All pay auctions and group size: Grading on a curve and other applications

James Andreoni [a,b], Andy Brownback [c,*]

[a] Department of Economics, University of California, San Diego, La Jolla, CA 92093, United States
[b] NBER, United States
[c] Department of Economics, University of Arkansas, Fayetteville, AR 72701, United States

ABSTRACT

We model contests with a fixed proportion of prizes, such as a grading curve, as all-pay auctions where higher effort weakly increases the likelihood of a prize. We find theoretical predictions for the heterogeneous effect auction size has on effort from high- and low-types. We test our predictions in a laboratory experiment that compares behavior in two-bidder, one-prize auctions with behavior in 20-bidder, 10-prize auctions. We find a statistically significant 11.8% increase in aggregate bidding when moving from the small to large auction. The impact is heterogeneous: as the auction size increases, low-types decrease effort but high-types increase effort. Additionally, the larger auction provides a stronger rank-correlation between effort and ability, awarding more prizes to the higher-skilled and improving the efficiency of prize allocation.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

We study the impact of changing group size in all-pay auctions with a fixed proportion of prizes. We derive theoretical predictions and confirm these with experimental evidence. While our findings are broadly applicable, to fix ideas and focus intuition we present the issues in terms of one of the most interesting applications of our result: grading on a curve. Students of heterogeneous skills choose how much effort to invest in a class in order to earn one of a fixed proportion of "A" grades. In this application, the skill level is the private value and the amount of effort invested in the course is the bid. The all-pay auction format suits the application since effort, the bid, is non-refundable. What makes the question of group size more interesting here is that the objectives of the educator as the auction designer may not be as simple as maximizing revenue, that is, total effort. They may also care about the distribution of effort among different students, or on accurately learning the relative skills of students, that is, the class ranking. Our analysis can help sort out the effects of class size on this array of possible objectives.

We show that as the enrollment of a course decreases, a given student becomes less sure of her percentile ranking in the course, and this increased uncertainty has asymmetric effects on those at the high and low ends of the population distribution of skills. For low skilled it means working harder is more likely to generate a top grade, while for the high skilled it means slacking off will be less likely to cost the student the coveted "A" grade. For the auction designer, the optimal class size will depend on the objectives. If the objective is to maximize total effort, then larger class sizes will produce more effort. If,

---

* Corresponding author.
  E-mail addresses: andreoni@ucsd.edu (J. Andreoni), abrownback@walton.uark.edu (A. Brownback).

however, the objective is to increase the effort of the lower skilled students, and thus improve their outcomes, then smaller class sizes will be optimal. Finally, if the objective is to accurately certify the ranking of students by the quality of their effort, then our experimental results suggest that larger class sizes are preferred.

While the application to classroom effort is natural and important, our results can be applied to many situations. Should an academic department set aside tenured slots for specific fields, such as a "macro" position and a "micro" position, or should it award the slots to the best candidates regardless of field? How should a large law firm structure it's awarding of partner to a cohort of associates? How does promotion based on seniority as well as effort compare to promotion based on effort alone? These and others are all interesting applications of our results.

We are, of course, not the first to examine the effect of the number of bidders in an all-pay auction.[1] We distinguish our study from these by examining the effects of increasing the group size while maintaining the *proportion* of prizes, not simply the *number* of prizes. In addition, we focus not simply on the effects on auction profits, but our theory and our within-subject experimental design allow us to explore the heterogeneous impacts of group size for a continuum of types.

2. Background

The literature on all-pay auctions and contests dates back several decades[2] and has been adapted to explore a wide variety of settings such as labor contracts (Lazear and Rosen, 1981), classroom grading (Becker and Rosen, 1992), and political lobbying (Hillman and Riley, 1989; Baye et al., 1993).

Experimental papers have studied changes in the structure of contests similar to ours. For example, Potters et al. (1998) and Davis and Reilly (1998) test bidding in Tullock contests under different structures. Other papers explore the effects of changes in the proportion of contestants awarded prizes (Barut et al., 2002; Gneezy and Smorodinsky, 2006; Lim et al., 2014; Harbring and Irlenbusch, 2005; Orrison et al., 2004). Still others have focused on the heterogeneity of bidding behavior in contests with a fixed size, noting a "bifurcation" between low- and high-ability types (Müller and Schotter, 2010; Noussair and Silver, 2006). Minor (2012) investigates this "coarse" behavior, where bidders choose between a small number of disparate strategies.[3]

Our investigation of the impact of changes in contest design on bidder heterogeneity, therefore, fills a natural and important gap in this literature. A focus on average effects of design changes has obscured the tension present between the interests of the high- and low-skilled contestants. We identify this tension theoretically, and then confirm its existence in our experimental results. In addition to the first test of the effects of a large-scale change in contest size, we employ a powerful design that identifies within-subject effects resulting from this change. By addressing this aspect of contest design, we hope to inform policy makers about the impact that changes in group size will have on people with different skill-levels when they are subject to relative evaluations.

3. Theoretical model

Consider an independent private value all-pay auction. Values are drawn from a known distribution and represent the money earned when winning a prize. Bidders compete for a limited number of such prizes by simultaneously placing irreversible bids that must be paid regardless of the outcome of the auction. Prizes are awarded to the highest bidders, meaning an increase in bidding weakly increases the probability of receiving a prize. Uncertainty enters this model through an unknown draw of opponents. Bidders know the distribution from which all valuations are drawn, but do not know each other's realized values.[4]

3.1. Basics of the model

Suppose a given auction has N bidders. Call this the auction's size. Bidders compete for one of P prizes with $P < N$. Each bidder independently draws a private value, v_i, with the commonly known cumulative distribution function, $F(v_i)$. These values can be thought of as a bidder's ability.[5] Higher ability bidders gain more surplus from winning an auction than do lower ability bidders. This framework is isomorphic to an independent private cost of effort with the higher "ability" bidders having either higher valuations or lower costs of effort in the respective environments.[6]

---

[1] Moldovanu and Sela (2001, 2006) theoretically explore the allocation of prizes and the optimal "architecture" of a contest for a fixed *number of contestants*, and Olszewski and Siegel (2013) look at how the number of bidders affects bids for a fixed *number of prizes*, rather than a fixed *proportion of prizes* as is done here. Experimentally, Harbring and Irlenbusch (2005) and Orrison et al. (2004) test the effects of small changes in a contest's size on bidder behavior. Harbring and Irlenbusch (2005) employ homogeneous types and Orrison et al. (2004) only address heterogeneity among two types of players, one type advantaged by the addition of a constant to their output.

[2] See Tullock (1967) and Krueger (1974). Baye et al. (1996) provide the full characterization of the equilibria in the complete information all-pay auction, while Amann and Leininger (1996) and Krishna and Morgan (1997) expand the model to incorporate incomplete and possibly correlated information.

[3] Dechenaux et al. (2012) provide a complete review of the contest literature.

[4] Using a Tullock contest could have further integrated uncertainty into the model, though the two environments are qualitatively similar. Experimentally, the all-pay auction is simpler for subjects and we did not believe that the additional layer of complexity of the Tullock contest was worth the risk of confusion.

