$ earn significantly more than high types who play *safe*. The prevalence of *safe* bidding, however, suggests that for most high types even a small chance of a positive bid from a weak type is enough to induce *safe* bidding. Using data from NASCAR races, Bothner, Kang, and Stuart (2007) demonstrate that lower ranked drivers can induce higher ranked drivers to take more risk to avoid losing the race or to fall in the ranking.[25] A similar motivation may lead to more *safe* bidding and to thus securing

[24]Regret has been previously analyzed in symmetric or incomplete information auctions (see e.g., Filiz-Ozbay and Ozbay 2007, Baye, Kovenock, and de Vries 2012 or Hyndman, Ozbay, and Sujarittanonta 2012). Unlike in our setting, in symmetric auctions there is no possibility for the bidders to generate a secure positive payoff and therefore the amount of regret a bidder experiences in the case of a loss depends on the winning bid of their opponent.

[25]The NASCAR race series is a multiple-round tournament, where rankings are publicly available and

the prize in our setup.

The composition of teams is critical for the performance of firms and organizations in general (Mathieu, Tannenbaum, Donsbach, and Alliger 2014) and for contest designers in particular. From a theoretical perspective installing a contest with homogeneous bidders leads to full rent dissipation (see e.g., Hillman and Riley 1989, Baye, Kovenock, and de Vries 1993), whereas revenues decrease with the heterogeneity of bidders. In practice, firms typically face a heterogeneous workforce, and still scholarly research indicates that relative performance schemes are widely applied in the firm context (e.g., Connelly, Tihanyi, Crook, and Gangloff 2014). Our paper focuses on heterogeneous contests with the presence of a superstar and our results may have several managerial implications in terms of designing contests.

That superstars prefer a *safe* bid over a mixed strategy even though this is less profitable suggests that our setup renders non-material factors important. In addition to the utility from the monetary outcome, the strongest bidder may, for example, incur disutility from feelings of regret in case of losing the auction or, more generally, suffer a psychological loss as they expect to win owing to their superiority. If this is the case, managers or contest designers may use such non-pecuniary motives of the strongest bidder in their advantage. For example, by making sure that the bidders' valuations are public knowledge as this may provide a fertile soil for regret that may induce superstars to bid *safe*. Although this is not always feasible, it is often possible to provide sufficiently correlated signals about the valuations of participating bidders or to provide updates of the performance ranking of contestants (Casas-Arce and Martínez-Jerez 2009).

There are also good reasons for heterogeneous contests. Indeed, the presence of a superstar can be, but need not be, detrimental to overall revenues. While we find evidence of the latter, there are situations where the overall performance in a contest with a

are updated after each race. This rank updating naturally leads to a heterogeneous contest as previous performance places some drivers in an advantageous and others in a disadvantageous position (see also Casas-Arce and Martínez-Jerez 2009).

superstar might be more superior than in a contest without a superstar because the valuation gap between the weaker bidders is too wide to ramp up competition in the smaller contest.[26] In other cases, managers may want to benefit from the presence of a superstar as they typically add status, popularity or glamor to the contest. The response of the organizers of the Formula One racing series to the dominance of Michael Schumacher discussed in the introduction illustrates such a case. Similarly, managers may implement heterogeneous contests because they may reveal more information about bidders' ability (Gürtler and Gürtler 2015). This can be, for example, a relevant issue in law firms, consultancies or investment banks where typically many associates compete for a few spots as a partner and the company only wants to promote the most able associates. If managers have a preference for a heterogeneous contest, they may boost overall performance, for example, by handicapping bidders with the highest valuation (e.g., Lazear and Rosen 1981, Che and Gale 1998, Llorente-Saguer, Sheremeta, and Szech 2016) or by changing the prize allocation rule (e.g., Cason, Masters, and Sheremeta 2010).[27]

Our findings suggest that admitting a superstar to the contest does not hurt overall performance in the contest, even though there is no handicapping or other measure to increase participation of weaker bidders in place. In fact, our results on individual behavior are mostly in line with the qualitative theoretical predictions: weaker bidders are discouraged in the presence of a superstar, but they substantially increase their participation and average effort (bids) in the absence of a superstar (see also Brown 2011 for field evidence from the PGA Tour). However, the increased effort of the weaker bidders in the absence of the strongest bidder cannot compensate for the lost effort of the superstar and therefore we find little support for the *exclusion principle*.

[26]In Appendix B we present a treatment (*EXUP*) where the valuations of the weaker bidders are sufficiently heterogeneous and find evidence that, as illustrated in the model in section 2.1, aggregate effort of competitors in the absence of the superstar does not compensate for the lost superstar effort, i.e., overall performance is superior in a contest with a superstar.

[27]Cason, Masters, and Sheremeta (2010) compare endogeneous entry of bidders in an all-pay auction and proportional prize contest and find that the all-pay auction primarily attract high-ability contestant, whereas the proportional prize contest also attracts low-ability contestants. Thus, there is a similar discouragement of low-ability contestants in their all-pay auction as in our study.

References

Anderson, S., J. K. Goeree, and C. A. Holt (1998). Rent seeking with bounded rationality: An analysis of the all-pay auction. *Journal of Political Economy 106*(4), 828–853.

