NBER WORKING PAPER SERIES

PARETO IMPROVEMENTS IN THE CONTEST FOR COLLEGE ADMISSIONS

Kala Krishna
Sergey Lychagin
Wojciech Olszewski
Ron Siegel
Chloe Tergiman

Working Paper 30220
http://www.nber.org/papers/w30220

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2022

We are indebted to Bela Szabadi for outstanding research assistantship. We thank Dawei Fang, Alex Frankel, Navin Kartik, Emir Kamenica, Shengwu Li, Jingfeng Lu, Nick Netzer, Philipp Strack, Jun Xiao, and seminar participants at Bocconi, Cambridge, Carlo Alberto, Cincinnati, European University Institute, Hong Kong, Johns Hopkins, Kyoto, Leicester, Northwestern, NUS, Oxford, Penn State, SKKU, Virginia, Wisconsin, Yonsei, the AEA Winter Meetings in Chicago, the Bern Workshop on Contest Theory, the Global Seminar on Contests and Conflict, the Pennsylvania Economic Theory Conference, the SAET Annual Conference in Taipei, and the Southampton WinterWorkshop in Economic Theory for very helpful comments and suggestions. Olszewski and Siegel gratefully acknowledge financial support from the NSF (grant SES-1325968). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Pareto Improvements in the Contest for College Admissions
Kala Krishna, Sergey Lychagin, Wojciech Olszewski, Ron Siegel, and Chloe Tergiman
NBER Working Paper No. 30220
July 2022
JEL No. D02,I23,L5

ABSTRACT

College admissions in many countries are based on a centrally administered test. Applicants invest a great deal of resources to improve their performance on the test, and there is growing concern about the large costs associated with these activities. We consider modifying such tests by introducing performance-disclosure policies that pool intervals of performance rankings, and investigate how such policies can improve students' welfare in a Pareto sense. Pooling affects the equilibrium allocation of studentso colleges, which hurts some students and benefits others, but also affects the effort students exert. We characterize the Pareto frontier of Pareto improving policies, and also identify improvements that are robust to the distribution of college seats.

We illustrate the potential applicability of our results with an empirical estimation that uses data on college admissions in Turkey. We find that a policy that pools a large fraction of the lowest performing students leads to a Pareto improvement in a contest based on the estimated parameters. We then conduct a laboratory experiment based on the estimated parameters to examine the effect of such pooling on subjects' behavior. The findings generally support our theoretical predictions. Our work suggests that identifying and introducing Pareto improving performance-disclosure policies may be a feasible and practical way to improve college admissions based on centralized tests.

Kala Krishna
Department of Economics
523 Kern Graduate Building
The Pennsylvania State University
University Park, PA 16802
and NBER
kmk4@psu.edu

Sergey Lychagin
Department of Economics
Central European University
Nador u. 9
Budapest 1051
Hungary
lychagin@gmail.com

Wojciech Olszewski
2001 Sheridan Road
Evanston, IL 60208
wo@northwestern.edu

Ron Siegel
The Pennsylvania State University
412 Kern Graduate Building
University Park, PA 16801
ron.siegel@gmail.com

Chloe Tergiman
cjt16@psu.edu

1 Introduction

College and university admissions in many countries are determined by students' performance on a centrally-administered test. This is the case, for example, for most colleges and universities in Brazil, China, Russia, South Korea, and Turkey. The students with the highest performance are admitted to the best colleges, those ranked slightly below them are admitted to the next best colleges, and so on. In many other countries factors such as high-school grades are also considered, but centralized test results are still of paramount importance in the college admissions process.

Consequently, students invest a great deal of time and effort preparing for these tests. In many Asian countries, including China, Japan, South Korea, and Taiwan, students attend specialized "cram schools,"[1] which focus on improving students' performance on the tests. This often consists of rote learning, solving a large number of practice problems, and practicing test-taking strategies tailored to the specific test. In other countries, students hire tutors, buy books, and take specialized courses, all geared entirely toward improving their test scores. These activities likely improve students' performance on the test, but are far less likely to generate substantial long-term improvements in students' productive human capital. These activities do, however, carry significant costs in terms of time, money, and effort. In South Korea, for example, it is not uncommon for high school students to spend several hours a day in cram schools, and the high stakes competition for college admissions is seen as one of the main causes for the high rates of unhappiness and suicide among teenagers.[2] Similar concerns have also been raised in the United States.[3]

Addressing this issue is more difficult than it might initially appear. Passing laws to pro-

[1]The term comes from the word "cramming," which also attests to the style and content of the instruction.

[2]See, for example, Matthew Carneys' discussion "South Korean education success has its costs in unhappiness and suicide rates" from June 15, 2015 on the Australian Broadcasting Corporation. The high suicide rate among teenagers is frequently attributed to their and others' expectations for them to do well in the competition for college admissions (Ahn and Baek (2013)).

[3]For example, Hsin and Xie (2014) report that the high academic effort Asian-Americans exert leads to lower subjective well being and to psychological and social difficulties. Schwartz (2016) discusses (without modeling formally) the psychological and other costs to students, and suggests that a lottery among those passing a threshold might increase welfare relative to the current admissions process. Bodoh-Creed and Hickman (2019) provide additional references, and also estimate the cost of competition in monetary terms.

hibit or limit test-preparation activities may be both difficult and ineffective.[4] Changing the admissions process may also be impractical. First, it is not clear what a better system would look like. For example, accurate tests lead to better students being admitted to better colleges, and other systems may lead to different outcomes, which may or may not be preferred. Second, implementing a new system may be expensive and technically difficult. Third, a new system that makes some students worse off would likely face significant resistance, even if it made other students better off.

This paper investigates simple modifications to admissions processes based on centralized tests that can make all students better off. We model college admissions as a contest with many players (students) and many prizes (college seats). Students exert costly effort and are admitted to colleges based on the rank order of their performance.[5] We consider performance-disclosure policies, which coarsen students' rank order by pooling together intervals of performance and assigning the same score to all performances in an interval. If many students obtain the same score, they are randomly admitted to the corresponding fraction of colleges.[6] For example, a "top pooling" policy that pools some fraction (an interval) of the top performing students leads to these students being randomly assigned to the same fraction of the top college seats.[7] An attractive property of performance-disclosure policies is that they do not require changing the tests or introducing new components to

[4]For example, In a 2014 New York Times article, (https://www.nytimes.com/2014/08/02/opinion/sunday/south-koreas-education-system-hurts-students.html), Se-Woong Koo reports that many South Korean presidents tried to limit cram schools' activities, including passing a 10 p.m. mandatory closure time. But even this restriction was circumvented "by operating out of residential buildings or blacking out windows so that light cannot be seen from outside."

[5]While our focus is on college admissions, our model may also be applied to other settings. One example is large corporate promotion contests, in which effort could correspond to "gaming" by smoothing store level sales numbers, distorting subordinate activity, or redirecting marketing resources. We thank a referee for suggesting this example.

[6]Such coarsening can be viewed as a particular way of making performance on the test noisier. Morgan et. al. (2022), in closely related independent research, suggest that other forms of noise can also be socially beneficial.

[7]A similar policy was proposed by Barry Schwartz in a 2007 LA times article (http://articles.latimes.com/2007/mar/18/opinion/op-schwartz18/2) in order "to dramatically reduce the pressure and competition that our most talented students now experience." "Then the names of all the 'good enough' students could be placed in a metaphorical hat, with the 'winners' drawn at random for admission. Though a high school student will still have to work hard to be 'good enough' for Yale, she won't have to distort her life in the way she would if she had to be the 'best.' The only reason left for participating in all those enrichment programs would be interest, not competitive advantage." We provide a framework to formalize, evaluate, and compare top pooling and other performance-disclosure policies.

the admissions process. They also respect the property that a higher score leads to a better college assignment than a lower score. These two properties may help make such policies appealing to policy makers.

A real-world example of a coarse performance-disclosure policy is the one recently adopted by the South Korean Ministry of Education for the College Scholastic Ability Test (CSAT), which determines college admissions in South Korea. Until 2018, each part of the test was graded on a 0-100 or 0-50 scale. Starting in 2018, scores of 90-100 in the English component of the CSAT are reported as one grade, scores of 80-89 as another grade, etc. "Students in the same graded classification will all be considered on an equal playing field in the college admissions process, regardless of their numerical scores."[8] The goal is to reduce costly competition between students,[9] while recognizing that the assortativity of the admissions process will be reduced as well.[10] One possible concern with such a policy, however, is that even if its overall effect on student welfare is positive, it may be that some students are harmed while others benefit.

We are interested in performance-disclosure policies that benefit all students, and refer to such policies as Pareto improving. In particular, we do not need to consider welfare trade-offs across students. A key finding of our analysis is that Pareto improving policies often exist. This may seem surprising, since a fixed set of college seats implies that a student can be admitted to a better college only if another student is admitted to a worse college. The crucial element that makes Pareto improvements possible is that test preparation is costly. The costs students incur, as well as the resulting college assignment, are determined in equilibrium, and the equilibrium is affected by the performance-disclosure policy. Relative to the baseline contest with no coarsening, introducing a performance-disclosure policy leads to some students being admitted to better colleges; this makes them better off even if they incur higher costs, as long as the cost increase is not too large. Other students are admitted

[8]Korea JoongAng Daily, October 10, 2015. "CSAT English section to take absolute grade scale."

[9] This "... grading system aims to reduce excessive competition among test-takers." "'We are trying to alleviate unnecessary and exorbitant competition between students who are competing with one another to gain one or two points more,' said Kim Doo-yong, a ministry official." Ibid.

[10]"In the last mock exam..., 23 percent of total examinees scored in the first grade... but only 4.64 percent of total examinees received perfect scores, which by current standards means they would have been the only ones to classify in the first grade." Ibid.

to worse colleges; if they also incur lower costs they are made better off as long as the reduction in the costs is large enough.

We characterize the Pareto improving policies and rank them in a Pareto sense. We also characterize the Pareto frontier of such policies. We do this initially for top pooling policies, which are the easiest to analyze, and then for policies that pool any single interval of performance. The characterization shows that pooling a performance interval is Pareto improving if and only if the student with the highest performance in the interval benefits from the pooling. This in turn happens if the population distribution of student ability conditional on the same interval (in percentile terms) first-order stochastically dominates (FOSD) the uniform distribution. We then generalize this condition to policies with multiple pooling intervals.