[5] For a detailed discussion of the differences between ability, value, and cost of effort, see Baik (1994).

[6] This isomorphism will not hold under preferences that do not maximize expected utility.

A bidder competes for a prize by casting a bid, b_i, the cost of which is paid regardless of the outcome. Auctions are one-shot with simultaneous bids, so bidders cannot condition bids on the behavior of their competitors. Bidders with the P highest bids receive one prize each. Let $Pr_{N,P}(b_i)$ be the probability that bid b_i will win a prize in an auction with N participants and P prizes.

Assume risk neutral, expected-utility maximizing bidders.[7] A bidder's utility is then:

$$
\begin{aligned}
U(b_i; v_i, N, P) &= Pr_{N,P}(b_i)(v_i - b_i) + (1 - Pr_{N,P}(b_i))(-b_i), \\
&= v_i \times Pr_{N,P}(b_i) - b_i
\end{aligned}
\tag{1}
$$

3.2. Equilibrium bidding

Athey (2001) demonstrates the existence of a monotone Bayesian Nash equilibrium bidding function in pure-strategies. This equilibrium in pure strategies is possible as a result of the incomplete information about the valuations of other bidders. The precision of inference about the distribution of valuations among other bidders changes with the auction size. Equilibrium bidding, therefore, changes with the auction size as well. We will now construct equilibrium bidding functions for an arbitrary auction size.

Call the monotonic equilibrium bidding function $B(v_i; N, P)$. This maps from the space of valuations onto the set of possible bids: $B : v_i \to b_i$, where $v_i \in [0, 1]$ and $b_i \in [0, 1]$.[8]

Under monotonicity, there exists a function, $Z_{N,P} : v_i \to [0, 1]$, that maps valuations to the equilibrium probabilities of winning a prize in an auction of size N with P prizes. Specifically, this function returns the probability that a valuation, v_i, is higher than at least $N - P$ of the $N - 1$ other valuations in the auction. This is expressed as an order-statistic:

$$
Z_{N,P}(v_i) = \sum_{k=N-P}^{N-1} \left(\frac{(N-1)!}{(N-1-k)!k!} \right) F(v_i)^k (1 - F(v_i))^{N-1-k}.
\tag{2}
$$

Recall, $F(v_i)$ is the cumulative distribution function for v_i.

It is straightforward to show that, as a well-defined, symmetric, strictly monotonic bidding function, $B : v_i \to \mathbb{R}^+$ must also be continuous.[9] Continuous, monotone functions are invertible implying that there exists a function, $B^{-1}(b_i)$, that maps bids cast onto the valuations implied by those bids. Denote this inverse function $V(b_i)$.

Incorporating $Z_{N,P}(V(b_i))$ into the bidder's utility captures the incentive to deviate from equilibrium by assuming the bid associated with a different valuation. The first order condition of this utility function must be satisfied by any equilibrium bidding function. The modified utility and its first order condition are given by (3) and (4), respectively.

$$
U(b_i; v_i) = v_i \times \sum_{k=N-P}^{N-1} \left(\frac{(N-1)!}{(N-1-k)!k!} \right) F(V(b_i))^k (1 - F(V(b_i)))^{N-1-k} - b_i
\tag{3}
$$

$$
\begin{aligned}
\frac{1}{V'(b_i)} = v_i \times \sum_{k=N-P}^{N-1} &\left\{ \left(\frac{(N-1)!}{(N-1-k)!k!} \right) \right. \\
&\left. \times \left[k \times f(V(b_i))F(V(b_i))^{k-1}(1 - F(V(b_i)))^{N-1-k} - (N-1-k) \times f(V(b_i))F(V(b_i))^{k}(1 - F(V(b_i)))^{N-2-k} \right] \right\}
\end{aligned}
\tag{4}
$$

Since $V(b_i) \equiv B^{-1}(b_i)$, it follows that, $\frac{1}{V'(b_i)} = B'(v_i)$. Additionally, equilibrium requires that $Z_{N,P}(V(b_i)) = Z_{N,P}(v_i)$. We thus arrive at a first order differential equation from which we can solve the general form of the equilibrium bidding function,

$$
B'(v_i) = v_i \times \sum_{k=N-P}^{N-1} \left\{ \left(\frac{(N-1)!}{(N-1-k)!k!} \right) \left[k \times f(v_i)F(v_i)^{k-1}(1 - F(v_i))^{N-1-k} - (N-1-k) \times f(v_i)F(v_i)^{k}(1 - F(v_i))^{N-2-k} \right] \right\}.
\tag{5}
$$

---

[7] In the appendix, we demonstrate that the qualitative predictions of the model generalize to risk aversion when risk aversion parameters are shared and common knowledge.

[8] Strict monotonicity follows naturally from the model. Parreiras and Rubinchik (2010) rule out weak monotonicity by contradiction: Suppose bidding were flat at a bid of $\hat{b}$ for a non-degenerate interval of valuations, then there would be a discontinuous jump in the probability of winning by bidding just above $\hat{b}$. This profitable deviation completes the contradiction by showing that it would never be optimal to bid $\hat{b}$. Strict monotonicity also satisfies intuition, as subjects likely understand that higher-valuation bidders will cast higher bids.

[9] Suppose a discontinuity were present at equilibrium. Bidders casting bids at the upper bound of the discontinuity would profit from decreasing their bid to just above the lower bound of the discontinuity. This would not change their probability of winning a prize, but would decrease their cost of bidding. This profitable deviation violates our assumption of equilibrium.
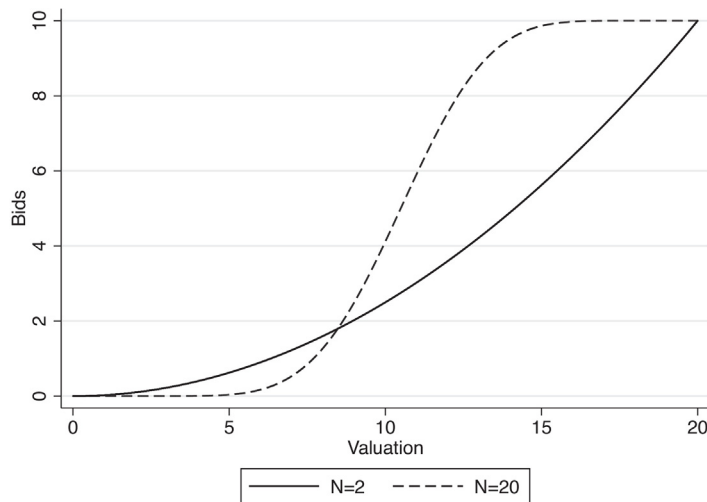
Fig. 1. Equilibrium bidding functions.