Barut, Y., D. Kovenock, and C. Noussair (2002). A comparison of multiple-unit all-pay and winner-pay auctions under incomplete information. *International Economic Review 43*, 675–707.

Baye, M. R., D. Kovenock, and C. G. de Vries (1993). Rigging the lobbying process: An application of the all-pay auction. *American Economic Review 83*, 289–294.

Baye, M. R., D. Kovenock, and C. G. de Vries (1996). The all-pay auction with complete information. *Economic Theory 8*, 1996.

Baye, M. R., D. Kovenock, and C. G. de Vries (2012). Contests with rank-order spillovers. *Economic Theory 51*(2), 315–350.

BBC (2002, February). F1 viewing figures drop. http://news.bbc.co.uk/sport2/hi/motorsport/ formula_one/1842217.stm.

Berger, J. and D. Pope (2011). Can losing lead to winning? *Management Science 57*(5), 817–827.

Bothner, M. S., J.-h. Kang, and T. E. Stuart (2007). Competitive crowding and risk taking in a tournament: Evidence from nascar racing. *Administrative Science Quarterly 52*(2), 208–247.

Brown, J. (2011). Quitters never win: The (adverse) incentive effects of competing with superstars. *Journal of Political Economy 119*(5), 982–1013.

Casas-Arce, P. and F. A. Martínez-Jerez (2009). Relative performance compensation, contests, and dynamic incentives. *Management Science 55*(8), 1306–1320.

Cason, T. N., W. A. Masters, and R. M. Sheremeta (2010). Entry into winner-take-all and proportional-prize contests: An experimental study. *Journal of Public Economics 94*(9), 604–611.

Che, Y.-K. and I. L. Gale (1998). Caps on political lobbying. *American Economic Review 88*(3), 643–651.

Chen, H., S. H. Ham, and N. Lim (2011). Designing multiperson tournaments with asymmetric contestants: An experimental study. *Management Science 57*(5), 864–883.

Connelly, B. L., L. Tihanyi, T. R. Crook, and K. A. Gangloff (2014). Tournament theory: Thirty years of contests and competitions. *Journal of Management 40*(1), 16–47.

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*(1-2), 89–115.

Dechenaux, E., D. Kovenock, and R. M. Sheremeta (2012). A survey of experimental research on contests, all-pay auctions and tournaments. Technical Report Discussion Paper, Chapman University.

Dohmen, T. and A. Falk (2011). Performance pay and multi-dimensional sorting: Productivity, preferences and gender. *American Economic Review 101*(2), 556–590.

Ernst, C. and C. Thöni (2013). Bimodal bidding in experimental all-pay auctions. *Games 4*, 608–623.

Filiz-Ozbay, E. and E. Y. Ozbay (2007). Auctions with anticipated regret: Theory and experiment. *American Economic Review 97*, 1407–1418.

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

Frank, R. H. and P. J. Cook (1995). *The Winner-Take-All society: Why the Few at the Top Get So Much More Than the Rest of Us*. Penguin Books.

Gill, D. and V. Prowse (2012). A structural analysis of disappointment aversion in a real effort competition. *The American economic review*, 469–503.

Gillen, B., E. Snowberg, and L. Yariv (2015). Experimenting with measurement error: techniques with applications to the caltech cohort study.

Gneezy, U. and R. Smorodinsky (2006). All-pay auctions - an experimental study. *Journal of Economic Behavior and Organization 61*, 255–275.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association* (1), 114–125.

Gürtler, M. and O. Gürtler (2015). The optimality of heterogeneous tournaments. *Journal of Labor Economics 33*(4), 1007–1042.

Hammond, R. G. and X. Zheng (2013). Heterogeneity in tournaments with incomplete information: An experimental analysis. *International Journal of Industrial Organization 31*(3), 248–260.

Hillman, A. L. and J. G. Riley (1989). Politically contestable rents and transfers. *Economics and Politics 1*(1), 17–39.

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

Hyndman, K., E. Ozbay, and P. Sujarittanonta (2012). Rent seeking with regretful agents: Theory and experiment. *Journal of Economic Behavior and Organization 84*(3), 866–878.

Konrad, K. (2009). *Strategy and Dynamics in Contests*. New York, NY: Oxford University Press.

Lazear, E. P. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy 89*(5), 841–864.

Llorente-Saguer, A., R. M. Sheremeta, and N. Szech (2016). Designing contests between heterogeneous contestants: An experimental study of tie-breaks and bid-caps in all-pay auctions. *mimeo*.

Lugovskyy, V., D. Puzello, and S. Tucker (2010). An experimental investigation of overdissipation in the all-pay auction. *European Economic Review 54*, 974–997.

Mathieu, J. E., S. I. Tannenbaum, J. S. Donsbach, and G. M. Alliger (2014). A review and integration of team composition models: Moving toward a dynamic and temporal framework. *Journal of Management 40*(1), 130–160.

Müller, W. and A. Schotter (2010). Workaholics and drop-outs in organizations. *Journal of the European Economic Association 8*(4), 717–743.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review 85*(5), 1313–1326.

Noussair, C. and J. Silver (2006). Behavior in all-pay auctions with incomplete information. *Games and Economics Behavior 55*, 189–206.

Sheremeta, R. M. (2013). Overbidding and heterogeneous behavior in contest experiments. *Journal of Economic Surveys 27*(3), 491–514.

Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior 10*(1), 218–254.

A Additional Figures and Tables – For Online Publication

Figure 2 presents the cumulative distribution of the bids of the strongest bidder in the exclusion and no-exclusion condition when not using the *safe* strategy along with the 45-degree line resulting from the equilibrium predictions. In both conditions, we observe that the distribution of bids is close to the theoretical prediction.

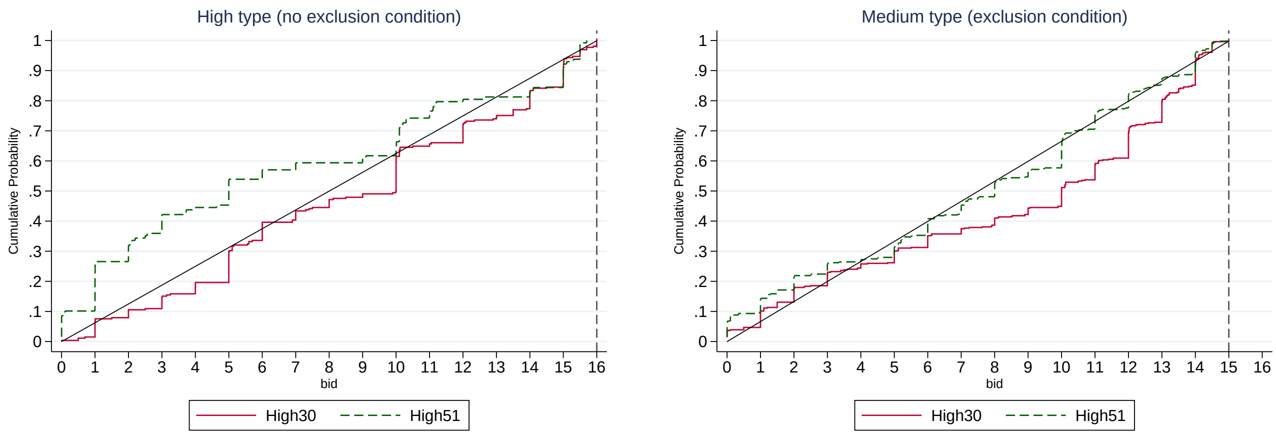

Figure 2: Cumulative distribution of bids when not using the *safe* strategy.

For a subset of participants we have a second risk measure ($N = 72$). More precisely, in the last three sessions we asked subjects to assess their general willingness to take risk on an 11-point scale. Following the suggestion of Gillen, Snowberg, and Yariv (2015) we include both measures in our regressions in Table 5 below. First, we include dummy variables for both measures. In the choice-list procedure we classify subjects as risk averse if they switch from the lottery to the safe option as long as the safe option is smaller than or equal to the expected value of the lottery. Similarly, we code subjects as risk averse if they indicate a willingness to take risk below 6 on the 11-point Likert scale of the risk question. As a second set of measures we use standardized values of the switching point in the choice list and the answer in the risk question. The results indicate that the coefficient for the treatment is remarkably stable across all specifications. While a measurement error is still an issue, the results lend further credibility to our results reported in Table 4.

Table 5: Regression: Risk aversion and Choice of the *safe* strategy in the no-exclusion condition

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
High51	0.176∗∗	0.174∗∗	0.174∗∗	0.172∗∗	0.167∗∗	0.171∗∗	0.165∗∗
	(0.067)	(0.063)	(0.068)	(0.064)	(0.071)	(0.072)	(0.074)
Period	−0.004	−0.004	−0.004	−0.004	−0.004	−0.004	−0.004
	(0.003)	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
Risk averse MPL(D)		0.085		0.092			
		(0.090)		(0.093)			
Risk averse Q(D)			0.062	0.072			
			(0.088)	(0.086)			
Std Risk MPL					−0.032		−0.026
					(0.028)		(0.031)
Std Risk Q						−0.042	−0.037
						(0.045)	(0.047)
N	624	624	624	624	624	624	624

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Linear Probability regressions with standard errors clustered on matching group level (in parentheses). Regressions use data from the last three sessions in which we elicited two risk measures. *High51* is a dummy variable indicating that the $v_1 = 51$. The variable "Period" captures a linear time trend. "Risk averse MPL" is a dummy variable, which equals one if a subject prefers safe options smaller or equal to the expected value of the lottery. Similarly, "Risk averse Q" is a dummy variable, which equals one if a subject's response is below six on the 11-point scale of the hypothetical question about general willingness to take risks. "Std Risk MPL" and "Std Risk Q" are standardized variables with zero mean and standard deviation one of subjects' responses to each question. (D) denotes dummy variable.

B Additional Treatments – For Online Publication

B.1 Setup

We conducted two further treatments to explore the *exclusion principle* in a richer environment that is not idiosyncratic to a specific choice of valuations. In these additional treatments we assigned new valuations to bidders in each period, which allowed us to explore the *exclusion principle* for a broad range of parameters. The results of these additional treatments corroborate our findings from the main treatments.

