Received: 10 July 2015 | Revised: 23 November 2017 | Accepted: 18 December 2017

DOI: 10.1111/jems.12243

ORIGINAL ARTICLE

WILEY Journal of Economics & Management Strategy

Exclusion in all-pay auctions: An experimental investigation

Dietmar Fehr[1] | Julia Schmid[2]

[1]University of Heidelberg, Alfred-Weber-Institute for Economics, Heidelberg, Germany (Email: dietmar.fehr@awi.uni-heidelberg.de)

[2]DIW Berlin, Berlin, Germany (Email: jschmid@diw.de)

Abstract

Contest designers and managers who wish to maximize the overall revenue of a contest 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 strongest bidder in order to promote greater homogeneity among the remaining bidders, even though the theoretical *exclusion principle* predicts otherwise. This is because the strongest bidders are willing to give up a substantial portion of their expected rent in order to minimize the chance of losing the contest.

1 | INTRODUCTION

Relative performance schemes are considered an important tool for motivating the performance of agents in organizational settings, sports, and many other domains of our society (e.g., Frank & Cook, 1995). For example, firms often make use of promotion tournaments and sales competitions, lobbyists seek for influence in the political domain, athletes compete for medals, and researchers strive 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 on the overall performance of the bidders. The closeness of competition, and, by extension, the composition of contestants are critical design parameters for a contest designer, as a heterogeneous contest may have adverse effects on agents' performance.

In recent years, many sports have seen the presence of dominant athletes, such as Roger Federer and Novak Djokovic on the Tennis ATP Tour, or Tiger Woods on the Golf PGA Tour. These superstars garner extensive publicity and serve as the face of their sport. However, excessive dominance by one athlete might also lead audiences to lose interest while encouraging a lower level of competition. For example, due to Michael Schumacher's dominance in Formula One racing, viewing figures dropped and, consequently, the FIA changed several of their rules to make the races more closely contested (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 a theoretical foundation 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 principle*). 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] This can lead to a lower overall performance. The idea behind the *exclusion principle* is thus to increase the bids of the remaining bidders by creating a smaller but more homogeneous contest (without a superstar).

This paper presents an experimental test of the *exclusion principle*. Specifically, 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, that is, 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 thank a coeditor, two anonymous referees, as well as Dorothea Kuebler, Joerg Oechssler, and Christoph Vanberg for very detailed and useful comments. We are grateful for financial support from the Deutsche Forschungsgemeinschaft (DFG) through the SFB 649 "Economic Risk."

 © 2018 Wiley Periodicals, Inc.

We find little support for the theoretical predictions. In homogeneous contests, that is, 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 when the strongest bidder participates than when the strongest bidder is excluded. This remains true independent of the strength of the strongest bidder. That is, revenues are comparable when the strongest bidders' valuation of the prize is almost twice as high as the valuation of the second-strongest bidder and when it is more than three times higher. Therefore, the presence of superstars is not detrimental to contest revenues in our setting.

The main reason for the failure of the *exclusion principle* is the bidding behavior of the strongest bidders. Although the weaker bidders increase their bids significantly in the exclusion condition, they cannot compensate for the lost revenue of the strongest bidder. If the strongest bidders participate, they frequently choose bids equal to or higher than the valuation of the second-strongest bidder. This strategy is clearly not in line with standard theory, as the unique equilibrium is in mixed strategies with no mass points for the strongest bidder. However, this strategy guarantees in theory that the strongest bidder will win the auction and obtain approximately the same profit as expected in equilibrium, because weaker bidders never bid above their own valuation. In the experiment, we find that this "*safe*" strategy wins the auction in 94% of cases as, in contrast to the theoretical prediction, weaker bidders do occasionally bid above their valuation. As a consequence, the profits of the strongest bidders from using the *safe* strategy are, on average, considerably lower than the expected profit from playing the equilibrium mixed strategy. Subjects are more likely to choose the *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 bids equal to or above the valuation of the second-strongest bidder.[4]

The results presented in this paper are linked to a large experimental literature on contests (for a comprehensive survey, see Dechenaux, Kovenock, & Sheremeta, 2015). Although this literature puts much emphasis on tournaments, Tullock contests, and incomplete information all-pay auctions, a smaller number of papers discuss complete information all-pay auctions (e.g., Cason, Masters, & Sheremeta, 2010; Davis & Reilly, 1998; Ernst & Thöni, 2013; Gneezy & Smorodinsky, 2006; Llorente-Saguer, Sheremeta, & Szech, 2016; Lugovskyy, Puzello, & Tucker, 2010). 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 Cason et al., 2010; Davis & Reilly, 1998; Llorente-Saguer et al., 2016).[5]

There are two noteworthy observations that emerge from these studies. First, subjects tend to overbid in comparison to the Nash equilibrium. Second, bidding behavior is bimodal. That is, while subjects seem to randomize their bids, they typically place too much weight on zero or low bids and on high bids.[6] Our results provide further support for these two observations. We find significant overdissipation by the strongest bidder (similar to Davis & Reilly, 1998) as well as evidence that weaker bidders frequently drop out of the auction. The latter is consistent with the finding that the presence of a strong bidder can, for example, translate into a lower entry rate of weaker contestants in all-pay auctions (Cason et al., 2010).