We impose the initial condition $B(0) = 0$ to solve this for given values for N and M.[10]

4. Experimental model

Before deriving predictions for this environment, it is important to address the selection of our experimental parameters. We believe our parameters accomplish a careful balance of salience for subjects, policy applicability, and ease of analysis and interpretation. To maximize salience, we chose a simple, linear prize structure with half of the participants receiving prizes. These choices may not perfectly reflect the classroom environment, but they capture the qualitative features of grading on the curve while allowing subjects to easily understand the differences in treatments.[11] We also chose nearly the largest change in the group size allowed by the size of the laboratory.[12] This served to make changes in the treatment clear and provide a strong test of the theory by increasing the differences between the equilibrium effort functions. By identifying a sensitivity to group size, our results can make clear contributions to policy. Even if the sensitivity to group size is magnified by the stylized and artificially salient environment, we can confirm that predictable responses to changes in group size exist and should be accounted for. Finally, our parameter choices allow for a simplified analysis of our environment and interpretation of our data. This is particularly important for our subjects attempting to best respond to their environment.

Imposing our experimental parameters of $(N, P) \in \{(2, 1), (20, 10)\}$ and $v_i \sim U[0, 1]$ on Eq. (5), we get[13]

$$B_{2,1}(v_i) = \frac{v_i^2}{2},$$ (6)

and

$$B_{20,10}(v_i) = 83980v_i^{11} - 692835v_i^{12} + 2558160v_i^{13} - 5542680v_i^{14} + 7759752v_i^{15} - \frac{14549535}{2}v_i^{16}$$

$$+ 4564560v_i^{17} - 1847560v_i^{18} + 437580v_i^{19} - 46189v_i^{20}.$$ (7)

Our subjects draw valuations uniformly from the set $\{\$0.01, \$0.02, \ldots, \$20.00\}$. Thus, we rescale the bidding functions to generate experimental predictions. Fig. 1 plots the equilibrium bidding functions for the experimental parameters.

To understand the intuition behind the shapes of the bidding functions, consider the limit case for an all-pay auction with $P = N/2$ as $N \to \infty$. The infinitely large sample eliminates uncertainty about the distribution of opponents' valuations, and the bidding function approaches a step function. Here, only the top half of bidders receive prizes, so any bidder below the median (i.e. $v_i < \$10$) bids $b_i = \$0$, while bidders above the median (i.e. $v_i > \$10$) bid $\tilde{b}$, a bid that makes the marginal

---

[10] This initial condition is justified through dominance. Bids of $b_i < 0$ are not possible, and, if $v_i = 0$, bids of $b_i > 0$ guarantee negative profits.

[11] If, for example, we had used a non-linear prize structure, subjects would have required calculators to make their decisions.

[12] The UCSD EconLab can hold 24 subjects.

[13] Eqs. (6) and (7) allow us to check the second-order conditions for utility maximization computationally. We can see that the derivative of the expected utility function is positive for any possible alternate bid, $\underline{b} < B(v_i)$, and negative for any alternate bid, $\tilde{b} > B(v_i)$. Thus, our bidding function maximizes the expected utility function. For a general proof of this, see Moldovanu and Sela (2001).
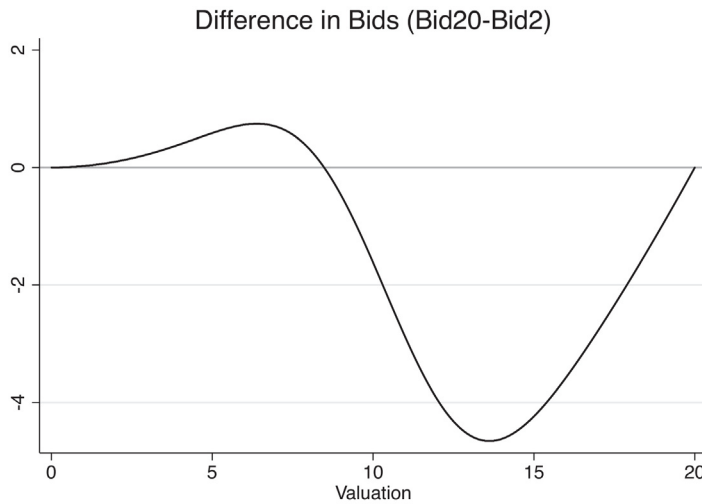
Fig. 2. Predicted difference in equilibrium bidding.

bidder indifferent between winning the auction with a bid of $\tilde{b}$ and losing the auction with a bid of zero. [14] In the case of $P = N/2$, the marginal bidder has $v_i = \$10$ and the bid that makes the marginal bidder indifferent is $\tilde{b} = \$10$.

Extending this intuition to our experiment, we can see how a change in the size of the auction affects the equilibrium bidding functions. The large auction in our experiment is ten times the size of the small auction, drawing its distribution of valuations closer to the underlying probability distribution than the small auction. For this reason, equilibrium bids in the 20-Person Auction more closely reflect bidding in the limit case.

Now, consider a bidder with $v_i \geq \$10$ choosing how far below the limit-case bid of $b_i = \$10$ he wants to bid in each auction. The marginal *benefit* of lowering a bid is constant across auctions. Bids are paid with certainty, so foregone bids are recovered with certainty. The marginal *cost* of lowering a bid—paid by reducing the probability of winning a prize—varies across auctions. This decrease in probability is greater in the large auction than in the small auction, thus bids cast by high-valuation bidders are larger in the large auction.

Consider the same marginal analysis for a bidder with $v_i < \$10$ choosing how far above the limit-case bid of $b_i = \$0$ he wants to bid in each auction. The marginal cost of increasing bids is constant, while the marginal benefit of increasing bids is probabilistic. At sufficiently low valuations, the marginal increases in the probability of winning a prize are higher in the small auction, causing bids in the small auction to rise above bids in the large auction.[15]

Our experimental design collects bids from both auction sizes. Thus, we will test the model's ability to predict *differences* in bids across the two auctions. Eq. (8) begins to generate predictions about bidding differences by subtracting the predicted bidding functions. This difference is captured graphically in Fig. 2.

$$B_{2,1}(v_i) - B_{20,10}(v_i) = \frac{v_i^2}{2} - 83980v_i^{11} + 692835v_i^{12} - 2558160v_i^{13} + 5542680v_i^{14} - 7759752v_i^{15} + \frac{14549535}{2}v_i^{16}$$
$$- 4564560v_i^{17} + 1847560v_i^{18} - 437580v_i^{19} + 46189v_i^{20}. \tag{8}$$

There are two features of this predicted difference that we will specifically test. The detailed proofs of each prediction can be found in the appendix.

4.1. Prediction 1: aggregate bidding increases in the size of the auction

We begin by testing for differences in aggregate bidding. This is visible through the relative size of the positive and negative area under the difference curve. Specifically, the area under the predicted bidding function for the 2-person auction

---

[14] At equilibrium in the limit case, bidders have probability measure zero. Thus, the rules must change to reflect the nature of the contest. Now, prizes will be awarded to bidders in descending order of their bids until half of the bidders have received prizes. Otherwise, there are profitable deviations from equilibrium.