As in our main treatment, bidding groups consist of three bidders with valuations $v_{H(igh)} > v_{M(edium)} > v_{L(ow)}$, which were drawn before the experiment. Specifically, two valuations v_M and v_L were drawn from the discrete uniform distribution over the interval $[11, 20]$ and the third valuation v_H was drawn from a discrete distribution over the interval $[15, 55]$. Based on these valuations we constructed two treatments. In one treatment the configuration of valuations is such that the exclusion of the high type v_H

is always profitable for the contest designer (in the following treatment EXP – *EX*clusion *P*rofitable). In the second treatment, the composition of groups is such that excluding the high type should result in lower revenues than letting all bidders participate in the auction (in the following treatment *EXUP* – *EX*clusion *U*n*P*rofitable). In *EXP*, the average valuations are $v_H = 35.3$, $v_M = 16$ and $v_L = 14.7$. In *EXUP*, these averages are $v_H = 30.9$, $v_M = 17.9$ and $v_L = 13$. Note that this setup presents a more complex strategic situation where finding the optimal bidding strategy is a rather difficult task as bidders face different sets of valuations in each period. To facilitate the understanding of the strategic aspects of the auction, subjects experienced each bidder role (v_H,v_M and v_L) over time.

Again, we randomized whether the bidder with valuation v_H is excluded from the auction or not and our aim was primarily to compare the revenue of an auction with two "*homogeneous*" bidders with valuation v_M and v_L (exclusion condition) to the revenue of an auction with all three bidders with valuations $v_H > v_M > v_L$ (no-exclusion condition) within a treatment. We conducted four computerized sessions at the experimental laboratory at the TU Berlin using the software tool kit *z-Tree* (Fischbacher 2007) with students recruited from a large database where they can voluntarily register for participating in experiments (ORSEE, Greiner 2015). The course of action was identical to our main treatments. In particular, subjects first completed a lottery choice task and then the all-pay auction was repeated 51 times (including one trial period) in both treatments. We randomly assigned subjects to a six-person group (matching group) and randomly matched subjects into two bidding groups of three in each period within a matching group. At the beginning of each period, the subjects in each bidding group were randomly assigned a valuation. The valuations were made public knowledge before bidding started and subjects learned whether the high type was participating in a particular period, which was randomly determined by the computer with probability $p = 0.5$. Subjects could place bids with up to three decimal places and they were informed of their earnings and the winning bid after each period. At the end of the second part of the experiment we publicly and randomly drew eight out of the 50 periods to determine subjects' earnings. The sum of points in these eight periods plus the earnings from the lottery choice task were exchanged at a rate of 10 points = €1. Additionally, the subjects received an initial endowment of €10 to cover potential losses. In total, 72 students (40 males and 32 females) from various disciplines participated in the experiment. Sessions lasted about 90 minutes and subjects' average earnings were approximately €15.

B.2 Results

In the following we focus our analysis on (i) aggregate behavior and (ii) individual bidding behavior. In particular, we provide evidence that underpin the key results from our main treatments. First, we confirm that excluding the bidder with the highest valuation is on average not beneficial in the treatment in which exclusion should be profitable for the contest designer (treatment EXP). Second, we show that this is due to the behavior of the strongest bidders when they participate in the auction. Specifically, they predominantly use the *safe* strategy, which in most cases guarantees them to win the auction. Moreover, this strategy is more prominent the larger the difference in the valuations of the strongest and second-strongest bidder is. Finally, the results represent a methodological contribution by showing that our results are not affected by how we assign valuations, i.e., whether subjects face the same set of valuations in each period or whether they get new valuations in each period.

Table 6 presents the summary statistics of observed behavior along with the theoretical predictions for both treatments broken down into the exclusion and no-exclusion condition.[28] The table shows that revenues are lower when the high type v_H is excluded from participation in both *EXUP* and *EXP*. While this is in line with the qualitative prediction for *EXUP*, it is in strong contrast to the prediction for *EXP* that exclusion increases revenues relative to no-exclusion. In *EXUP*, the sum of bids is, on average, 21.55 when all three bidders participate compared to 13.63 when only the two weaker bidders participate. The difference in the sum of bids in the two conditions (exclusion and no-exclusion) is statistically significant according to a Wilcoxon signed-rank test ($z = 2.201$, $p < 0.03$, $n = 6$). In contrast to the theoretical prediction, in *EXP* the sum of bids is larger when all bidders participate in the auction (18.48) than when the high type is excluded (14.02). We can reject the hypothesis of equal revenues in the two conditions (Wilcoxon signed-rank test $z = 2.201$, $p < 0.03$, $n = 6$), but not in favor of our alternative hypothesis.

It is apparent that, on average, the sum of bids is always higher than predicted (overbidding) except in the exclusion condition in *EXP*. When the high type v_H is excluded, we observe that the sum of bids in *EXUP* is about 1.2 times higher than predicted, whereas the sum of bids is about 1.5 times higher than predicted in the no-exclusion condition of both treatments (158% in *EXP* and 148% in *EXUP*).

Conducting the same thought experiment as in *High30* and *High51* in Section 3.1

[28] Due to a programming mistake we implemented the same valuation for v_M and v_L in 20 percent of cases in treatment *EXP*. Recall that the theory requires $v_H > v_M > v_L$. Excluding these cases yield qualitatively the same results and thus we include this data throughout our analysis. Note also that five out of 3,600 individual bids are significantly larger than 55. We exclude the data of the whole bidding group from these periods throughout our analysis.