We then consider robust Pareto improving policies, which are Pareto improving for any distribution of college seats. We characterize the robust Pareto improving policies and show that the Pareto optimal policy among them is unique. This policy consists of pooling each of the maximal intervals on which the conditional distribution of student ability FOSD the uniform distribution. This characterization is useful because it only involves students' ability distribution. Given this distribution, it is straightforward to derive the Pareto optimal policy, which is robust to the distribution of college seats.

We illustrate the potential applicability of our results with an empirical estimation that uses data on college admissions in Turkey. We use the framework of Krishna et al. (2018) to calibrate the model and estimate applicants' ability distribution and the distribution of college seats. We then simulate a college admissions contest with these distributions and find a maximal Pareto improving "bottom pooling" interval, which pools together a fraction of the lowest test scores.[11] Finally, we design a laboratory experiment based on the calibrated distributions and Pareto improving bottom pooling policy. We evaluate subjects' behavior in the baseline contest and with bottom pooling, and find that the behavior is in broad agreement with the theory. A small set of subjects, those with the lowest ability among the subjects who should not be affected by the bottom pooling policy, behave in a way that slightly decreases their monetary payoffs. We provide possible explanations for

[11] Appendix F describes the Pareto optimal robust Pareto improving policy.

this and argue that these subjects may be better off as well. Taken together, our theory, empirical estimation, and experiment suggest that the simple performance-disclosure policies we investigate have the potential to improve the welfare of millions of college admissions applicants.

The rest of the paper is organized as follows. Section 1.1 reviews the related literature. Section 2 introduces the model. Section 3 presents the equilibrium and the notion of Pareto improvements. Section 4 investigates top pooling. Section 5 investigates policies with a single pooling interval and with multiple pooling intervals. Section 6 derives the conditions for robust Pareto improvements. Section 7 conducts the empirical estimation using data on Turkish college admissions to show what a Pareto improving policy might look like. Section 8 uses the results from Section 7 to conduct a laboratory experiment and evaluate whether such a policy might lead to better outcomes in practice. Section 9 concludes. The appendices contain proofs, examine peer effects, extend our results for top pooling to more general student utility functions, and provide additional material regarding the empirical estimation and the laboratory experiment.

1.1 Relation to the literature

College admissions feature prominently in the matching literature, beginning with Gale and Shapley's (1962) seminal contribution. The focus of much of this work is on stability and efficiency in the presence of heterogeneous student preferences, while abstracting from the effort students exert. Since we are interested in Pareto improvements, endogenous effort choice is an important feature of our framework.

Condorelli (2012) characterizes the ex-ante efficient allocations of heterogeneous objects to heterogeneous agents with private valuations. We are interested in ex-post Pareto improvements. The difference between ex-ante and ex-post Pareto is important in the context of college admissions, because the former allocations may treat some agents better at the expense of other agents and consequently be controversial. And the set of ex-ante optimal allocations can be completely different from those that are ex-post Pareto improving, as we discuss at the end of Section 2.1 below. In addition, although Condorelli (2012) provides an elegant general solution to the problem he studies, his solution delivers little insight regarding pooling intervals of performance in the context of college admissions, which is the focus of our work. His main take-away insights apply when all players' type distributions have monotone hazard rates. Our results do not require such a condition.

Other authors studied the optimality of coarse partitions with random lotteries within elements of the partitions in various settings. The most closely related paper is the one by Che et al. (2018), who study the possibility of collusion among privately-informed bidders in an auction for a single object. They are interested in auction formats that are immune to collusion, as well as in optimal cartels. Our paper contains several results that look similar to theirs, and which are derived by referring to the same tools from the theory of mechanism design. The results of the two papers cannot be directly compared, however, because they are derived in different settings. The models are different and so are details of the analyses. For example, a cartel in their model is free to choose the bids for colluding bidders (as long as this choice is incentives compatible), and prizes are determined according to the rules of an exogenous auction. In contrast, our disclosure policies determine the rules, and bids and prizes are determined endogenously given the rules.

Chao and Wilson (1987) study priority services of scarce supplies. They show in a model with an interval of customers' valuations that few priority classes suffice to capture most of the gains from priority service, and more generally that the surplus not realized due to using only a finite number n of priority classes (with customers within each class served in a random order) is only of order $1/n^2$. In addition, McAfee (2002) shows that in a general matching setting two priority classes capture a large share of the payoffs produced by perfect matching. Similar findings were reported by Wilson (1989).

By studying a monopolist matchmaker, Damiano and Li (2007) show that perfect sorting may not be optimal. Instead, the monopolist can increase revenue by pooling a small interval of types, within which matching is random. However, unlike in our model, the reason for the optimality of pooling intervals in their model is the tradeoff between efficiency of matching and rent extraction.[12] Similarly, in Rayo's (2013) analysis of the design of positional goods,

[12] A less closely related paper by Hoppe et al. (2011) explores the relative performance of coarse matching versus assortative matching.

a monopolist can restrict the variety of positional goods that consumers use to signal their types in order to extract more surplus from consumers. Moldovanu et al. (2007) show that the designer of a contest for status may prefer to pool contestants into status categories in order to increase the aggregate performance.

Several papers compare allocating objects via contests[13] and lotteries from the perspective of contestants' welfare (see Taylor et al. (2003), Koh et al. (2006), Hoppe et al. (2009), and Chakravarty and Kaplan (2013)). The most closely related work is by Hoppe et al. (2009). They consider a two-sided matching model with ex-ante symmetric agents on each side, in which assortative matching takes place based on costly signals. They provide conditions (expressed in terms of monotone failure rates) under which random matching leads to ex-ante higher welfare than assortative matching, and show that random matching is Pareto improving for agents on one side if the distribution of types of that side first-order stochastically dominates the uniform distribution.

Hafalir et al. (2018) investigate a model of college admissions with entrance exams and two colleges with different qualities. They compare centralized admissions, in which students can apply to both colleges, and decentralized admissions, in which each student can apply to only one college. They show that lower ability students prefer the decentralized setting and higher ability students prefer the centralized setting. Fang and Noe (2018) consider a selection contest with identical prizes, and show that pooling a larger number of the top performers than the number of prizes can sometimes lead to lower risk taking without reducing winner quality. Fang, Noe, and Strack (2018) consider a large contest framework similar to ours to investigate the effect of different university grading curves when post-graduation salaries depend on inferences employees make from grades about student ability and human capital accumulation.

Ostrovsky and Schwarz (2010) investigate information disclosure policies by schools when students are passive and exert no effort, and focus on the amount of information schools reveal in equilibrium. In our analysis, performance disclosure policies affect students' efforts, and this determines which policies are Pareto preferred. A more recent contribution by Boleslavsky and Cotton (2015) considers schools' incentives to invest in quality when they

[13] Contests in these papers typically have the form of waiting in line.

can choose imperfectly informative grading policies. As a result of this strategic choice, schools have a greater incentive to invest in quality, which can increase welfare. Gottlieb and Smetters (2014) investigate why MBA students vote for grade non-disclosure policies when employers make inferences about students' abilities based on the disclosed information. Frankel and Kartik (2019) point out that the value of standardized tests can be diminished by students muddling the ability signals contained in these tests by engaging in preparation that is not available to all students. An important common feature of our and their model is that wasteful effort can distort signaling in contests.

Dubey and Geanakoplos (2010) consider a game of status between students. A student's status is equal to the difference between the number of students with a lower grade and the number of students with a higher grade. In particular, the aggregate allocation value of status is always 0. A student's performance is a noisy measure of his costly effort, and, similarly to our model, a grading policy pools intervals of performance. The focus is on characterizing grading policies that maximize effort. Such policies involve some pooling, and with heterogeneous students necessary conditions for such policies are derived. Coarse grades also arise in the setting of Harbaugh and Rasmusen (2018), in which a sender can choose whether to certify his privately-known quality. Certification schemes with coarse grades can result in more information by inducing the sender to certify a larger set of qualities.

Our paper also belongs to the literature on all-pay contests. Most papers in this literature focus on settings with two players, ex-ante symmetric players, or identical prizes. Olszewski and Siegel (2016) introduced the approximation approach to large contests, which makes it possible to study contests with many ex-ante asymmetric players and heterogeneous prizes, as we do here. Olszewski and Siegel (2020) use this approach to study performance-maximizing contests.[14] Bodoh-Creed and Hickman (2018) use a similar (and independently developed) approach to study quotas and affirmative action in college admissions. While there are several technical differences between their model and ours,[15] the main differences are in the

[14] Fang, Noe, and Strack (2020) study the effect of different prize structures on aggregate effort in symmetric all-pay auctions with complete information.

[15] They consider a more general utility function but assume that the limit distribution of college seats is atomless, and consider two groups of students, the minority and the majority, such that students within each group are ex-ante symmetric.

design instruments they consider (quotas and affirmative action) and their focus on aggregate welfare as opposed to our focus on Pareto improvements. They also allow for productive effort, which is potentially important in college admissions settings that take into account factors like high school performance. One of their findings is that using a lottery to assign students to colleges would generate higher aggregate student welfare than a contest for college admissions. Our investigation of optimal category rankings shows that a pure lottery can be improved upon for all students by partitioning the set of students into several categories based on their performance and using a separate assignment lottery for each category.

2 The baseline contest

A large number of players (students) compete for prizes (college seats) by taking a test. Each prize is characterized by its known value $y \in [0,1]$, and each player is characterized by her ability (type) $x \in [0,1]$, which affects her cost of performance on the test and/or her prize valuation. Each player's type is drawn from a player-specific distribution, independently across players. This accommodates ex-ante asymmetry across players. After privately observing her type, each player exerts costly effort to achieve her desired performance $t \geq 0$ on the test. The test may have several parts or be comprised of several examinations, provided that they are weighted in a way that produces a single number (the performance) according to which players are ranked. The player with the highest performance obtains the highest prize, the player with the second-highest performance obtains the second-highest prize, and so on. Some prizes may be identical, which allows for multiple seats in a given college (or tier of colleges). Ties are resolved by a fair lottery. The utility of a type x player who chooses performance t and obtains prize y is

$$g_1(x)\, y - \frac{c(t)}{g_2(x)}, \tag{1}$$

where c is strictly increasing and twice continuously differentiable, and $\lim_{t\to\infty} c(t) = \infty$.[16] Function c captures the cost of performance, function $g_1 \geq 0$ captures the effect of the player's type on her prize valuation, and function $g_2 \geq 0$ captures the effect of the player's type on her cost of performance. We order types so that $g_1(x)\, g_2(x) = x$. Two special cases (which are assumed in most of the contest literature) are

$$xy - c(t), \tag{2}$$

in which the player's type only affects her prize valuation, and

$$y - \frac{c(t)}{x}, \tag{3}$$

in which the player's type only affects her performance cost. Utilities (1) for different functions g_1 and g_2 are strategically equivalent, because for each type x multiplying (1) by $g_2(x)$ gives (2). Throughout our analysis we will assume utility (2). This is for convenience only. As we now discuss, all our results hold without change for any utility (1) (and for the special case (3)).