More generally, 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.

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 et al. (1993) with one prize and up to three bidders. Participants in the auction value the prize differently and they are risk neutral. The valuations $v_{i \in \{1,2,3\}}$ are commonly known and heterogeneous in our setup, such that they can be ordered 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, there exists a unique mixed strategy equilibrium. 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 her 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 bid of the bidder with the second-highest valuation v_2 is also uniformly distributed conditional on submitting a positive bid. She remains inactive, that is, 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 $G_2(x) = 1 - v_2/v_1 + x/v_1$. Accordingly, 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 $\mathbb{E}[x_2] = (v_2)^2/2v_1$.

The expected payoff of the strongest bidder in the auction is $v_1 - v_2$, whereas the expected payoff of the second-strongest bidder is zero. The expected sum of bids, that is, the revenue of the auction, adds up to $W(v_1, v_2) = (1 + v_2/v_1)v_2/2$. 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 $v_1 >> v_2 \geq v_3$, that is, 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. The revenue benefit of excluding the strongest bidder depends on v_1 and on how small the difference $v_2 - v_3$ is, as the expected revenue from exclusion is increasing in v_1 and v_3; see inequality (1). 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 less intense competition.

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, as described above.

2.2.1 | Risk elicitation task

The theoretical model assumes risk-neutral players, but risk aversion is often proposed to explain behavior in auctions (see Chen, Ong, & Segev, 2017 for a theoretical discussion). In order to have a measure of subjects' risk attitudes, we directly elicit risk preferences using a binary lottery procedure (see, for example, Dohmen & 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% chance of each outcome, 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.[7]

2.2.2 | All-pay auction

After the first task, subjects repeatedly play the all-pay auction for 50 periods. In the beginning, we randomly assign subjects to matching groups consisting of six persons each, which are 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. Each bidding group consists of a high (v_1), medium (v_2), and low type (v_3) with $v_1 > v_2 > v_3$. Although 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 run sessions with valuations $v \in \{30, 16, 15\}$, hereafter denoted as treatment *High30*, and sessions with $v \in \{51, 16, 15\}$, hereafter denoted as treatment *High51*.

In each period, the bidder with valuation v_1 is excluded from the auction with probability $p = 0.5$.[8] Subsequently, subjects learn the valuation of the other bidders and whether the auction is run among two or three bidders before placing their bids. 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 experience each bidder role—high, medium or low type—over time, that is, the valuations $v \in \{30, 16, 15\}$ or $v \in \{51, 16, 15\}$ are randomly assigned to subjects in a bidding group in each period.[9]

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.[10] High types always face two bidders with valuations $v_2 = 16$ and $v_3 = 15$. Thus, they should bid in both treatments the same in expectation, as their behavior depends only on v_2. In contrast, the behavior of a medium type depends on v_1. Therefore, the share of zero bids should increase in v_1 while revenues decrease in v_1. 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 valuations v_2 and v_3 (exclusion condition) to the revenue of an auction with all three bidders with the 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.

2.2.3 | Procedures

We conducted seven computerized sessions with 18–24 participants each at the experimental laboratory at the Technical University of Berlin using the software tool kit *z-Tree* (Fischbacher, 2007). Because we randomly match subjects in groups of three, the behavior over time is likely to depend on previous interactions between bidders within matching groups. Thus, we treat a matching group as a unit of observation, which results in 24 independent groups for our statistical analysis below.

Subjects were recruited from a large database where students can voluntarily register for participation 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 second 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 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 auction. Table 1 presents summary statistics of the 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 (high type) 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 the exclusion of the strongest bidder 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 separately confirms that it does not pay off to exclude the strongest bidder in our setup. In fact, the revenues in the no-exclusion condition are higher than the revenues with exclusion in both *High30* and *High51*, but the difference is less pronounced when the strongest bidder is weaker, that is, when $v_1 - v_2$ is smaller as in *High30*. The difference in the

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
	[16.5]	[15.0]	[16.2]	[15.0]	[17.0]	[14.9]
	(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	1,203	1,193	707	692	496	501
Minimum bid	0	0	0	0	0	0
Maximum bid	57.5	61.6	57.5	61.6	53	31

Note: Medians are in brackets and 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.