Table 6: Sum of bids in *EXP* and *EXUP*

	EXUP		***EXP***	
	no exclusion (3 bidders)	exclusion (2 bidders)	no exclusion (3 bidders)	exclusion (2 bidders)
avg. sum of bids (observed)	21.55	13.63	18.48	14.02
	(11.58)	(8.27)	(10.88)	(8.72)
avg. sum of bid (predicted)	14.60	11.43	11.72	14.27
	(2.22)	(2.24)	(2.39)	(2.63)
minimum bid	0	0	0	0
maximum bid	50	40	55	40
N	*284*	*312*	*299*	*300*

Notes: Standard deviations in parentheses. We excluded the sum of bids if $x_i > 55$, as occurred in 5 out of 3600 individual bids.

illustrates again that the larger revenues in the no-exclusion condition in *EXP* is due to the behavior of the high type. Using the actual bids of the two weaker bidders and the theoretical bids of the strongest bidders to calculate the revenues shows that revenues would be lower than with exclusion in *EXP* (12.82 vs. 14.02). However, the difference in revenues is not statistically different (Wilcoxon signed-rank test $z = 0.943$, $p > 0.34$, $n = 6$).

Intuitively, one would expect that exclusion is more likely profitable the more heterogeneous the group is. However, our results suggest that exclusion is not profitable for a wide range of valuations in *EXP*, because of the prevalent overbidding in the no-exclusion condition. When all three bidders participate in the auction, the sum of bids is in about 80 percent of cases higher than the predicted sum of bids. While in these cases the sum of bids is on average 22.23, it is, on average, only 6.51 when no overbidding occurred. Overbidding occurs in particular when the valuation of the strongest bidder is high. In groups with overbidding the valuations are on average $v_H = 36.3$ and $v_M = 15.7$ compared to $v_H = 33.3$ and $v_M = 16.1$ in groups where the sum of bids that are equal or lower than predicted. Accordingly, there is a significant and positive correlation between overbidding and the distance in valuations between the strongest and second-strongest bidder ($\rho = 0.17$, $p < 0.01$). That is, the stronger the high type, the more likely overbidding takes place. This suggests that the *exclusion principle* may be profitable if the strongest bidder is not too strong.

In summary, these results corroborate the findings from the main treatment that the *exclusion principle* is not profitable and, in addition, indicates that this is likely more often the case when the strongest bidder is far more superior to the other bidders. Next, we

Table 7: Summary statistics of individual bids of bidder types

	EXP			***EXUP***		
bidder type	High	Medium	Low	High	Medium	Low
avg. bid	13.56	2.24	2.67	15.84	3.50	2.18
no exclusion	(6.43)	(5.30)	(6.81)	(5.87)	(7.66)	(6.29)
avg. predicted bid	7.92	3.81	0.00	8.98	5.58	0.00
no exclusion	(1.19)	(1.42)	(0.00)	(0.84)	(1.70)	(0.00)
avg. bid	-	8.40	5.63	-	9.69	3.94
exclusion		(6.06)	(6.47)		(5.55)	(5.77)
avg. predicted bid	-	7.43	6.84	-	6.54	4.88
exclusion		(1.26)	(1.39)		(0.96)	(1.30)
minimum bid	0	0	0	0	0	0
maximum bid	47	40	40	32	50	50

Notes: Standard deviations in parentheses. We excluded bids $x_i > 55$. This was the case in 5 out of 3,600 individual bids.

look at individual bidding behavior. Table 7 provides a first overview of the average bids of each bidder type in the no-exclusion condition (top panel) and the exclusion condition (bottom panel) for each treatment.

We find that in both treatments high types bid, on average, almost twice as much as predicted by theory if they participate in the auction (Table 7, top panel). The difference between the actual bids and predicted bids is statistically significant in both treatments (Wilcoxon signed-rank test, $z = 2.201$, $p < 0.03$, $n = 6$ for both treatments). As a consequence, the strongest bidders forgo a substantial part of their rent in order to increase their chance of winning. In particular, they win about 82% of the auctions in both treatments, which is about 10 percentage points more often than predicted. But they only earn about 81 (74) percent of the expected profits in *EXP* (*EXUP*).[29]

Weaker types often drop out of the bidding process in the no-exclusion condition (discouragement effect). For example, medium types abstain from bidding (placing a zero bid) in 64 percent of cases in *EXP* and *EXUP*. On the other hand, low types who should never place a positive bid in the no-exclusion condition, place zero bids in only 71 percent of cases. Accordingly, we observe on average strictly positive bids of low types (2.67 and 2.18, see Table 7) and similar bids of the medium types (2.24 and 3.50).

We now focus on the behavior of the strongest bidder in a bidding group, which is either the bidder with v_H in the no-exclusion condition or the bidder with second-highest

[29] Subjects earn 82 (90) percent of the expected profits in *EXP* (*EXUP*) when they place a bid in the interval $(0, v_2]$.