[15] We show in the appendix that (surprisingly) the equilibrium is not strongly affected by imposing risk-aversion on the bidders' utilities. This is because bidding is driven by changes in probability—which is unaffected by risk aversion, not changes in surplus—which is affected by risk aversion. See the appendix for a graph of the equilibrium bidding functions under common risk aversion parameters. We also include a graph of the equilibrium under a common joy-of-winning value in the appendix. Trivially, this added degree of freedom improves the model's fit. We do not include this specification in the discussion, because it provides identical predictions about the effect of changing the auction size.

is $\int_{v=0}^{20} B_{2,1}(v) = \$66.67$, while the area under the predicted bidding function for the 20-person auction is $\int_{v=0}^{20} B_{20,10}(v) = \$95.25$. Since we will have exactly 20 bidders each round drawing values uniformly from $v_i = \$0$ to $v_i = \$20$, the area under each curve represents the expected revenue from all bids cast in each auction size.[16] In absolute terms, we expect that the larger auctions capture $28.58 more revenue per period or 1.43 times as much revenue. The prediction that revenues increase as the size of the auction grows—holding fixed the proportion of winners—is proven in Moldovanu and Sela (2006).[17]

4.2. Prediction 2: predicted bids from the larger auction cross predicted bids from the smaller auction from below exactly once

The second feature of this predicted difference we will test is the single-crossing property where predicted bids in the 20-person auction meet those of the 2-person auction. We derive this prediction by setting Eq. (8) equal to zero and solving for v_i. At $v_i \approx 8.477$, bids in the two auctions are equal with $B_{2,1}(8.477) = 1.796 = B_{20,10}(v)$. We predict that bids in the larger auction cross from below. For $v_i < 8.477$, predicted bids in the 2-person auction exceed those of the 20-person auction. Conversely, for $v_i > 8.477$, predicted bids in the 20-person auction exceed those of the 2-person auction. Following the framework of Moldovanu and Sela (2006), we prove this single-crossing property for generic auction sizes in the appendix.

5. Experimental hypotheses

The model provides strong predictions about how bidders will behave in each auction. A strict test of the model's fit compares observed bidding *levels* to the Bayes–Nash equilibrium predictions shown in Fig. 1. Our paired auction design, which compares bidding in different auction sizes, opens up a novel second channel through which we can evaluate the model—its accuracy in predicting *changes* in bidding.

In addition to these specific tests, we consider three general predictions of the model. We state these general predictions in terms of a teacher's objectives and use them to highlight tensions between different objectives a mechanism designer may have. In the event that the equilibrium prediction holds, the general predictions will follow. But, even if the precise equilibrium prediction is rejected by statistical tests, tests of the more general effects of contest size will still be informative for the policy debate about the merits of different contest sizes. Thus, we test the equilibrium predictions and the qualitative predictions independently.

Given the complexity of the equilibrium bidding functions, strict hypotheses about equilibrium bidding levels are rather restrictive. Even with experience, the complexity of the bidding functions and the lack of clear heuristics makes a tight fit between the data and the theory unlikely. Our design allows us to consider a less restrictive test of the changes in bidding between auctions. These tests of equilibrium behavior will form a starting point for our hypotheses before we move on to the more general predictions identified below.

5.1. Qualitative predictions: the designer's objectives

The objectives we consider are designed around encouraging effort from students, since educators generally pursue classroom policies that increase rather than decrease effort. While auction bids accrue solely to the designer, it need not be the case that the designer alone benefits from effort. Our results hold implications for any environment where the rewards for effort are more salient—for example, psychologically or materially—to the designer than to the agents. For instance, exerting effort may feel like a wasteful action to students in principles of economics, but educators believe that there are benefits to this effort that are not yet clear to the students.

One typically thinks of an auction designer as having the unique objective of maximizing revenues but there may be other objectives competing for attention in the minds of policy-makers applying these insights to education. We consider three such potential objectives. While we will only develop clear predictions for two of them, our experiment cleanly tests all three.

5.1.1. Objective 1: maximize aggregate effort from students

Holding fixed the proportion of prizes, the model predicts that aggregate bidding increases in the size of the auction. These increases reduce the uncertainty about the distribution of valuations among other bidders. With tighter predictions about the distribution of valuations a bidder faces, high-valuation bidders face greater probabilistic costs associated with casting low bids. Increased bidding from high-valuation bidders more than offsets decreased bidding from low-valuation bidders.

---

[16] The expected revenue from any one 2-person auction is proportionally smaller: $6.67.

[17] As part of their proof of Proposition 2, they show that the expected bid in an isomorphic contest with *N* bidders and *p* winners is proportional to the order statistic of the $p + 1$-th highest ability bidder. For $p = N/2$, as $N \to \infty$, this order-statistic monotonically increases towards the median of the distribution of abilities.

5.1.2. Objective 2: differentially weight effort from low- and high-ability students

It is reasonable to expect, especially in the arena of education, that policy makers may not simply wish to maximize total effort, but might care from what type of students that effort comes. A policy may be designed to help the best students "realize their full potential" or help prevent the most challenged students from being "left behind."

As discussed under Objective 1, we predict that high-valuation bidders will cast higher bids in the larger auction. For low-valuation bidders, however, the decreased uncertainty about the distribution of bidders in the larger auction drives down bidding. Therefore, we expect to see higher bids in the small auction for low-valuation bidders. It follows that policies designed to differentially favor effort from low-ability students would promote smaller class sizes, while policies designed to favor effort from high-ability students would promote larger classes.

Fig. 1 plots the model's specific prediction that the bidding functions intersect at $v_i = \$8.48$. We test this prediction in two ways. First, we compare bids in the small auction to those from the large auction for bidders with valuations on either side of $v_i = \$8.48$. Second, we estimate the point where observed bids in the large auction begin to dominate bids in the small auction and compare it to the model's prediction of $v_i = \$8.48$.

5.1.3. Objective 3: accurately order student ability based on effort

A mechanism designer may want to use the bids cast to infer the ranking of valuations. For instance, a teacher may want to use her students' test scores to estimate the relative intelligence of her students. Deriving predictions in this domain from a symmetric equilibrium model provides little traction because symmetric bidding functions produce a one-to-one, monotonic matching of bids to valuations. As a result, relative bidding perfectly reflects relative valuations.

Since the ability to rank valuations based on bids holds important policy-relevance and our design provides a clean test of this relative ability, we will explore this objective empirically. In particular, we will empirically explore the role of changes in the size of the auction on (1) the overall ordering of bids, (2) the ordering of bids within low-valuation ($v_i < \$10$) and high-valuation ($v_i > \$10$) bidders, and (3) the ordering of bids between low-valuation and high-valuation bidders—comparisons where valuations are on opposite sides of $v_i = \$10$.

6. Experimental procedures

To test the sensitivity of effort to a change in the size of a contest, we employ an independent private value all-pay auction with multiple prizes. We use a paired bidding design similar to Kagel and Levin (1993) and Andreoni et al. (2007) in which we elicit two bids from each bidder in each round—one bid for the large auction and one for the small auction. This is a powerful design that allows us to pair data perfectly and test our hypotheses about the difference in bidding within-subject while holding valuation constant.