average sum of bids between the two conditions (exclusion and no exclusion) is not statistically significant in *High30* (Wilcoxon signed-rank test, $z = 0.97$, $p = 0.33$, $n = 14$), but is in *High51* (Wilcoxon signed-rank test, $z = 2.4$, $p < 0.017$, $n = 10$).[11]

From Table 1 it is apparent that, on average, revenues are always higher than predicted, 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% of cases revenues are higher than predicted when all three bidders participate in the auction (*High30*). This share is slightly higher in *High51* at 85% . 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? Having observed substantial overbidding in the presence of three bidders, we can ask whether exclusion would have been profitable if the strongest bidders had behaved as 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

To get a deeper insight into the underlying reasons for the failure of the *exclusion principle*, 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) along with the theoretical benchmark (long-dashed lines).

It is striking that the strongest bidders bid more than predicted, as evidenced in Table 2 and Figure 1. 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).[12] 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 medium types in the no-exclusion condition 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, that is, 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. Figure 1 reveals that a substantial share of bids are equal to or above the valuation of the second-strongest bidder, that is, $x_1 \geq v_2$. This is particularly the case in the no-exclusion condition (top-left panel of Figure 1). When high types participate in the auction, we observe bids that are equal to or higher than the valuation of the second-strongest bidder in about 63% of cases in *High30*. Although the strategic situation for high types does not differ between treatments, we observe more mass at the second-highest valuation when the contest is more heterogeneous, as in *High51*. More precisely, the share of bidding at

TABLE 2 Summary statistics of individual bids of bidder types

		High30			*High51*		
		High	**Medium**	**Low**	**High**	**Medium**	**Low**
		v_1	v_2	v_3	v_1	v_2	v_3
No exclusion	Avg. bid	13.74	2.36	1.71	15.07	2.15	1.09
		[16]	[0]	[0]	[16]	[0]	[0]
		(4.92)	(5.19)	(4.57)	(6.71)	(5.04)	(3.58)
	Avg. predicted bid	8	4.27	0.00	8	2.51	0.00
Exclusion	Avg. bid	–	10.39	6.49	–	9.09	5.45
			[12]	[5]		[10]	[2.33]
			(5.44)	(6.59)		(5.23)	(5.81)
	Avg. predicted bid	–	7.5	7.03	–	7.5	7.03

Note: Medians are in brackets and standard deviations in parentheses. We excluded bids $x_i > 55$. This was the case in four of 7,200 individual bids.