2.1 Discussion of the model

Our baseline contest accommodates heterogeneity in college quality and student ability, and models costly test preparation as a strategic choice. Like any other model, it abstracts from certain realistic and potentially important aspects. First, the model stipulates a common ordinal ranking of college quality across students.[17] Second, the model abstracts from factors that are not controlled by the players and may affect their performance ("noise"). These two assumptions are made for tractability (but they may also be fairly realistic in some settings). Third, the model assumes that test preparation is costly, as in Spence (1973).

[16]The linearity of y is a normalization; we can replace y in players' utility with $h(y)$, where h is strictly increasing and twice continuously differentiable and $h(y) = 0$, without affecting any of the results. We can also replace the assumption that $\lim_{t\to\infty} c(t) = \infty$ with the assumption that $\lim_{t\to\bar{t}} c(t) = \infty$ for some positive $\bar{t}$ that represents a cap on students' maximal effort.

[17]Homogeneous ordinal preferences are also assumed in some matching papers on school choice (for example in Lien, Zheng, and Zhong (2017)).

This cost is captured by function c, and should be interpreted as net of any direct benefit from the preparation activities. This is most appropriate for activities specifically geared toward improving students' performance on the test, as discussed in the introduction. The model is less suitable when other activities, such as taking AP classes, play an important role in college admissions and may have significant direct benefits at moderate levels of investment. But even there the costs may exceed the benefits past a certain point. Bodoh-Creed and Hickman (2019) provide support for this in the context of college admissions in the United States.[18] In such cases, function c can be thought of as a simplification that assumes that all preparation activities are costly.

Fourth, similarly to much of the matching literature, peer effects are absent: a student's valuation for being admitted to a college does not depend on which other students are admitted to the same college. In fact, peer effects can be accommodated without changing the substance of any of our results. This is done in Appendix C. The idea is that in a large contest each student is fairly certain about the equilibrium distribution of student types admitted to the various colleges. We can therefore replace the value y of being admitted to a specific college with another value that includes the peer effects generated by the set of students admitted to that college. The rest of the analysis is unchanged. Finally, in Appendix D we show that our results for top pooling generalize to separable utility functions of the form $h(x, y) - c(t)$ and $h(y) - c(x, t)$ that satisfy some conditions.

It is also important to point out that while utilities (1) for different functions g_1 and g_2 are strategically equivalent, different functions g_1 have different implications for aggregate welfare. Setting aside players' performance, if $g_1(x)$ increases in x (e.g., utility (2)), then the allocation of prizes that maximizes aggregate welfare allocates the highest prize to the player with the highest type, the second highest prize to the player with the second highest type, etc. But the opposite is true if $g_1(x)$ decreases in x. And if $g_1(x)$ is independent of x (e.g., utility (3)), then all prize allocations generate the same aggregate welfare. These different functional forms make no difference for our analysis, however, because we focus on performance-disclosure policies that make all students better off, and in fact all types of

[18] They study a rich data set and a contest model in which effort can be productive, but show that for most students most of the effort is in fact wasteful (and the wasted effort is three times higher than the productive effort for the middle 50 percent of the learning cost distribution).

students better off (we provide a precise definition in Section 3.1 below). Since each player's ranking of prize-performance pairs is the same for all utility functions (1), our results hold for all utility functions (1). In particular, we do not need to take a stand on aggregate welfare or on whether awarding higher prizes to players with higher types is desirable from a welfare perspective. In contrast, the ex-ante efficient allocations (Condorelli (2012)), which maximize the sum of expected utilities, differ across the functional forms. For example, under utility (3) the ex-ante efficient allocations are those that induce a performance of 0, that is, all the possible lotteries over prizes (including deterministic allocations), which do not depend on players' performance.

3 Equilibrium

A direct equilibrium analysis of the baseline contest described in Section 2 is intractable because the equilibria generally involve mixed strategies and are not symmetric. Since the contests we consider have many players and prizes, we can make use of the tractable approach to studying the equilibria of large contests, which was developed in Olszewski and Siegel (2016). They show that all the equilibria of such large contests are closely approximated by the unique single-agent mechanism in a specific environment that implements the assortative allocation of prizes to agent types and gives the lowest type a utility of 0. More precisely, denote by F the average distribution of players' types and suppose that it has a continuous, strictly positive density f, and denote by G the empirical distribution of prizes, which need not be continuous or have full support. For example, G may consist of atoms that represent colleges (or tiers of colleges). The size of each atom represents the fraction of overall seats offered by the corresponding college. The *assortative allocation* assigns to each type x prize

$$y^A(x) = G^{-1}(F(x)),$$

where

$$G^{-1}(z) = \inf\{y : G(y) \geq z\} \text{ for } 0 \leq z \leq 1.$$

That is, the quantile in the prize distribution of the prize assigned to type x is the same as the quantile of type x in the type distribution. It is well known (see, for example, Myerson (1981)) that the unique incentive-compatible mechanism that implements the assortative allocation and gives type $x = 0$ utility 0 specifies for every type x performance

$$t^A(x) = c^{-1}\left(xy^A(x) - \int_0^x y^A(\tilde{x})\, d\tilde{x}\right). \tag{4}$$

This implies that type x obtains utility

$$U(x) = xy^A(x) - c(t^A(x)) = \int_0^x y^A(\tilde{x})\, d\tilde{x}. \tag{5}$$

Roughly speaking, the approximation shows that in any equilibrium of a large contest a player with type x with high probability chooses a performance close to $t^A(x)$ and obtains a prize close to $y^A(x)$, which gives her a utility close to $U(x)$. See Olszewski and Siegel (2016) for a precise statement and additional details.

The intuition for why this single-agent mechanism approximates the equilibria of large (finite) contests is that, given players' equilibrium strategies, with a large number of players the law of large numbers implies that each bid leads to an almost deterministic rank-order quantile (in the distribution of bids) and thus to an almost deterministic prize. In the limit we obtain an "inverse tariff" that maps bids to prizes. Utility (1) implies that higher types choose higher bids from any tariff, so the mechanism induced by the inverse tariff implements the assortative allocation. Any player can bid 0 and obtain the lowest prize, so the utility of type 0 is 0.

In the rest of paper we focus on the approximating single-agent mechanism to investigate how different performance-disclosure policies affects students' welfare in a Pareto sense, which we define below. As discussed in the introduction, the potential for Pareto improvements exists because performance is costly: if pooling reduces students' performance, all students could be made better off even though the allocation of college seats changes.

3.1 The notion of Pareto improvements

We study the approximating mechanisms under various performance-disclosure policies, and use the term "Pareto-improving" in reference to the utility of the types in these approximating mechanisms. A performance-disclosure policy is Pareto improving if all types are better off and there is a positive measure of types that are strictly better off. Such an improvement implies that in a sufficiently large contest some players are strictly better off and no player is worse off by more than an arbitrarily small amount; moreover, the sum of these small amounts across all players who are worse off is arbitrarily small compared to the gains of the players who are strictly better off. We point out, however, that these gains and losses are in terms of expected utilities. Since pooling in our performance-disclosure policies leads to lotteries over prizes, by "gains" for a player we mean that the player prefers the lottery to the original disclosure policy, but she may or may not prefer the outcome once the lottery is realized.

4 Top pooling

We begin by considering "top q pooling," in which a fraction q of the highest performing students are pooled. These students still obtain the best college seats, but the allocation of these seats to the students is random. Thus, to study the effect of top pooling we can simply consider a contest in which the top fraction q of prizes are replaced with mass q of identical prizes whose value is equal to the average value of the top prizes. To do this, let $x^* = F^{-1}(1-q)$ be the type whose quantile in the average type distribution is $1-q$, and let G^q be the prize distribution that results from replacing the top mass q of prizes in distribution G with a mass q of prize

$$y(q) = \frac{\int_{1-q}^{1} G^{-1}(z)\, dz}{q} = \frac{\int_{x^*}^{1} y^A(x) dF(x)}{1-F(x^*)}.$$

That is, $(G^q)^{-1}(F(x)) = G^{-1}(F(x))$ for $x \leq x^*$, and $(G^q)^{-1}(F(x)) = y(q)$ for $x > x^*$. The corresponding assortative allocation $y^{A,q}$ satisfies

$$y^{A,q}(x) = (G^q)^{-1}(F(x)). \tag{6}$$

The unique single-agent mechanism that implements this allocation and gives type $x = 0$ a utility of 0 specifies performance

$$t^{A,q}(x) = c^{-1}\left(xy^{A,q}(x) - \int_0^x y^{A,q}(\tilde{x})\,d\tilde{x}\right). \tag{7}$$

Consider how this mechanism compares with the one in Section 3, which implements the assortative allocation y^A and in which the performance t^A is given by (4). By definition of G^q and $y^{A,q}$, we have that $y^{A,q}(x) = y^A(x)$ and $t^{A,q}(x) = t^A(x)$ for $x \leq x^*$, and $y^{A,q}(x) = y(q)$ and $t^{A,q}(x) = M$ for $x > x^*$, where

$$M = c^{-1}\left(x^*y(q) - \int_0^{x^*} y^A(\tilde{x})\,d\tilde{x}\right) = c^{-1}\left(x^*\frac{\int_{x^*}^1 y^A(\tilde{x})dF(\tilde{x})}{1-F(x^*)} - \int_0^{x^*} y^A(\tilde{x})\,d\tilde{x}\right). \tag{8}$$

Type x^* is a threshold type, above which pooling occurs: all higher types choose the same performance M and obtain the same lottery over prizes. Since there is a one-to-one correspondence between q and x^*, in what follows we also refer to top q pooling as "top pooling with threshold x^*." To gain some intuition for performance M, notice that (5) and (8) imply that

$$x^*y^A(x^*) - c(t^A(x^*)) = x^*\frac{\int_{x^*}^1 y^A(x)dF(x)}{1-F(x^*)} - c(M), \tag{9}$$

that is, type x^* is indifferent between choosing performance $t^A(x^*)$ and obtaining prize $y^A(x^*)$ and choosing performance M and obtaining a prize randomly from the mass $1-F(x^*)$ of the highest prizes. (Note that $t^A(x^*) < M$.)