6.1. Recruitment and participation

60 undergraduate students from the University of California, San Diego (UCSD) participated in our experiment. The experiment was conducted in 3 sessions in February of 2013 in the EconLab at UCSD. Each session required exactly 20 subjects and lasted approximately 90 min. All subjects received a $20 participation payment. Subjects took home between $15 and $45 after adding any gains or deducting any losses.

6.2. Instructions

We first read the instructions aloud to all subjects and then administered a quiz to test their comprehension of the auction formats and the payoffs from different outcomes.[18]

Our instructions made clear the means by which we awarded prizes and the payoffs conditional on receiving or not receiving a prize. We carried out the entire experiment on computers using the Z-Tree Economics Software (Fischbacher, 2007). Key instructions remained posted at the top of every decision screen. This feature helped us remind subjects that bids would be deducted from their winnings regardless of the outcome of the auction, that their final payment for the experiment would be determined by one auction selected at random, and that they would see two auctions for each valuation they drew—the 2-Person Auction, then the 20-Person Auction. For all auctions, subjects were reminded of the number of opponents they faced and the number of prizes available.

6.3. Auction design

Every round, subjects received valuations drawn uniformly in $0.01 increments between $.01 and $20.00. That is, $v_i \sim U[\$.01, \$20.00]$. Using this valuation, subjects participated in both a "2-Person Auction," for which they were randomly and anonymously paired, and a "20-Person Auction," which consisted of all participants in the session.[19]

---

[18] The instructions and quiz can be found in the appendix.

[19] This structure gives us the largest change from one size of contest to the other. Since the 20-Person Auction features a stable cohort, issues with collusion could arise. If this were the case, bids in the 20-Person Auction would decrease with experience. We find that bids do drop in later rounds, though bids drop
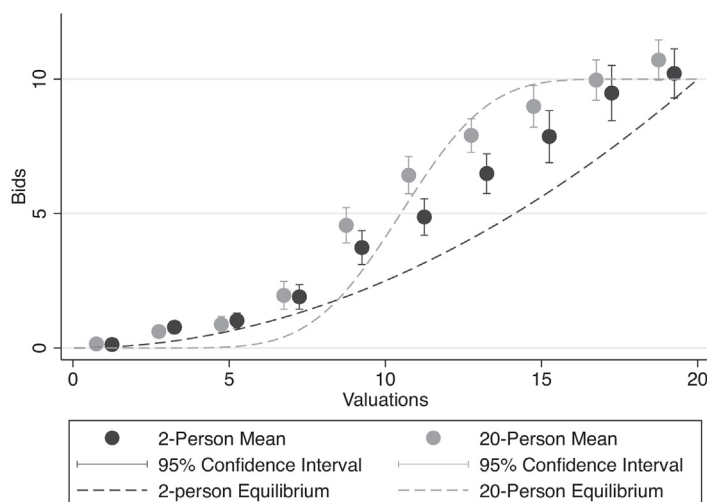
Fig. 3. Mean bidding across valuations.

Subjects cast bids for both auctions in every round. After all subjects cast bids, we revealed all winning and losing bids cast by all bidders in all auctions that round. Each subject was also informed of his own payoffs from both auctions that round. All bids were cast without knowledge of opponents' bids or valuations, so we consider them simultaneous bids under incomplete information. No communication was allowed, and we never revealed the valuations associated with any bids by other subjects.

Subjects then drew new valuations and new partners for the 2-Person Auction before repeating the auction round. For each new valuation they drew, they cast bids in the 2- and 20-Person Auctions. Subjects in session 1 repeated the auction rounds 20 times, while subjects in sessions 2 and 3 repeated them 15 times. To be conservative and uniform in our analysis, we restrict our sample to the first 15 rounds of each session.[20] At the end of the experiment, one auction in one round was selected at random to determine each subject's payments.

After completing the auction rounds, we used a discretization of the method developed in Andreoni and Harbaugh (2016) to elicit subjects' risk preferences.[21] We will use our estimates of risk preference and demographic data as controls in our analysis.

7. Results

Subjects drew valuations and cast bids in two different auctions for 15 rounds. From this, we collected a total of 1800 different bids for 900 valuations. 12 of these bids were strictly dominated—the bid exceeding the valuation of the bidder—and guaranteed the bidder a negative payoff.[22] Dominated bids were evenly distributed across periods and auction sizes. We omit only one of these dominated bids, a $99 bid in the final 20-Person Auction of session 3. This bidder had a valuation of $0.38, guaranteeing a loss of -$98.62 from this auction. Since we cannot fathom that this choice represents true preferences, we omit it in order to improve the accuracy of our point estimates, although none of our significant results are rendered insignificant by its inclusion.

Fig. 3 compares the aggregate patterns of bidding across different valuations to their predictions. We split the data by valuation into 10 equally-sized bins. Adjacent dots show the mean and 95% confidence intervals of the bids in each auction size. Fig. 3 shows that all bids are near-zero for low valuations regardless of auction size. The larger auction begins to dominate over middle-valuations, as predicted. For high valuations, bids are close to the equilibrium predictions with little difference between the two auction sizes.[23]

---

more in the smaller auction than the larger auction, contradicting the collusion hypothesis. To the extent that collusion remains a possibility, it suggests bidding should *decrease* in the larger auction, which biases us away from our hypotheses.

[20] See the appendix for the analysis including all 20 rounds of session 1. Our results increase in magnitude and significance. Our initial experimental design called for 20 rounds in 90 minutes. During the first session we discovered that this time allowed for only 15 rounds, which is how we completed the experiment.

[21] This elicitation can be found in the appendix.

[22] Overbidding in this way has been documented in several studies using common-value all-pay auctions (Gneezy and Smorodinsky, 2006, for example), but does not represent a large fraction of bids in independent private value all-pay auctions (Noussair and Silver, 2006).

[23] For a finer resolution display of subject bidding, we separate individual bids by auction size and session and plot them against the bidder's valuation. We include these graphs in the appendix rather than here, since the volume of bids—300 per graph—makes the graphs difficult to interpret parsimoniously.
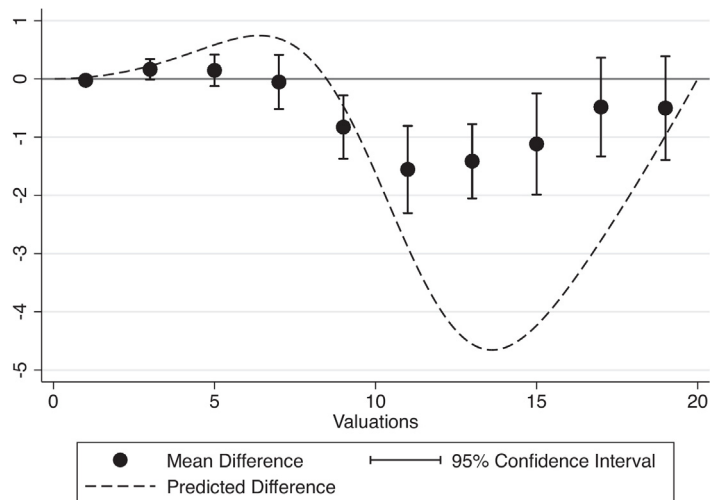
Fig. 4. Bidding differences: 2-person auction bid minus 20-person auction bid.