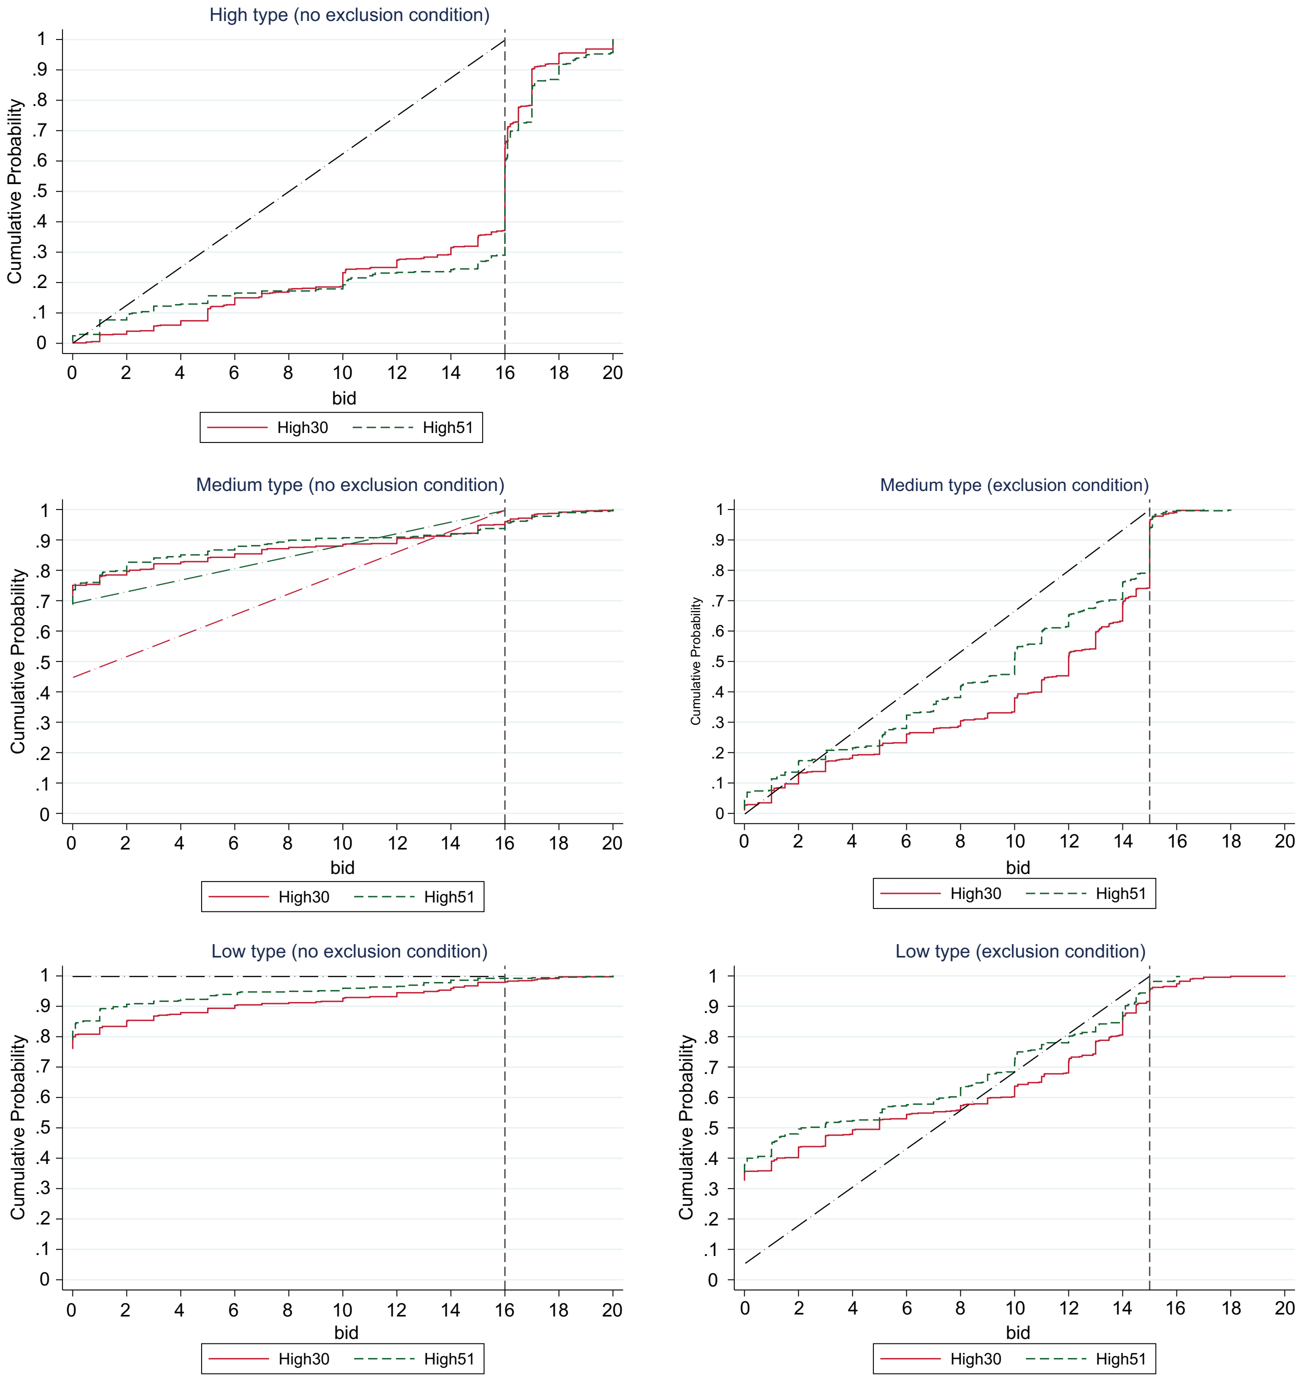

FIGURE 1 Cumulative distribution of bids in the exclusion and no-exclusion condition along with the theoretical benchmark (long-dashed lines) [Color figure can be viewed at wileyonlinelibrary.com]

or above the valuation of the second-strongest bidder is 11 percentage points higher in *High51* than in *High30*. 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% of bids in *High30* and 21% of bids in *High51* are equal to or larger than the valuation of the second-strongest bidder.[13]

There are two major differences between the exclusion and no-exclusion condition that may account for the pronounced difference in bidding at or above v_2 in these two conditions. First, the difference in valuations $v_1 - v_2$ (and thus the "certain" profit) is obviously larger in the no-exclusion than in the exclusion condition. Thus, high types have more to lose and may have more to regret if they place a bid lower than v_2 and lose the auction than medium types in the exclusion condition. Second, in the no-exclusion condition high types face two other bidders who are almost equally strong, whereas medium types in the exclusion condition are confronted with only one opponent. Although the number of bidders should not matter from a theoretical perspective, as only the two strongest bidders participate in the auction, it may make a difference if high types believe that they compete against both weaker bidders.