4.1 Welfare comparisons

Our first result compares each type's utilities in the approximating mechanisms with and without top pooling. The proof of this result, as well as those of other results, are in Appendix B. For the result, we assume that not all prizes are identical in the top mass $1 - F(x^*)$ of prizes.[19]

Proposition 1. *Consider top pooling with threshold x^*.*

(a) The utility of types $x < x^*$ is not affected.

(b) The utility of type $x > x^*$ increases if and only if

$$\frac{\int_{x^*}^{1} y^A(\tilde{x})dF(\tilde{x})}{1 - F(x^*)} \geq \frac{\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}}{x - x^*}. \tag{10}$$

(c) The gain in utility for types $x > x^*$ first increases and then decreases in type. Thus, there is a type x^{**} in $(x^*, 1]$ such that the utility of types x in (x^*, x^{**}) is higher with pooling than without pooling, and the utility of types $x > x^{**}$ is lower with pooling than without pooling.

(d) Top pooling is Pareto improving if and only if it increases the utility of type 1, that is,

$$\frac{\int_{x^*}^{1} y^A(\tilde{x})dF(\tilde{x})}{1 - F(x^*)} \geq \frac{\int_{x^*}^{1} y^A(\tilde{x})\, d\tilde{x}}{1 - x^*}. \tag{11}$$

Proposition 1 shows that the effect of top pooling on players' welfare depends on their types. Players with types lower than x^* are unaffected because their performance and the prize they obtain do not change. Players with types in (x^*, x^{**}) benefit, but the reason for this may vary across the players. Players with types higher than but close to x^* obtain a prize lottery that is better than the prize they obtain without top pooling (because prize $y^A(x^*)$ is the lower bound of the support of the prize lottery). Since type x^* is indifferent between the contests with and without top pooling, his performance with top pooling must be higher, that is, $t^A(x^*) < M$. Thus, players with types close to x^* benefit from top pooling because they obtain (with high probability) a better prize, even though they choose a higher performance. On the other hand, players with types lower than but close to 1 obtain a prize

[19] If they are identical, then top pooling has no effect.

lottery that is worse than the prize they obtain without top pooling (because prize $y^A(1)$ is the upper bound of the support of the prize lottery). If top pooling is Pareto improving, therefore, these players must choose a sufficiently lower performance with top pooling that offsets the loss from the prize lottery. In particular, $t^A(1) > M$.[20]

Note that the left-hand side of (10) is the average of $y^A(\tilde{x})$ across types $\tilde{x}$ that choose performance M. This average is taken with respect to the actual (truncated) distribution of types. The left-hand side is independent of x. The right-hand side is the average of $y^A(\tilde{x})$ across all types lower than x that choose performance M taken with respect to the (truncated) uniform distribution. The right-hand side increases in x. To understand (10), which is the key condition in Proposition 1, multiply each side of (10) by $x - x^*$. Then, the left-hand side of (10) is the difference between the utilities of type $x > x^*$ and type x^* in the contest with top pooling. The right-hand side of (10) (after multiplying it by $x - x^*$) is

$$\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x} = \int_0^x y^A(\tilde{x})\, d\tilde{x} - \int_0^{x^*} y^A(\tilde{x})\, d\tilde{x},$$

which is the difference between the utilities of type $x > x^*$ and type x^* in the contest without top pooling. The fact that $\int_0^x y^A(\tilde{x})\, d\tilde{x}$ is the utility of type x in the contest without top pooling is well known from standard mechanism design because a single-agent mechanism approximates the equilibria of large finite contests(see Section 3). Intuitively, type $\tilde{x} + d\tilde{x}$ can pretend to be type $\tilde{x}$, obtain prize $y^A(\tilde{x})$, and enjoy a utility increase of $y^A(\tilde{x})\, d\tilde{x}$ relative to type $\tilde{x}$.

Proposition 1 shows that types slightly higher than x^* benefit from top pooling,[21] but high types may or may not benefit. This depends on whether type 1 benefits, in which case all types higher than x^* do. The two possibilities are depicted in Figure 1, which illustrates the utility gain resulting from top pooling as a function of type. The left-hand side corresponds to top pooling with $x^{**} < 1$, so it is not Pareto improving, and the right-hand side corresponds to top pooling that is Pareto-improving, so $x^{**} = 1$.

[20]If top pooling is Pareto improving, then it reduces the aggregate cost of performance because the set of prizes is unchanged, and all players are made (at least weakly) better off.

[21]This is because the marginal equilibrium utility, or the marginal information rent, is equal to a type's prize, which is higher with top pooling for types slightly higher than x^*.

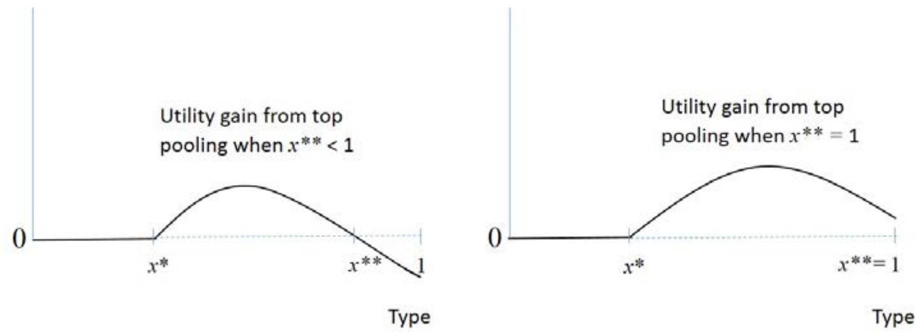

Figure 1: Utiltiy gain from top pooling.

The following example illustrates the results from Proposition 1.

Example 1. *Suppose that F and G are uniform. Consider utility (2) with $c(t) = t$.*[22] *The assortative allocation is $y^A(x) = x$, and the approximating mechanism specifies performance $t^A(x) = x^2/2$. The payoff of type x is $x^2/2$.*

Under top pooling with threshold x^*, every type $x < x^*$ chooses performance $\frac{1}{2}x^2$ and obtains prize x, and every type $x > x^*$ chooses performance M and obtains a prize drawn uniformly from interval $[x^*, 1]$. Performance M is given by

$$\frac{1}{2}(x^*)^2 = x^*\frac{1+x^*}{2} - M,$$

so $M = x^*/2$. Thus, the payoffs are $\frac{1}{2}x^2$ for $x < x^*$ and

$$x\frac{1+x^*}{2} - \frac{x^*}{2} \geq \frac{1}{2}x^2$$

for $x > x^*$. All top pooling thresholds are Pareto-improving.

Example 1 shows that there may exist multiple Pareto-improving pooling thresholds. It is therefore reasonable to ask whether these thresholds can be Pareto ranked. In Example 1, the derivative of the payoff of every type $x > x^*$ with respect to the threshold type x^* is $(x-1)/2$, so all types prefer $x^* = 0$, i.e., the Pareto preferred top pooling is a lottery over all prizes. In general, however, Pareto-improving pooling thresholds are not Pareto ranked, as the following example shows.

Example 2. *Let $F = G$ have density $f = g = 7/4$ on intervals $[0, 1/4]$ and $[3/4, 1]$, and*

[22]Recall that all our results hold for utility (1).

density $f = g = 1/4$ on interval $(1/4, 3/4)$. Consider utility (2) with $c(t) = t$. The assortative allocation is $y^A(x) = x$, and the approximating mechanism specifies performance $t^A(x) = x^2/2$. The payoff of type x is $x^2/2$.

Top pooling with threshold $x^* = 0$, which is a lottery, is Pareto improving, since the expected prize of $1/2$ at performance 0 gives each type x utility $x/2$, which exceeds $x^2/2$. Now consider top pooling with threshold $x^* = 1/2$. For this threshold we have $M = 19/64$, since

$$\frac{1}{2}\left(\frac{1}{2}\right)^2 = \frac{1}{2}\left(\frac{1}{8}\frac{5}{8} + \frac{7}{8}\frac{7}{8}\right) - \frac{19}{64}.$$

Type $x = 1$ benefits from this top pooling, since $((1/8)(5/8) + (7/8)(7/8)) - (19/64) > 1/2$. It is therefore Pareto improving, by part (d) of Proposition 1. Type $x = 1$ (as well as slightly lower types) also prefer this top pooling to a pure lottery. However, types in the interval $(0, 1/2)$ have the opposite preference, because a lottery gives each of them an expected utility of $x/2$, and top pooling with threshold $x^* = 1/2$ gives each of them an expected utility of $x^2/2$.

The following corollary of Proposition 1 clarifies when Pareto-improving pooling thresholds are Pareto-ranked.

Corollary 1. *Suppose that $x_1^* < x_2^*$ are top-pooling thresholds and not all prizes are identical in the top mass $1 - F(x_1^*)$ of prizes.*

(a) If type $x = 1$ weakly prefers x_1^* to x_2^*, then types x in $(x_1^*, 1)$ strictly prefer x_1^* to x_2^*, so x_1^* is Pareto preferred to x_2^*.

(b) If type $x = 1$ strictly prefers x_2^* to x_1^*, then x_1^* and x_2^* are not Pareto ranked. There is an x^{**} in $(x_2^*, 1)$ such that types x in (x_1^*, x^{**}) strictly prefer x_1^*, and types $x > x^{**}$ strictly prefer x_2^*.

Figure 2 illustrates the two parts of Corollary 1. The left-hand side corresponds to part (a), and the right-hand side corresponds to part (b).