To leverage our paired design, we evaluate the experimental hypotheses based on their predictions about the difference between the 2- and 20-Person Auction bids at a given valuation. Fig. 4 compares the observed differences in bids to the model's predicted difference. The observed differences are averaged within the same 10 equally-sized bins as before.

7.1. Nash equilibrium predictions

We begin with our strict test of equilibrium bidding.

Statistical tests reveal that there are systematic deviations between individual bids and their equilibrium predictions both in the 2-Person Auction ($t = 7.00$, $p < 0.001$, $n = 900$) and the 20-Person Auction ($t = 4.73$, $p < 0.001$, $n = 893$). Comparing differences in bidding with the difference in the equilibrium predictions yields a similar rejection of the theoretical predictions ($t = 4.87$, $p < 0.001$, $n = 899$).[24]

Of course, one should not be surprised that in this unique environment that a model with complex and precise predictions fails these statistical tests. As the literature on auction experiments reveals, the theory often provides a good benchmark but seldom predicts precise outcomes. Indeed, Fig. 4 makes it clear that the differences between bids in the 2- and 20-Person Auctions evolve in a way that is qualitatively similar to the model's predictions. Therefore, the question for the following objectives becomes whether, despite the lack of a precise fit, the qualitative predictions and intuitions gleaned from theory are useful in predicting behavior.

7.2. Objective 1: Maximize aggregate bidding

In 34 of the 45 rounds, revenues were greater in the larger auction, consistent with our theory. Using each round as an observation, we perform a paired t-test on the difference in total revenue between auctions and find that the larger auctions generate significantly more revenue ($t = 4.92$, $p < .001$, $n = 45$). Our paired design allows us to test this hypothesis at the individual level and perform a paired t-test matching each bid in the large auction with its counterpart in the small auction. We find a mean increase in bidding of \$0.56 when moving from the small to the large auction ($t = 2.85$, $p = .006$, $n = 899$). While this difference is significantly different from zero, it is also significantly smaller than the predicted average difference of \$1.43 ($t = 4.87$, $p \approx .000$, $n = 899$).

We also derived predictions about the proportional change in bidding across auction sizes. The model predicts that the average bid in the 20-person auction is 1.43 times as large as the average bid in the 2-person auction. We observe that the average bid in the 20-person auction is only 1.11 times as large as its counterpart in the 2-person auction.

---

[24] These and all other within-subject tests use standard errors clustered at the subject level. Session level clusters are not advised. With only 3 clusters, we introduce the potential for bias (Bertrand et al., 2004). While there may be cause to believe that session characteristics influence bidding *levels*, our unit of analysis is within-subject, thus clustering at the session level is only justified if there is reason to believe that session characteristics somehow drive bidding *differences* within an individual bidder.
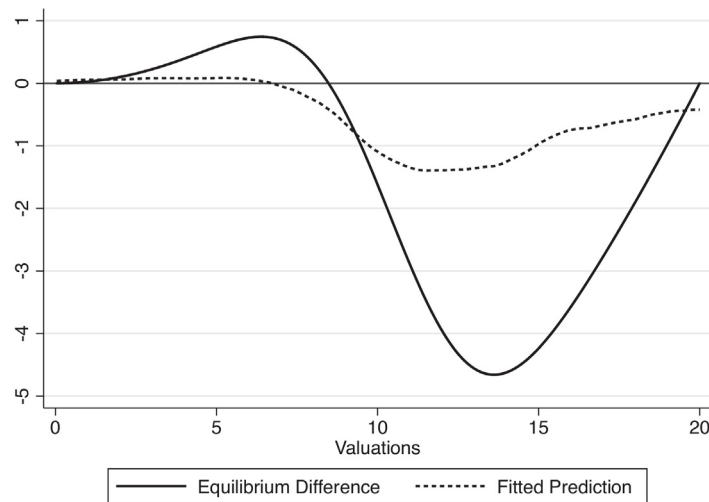
Fig. 5. Predicted difference in bids: 2-person auction bid minus 20-person auction bid.

7.3. Objective 2: Differentially weight effort from low- and high-ability students

The model predicts that bids in the small auction will dominate bids in the large auction for low valuations, while bids in the large auction will dominate for high valuations. Our specific theoretical prediction places this crossover at $v_i = \$8.48$.

Fig. 4 confirms this prediction of relative dominance for high-valuation bidders, but not low-valuation bidders. The behavior of low-valuation bidders appears to be invariant to auction size, while high-valuation bidders seem to show sensitivity to the auction size.

We test this prediction in two ways. First, we estimate a fitted crossover point where bids in the large auction begin to dominate bids in the small auction. To estimate the valuation at which bids in the large auction begin to dominate those of the small auction, we use a locally-linear polynomial smoothing function to fit the difference in bids across valuations. We plot this function in Fig. 5, which shows the fitted crossover point at $v_i \approx \$6.71$. While this diverges from the precise theoretical prediction about the location of the single-crossing point, it confirms the single-crossing nature of the bidding functions with bids in the small auction dominating at first, before bids in the large auction begin to dominate.

The second way we test predictions about relative dominance is by comparing relative bidding levels on either side of the theoretical and fitted crossover points. We use the theoretical crossover point of $v_i = \$8.48$ to split our data into two bins, *LowerBin* and *UpperBin*, and test our theory's predictions about relative bidding in each bin. To best exploit our paired data, our dependent variable will be the bid in the small auction minus the bid in the large auction. Our theory predicts that the difference will be positive until $v_i = \$8.48$ and negative afterwards. In column 1, we regress this difference onto *LowerBin* and *UpperBin* with standard errors clustered at the subject level. In column 2, we repeat this regression with controls for

Table 1
The difference in bids (2-person auction bid – 20-person auction bid).

Periods:	Predicted cutoff: $8.48			Fitted cutoff: $6.71		
	1–15	1–15	9–15	1–15	1–15	9–15
LowerBin ($V_i \leq \$8.48$)	0.0816 (0.10)	0.0857 (0.09)	−0.0920 (0.14)			
UpperBin ($V_i > \$8.48$)	−0.997[***] (0.31)	−0.957[***] (0.29)	−1.151[***] (0.40)			
LowerFitBin ($V_i \leq \$6.71$)				0.0598 (0.08)	0.0576 (0.08)	−0.0516 (0.09)
UpperFitBin ($V_i > \$6.71$)				−0.849[***] (0.28)	−0.835[***] (0.26)	−1.081[***] (0.37)
R.A. controls	No	Yes	Yes	No	Yes	Yes
N	899	899	419	899	899	419

Standard errors clustered by subject.
See the appendix for analysis including the $99 bid. No significant results are affected.
[*] $p < 0.10$.
[**] $p < 0.05$.
[***] $p < 0.01$.

Table 2
Measuring the well-ordering of bids.