Observing bids $x_1 \geq v_2$ is certainly not in line with theory, which involves no mass points in the distribution of bids. In the unique mixed strategy equilibrium, the strongest bidder randomizes bids in the interval $[0, v_2]$. The second-strongest bidder enters the competition only with probability v_2/v_1 and then randomizes bids in the interval $[0, v_2]$ as well. As weaker bidders theoretically never place a bid above their valuation, the strongest bidder can win the auction with certainty when bidding v_2 (or $v_2 + \epsilon > 0$). Such bidding behavior is not optimal, but it gives rise to profits that are approximately the same as the expected profit in equilibrium, that is, $v_2 - v_1$.[14] In the following, we refer to bidding at or above v_2 by the strongest bidder as "*safe*" strategy. Empirically, the likelihood of winning the auction when bidding "*safe*" is, however, below one, as the weaker bidders ($j = 2, 3$) occasionally bid above their own valuation ($x_j > v_2$). We examine the use of the "*safe*" strategy in more detail in the next section.

3.3 | Anatomy of *safe* bidding

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 nearly resembles a uniform distribution, as predicted (see Figure 1 in Appendix A in Supporting Information). 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 8.7 in both conditions, 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 bids is strictly greater than v_2. However, we have to note that the overwhelming majority of these bids (74% in the no-exclusion and 97% in the exclusion condition) are in a comparatively small interval $(v_2, v_2 + 1]$.

In the following, we confine our analysis 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. As pointed out in the previous section, playing the *safe* strategy is clearly not in line with theory. Yet, a subject's profit from playing *safe* should be close to the expected equilibrium profit, $v_1 - v_2$, provided that the bid is equal to or infinitesimally larger than v_2.

Although we have seen that the large majority of *safe* bids is indeed close to v_2, implying profits not far from $v_1 - v_2$, we find that the profits from playing the *safe* strategy are substantially lower than this theoretical benchmark. Table 3 illustrates that in the no-exclusion condition bidding in $(0, v_2)$ results in average profits that are close to $v_1 - v_2$ (with 14.9 in *High30* and 35.9 in *High51*), whereas the *safe* strategy leads to lower average profits, with 11 in *High30* and 30.4 in *High51*. 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*).[15]

That the *safe* strategy performs worse than bidding in $(0, v_2)$ is mainly due to weaker bidders who occasionally bid above their valuation (see also Figure 1). In fact, high types using the *safe* strategy lose the auction in about 7% and 5% of cases in *High30* and *High51*, respectively.[16] Consequently, by choosing the *safe* strategy the strongest bidders forgo a substantial part of their rent in order to increase the likelihood of winning. More precisely, they earn only about three quarters of the average profits that would accrue from placing a bid in the interval $(0, v_2)$ in *High30*. Similarly, the strongest bidders earn about 85% of the expected profit of a bid in the interval $(0, v_2)$ in *High51*.

Is the prevalence of the *safe* strategy and the associated lower profits driven by 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,

TABLE 3 Bidding behavior of the strongest bidder

		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

Note: 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*.

that is, bid zero (see left-panel of Figure 1). However, while medium types tend to drop out too much, low types drop out too little.[17] In fact, low types should never place a positive bid in the no-exclusion condition.

In *High30,* we observe that medium types place a zero bid in about 72% of cases although they should only abstain in 47% of cases. If they participate, the average bid is 8.4 (17% of bids are above v_2). At the same time, we observe that low types place a positive bid in 24% of cases and the average bid is 7.1 (12% of bids are above v_3). Thus, it seems that low types' behavior compensates in a large part for the less frequent bids of medium types. This bidding behavior gives rise to winning the auction in 26% (20%) of cases for medium (low) types, conditional on participation. Overall the profits of the weaker bidders are negative, but close to zero for both types (−1).

In *High51,* we see that medium types place a zero bid in about 68% of cases, which is close to the theoretical prediction. If they place a positive bid, it is on average 6.8 (17% of bids are above v_2). Again, we observe that low types participate in 22% of cases, with an average bid of 5 (7% of bids are above v_3). Owing to this behavior of weaker bidders, high types in *High51* face an active competitor more often than predicted. Again, this bidding behavior results in a fairly high share of wins (17% for medium types and 14% for low types) and to profits that are, on average, close to zero (−1).

The distribution of winning bids, which is observed by subjects, reflects these differences in the participation of weaker types. There are significantly fewer winning bids $x_i \in [0, 16)$ in *High51* than in *High30* (24% vs. 36%) and, consequently, significantly more winning bids $v_2 \geq 16$ in *High51* (Fisher's exact test, $p < 0.01$).