Corollary 1 follows from Proposition 1 by interpreting top pooling with threshold x_1^* as the composition of two top poolings: top pooling with threshold x_2^* followed by top pooling with threshold x_1^*. More precisely, top pooling with threshold x_2^* is equivalent to a

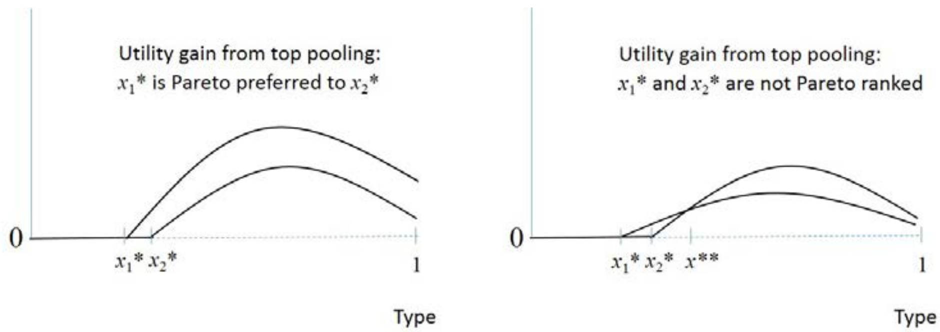

Figure 2: Comparing the utility gain from pooling thresholds $x_1^* < x_2^*$

baseline contest with a modified prize distribution G^q in which the prizes allocated to types $(x_2^*, 1]$ are replaced by a mass $q = 1 - F(x_2^*)$ of the average (according to the original prize distribution G) of the prizes in G. Top pooling with threshold x_1^* applied to this modified prize distribution G^q is clearly equivalent to top pooling with threshold x_1^* applied to the original prize distribution. Proposition 1 for top pooling with threshold x_1^* applied to prize distribution G^q immediately implies the corollary: part (a) of the corollary follows from Part (d) of Proposition 1, and part (b) of the corollary follows from part (c) of Proposition 1.

Corollary 1 leads to a simple description of the Pareto frontier of top-pooling thresholds. To see this, consider the function ϕ that assigns to any threshold x^* the utility of type $x = 1$ in the approximating mechanism with this threshold. Denote by ψ the lowest monotone non-decreasing function that is pointwise weakly higher than ϕ. In Figure 3, ψ coincides with ϕ on $[0, x_2]$, $\psi \equiv \phi(x_2)$ is constant on $[x_2, x_3]$, it again coincides with ϕ on $[x_3, x_4]$, and $\psi \equiv \phi(x_4)$ is constant on $[x_4, 1]$. The Pareto frontier of top-pooling thresholds consists of all the thresholds x^* at which $\psi(x^*) = \phi(x^*)$ except that only the lower endpoints of the intervals on which ψ is constant belong to the Pareto frontier. In Figure 3, the Pareto frontier consists of the closed interval $[0, x_2]$ and the left-open interval $(x_3, x_4]$.

Indeed, if $\phi(x^*) \leq \phi(x)$ for some $x < x^*$, which means that $\psi(x^*) = \psi(x')$ for some $x' < x^*$, then x^* does not belong to the Pareto frontier by part (a) of Corollary 1. Otherwise, that is, if $\phi(x^*) > \phi(x)$ for every $x < x^*$, then type 1 prefers x^* to any $x < x^*$ so x^* is not Pareto dominated by any $x < x^*$. In addition, x^* is not Pareto dominated by any $x > x^*$ by parts (a) and (b) of Corollary 1. Finally, note that not all top poolings from the Pareto frontier are necessarily Pareto improving. The Pareto improving top poolings x^* satisfy the necessary and sufficient condition from part (d) of Proposition 1, which can be expressed as

$\psi(x^*) \geq \phi(1)$. In Figure 3, the top poolings from interval $[0, x_1)$, which is part of the Pareto frontier, violate this condition.

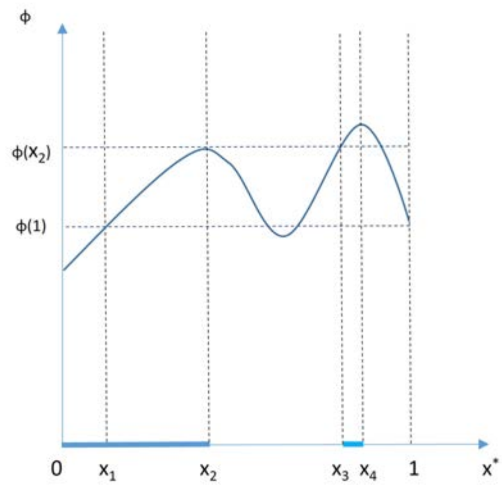

Figure 3: Pareto frontiers of top pooling thresholds.

When we consider only a finite set X^* of top pooling thresholds, Corollary 1 implies that the Pareto frontier of X^* consists of the threshold x^* that is most preferred by type 1 among all the thresholds in X^*, the threshold that is most preferred by type 1 among the thresholds in X^* that are lower than x^*, and so on.

5 Category Rankings

We now consider more general performance disclosure policies, which may include one or more pooled intervals of performance ranking. We investigate how different policies affect students' welfare, and identify the Pareto improving ones. We will use the term "category rankings" to describe such policies. A category ranking is a monotone partition of the players according to the ranking of their performance. One example is partitioning them above and below the median performance. Another example is partitioning them according to whether their performance is below the 10-th percentile, between the 10-th percentile and the 20-th percentile, etc. A category ranking induces a partition of the set of prizes, and the prizes within each element of the partition are randomly assigned to the players in the corresponding element of the category ranking.

Formally, a category ranking is a monotone partition $\mathcal{J}$ of the set $[0,1]$ of quantiles into singletons and left-open intervals. The intervals are $J_k = \left(q_k^l, q_k^h\right]$ for $1 \leq k \leq K \leq n$, where $0 \leq q_1^l < q_1^h \leq \cdots \leq q_K^l < q_K^h \leq 1$. The interpretation is the fraction $q_k^h - q_k^l$ of players whose performance quantile rankings lie in J_k are grouped together (any rule can be used to break ties in the ranking of two or more players who choose the same performance). Prizes are assigned in decreasing value to the partition elements, and distributed according to a fair lottery among the players in each partition element. To describe the approximating mechanism, denote by $G^{\mathcal{J}}$ the distribution of prizes when the prizes in each interval J_k are replaced with an equal mass of prize

$$y(J_k) = \frac{\int_{q_k^l}^{q_k^h} G^{-1}(z)dz}{q_k^h - q_k^l} = \frac{\int_a^b y^A(x)dF(x)}{F(b) - F(a)} \tag{12}$$

for $a = F^{-1}\left(q_k^l\right)$ and $b = F^{-1}\left(q_k^h\right)$. The corresponding assortative allocation $y^{A,\mathcal{J}}(x)$ satisfies

$$y^{A,\mathcal{J}}(x) = \left(G^{\mathcal{J}}\right)^{-1}(F(x)). \tag{13}$$

The unique incentive-compatible mechanism that implements this allocation and gives type $x = 0$ a utility of 0 specifies performance

$$t^{A,\mathcal{J}}(x) = c^{-1}\left(xy^{A,\mathcal{J}}(x) - \int_0^x y^{A,\mathcal{J}}(\tilde{x})\, d\tilde{x}\right). \tag{14}$$

Note that a category ranking induces a partition $\mathcal{I}$ of the set of types $X = [0,1]$ into singletons and K intervals $I_k = \left(F^{-1}\left(q_k^l\right), F^{-1}\left(q_k^h\right)\right]$, such that all types in interval I_k choose the same performance and obtain the same prize $y(J_k)$ in the approximating mechanism, and singleton types obtain the prize they did in the original approximating mechanism. Thus, the assortative allocation and approximating mechanism can be equivalently defined from the partition $\mathcal{I}$ of types (instead of the partition $\mathcal{J}$) by letting $G^{\mathcal{I}}$ coincide with $G^{\mathcal{J}}$ and defining $y^{A,\mathcal{I}}(x)$ and $t^{A,\mathcal{I}}(x)$ as in (13) and (14) with $\mathcal{I}$ instead of $\mathcal{J}$.

Thus, from the perspective of the approximating mechanism, a category ranking $\mathcal{J}$ corresponds to a partition $\mathcal{I}$ of the set of types into singletons and a finite number of left-open intervals. In what follows, it will be convenient to consider such partitions of the set of types and the corresponding approximating mechanisms. We will abuse terminology slightly by also referring to such partitions $\mathcal{I}$ of the type interval $[0,1]$ as category rankings.[23]

5.1 The added value of category rankings

Top pooling is a particular kind of category ranking: top pooling with threshold x^* is the category ranking $\mathcal{I} = \{(x^*,1]\} \cup \{\{x\} : x \leq x^*\}$. The richer set of outcomes that can be generated by category rankings may include outcomes that are Pareto preferred to all outcomes that can be generated by top poolings. This is what the following example demonstrates.

Example 3. *Let $F = G$ have density $f = g = 5/4$ on interval $[0,3/4]$, and density $f = g = 1/4$ on interval $[3/4,1]$. Consider utility (2) with $c(t) = t$. The assortative allocation is $y^A(x) = x$, and the approximating mechanism specifies performance $t^A(x) = x^2/2$. The payoff of type x is $x^2/2$.*

Top pooling with threshold $x^* = 3/4$ is Pareto improving. Indeed, the corresponding performance M is given by

$$\frac{1}{2}\left(\frac{3}{4}\right)^2 = x^*\left(\frac{1+\frac{3}{4}}{2}\right) - M,$$

which gives $M = 3/8$. Types in $(3/4,1]$ choose performance M, and each of them obtains a prize drawn uniformly from interval $(3/4,1]$. The utility of type $x = 1$ is equal to $1/2$ both with and without top pooling. So, by part (d) of Proposition 1, top pooling with threshold $x^* = 3/4$ is Pareto improving. One can readily check that top pooling with any threshold $x^* > 3/4$ is also Pareto improving and gives type $x = 1$ utility $1/2$. We will show that top pooling with any threshold $x^* < 3/4$ is not Pareto improving. So, by part (a) of Corollary 1, the threshold $x^* = 3/4$ is the Pareto preferred one.

[23]The fact that category rankings can be analyzed using the large contest approach makes them tractable. Of course, many other forms of noisy ranking exist, but their analysis may require different techniques. Pooling non-continuous intervals, for example, violates the property that a higher score leads to a higher expected prize, and therefore invalidates the large contest approach. While we suspect that such pooling will not be optimal, a formal analysis goes beyond the scope of this paper. It may also be difficult to convince policy makers to adopt such "non-monotonic" policies.