		Kendall–Tau coefficient of well-ordering		
	Obs	2-Person auction	20-Person auction	Difference
Pooled[a]	900	0.546 (0.014)	0.585 (0.012)	−0.040** (p = 0.11)
Pooled[a,b]	899	0.546 (0.014)	0.589 (0.012)	−0.043*** (p = 0.004)
Per period[b]	45	0.566 (0.014)	0.612 (0.014)	−0.042*** (p = 0.005)

[a] Standard errors clustered at the period level.
[b] $99 bid omitted.
*p < 0.10.
** p < 0.05.
*** p < 0.01.

the elicited risk preferences.[25] The results can be seen in the left-hand side of Table 1. In column 3, we restrict the sample to experienced bidders.

We repeat this analysis with the fitted crossover point, generating new dummy variables, *LowerFitBin* and *UpperFitBin*, for valuations below the fitted crossover point, $v_i = \$6.71$. The results are found on the right-hand side of Table 1.

The difference in bids in the upper bin is large in magnitude and highly statistically significant. This difference in bids is also statistically distinguishable from the same difference in the lower bin. However, the difference in bids in the lower bin is not statistically distinguishable from zero. Recall that the predicted difference in bids is small over lower valuations, so it is not surprising that the observed differences are not statistically significant. For valuations just above the median, however, we predict large differences between bids cast in the large and small auction. This separation is clearly displayed in statistical tests and graphical analysis. Experience appears to favor the large auction, with the difference in bids between the 2- and 20-Person Auctions growing in the later rounds.

7.4. Objective 3: Accurately order bidders based on their bids

We test which auction best accomplishes Objective 3 with Kendall's (1938) Tau rank correlation coefficient (Hereafter, K-T). The K-T coefficient measures the ratio of concordant pairs to discordant pairs. A pair of bids is concordant if the ranking of the two bids is consistent with the ranking of the valuations of the bidders ($b_i < b_j \Leftrightarrow v_i < v_j$), while a pair of bids is discordant if it is inconsistent with the rankings of the bidders' valuations ($b_i > b_j \Leftrightarrow v_i < v_j$).[26]

The K-T coefficients and their significance tests for the pooled data are reported in the first two rows of Table 2. In row 3, we generate one K-T coefficient for each auction in each period before running a paired *t*-test of the two auctions from a given period.[27]

Our results show that neither auction is able to generate a perfect ordering of valuations based on bids. Observing a concordant pair of bids is 55% more likely than observing a discordant pair of bids in the small auction and 59% more likely in the large auction. This difference in relative likelihood is statistically significant and economically meaningful, as it leads to 365 more discordant pairs of bids.

Digging deeper into the K-T coefficients reveals some of the mechanisms that drive bids out of their well-ordered ranking. For bidders with similar valuations, Fig. 1 shows that bids in the 2-Person Auction have a more gradual and consistent differences. On the other hand, expected bids from bidders with very distinct valuations show much starker differences in the 20-Person Auction.

We divide bidders based on their valuations into ex-ante expected winners and losers. Those with valuations above the median should expect to win ex-ante, and those below should expect to lose ex-ante. Rows 1 and 3 of Table 3 report the K-T coefficients within expected losers and winners, respectively. These within-group coefficients differ meaningfully from the pooled coefficients in Table 2, and are relatively more favorable towards the 2-Person Auction. Row 2 of Table 3 compares sorting between expected winners and losers while maintaining the same interval length. This between-group sorting strongly favors the 20-Person Auction, indicating that the stark probabilistic distinction between expected winners and losers in the 20-Person Auction may have motivated a stronger separation in bids.

---

[25] Risk aversion does not have a clear prediction, since it affects high-valuation and low-valuation bidders differently. Thus, we interact these controls with the bins for different valuations.

[26] Mathematically, the K-T coefficient is $K = \frac{C-D}{((N \times (N-1))/2)}$, where N is the total number of bids cast. We can also correct for ties in bids or valuations according to Kendall (1970). The signs and significance of our results are maintained or increased with this correction. We chose Kendall's Tau over a possibly more recognizable test of rank-correlation, the Spearman (1904) coefficient, because Spearman's coefficient is more heavily influenced by outliers, biasing our results in the direction of the large auction, and because Kendall's Tau has a more straightforward interpretation. Our results have the same signs and larger magnitudes using Spearman's coefficient.

[27] Given that we have 45 pairs of K-T statistics, we invoke the Central Limit Theorem and perform a standard *t*-test on the difference between the pairs.

Table 3
Kendall's Tau coefficient for ordering among expected winners and losers.

	Obs	2-Person auction	20-Person auction	Difference
Valuation$_i$ < 10	422	0.365	0.289	0.076***
		(0.028)	(0.028)	(p = 0.006)
5 < Valuation$_i$ < 15	477	0.401	0.479	-0.078***
		(0.030)	(0.021)	(p = 0.004)
10 < Valuation$_i$	477	0.333	0.342	−0.008
		(0.031)	(0.027)	(p = 0.805)

$99 bid omitted and standard errors clustered at the period level.
* p < 0.10.
** p < 0.05.
*** p < 0.01.

The two observed sorting mechanisms follow naturally from the model, which predicts that the large auction should see pooling among high- and low-valuation bidders but strong separation between the two groups, while the small auction should better separate bidders within those same groups. Interpreting Table 2 in light of the results from Table 3 indicates that the superior sorting *between* expected outcomes in the 20-Person Auction dominates the diminished sorting *within* expected outcomes, resulting in a better-ordered set of bids.

8. Discussion

Our results show that subjects are sensitive to changes in the number of competitors even when the proportion of winners remains constant. Moreover, the model provides useful benchmark predictions about the way subjects will adjust their behavior to the contest size.

Before discussing policy recommendations, it is important to note that our analysis explores the effect of contest size while holding all other inputs fixed. We abstract away from systematic correlations between the quality of inputs and the class size. Of course, in any discussion of policy these correlations will need to be weighed against our results about the strategic effects of class size on effort. This paper only identifies the potential for a strategic effect of class size and leaves it to further research to weave these strategic effects into the broader discussion of the effects of class size on performance. It is also important to note that we explore effort responses to large changes in the contest environment, thus our results should be adjusted to account for the actual change in class size being considered.

Looking only at these strategic effects, for classroom designers concerned with maximizing aggregate effort, we recommend larger class sizes, as our larger contests yielded greater aggregate bidding. Thus, we expect that students in classes graded on the curve will increase effort as the class size grows. In smaller contests, we find that the greater uncertainty about the effort required for a given grade causes participants to strategically lower their effort in an attempt to exploit the randomness of the environment. Larger classes hold less uncertainty about the effort required for a given grade, and we expect students to respond with increases in effort just large enough to meet the effort required for their desired grade.

A designer may also have preferences over the distribution of effort across students with different abilities. We found that this distribution of effort is strongly affected by the size of the contest. From our results, we expect that high ability students will exert greater effort in larger classes, while students with lower-middle abilities may exert slightly more effort in smaller classes. Our results also suggest that the lowest ability students will not be sensitive to changes in the size of the strategic environment.

Larger classes may also have an advantage with respect to the ability to sort students by their classroom performance. Since student abilities are unobservable, classroom performance must serve as a sorting mechanism that distinguishes relative ability. We discovered that the ability of a relative grading scheme to accurately sort students may be systematically undermined in courses with low enrollment, creating an environment where a student's ability is less correlated to her relative performance.

