Journal of Economic Behavior & Organization 137 (2017) 361–373

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

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

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

ARTICLE INFO

Article history:
Received 24 April 2016
Received in revised form 17 March 2017
Accepted 20 March 2017
Available online 31 March 2017

Keywords:
All-pay auctions
Contests
Group size
Education
Grading

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.

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).

http://dx.doi.org/10.1016/j.jebo.2017.03.017

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.