To see that no threshold $x^* < 3/4$ is Pareto improving, recall that the performance $M < 3/8$ satisfies

$$\frac{1}{2}(x^*)^2 = x^* E[y \mid t = M] - M,$$

where $E[y \mid t = M]$ is the expected prize contingent on choosing performance M. The utility of type $x = 1$ is thus

$$E[y \mid t = M] + \underbrace{\frac{1}{2}(x^*)^2 - x^* E[y \mid t = M]}_{-M} < \frac{1}{2},$$

because $E[y \mid t = M] < (1 + x^*)/2$ for $x^* < 3/4$.

Top pooling with threshold $x^* = 3/4$ is the category ranking that pools together the top $1/16$ of the types and leaves the other types as singletons. However, this category ranking is Pareto inferior to the category ranking that pools together the top $1/16$ of the types, and pools together the bottom $15/16$ of the types. Indeed, under this category ranking, the bottom $15/16$ of the types exert no effort and obtain an expected prize of $3/8$, while the top $1/16$ of the types choose performance $3/8$ and obtain an expected prize of $7/8$. Under the former category ranking, the top $1/16$ of the types also choose performance $3/8$ and obtain an expected prize of $7/8$, but the bottom $15/16$ of the types x obtain a lower utility of $x^2/2$.

5.2 Welfare comparisons

Consider first single-interval category rankings, that is, category rankings of the form $\mathcal{I} = \{(x^*, x^{**}]\} \cup \{\{x\} : x \leq x^* \text{ or } x > x^{**}\}$ for some types $0 \leq x^* < x^{**} \leq 1$; top pooling is a special case in which $x^{**} = 1$. As in the case of top pooling, we assume that not all of the prizes in quantiles $[F(x^*), F(x^{**})]$ are identical.[24] The following result generalizes Proposition 1.

Proposition 2. *(a) The utility of type $x \in (x^*, x^{**}]$ increases as a result of the single-interval*

[24] If they are identical, then the category ranking has no effect.

category ranking $\mathcal{I}$ if and only if

$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})dF(\tilde{x})}{F(x^{**}) - F(x^*)} \geq \frac{\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}}{x - x^*}.$$

(b) The category ranking $\mathcal{I}$ is Pareto improving if and only if

$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})dF(\tilde{x})}{F(x^{**}) - F(x^*)} \geq \frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\, d\tilde{x}}{x^{**} - x^*}. \tag{15}$$

The intuition for Proposition 2 is similar to the one underlying Proposition 1, applied to types in the interval $(x^*, x^{**}]$. In particular, if a type $x \in (x^*, x^{**}]$ benefits from the category ranking, then all types in the interval $(x^*, x]$ benefit as well. Types $x \leq x^*$ are clearly not affected by the category ranking, and the derivative of the utility of types $x > x^{**}$ is equal to $y^A(x)$ both in the original contest and under the category ranking. Thus, if type x^{**} is better off under the category ranking, then so are all types higher than x^{**}, which gives part (b).

For category rankings that include more than one interval, a generalization of the conditions in Proposition 2 provides sufficient conditions for a category ranking to increase the utility of a type and to be Pareto improving, but these conditions are no longer necessary. This is because pooling on an interval may increase the utility of types above the interval to such a degree that even if these types are pooled in a way that lowers their utility, the overall effect may be to increase their utility relative to the baseline contest.

To obtain the sufficient conditions, consider a category ranking $\mathcal{I}$ that includes precisely the $K \geq 2$ intervals $I_1, \ldots, I_K$, where $I_k = (x_k^*, x_k^{**}]$ and $x_k^{**} \leq x_{k+1}^*$ for $k < K$. The effect of the category ranking can be described as follows. For each $k < K$ let G^k be the distribution of prizes when the prizes corresponding to intervals $I_1, \ldots, I_k$ are replaced by their averages. Then, the contest with the category ranking that pools only intervals $I_1, \ldots, I_{k+1}$ is the same as the contest with the single-interval category ranking that pools only interval I_{k+1} but starts with prize distribution G^k. Proposition 2 describes the effect of this single-interval category ranking on a baseline contest with prize distribution G^k. By induction on k we immediately obtain the following corollary of Proposition 2.

Proposition 3. *(a) The utility of type $x \in (x_k^*, x_k^{**}]$ increases as a result of the category ranking $\mathcal{I}$ if*

$$\frac{\int_{x_k^*}^{x_k^{**}} y^A(\tilde{x}) dF(\tilde{x})}{F(x_k^{**}) - F(x_k^*)} \geq \frac{\int_{x_k^*}^{x} y^A(\tilde{x})\, d\tilde{x}}{x - x_k^*} \text{ and } \frac{\int_{x_j^*}^{x_j^{**}} y^A(\tilde{x}) dF(\tilde{x})}{F(x_j^{**}) - F(x_j^*)} \geq \frac{\int_{x_j^*}^{x_j^{**}} y^A(\tilde{x})\, d\tilde{x}}{x_j^{**} - x_j^*} \text{ for all } j < k.$$

(b) The category ranking $\mathcal{I} = \{I_1, ..., I_K\}$ is Pareto improving if

$$\frac{\int_{x_j^*}^{x_j^{**}} y^A(\tilde{x}) dF(\tilde{x})}{F(x_j^{**}) - F(x_j^*)} \geq \frac{\int_{x_j^*}^{x_j^{**}} y^A(\tilde{x})\, d\tilde{x}}{x_j^{**} - x_j^*} \text{ for all } j \leq K.$$

Characterizing the Pareto frontier of category rankings is more complicated than for top poolings. In Appendix A we provide a method for checking whether a category ranking belongs to the Pareto frontier of category rankings, and we illustrate its usefulness with an example.

6 Robust Pareto improvements

The results in the previous sections suggest that Pareto improvements exist in some college admissions settings. To apply these results we need to construct estimates of the type distribution F and the prize distribution G (which affects the results via the assortative allocation y^A).

We now present simpler results that rely only on properties of the type distribution F and correspond to Proposition 1, Corollary 1, and Proposition 2. We will use the term "robust Pareto improvement" as shorthand for "weakly better for every type, for any functions c and G, and a Pareto improvement for some functions c and G."[25] Robustness is useful because there are many aspects of the college experience that students may value, including the quality of faculty, the location and amenities of the campus, and the alumni network. The various aspects are aggregated in our model into a unidimensional prize value. Robustness

[25]Note that this robustness is different from the one usually studied in the mechanism design literature, which typically concerns the underlying information structure.

frees the analyst from making assumptions about the details of the aggregation.[26] These results may therefore be particularly useful for empirical work.[27] However, our model still entails the important assumption that students agree on the ranking of colleges.[28]

Our main robustness result characterizes the Pareto frontier of robust Pareto improving category rankings and shows that it is a singleton. We begin with a definition and some preliminary results.

Definition 1. *Distribution F truncated below x^* and above x^{**} FOSD (first-order stochastically dominates) the uniform distribution truncated below x^* and above x^{**} if*

$$\frac{F(x) - F(x^*)}{F(x^{**}) - F(x^*)} \leq \frac{x - x^*}{x^{**} - x^*} \tag{16}$$

for every x in $[x^*, x^{**}]$.

Distribution F truncated below x^* (truncated above x^{**}) FOSD the uniform distribution truncated below x^* (truncated above x^{**}) if (16) holds for every x in $[x^*, 1]$ (for every x in $\in [0, x^{**}]$).[29]

The first result characterizes robust Pareto improving top pooling thresholds.

Corollary 2. *Type x^* is a robust Pareto improving top pooling threshold if and only if the distribution F truncated below type x^* FOSD the uniform distribution truncated below this x^*.*

Corollary 2 follows from part (d) of Proposition 1. Since $y^A(\tilde{x}) = G^{-1}(F(x))$ can be an arbitrary increasing function with values in $[0, 1]$ for an appropriate G, (11) is one of the equivalent conditions that define (conditional) FOSD.

Consider two robust Pareto improving top pooling thresholds $x_1^* < x_2^*$. By definition of top pooling, the effect of top pooling with threshold x_2^* is identical to the effect of using

[26]We are grateful to a referee for suggesting this important comment.

[27]An estimate of the distribution F is still required for all of our results.

[28]We also recall out that while one need not know c, it is assumed that c is strictly increasing, twice continuously differentiable, and $\lim_{t\to\infty} c(t) = \infty$; distribution G can be an arbitrary empirical distribution of prizes, which need not be continuous or have full support; and F is assumed only to have a strictly positive density on $[0, 1]$.

[29]This definition uses our assumption that F is distributed on $[0, 1]$. More generally, 0 and 1 must be replaced by the bounds of the support of F.

a baseline contest and modifying the prize distribution by changing the prizes allocated to types $(x_2^*, 1]$ from what is specified by G to a mass $1 - F(x_2^*)$ of the average (according to distribution G) of these prizes. Then, by Corollary 2, top pooling with threshold x_1^* leads to a further robust Pareto improvement. This proves the following result.

Corollary 3. *If $x_1^* < x_2^*$ are robust Pareto improving top pooling thresholds, then top pooling with threshold x_1^* is robust Pareto preferred to top pooling with threshold x_2^*. Thus, the Pareto frontier of robust Pareto improving top pooling thresholds is a singleton, which is the lowest robust Pareto improving top pooling threshold.*

Corollary 3 explains why in Example 1 lower top pooling thresholds are Pareto preferred to higher ones, and why a lottery is Pareto preferred to any positive top pooling threshold.[30]

The next result characterizes robust Pareto improving single-interval category rankings. It follows immediately from part (b) of Proposition 2 and the definition of FOSD.

Corollary 4. *Category ranking $\mathcal{I} = \{(x^*, x^{**}]\} \cup \{\{x\} : x \leq x^* \text{ or } x > x^{**}\}$ is robust Pareto improving if and only if distribution F truncated below x^* and above x^{**} FOSD the uniform distribution truncated below x^* and above x^{**}.*

Consider two robust Pareto improving single-interval category rankings $\mathcal{I}_1$ and $\mathcal{I}_2$, with corresponding pooled type intervals $I_1 = (x_1^*, x_1^{**}]$ and $I_2 = (x_2^*, x_2^{**}]$. Suppose $I_1 \subseteq I_2$. The effect of pooling on interval I_1 is identical to the effect of using a baseline contest and modifying the prize distribution by changing the prizes allocated to types $(x_1^*, x_1^{**}]$ from what is specified by G to a mass $F(x_1^{**}) - F(x_1^*)$ of the average of these prizes. By Corollary 4, applying the category ranking $\mathcal{I}_2$ leads to a further robust Pareto improvement. Thus, $\mathcal{I}_2$ is robust Pareto preferred to $\mathcal{I}_1$.