While this paper analyzes the motivating effects of class size when graded in a binary manner (e.g. "pass" or "fail"), the majority of courses have finer grade differentiation (e.g. A,B,C,D,F). These multi-tiered environments would simply break the larger contest into several smaller contests between grade levels. As such, we do not believe this would affect the qualitative predictions of the model. A student at the 90th percentile will compete with other high-ability students for A's in the same way that median students in our experiment compete for "pass" grades. With a finer grading scheme, low enrollment courses have an even larger effect on the uncertainty each student faces with regards to their draw of classmates and their relative ability. Our results indicate that this uncertainty may cause effort to diminish and further weaken the correlation between ability and performance. On the other hand, with large enough enrollment, we still expect students to just match the effort required for a given grade. For example, the top student in a class of 500 may pool with the other high-ability students just above the cutoff for A's.[28] The ordering of students under this finer grading scheme will depend more heavily on sorting

---

[28] For a discussion of the optimal fineness of a grading scheme, see Dubey and Geanakoplos (2010).

across expected outcomes, giving a further advantage to large courses. This is largely speculative, however, and needs to be verified by further studies.

9. Conclusions

We designed and executed an experiment to examine how the size of a contest affects the effort choices by contestants with different abilities. We used an all-pay auction format to measure effort in auctions with different sizes. We used a paired auction design to abstract away from the exacting question of equilibrium bidding *levels* and focus on how bidding *changes* with auction size. We broke these changes down by ability level and showed the model's ability to qualitatively predict how bidding differences would evolve with ability. Our results hold economically and statistically significant implications for a broad category of mechanism design.

The theoretical bidding functions we developed correctly predicted that larger auctions elicit greater aggregate bidding as well as higher bids from high-ability bidders. The model draws attention to the uncertainty that arises as the size of an auction decreases. This uncertainty not only drives down bids from high-ability bidders, but also limits the ability of the designer to accurately order bidders' valuations based on their bids.

As policy makers, administrators, and other mechanism designers continue to employ relative awarding schemes, our research sheds light on the effect that the number of participants has on the effort exerted in the mechanism. While this study considered only two sizes in a stylized contest, we believe it serves as a strong foundation for further research on the independent effect of a contest's size on the performance of its participants.

Funding

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.jebo.2017.03.017.

References

Amann, E., Leininger, W., 1996. Asymmetric all-pay auctions with incomplete information: the two-player case. Games Econ. Behav. 14, 1–18.
Andreoni, J., Che, Y.K., Kim, J., 2007. Asymmetric information about rivals' types in standard auctions: an experiment. Games Econ. Behav. 59 (2), 240–259.
Andreoni, J., Harbaugh, W., 2016. Unexpected Utility: Experimental Tests of Five Key Questions about Preferences over Risk. Mimeo.
Athey, S., 2001. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. Econometrica 69 (4), 861–889.
Baik, K.H., 1994. Effort levels in contests with two asymmetric players. South. Econ. J., 367–378.
Barut, Y., Kovenock, D., Noussair, C., 2002. A comparison of multiple-unit all-pay and winner-pay auctions under incomplete information. Int. Econ. Rev. 43, 675–707.
Baye, M., Kovenock, D., de Vries, C.G., 1996. The all-pay auction with complete information. Econ. Theory 8, 291–305.
Baye, M., Kovenock, D., de Vries, C.G., 1993. Rigging the lobbying process: an application of the all-pay auction. Am. Econ. Rev. 83, 289–294.
Becker, W.E., Rosen, S., 1992. The learning effect of assessment and evaluation in high school. Econ. Educ. Rev. 11 (2), 107–118.
Bertrand, M., Duflo, E., Mullainathan, S., 2004. How much should we trust differences-in-differences estimates? Q. J. Econ. 119 (1), 249–275.
Davis, D., Reilly, R., 1998. Do many cooks always spoil the stew? An experimental analysis of rent seeking and the role of a strategic buyer. Public Choice 95, 89–115.
Dechenaux, Kovenock, Sheremeta, 2012. A Survey of Experimental Research on Contests, All-Pay Auctions and Tournaments. mimeo.
Dubey, P., Geanakoplos, J., 2010. Grading Exams: 100, 99, 98,. . . or A, B, C? Games Econ. Behav. 69, 72–94.
Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10, 171–178.
Gneezy, U., Smorodinsky, R., 2006. All-pay auctions—an experimental study. J. Econ. Behav. Org. 61, 255–275.
Harbring, C., Irlenbusch, B., 2005. Incentives in tournaments with endogenous prize selection. J. Inst. Theor. Econ. 127, 636–663.
Hillman, A., Riley, J.G., 1989. Politically contestable rents and transfers. Econ. Polit. 1, 17–40.
Kagel, J.H., Levin, D., 1993. Independent private value auctions: bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. Econ. J., 868–879.
Kendall, M.G., 1938. A new measure of rank correlation. Biometrika 30, 81–93.
Kendall, M.G., 1970. Rank Correlation Methods, fourth ed. Charles Griffin & Co., London.
Krishna, V., Morgan, J., 1997. An analysis of the war of attrition and the all-pay auction. J. Econ. Theory 72, 343–362.
Krueger, A.O., 1974. The political economy of the rent-seeking society. Am. Econ. Rev. 64, 291–303.
Lazear, E.P., Rosen, S., 1981. Rank-order tournaments as optimum labor contracts. J. Polit. Econ. 89, 841–864.
Lim, W., Matros, A., Turocy, T., 2014. Bounded rationality and group size in Tullock contests: experimental evidence. J. Econ. Behav. Organ. 99, 155–167.
Minor, D., 2012. Coarse Strategy and Competition. mimeo.
Moldovanu, B., Sela, A., 2001. The optimal allocation of prizes in contests. Am. Econ. Rev. 91, 542–558.
Moldovanu, B., Sela, A., 2006. Contest architecture. J. Econ. Theory 126 (1), 70–96.
Müller, W., Schotter, A., 2010. Workaholics and dropouts in organizations. J. Eur. Econ. Assoc. 8, 717–743.
Noussair, C., Silver, J., 2006. Behavior in all pay auctions with incomplete information. Games Econ. Behav. 55, 189–206.
Olszewski, W., Siegel, R., 2013. Large Contests. Mimeo.
Orrison, A., Schotter, A., Weigelt, K., 2004. Multiperson tournaments: an experimental examination. Manag. Sci., 1–10.
Parreiras, S.O., Rubinchik, A., 2010. Contests with three or more heterogeneous agents. Games Econ. Behav. 68 (2), 703–715.
Potters, J.C., de Vries, C.G., van Winden, F., 1998. An experimental examination of rational rent seeking. Eur. J. Polit. Econ. 14, 783–800.
Spearman, C., 1904. The proof and measurement of correlation between two things. Am. J. Psychol. 15, 72–101.
Tullock, G., 1967. The welfare costs of tariffs, monopolies, and theft. West. Econ. J. 5, 224–232.