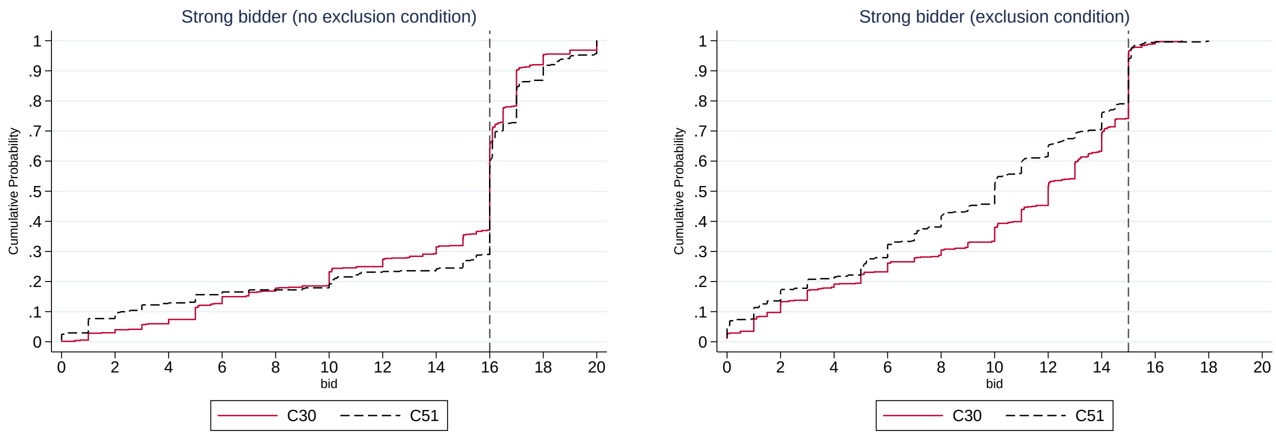

Figure 3: Distribution of cumulative bids relative to the second highest valuation in *EXP* and *EXUP*.

valuation v_M in the exclusion condition. Similar to our findings in the main treatments, we observe behavior that is completely distinct from the theoretical prediction. Figure 3 shows on the left-hand side the cumulative distribution of the bids of high types x_H relative to the valuation of the medium type v_M (no-exclusion condition) and on the right-hand side the bids of the medium type x_M relative to the valuation of the low type v_L (exclusion condition).[30] The figure also shows the theoretical benchmark for the strongest bidder (45-degree line).

It is apparent that in the no-exclusion condition of both *EXP* and *EXUP* (Figure 3, left-hand panel) almost two-thirds of the high types' bids (64 percent) are equal or above the valuation of the medium type, i.e., they apply a "*safe*" bidding strategy. The same behavioral regularity can be observed for medium types in the exclusion condition, who use the *safe* strategy, too. About 32 percent of medium types in *EXP* and about 48 percent of medium types in *EXUP* choose the *safe* strategy when they are the strongest bidder. The difference in using the *safe* strategy in *EXP* and *EXUP* is not statistically significant (Mann-Whitney test $z = 1.281$, $p > 0.20$). However, the difference in the fraction of the *safe* strategy between high and medium types is statistically significant in both *EXP* (Wilcoxon signed-rank test $z = 2.201$, $p < 0.03$) and *EXUP* (Wilcoxon signed-rank test $z = 1.992$, $p < 0.047$).

[30] Note that the predicted bid of the strongest bidder (either the high or medium type) depends on the valuation of the second-strongest bidder v_2 (either the medium or low type) and is uniformly distributed over the interval $[0, v_2]$. Since v_2 varies in each period, we transform the distribution such that the support is independent of v_2 in order to draw the cumulative distribution function (cdf) of the observed bids: the strongest bidder should never bid more than v_2 and thus the maximum ratio of her bid relative to v_2 is one. All bids lower than v_2 are chosen with equal probability. This implies that the strongest bidder's bid relative to v_2 is uniformly distributed over the unit interval: $(x_1/v_2) \sim Uni[0,1]$.

Table 8: Regression: Choice of the *safe* strategy

Dependent variable:	Safe strategy of the strongest bidder.			
	(1) (high/med type)	(1) (high/med type)	(2) (high/med type)	(3) (high/med type)
Three bidders (D)	0.231***	−0.011	0.020	
	(0.040)	(0.061)	(0.062)	
Ten period blocks	−0.008	−0.008	−0.008	−0.028*
	(0.014)	(0.014)	(0.011)	(0.017)
Distance in valuation		0.032***	0.032***	0.037**
		(0.006)	(0.006)	(0.015)
Squared distance		−0.001***	−0.001***	
		(0.000)	(0.000)	
Risk averse (D)			0.097	
			(0.103)	
Avg. bid as med/low type				0.019**
				(0.009)
N	1198	1198	888	141
Pseudo R^2	0.04	0.07	0.09	0.24

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Probit regressions (average marginal effects) with standard errors clustered on matching group level (in parentheses). The dummy variable "Three bidders" equals one in the no-exclusion condition. The variable "Ten period blocks" divides the 50 periods in five 10 period blocks and captures time effects. "Distance in valuations" and "Squared distance" denotes the difference between the strongest and the second strongest valuation and the squared difference, respectively. "Risk averse" is a dummy variable which equals one if a subject prefers safe options smaller or equal to the expected value of the lottery. The variable "Avg. bid as med/low type in period 1-5" captures subjects' own behavior as a weak type in the first five periods. (D) denotes dummy variable.