Similarly, if I_1 and I_2 are disjoint, then the two-interval category ranking with pooled type intervals I_1 and I_2 is robust Pareto preferred to $\mathcal{I}_1$ and $\mathcal{I}_2$. Finally, suppose that I_1 and I_2 intersect. In this case, the following lemma and our result for the case $I_1 \subseteq I_2$ imply that the single-interval category ranking with pooled type interval $I_1 \cup I_2$ is robust Pareto preferred to $\mathcal{I}_1$ and $\mathcal{I}_2$.

[30] Example 2 fails this condition, because F does not FOSD the uniform distribution on $[0, 1]$.

Lemma 1. *Consider two intervals $I_1 = (x_1, x_3]$ and $I_2 = (x_2, x_4]$ for $0 \leq x_1 < x_2 < x_3 < x_4 \leq 1$. If for each interval I_1 and I_2, F restricted to the interval FOSD the uniform distribution restricted to the interval, then F restricted to the union $I_1 \cup I_2$ of the intervals FOSD the uniform distribution restricted to $I_1 \cup I_2$.*

The next proposition, which is our main robustness result, is an immediate consequence of Corollary 4 and the purely statistical observation in Lemma 1.

Proposition 4. *The Pareto frontier of robust Pareto improving category rankings is a singleton $\mathcal{I}^{PF}$, which consists of the maximal intervals such that F restricted to each interval FOSD the uniform distribution restricted to the interval, along with singletons for all other types.*

The following example illustrates Proposition 4.

Example 4. *Let F have density $f = 4/3$ on interval $[0, 1/4]$, $f = 2/3$ on interval $(1/4, 1/2]$, $f = 1/2$ on interval $(1/2, 3/4]$, and $f = 3/2$ on interval $(3/4, 1]$. (Notice that for robust Pareto improvements we do not specify functions c and G.) Then, F restricted to interval $[1/2, 1]$ FOSD the uniform distribution restricted to the same interval, because the former has an increasing density and the latter has a constant density.*

On interval $[1/4, 1]$, the uniform distribution has density $4/3$, and F restricted to this interval has density 1 on interval $(1/4, 1/2]$, density $3/4$ on interval $(1/2, 3/4]$, and density $9/4$ on interval $(3/4, 1]$. So F also FOSD the uniform distribution, when both are restricted to interval $(1/4, 1]$.

Distribution F does not, however, FOSD the uniform distribution when both are restricted to a longer interval that contains $[1/4, 1]$, because their densities on interval $[0, 1/4]$ are $4/3$ and 1, respectively. In addition, F "weakly" FOSD the uniform distribution when both are restricted to interval $[0, 1/4]$. Thus, by Proposition 4, the unique robust Pareto improving category ranking on the Pareto frontier consists of two intervals: $[0, 1/4]$ and $(1/4, 1]$. It is easy to construct examples in which a single Pareto-frontier category ranking consists of any finite number of intervals.

7 Empirical Application

In this section, we demonstrate how to apply our theory in a college admissions setting to obtain Pareto improvements. Our goal is to provide a “proof of concept” using data in a concrete setting, while recognizing that many factors beyond our theoretical model will need to be considered when making actual policy recommendations. We consider Turkish college admissions, which are based on a centralized test, and use data on college applications of Turkish high school students as used by Krishna et al. (2018).[31]

Our empirical application includes two parts, calibration and simulation. The calibration uses the analysis of Krishna et al. (2018) to estimate the distribution of prizes (college seats) and student types. An important part of their analysis considers test re-taking behavior, which is very useful for our purposes because it is instrumental in estimating the payoff from obtaining any particular score in the exam. The second part of our exercise simulates students' behavior in a one-stage college-admissions setting based on the estimated distributions that corresponds to our theoretical model. We then derive the maximal Pareto improving “bottom-pooling” policy, in which some fraction of the lowest-performing students are pooled together.[32] An attractive feature of this maximal interval is that it does not change the equilibrium behavior and prize allocation of students with types above this interval. Section 8 conducts an experiment based on the estimated distributions and the Pareto improving bottom pooling policy to check its effect on actual subjects.

7.1 Calibration

For the calibration, we rely on the model estimated by Krishna et al. (2018). Players' payoff in that model share essential features with (3), where the net payoff is $y - c(t)/x$: players have the same prize valuations but different costs that depend on their ability.[33] In order to design a Pareto improving policy, we need two primitives: the distribution of prizes, $G(y)$

[31] See Krishna et al. (2018) for details regarding the data and the university entrance exam system in Turkey. We used the survey of exam applicants and the administrative data on exam performance (OSYM (2002c) in conjunction with the official publications of exam rules and results (OSYM (2002) and OSYM (2002b)).

[32] No significant top pooling policy is Pareto improving.

[33] Appendix E presents more details on the relationship between the two models.

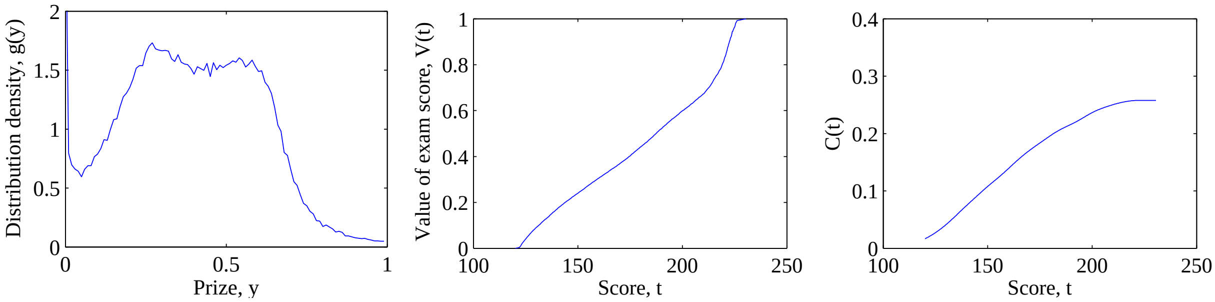

(a) Prize distribution density, $g(y)$. *(b) The value of score, $V(t)$.* *(c) Realized cost conditional on placement score, $C(t) = c(t)/x^A(t)$.*

Figure 4: Key inputs from Krishna et al. (2018) used to calibrate the model.

and the distribution of student abilities, $F(x)$.

We obtain $G(y)$ for the Turkish college applicants from Krishna et al. (2018), who estimate the value of admission to a college as a function of one's rank in the exam. The inverse of this value function is precisely $G(y)$, which maps the value of a particular placement, y, to one's percentile rank in the exam. Intuitively, Krishna et al. (2018) use test retaking decisions to pin down the shape of $G(y)$: given the potential improvement in score from retaking the test, the more sharply the value function rises with score, the more attractive retaking becomes. The estimated density of prizes, $g(y)$, is depicted in Figure 4a.

The second primitive, $F(x)$, is backed out from the fact that scores are produced by investment and ability so as to maximize the value of the score minus its cost. For this, we derive estimates of the cost of investment, $C(t)$, for each score t observed in the data. Let $x^A(t)$ denote the type of students who attain score t. Recall that, in equilibrium, each score is attained by a unique type so that

$$C(t) = \frac{c(t)}{x^A(t)}. \tag{17}$$

We also need the value, $V(t)$, of obtaining score t in the college entrance exam. Function $V(t)$ is the value of placement y, including the option of retaking the exam, net of any associated costs. Details on how $V(t)$ is obtained are in Appendix E. Figure 4b plots the estimate of $V(t)$ from Krishna et al. (2018).

Once we have these two functions, we note that optimality of effort implies that:

$$\frac{c'(t)}{x^A(t)} = V'(t). \tag{18}$$

Differentiating (17) and substituting for $\frac{c'(t)}{x^A(t)}$ using (18) gives

$$C'(t) = \frac{c'(t)}{x^A(t)} - \frac{c(t)}{x^A(t)^2}\frac{dx^A(t)}{dt} = V'(t) - \frac{C(t)}{x^A(t)}\frac{dx^A(t)}{dt}$$

After re-arranging terms, we obtain a differential equation with $x^A(t)$, the equilibrium mapping from score to ability, as the unknown function:

$$\frac{dx^A(t)}{dt} = x^A(t)\frac{V'(t) - C'(t)}{C(t)}$$

We integrate this equation numerically to obtain $x^A(t)$ using the fact that the highest-ability student gets the highest score. Once we have $x^A(t)$, we invert this function to get the distribution of ability since we have the distribution of scores in the data.[34]

The estimate of expended costs conditional on exam score, $C(t)$, is obtained in Krishna et al. (2018) using data on pre-exam schooling investments. In the Turkish context, each middle-school student chooses between public, private, and exam high school, with or without extra preparatory courses, which results in six possible investment levels. Investments are costly: selective schools require entrance exams of their own, while private schools and preparatory courses charge tuition.[35] The cost of choosing each option is modeled and estimated using a mixed logit model in which the cost has a random component with a specified distribution. Further details are in Appendix E.

Because we observe the investment choices for each student, we can impute his costs using the cost estimates. We also know his score, so averaging over imputed costs for each

[34] The distribution of ability F can be expressed via the observed c.d.f. of scores, $H(t)$, and the inverse function $t^A = \left(x^A\right)^{-1}$: $F(x) = \Pr\{X < x\} = \Pr\{t^A(X) < t^A(x)\} = \Pr\{t < t^A(x)\} = H(t^A(x))$.

[35] There is a clear relationship between pre-test investment and test outcomes. Higher scoring students are more likey to come from more selective schools and to have taken preparatory courses. Low scoring students overwhelmingly come from public schools and are unlikely to take preparatory courses. Students coming from private schools on average score in the middle of the distribution. Figure F2 in Appendix F illustrates these patterns.