Together, this suggests that the participation of low types triggers more *safe* bidding of high types in *High51* compared to *High30*. On the surface, it seems that weaker types bid too much on average and that they could improve their profits by abstaining more often from bidding (particularly low types), or at least by bidding below their own valuation in both treatments. However, the profits of both the medium and low types are already close to zero as predicted, such that there is little room for improvement in response to high types' behavior. In contrast, the strongest bidders 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 has provided some suggestive evidence for the differences in the use of the *safe* strategy between treatments. 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, that is, 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 $v_2 \geq 16$.

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 also 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, that is, 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. Columns (3) and (4) investigate how an individual's past profits affect *safe* bidding. Although 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 particularly less successful bidders 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 coefficient estimates. We should, however, keep in mind that our measure for risk aversion is naturally measured with error resulting in a downward bias.[18]

In the final three columns, we investigate how participants' behavior and experience in early periods affect their 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.[19] First, we investigate the relationship of initial bidding behavior as a weak type—either as a medium or low type—and subsequent bidding as a strong type in periods 21–50. Column (7) shows that higher average bids as a weak type in the first 20 periods are associated with 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

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

	Safe Strategy of the Strongest Bidder								
Dependent Variable:	**(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							0.024**		
in period 1–20							(0.011)		
Avg. bid in								0.038***	
period 1–20								(0.014)	
Avg. share of *safe* bidding									0.346***
in period 1–20									(0.116)
N	1,205	1,181	1,183	1,177	1,153	1,153	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 (7–9) only include data from period 21 to 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 euros in period 1. "Risk averse" is a dummy variable, which equals one if a subject prefers safe options equal or smaller to the expected value of the lottery. The variable "Avg. bid in period 1–20" ("Avg. bid as weak type in period 1–20") captures initial bidding behavior of an individual (as a weak type) 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.

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 encourages *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) relates to more *safe* bids in later periods as well, thus confirming our earlier finding.

In summary, the regressions reveal 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 also seems that a significant part of this behavior is related to past performance, initial behavior as one of the weak types, and to more exposure to *safe* bidding in the beginning.

4 | DISCUSSION

We have seen that the strongest bidders engage in frequent bidding at or above v_2, particularly in the no-exclusion condition. In the following, we discuss several possible explanations for such behavior.

It is conceivable that the prevalence of *safe* bidding is the result of revealing only the winning bids to subjects. Although we cannot rule out that this limited feedback slows down or hinders the learning of subjects, there are good reasons to think that such a lack of complete information about other bids is not solely responsible for the large share of *safe* bidding that we observe. First, subjects experience all bidder roles and thus have the possibility to learn bidding behavior over time. In fact, the average winning bid in both conditions and treatments is 13.2 (median 15) and in 69% of cases in the interval [0, 16], indicating that it is possible to win with bids that are (substantially) below v_2. Second, we observe that the extent of *safe* bidding declines over time. For example, in the no-exclusion condition it is 10 percentage points lower in the second half (periods 26–50) than in the first half (periods 1–25). These findings suggest that other reasons, which we will discuss below, play a role as well.[20]

The observed *safe* bidding is consistent with a nonequilibrium model of limited sophistication (Nagel, 1995; Stahl & Wilson, 1995). In the standard formulation of this model, players anchor their beliefs on a nonsophisticated level-0 player but differ with respect to their levels of reasoning.[21] 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 that guarantees approximately the same profit as expected from playing a mixed strategy. Consequently, engaging in more levels of reasoning or investing more cognitive resources may not be worthwhile. 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.

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, generating a guaranteed profit of $v_1 - v_2$.[22] This kind of regret is specific to our complete information environment since it presumes that the valuations of the competitors are public knowledge.[23] 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 high types' behavior. However, we observe that, on average, their profits are already close to the prediction of zero profits in both treatments, such that improvements are negligible. On the other hand, high types could clearly improve their profits by adopting a mixed strategy, given the behavior of the weaker types. 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 falling in the ranking.[24] A similar motivation may have led to *safe* bidding in our setting.

5 | CONCLUSION

Superstars can have a major impact on the attractiveness of contests, but at the same time their presence can detrimentally affect the performance motivation of competitors. In this paper, we experimentally investigated the effect of excluding superstars from the contest and thereby creating a more homogeneous participant pool. We find that the behavior of the weaker bidders is mostly in line with qualitative theoretical predictions: they are discouraged by the presence of a strong bidder ("superstar"), and they substantially increase their participation and average effort (bids) in the absence of a superstar (see, for example, Brown, 2011 for field evidence from the PGA Tour). However, in strong contrast to the theoretical predictions, excluding the strongest bidder is not beneficial for the contest designer. In fact, total revenues are significantly higher when superstars are allowed to participate.