Table 8 presents results from probit regressions with *safe* bidding as dependent variable. This variable equals one if the strongest bidder has chosen a bid that is at least as high as the respective second-highest valuation, i.e., in the no-exclusion condition *safe* equals one, if $x_H \geq v_M$, and in the exclusion condition *safe* is one if $x_M \geq v_L$. Note also that we pool the data across treatments, since we observe similar behavioral patterns in both treatments.

In column (1) we include a dummy variable "Three bidders," which captures the effect of facing two opponents versus one opponent. The coefficient estimate confirms the earlier result that the use of the *safe* strategy is more prevalent in three bidder groups, i.e., when high types participate in the auction. However, this effect does not persist if we include variables to capture the differences in valuation of the strongest and the second strongest bidder. The variable "Distance in valuation" is $(v_H - v_M)$ for the high type and $(v_M - v_L)$ for the medium type in the exclusion condition. The variable "Squared distance in valuation" is defined accordingly. Column (2) shows a positive and significant effect for the distance in the valuations and a small negative and significant effect for the squared distance. This not only indicates that with an increasing distance in valuations bidders are more likely to choose the *safe* strategy, but also that this effect diminishes after a certain point. A one-point higher difference in valuations is associated with a 3 percentage point increase in the likelihood of playing *safe*. In column (3) we additionally control for risk aversion by including a dummy variable which equals one if a subject prefers safe options that are smaller or equal to the expected value of the lottery in the lottery task. Keeping in mind that our control variable is plagued by a measurement error, the coefficient for risk aversion indicates that risk-averse subjects are more likely to play the *safe* strategy, but it is not significant. For all three specifications we find no evidence of a time trend. Overall, the regression results corroborate the earlier findings of *EX* and *EXUP* as well as of the main treatments.

In the last column we examine how *safe* behavior is driven by the bidding behavior of the other (weak) types. Note that in our setup subjects only learn the winning bid and, thus, they only get information about others' bidding behavior in case they lose the auction. However, it is likely that the strongest bidders draw inferences about the behavior of their competitors from their own behavior as a weak type. In order to look at this potential channel, we examine how own bidding behavior as a medium or low type in early periods affects subjects' inclination to play the *safe* strategy as a strongest bidder in later periods. Note that we concentrate here only on those subjects who had no experience as the strongest bidder in the first five periods, which results in a significantly smaller sample.[31] The variable of interest here is "Avg. bid as med/low type,"which indicates a positive and significant impact of own bidding behavior on playing the *safe* strategy. This suggests that early experience in the role of the weak types has an influence on playing *safe* as the strongest bidder in later periods. If we assume that subjects project their experience as a weak type onto others, this result provides support for the conjecture that, in part, the behavior of the weak types triggers the *safe* strategy of the strongest bidders.

To sum up, we have presented evidence that the *exclusion principle* is not predictive

[31]The choice of five periods reflects the trade-off between the number of observation (subjects) and a sufficient time to experience the different bidder roles. For example, extending the initial periods to 10 reduces the observations from 141 to 35.

for a wide range of parameter values and thus corroborates our findings from the main treatments. As in the main treatments this failure can be traced backed to the behavior of high types who predominantly use a strategy that ensures a secure profit. Moreover the use of this strategy is positively associated with the valuation of the high types, i.e., the larger the distance in valuations between the strongest and second-strongest bidder, the more likely it is that we observe the use of the *safe* strategy. Finally, our results from *EXP* and *EXUP* illustrate that how we assign valuations to subjects has little impact on our qualitative results.

C Appendix – For Online Publication

C.1 Instructions for the All-Pay Auction

General

The second part of the experiment consists of 50 periods in each of which you have to make a decision. Through your decision you can earn points. These points constitute your income which is exchanged to euros according to the conversion rate stated below. Your earnings from the first part of the experiment and from this part will be paid in cash to you at the end of the session.

In each of the 50 periods you will be randomly matched with two other participants to form a group. From now on we will label these two participants as group members. You and the other group members will not learn each other's identity at any point of time. In the following we explain the different decisions you have to make and the procedure of the experiment.

Decision in one period

In each period the computer randomly generates and assigns a number to you and the other group members. One of these numbers will be drawn from the set $\{15, 16, \ldots., 55\}$ and the other two numbers from the set $\{11, 12, \ldots.., 20\}$. In the beginning of each period you will learn your number and the two numbers of the other group members. In the remainder, we will refer to these numbers as "random numbers."

Before you make your decision, the computer randomly decides with a probability of 50% whether the group member with the highest random number is excluded from this period. This means that on average in five out of 10 cases the group member with the highest random number actively participates in that period. Also, in five out of 10 cases the group member with the highest random number is excluded and will not receive an income in that period. If it is not you who has the highest random number in a period you will definitely participate. You will learn in each period whether the group member with the highest random number is being excluded or not.

Every participating group member has to choose an arbitrary number. The number can have up to three decimal places and has to be non-negative (zero is possible). All group member will choose their number simultaneously. We denote this number "decision number."

Calculation of your income in one period

Your income depends on your decision number, as well as the decision number of the other group members and your random number.

After the decisions of all group members have been made, the computer compares and ranks the three decision numbers.

- If your decision number is the highest number, you earn your random number minus your decision number in this period.

  period income = random number – decision number

- If your decision number is not the highest number, you earn zero minus your decision number in this period.

  period income = 0 – decision number

In case of a tie, the highest number is determined randomly.