The main reason for this result is that the superstars often apply a strategy of bidding at least as much as their most powerful competitor. This strategy performs, on average, worse than bidding below the valuation of the second-strongest bidder. Thus, it seems that the strongest bidders prefer to give up a substantial portion of their rent (25% in *High30* and 15% in *High51*) in order to avoid losing the auction. As a consequence, the increased effort of the weaker bidders in the absence of the strongest bidder cannot compensate for the lost effort of the superstar. We therefore find little support for the *exclusion principle*.

Our finding that superstars prefer a *safe* bid over a mixed strategy even though this is less profitable suggests that psychological factors may play a role. Indeed, the strongest bidder may, for example, incur disutility from feelings of regret when losing the auction or, more generally, suffer a psychological loss, as they expect to win owing to their superiority. To the extent this is true, managers and contest designers can harness such behavior to their advantage. For example, by making sure the bidders' valuations are public knowledge, one can create preconditions for regret, thus encouraging *safe* bidding. Although providing such transparency is not always feasible, it is often possible to provide sufficiently correlated signals about the valuations of participating bidders or to provide updates concerning the performance ranking of contestants (Casas-Arce & Martínez-Jerez, 2009).

The composition of teams is critical for the performance of firms and organizations in general and for contest designers in particular (Mathieu, Tannenbaum, Donsbach, & Alliger, 2014). In practice, firms typically face a heterogeneous workforce. Nevertheless, research indicates that relative performance schemes are widely applied in the firm context (e.g., Connelly, Tihanyi, Crook, & Gangloff, 2014). Although some authors have called the use of such schemes into question by pointing to the disincentivizing effects that may be exerted by superstars (e.g., Brown, 2011), our findings highlight that the presence of a superstar is not detrimental to overall revenues.

That the results provide no support for the *exclusion principle* suggest that the composition of contestants is of second-order importance. Arguably, there are situations in organizational contexts where the presence of a superstar is desirable. For example, when the valuation gap between the weaker bidders might be too large to facilitate heightened competition.[25] In other cases, managers may implement heterogeneous contests because they may reveal more information about bidders' abilities (Gürtler & Gürtler, 2015). For example, this could be a relevant issue at law firms, consultancies or investment banks, where many associates typically compete for a few spots as a partner and the company only wants to promote the most able associates. Finally, managers may want to benefit from the presence of a superstar, as they typically add status, popularity or glamor to the contest.[26] In short, our research provides insights into contest design with heterogeneous workers.

ENDNOTES

[1] Another anecdote recounts that Tiger Woods' landslide victory at the 1997 Masters in Augusta led to deliberations about redesigning the Augusta National (Golf Digest, 2008).

[2] Likewise, U.S. professional sport leagues, such as the NBA, NFL, NHL, and MLB, put considerable 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 granting the weaker teams from the previous season the right to have first choice of the rookies from the pool of the best young prospects. This is in contrast to many sport leagues in Europe. Soccer leagues, for example, are usually dominated by the same two or three teams.

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

[4] In Supporting Information, we present additional evidence from two treatments where valuations vary across periods, which allows us to study the *exclusion principle* for a broad range of parameters. We show that in about 80% 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 of the strongest bidders and the associated prevalence of bidding at or above the valuation of the strongest competitor.

[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.

[6] 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, for example, Barut, Kovenock, & Noussair, 2002; Müller & Schotter, 2010, and Noussair & Silver, 2006).

[7] 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 of 144 subjects (12%) switched multiple times between the safe option and the lottery.

[8] The instructions stated that the computer will randomly decide with a probability of 50% whether the group member with the highest valuation is excluded in a period. For the detailed explanation in the instructions; see Appendix C in Supporting Information.

[9] 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 (random valuations). 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 more complex strategic situation and rich environment that is not idiosyncratic to a specific choice of valuations. We present the results of these two treatments in Appendix B in Supporting Information and show that our findings from the fixed-valuation treatments reported here are robust to the assignment of valuations (random or fixed valuations).

[10] 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 because the valuations of medium and low types are always the same. Although the absolute difference between the exclusion and no-exclusion conditions seem small, the relative differences are substantial (18% and 38%).

[11] Note that the average sum of bids in the exclusion condition is higher in *High30* than in *High51* despite an identical strategic situation in both treatments (Mann–Whitney test $z = 1.93$, $p < 0.06$). However, this difference is mainly due to excessive bidding in 2 of 14 groups in *High30*. Indeed, the median sum of bids is almost identical in the exclusion conditions of the two treatments; see Table 1.

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

[13] Bidding at or above v_2 is consistent with a level-k reasoning process (see, for example, Nagel, 1995; Stahl & Wilson, 1995). Assuming that a level-0 player i randomizes bids in the interval $[0, v_2]$, a level-1 *high type* best responds to this belief by bidding v_2, whereas a level-1 *medium type* best responds to this belief by playing the equilibrium mixed strategy. In turn, a level-2 *high type* believes that she is facing a level-1 *medium type* and thus best responds by playing the equilibrium mixed strategy, whereas a level-2 *medium type* assumes her opponent is a level-1 *high type* and best responds to this belief by placing a zero bid. Note that a level-k *low type* ($k > 0$) never places a positive bid.

[14] If the strongest bidder bids $x_1 = v_2$ (or slightly above v_2), then the second-strongest bidders' best reply is $x_2 = 0$. In this case, it is profitable for the strongest bidder to deviate to a substantially lower bid $x_1 = \epsilon > 0$. But then $x_2 = 0$ is not a best reply to $x_1 = \epsilon > 0$. More generally, as x_2 cannot be a best reply to x_1 if x_1 is a best reply to this x_2, there exists no pure-strategy equilibrium.

[15] 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 $z = 2.3$, $p < 0.025$, $n = 10$).

[16] The likelihood of losing with a bid $x_i = v_2$ or $x_i > v_2$ is about the same in both treatments (i.e., 10% if $x_i = v_2$ and 7% if $x_i > v_2$ in *High30*, and 7%, if $x_i = v_2$ and 5% if $x_i > v_2$ in *High51*). Average losses when bidding *safe* and losing the auction are −16.3 in *High30* and −16.7 in *High51*.

[17] This represents mixed evidence for a discouragement effect, that is, 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), and Chen, Ham, and Lim (2011) find no effect.

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

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

[20] There is very little evidence for the impact of feedback on bidding behavior in (heterogeneous) all-pay auctions (see, for example, Barut et al., 2002 and Hyndman, Ozbay, & Sujarittanonta, 2012) and it would be of interest to scrutinize this issue further.

[21] 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 that have higher expected payoffs with higher probability, which is consistent with observed overdissipation patterns found in all-pay auctions.

[22] Note that the observed overbidding may be explained by loss aversion, as playing a mixed strategy can involve a loss (e.g., Ernst & Thöni, 2013; Müller & Schotter, 2010). 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*.

[23] Regret has been previously analyzed in symmetric or incomplete information auctions (see, for example, Baye, Kovenock, & de Vries, 2012; Filiz-Ozbay & Ozbay, 2007, and Hyndman et al., 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.

[24] The NASCAR race series is a multiple-round tournament, where rankings are publicly available and 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 & Martínez-Jerez, 2009).

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

[26] Evidence from NFL teams suggests that a larger pay disparity within teams (which indicates the presence of superstars) is positively related to the franchise value of NFL teams (Mondello & Maxcy, 2009).

ORCID

Dietmar Fehr http://orcid.org/0000-0001-9735-390X

REFERENCES

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

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

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

Baye, M. R., Kovenock, D., & de Vries, C. G. (1996). The all-pay auction with complete information. *Economic Theory*, *8*, 291–305.

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

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

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

Bothner, M. S., Kang, J.-h., & Stuart, T. E. (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., & Martínez-Jerez, F. A. (2009). Relative performance compensation, contests, and dynamic incentives. *Management Science*, *55*(8), 1306–1320.

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

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

Chen, Z. C., Ong, D., & Segev, E. (2017). Heterogeneous risk/loss aversion in complete information all-pay auctions. *European Economic Review*, *95*, 23–37.

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

Davis, D. D., & Reilly, R. J. (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., Kovenock, D., & Sheremeta, R. M. (2015). A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics*, *18*(4), 609–669.

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

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

Filiz-Ozbay, E., & Ozbay, E. Y. (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., & Cook, P. J. (1995). *The winner-take-all society: Why the few at the top get so much more than the rest of us*. London, UK: Penguin Books.

Gill, D., & Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *The American Economic Review*, *102*, 469–503.

Gillen, B., Snowberg, E., & Yariv, L. (2015). *Experimenting with measurement error: Techniques with applications to the Caltech cohort study*. NBER Working paper.

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

Golf Digest (2008). *Tiger-proofing Augusta*. Retrieved from http://www.golfdigest.com/ story/gw19970418rosaforte.

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

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

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

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

Hyndman, K., Ozbay, E., & Sujarittanonta, P. (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*. Oxford, UK: Oxford University Press.

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

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

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

Mondello, M., & Maxcy, J. (2009). The impact of salary dispersion and performance bonuses in NFL organizations. *Management Decision*, *47*(1), 110–123.

Müller, W., & Schotter, A. (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., & Silver, J. (2006). Behavior in all-pay auctions with incomplete information. *Games and Economics Behavior*, *55*, 189–206.

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

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Fehr D, Schmid J. Exclusion in all-pay auctions: An experimental investigation. *J Econ Manage Strat*. 2018;1–14. https://doi.org/10.1111/jems.12243