Please note: The decision number you have chosen will be deducted from your period income independent of the rank of your decision number, i.e., your income will in any case be reduced by your decision number.

If you choose a high decision number, you increase the probability that your decision number is the highest. But a high decision number also reduces your income, since a higher number is deducted from your random number. If your decision number is not the highest, your income will also be reduced by your decision number. At the end of a period you will learn your income in this period. If your decision number was not the highest, you additionally learn the highest decision number. If your decision number was the highest number you will only learn your income in this period.

Example for calculation of the income in one period

Consider the following situation:

Your random number is 28 and you learn the random of the other group members. The computer decides that all group members will participate in this period. You choose 16 as your decision number.

a) If you have the highest decision number, you will earn your "random number" minus your decision number, i.e., your income in this period is 28 – 16 = 12

b) If your decision number is not the highest decision number, you will earn zero minus your decision number, i.e., your income in this period is 0 – 16 = -16

Please note, that your income depends on your random number, your decision number, and the decision numbers of the other two group members.

Consider now the following situation:

Your "random number" is 28 and you learn the random of the other group members. You find out that your decision number is not the highest number in the group. Thus, you participate in any case in this period. The computer decides that the group members with the highest "random number" will be excluded in this period. You choose 16 as your decision number.

a) If you have the highest decision number, you will earn your "random number" minus your decision number, i.e., your income in this period is 28 – 16 = 12

b) If your decision number is not the highest decision number, you will earn zero minus your decision number, i.e., your income in this period is 0 – 16 = -16

Please note, that your income depends on your random number, your decision number, and the decision numbers of the other two group members.

Consider now the following situation:

Your "random number" is 28 and you learn the random of the other group members. You find out that your decision number is the highest number in the group. The computer decides that the group member with the highest random number will be excluded in this period. For you, that means this period is finished and you will not get an income in this period.

After the first period, we repeat this procedure in period 2, period 3, through period 50. In each of the 50 periods you will be randomly matched with two other participants. You are assigned a random number and will learn the random numbers of the other two group members. Then the computer will decide whether the group member with the highest random number will participate in this period. All participating group members simultaneously choose their decision number and learn their income at the end of the period.

Calculation of the total income of the second part of the experiment

In the beginning you will receive a lump-sum payment of 100 points. At the end of the experiment the computer randomly draws 10 periods which determine your income. The points you earned in this period are then added up.

Your total income = 100 + sum of points in 10 randomly drawn periods

Your total income will be converted into euros at a rate of 10 points for one euro.

Trial period

Before we begin, we ask you to participate in a trial period that is not relevant for your earnings.

Quiz for the all-pay auction

Please answer the following questions and mark or fill in the correct answers.

1. Suppose your random number is 19 and your decision number is 12. Your decision number is the highest in your group. Your income in this period is:

   (a) 19

   (b) 12

   (c) 7

   (d) -12

2. Suppose your random number is 15 and your decision number is 6. Your decision number is not the highest in your group. Your income in this period is:

   (a) 9
   (b) – 6
   (c) – 9
   (d) – 15

3. Suppose your random number is 19 and your decision number is 12. All three group members participate in this period.

   (a) If your decision number is the highest in your group, you get ______ points minus ______ points. Your income in points in this period is _________.
   (b) If your decision number is the second highest in your group, you get ______ points minus ______ points. Your income in points in this period is _________.

4. What is your income in 3a) and 3b), when the group member with the highest "random number" is excluded and you participate in this period?

   (a) Income in situation 3a: ____________
   (b) Income in situation 3b: ____________

5. In each period you will be randomly matched with two other participants.

   (a) correct
   (b) wrong

6. If you participate in a period, is the decision number deducted from your income independent of the decision numbers of the other group members?

   (a) Yes
   (b) No

7. The probability of an exclusion of the group member with the highest random number in a period is 30%.

   (a) correct
   (b) wrong

8. A group member with the second or the third highest random number is not excluded in any period.

   (a) correct
   (b) wrong

9. In case two or more decision numbers are the highest number, the highest number is randomly determined.
    (a) correct
    (b) wrong

Discussion Papers of the Research Area Markets and Choice 2017

Research Unit: **Market Behavior**

Dorothea Kübler, Julia Schmid, Robert Stüber SP II 2017-201
Be a man or become a nurse: Comparing gender discrimination by employers across a wide variety of professions

Dietmar Fehr, Julia Schmid SPII 2017-202
Exclusion in the all-pay auction: An experimental investigation

Research Unit: **Economics of Change**

Jannis Engel, Nora Szech SP II 2017-301
The political economy of multilateral aid funds

Maja Adena, Jeyhun Alizade, Frauke Bohner, Julian Harke, Fabio Mesners SP II 2017-302
Quality certifications for nonprofits, charitable giving, and donor's trust: experimental evidence

Terri Kneeland SP II 2017-303
Mechanism design with level-k types: Theory and an application to bilateral trade

Dominik Duell, Justin Mattias Valasek SP II 2017-304
Social identity and political polarization: Evidence on the impact of identity on partisan voting trade

All discussion papers are downloadable:
http://www.wzb.eu/en/publications/discussion-papers/markets-and-choice