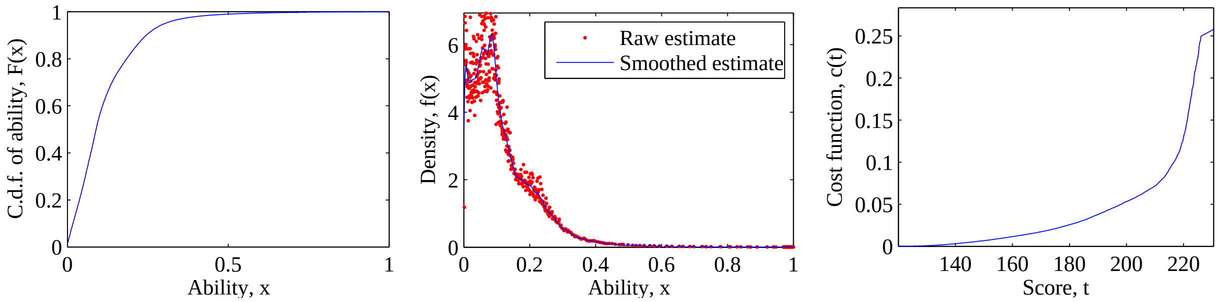

(a) Distribution function, $F(x)$.

(b) Density, $f(x)$.

(c) Cost function $c(t)$

Figure 5: Estimated elements of the model

score we can obtain the mean investment cost for each score in the data.[36] The resulting estimate of $C(t)$ is depicted in Figure 4c.

The calibrated primitives of the model are presented in Figure 5. Panels 5a and 5b plot the ability distribution $F(x)$ and its density, $f(x)$. As the estimated density is somewhat noisy, we plot its kernel-smoothed version in the same figure.[37] From the estimated $C(t)$ and the mapping from score to ability we obtain the estimate for $c(t)$ shown in Figure 5c.

7.2 Simulating a Pareto Improving Pooling Policy

We now simulate a college admissions contest based on the calibrated ability and prize distributions, and identify a Pareto improving policy. As in our theoretical model, we shut down test retaking and use the distribution of prizes $G(y)$, which does not include the value of retaking, so that players face the rank-to-prize mapping described by the density function in Figure 4a. Since shutting down test retaking alters the returns to effort relative to the setting of Krishna et al. (2018), we first use the characterization in Section 3 to simulate the equilibrium effort, allocation, and payoffs in the baseline contest without any pooling. This serves as a benchmark for the effects of pooling.

[36]This removes the extra dimensions of student heterogeneity that are present in Krishna et al. (2018), but are not included in this paper.

[37]The distribution of ability is skewed towards zero. This is driven by two patterns in the target moments in the calibration exercise. First, the steep slope of the value function $V(t)$ (see Figure 4b) implies that the cost of attaining top scores is also steeply sloped. A top student can drastically reduce the cost by performing slightly below the top score. At the same time, the expended cost profile $c(t)/x^A(t)$ is nearly flat for high-scoring agents as shown in Figure 4c: students at the top and those slightly below the top make similar pre-exam investments. These two patterns can only be explained together if ability $x^A(t)$ varies a lot in the top percentiles of the score distribution, which means that the ability distribution has a thin right tail as seen in Figure 5b.

We search for a maximal Pareto-improving bottom-pooling policy using the results in Proposition 2. Bottom-pooling policies are single-interval category rankings that pool some fraction of the lowest-performing players. We focus on a maximal bottom-pooling policy for several reasons. First, since our main goal is a proof of concept, we restrict attention to single-interval policies to keep the exercise relatively simple.[38] Second, the fact that the ability distribution is skewed toward zero and has a thin right tail makes top pooling unattractive as a Pareto improving policy. Finally, by choosing a maximal Pareto improving bottom pooling interval we guarantee that types above the interval are not affected by the policy (since the type at the top of the interval is indifferent). This simplifies the analysis and facilitates the experiment in Section 8.

We set the lower bound of the pooling interval, x^*, at zero and gradually increase the upper bound, x^{**}, in small increments until no further increase in x^{**} leads to a Pareto-improving policy. The resulting maximal Pareto improving bottom pooling policy pools all types below $x^{**} = 0.103$. This ability range encompasses roughly 58% of all applicants in the data, as evident from the ability distribution depicted in Figure 5a.

Figure 6a shows types' payoffs, $y(x)-c(t(x))/x$, under bottom pooling and in the baseline contest. Figure 6b shows the equilibrium scores under the two policies. Bottom pooling strictly increases the payoff of the types in the pooled interval because pooling induces these types to reduce their investment while still obtaining one of the pooled college seats. Higher types, those above the pooled interval, are not affected since they are indifferent between bottom pooling and the baseline contest. Overall, the mean payoff increases by 27%, while the pooled types gain 83% on average.

8 Experimental Evaluation

This section describes a laboratory experiment based on a discretized version of the calibration exercise and the Pareto improving bottom pooling policy identified in Section 7.

[38] We also search for the optimal robust Pareto improving category ranking characterized in Proposition 4. The policy we find resembles bottom pooling. While it prescribes using many pooling intervals, the biggest one lumps roughly a half of the student population at the lower end of the ability distribution. More details on the robust policy are available in Appendix F.

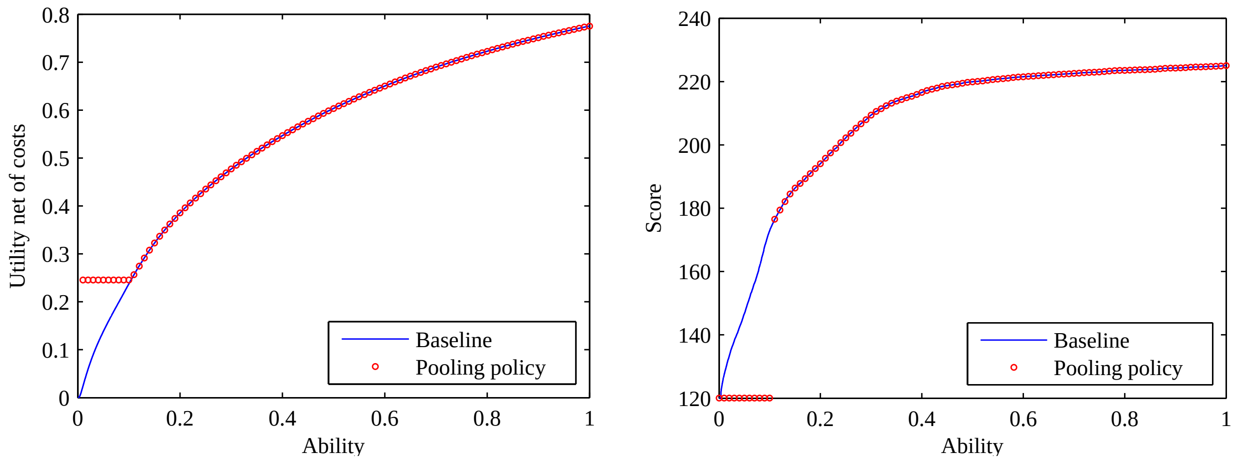

(a) *Net utility,* $y - c(t)/x$ (b) *Scores*

Figure 6: Equilibrium payoffs and effort under assortative matching and bottom pooling.

To conduct the experiment, we transformed the game into an individual decision-making problem without strategic uncertainly. We specified the costs of reaching particular exam scores and the benefit of enrolling in various colleges (the prizes). These costs and benefits, exogenous from the subjects' point of view, correspond to a situation in which the admission criteria are known in advance, as is often the case in college admissions settings that involve a large number of applicants and therefore entail little uncertainty.[39] The goal of the experiment was to investigate whether the Pareto improving policy from Section 7 would lead to improvements in practice. We also wanted to see when deviations from the theoretical predictions might occur and what would be the welfare implications of such deviations.

8.1 Experimental Design

This section describes the experiment. Appendix I contains all the experimental materials subjects faced. The experiment had two parts, each with multiple rounds. Subjects' decisions in any round did not affect the choices they faced in other rounds. Payment was based on subjects' decision in a randomly chosen round.

Part 1: the college admissions task. Part 1 was the main part of the experiment and consisted of two rounds. Each round corresponded to a college admissions setting, one

[39]This is the case, for example, in Turkey, Israel, and many colleges in the United States.

round with a bottom pooling policy (the "pooling policy" round) and one round without any pooling (the "discrete policy" round). Each subject participated in both rounds, and the order of the rounds was randomly determined.

In each round, subjects made an investment choice that determined which college they would enroll in. Each round had ten colleges, labeled College A (best), College B, etc. up to College J (worst). The payoff associated with enrolling in each college was fixed for both rounds, with that of College A being the highest, followed by that of College B, etc. Prior to the first round, each subject was assigned an "ability" in the form of an investment cost for each college. This ability remained fixed for both rounds. The investment costs were denominated in "tokens," did not exceed 100, and subjects were given a 100-token endowment in each round. The payoffs, costs, and abilities, were derived from a discrete version of the calibration and equilibrium results of Section 7.[40]

In each round, subjects decided how much to invest in "virtual study materials." A subject's investment determined the college that subject enrolled in, which determined payment if that round was chosen for payment. In the pooling round, a subject who chose to invest zero tokens in study materials participated in a lottery that randomized among the bottom six colleges (E to J) to determine the college in which the subject enrolled, with the associated payment.[41] Figure 7a presents the expected profits predicted by the theory in both rounds for each ability level. The overall weighted profits, using the appropriate weights for each ability level, are predicted to be 20% higher under the pooling policy than under the discrete policy. The predicted increase for low ability subjects is 65.7%.

Instructions for the second round were given after the first round was completed. After reading the instructions for any given round, subjects had to answer three quiz questions that tested their understanding of the instructions as well as their ability to calculate payoffs. Only subjects who answered the quiz questions correctly could move on to the the decision-

[40]We included eleven ability levels. Ability levels and cost functions, as well as a description of how the empirical data from Section 7 was used to determine the experimental parameters, are in Appendix G.

[41]To isolate the possible effects of the lottery, we also ran sessions in which a subject who invested zero tokens in study materials got a fixed payment equal to the expected value of the lottery. This corresponds to a pooling round in which the lottery is replaced with its expected value. The results from these sessions are in Appendix H.

making component of that round.[42]

Part 2: risk elicitation. Following the main part of the experiment, subjects participated in a series of ten rounds in which they had to choose between a fixed amount and a risky gamble.[43] This task was included in order to identify subjects whose choices were consistent with risk neutrality, an important element for the analysis of the pooling-induced lottery.

Procedures. We recruited subjects using Prolific, an online platform that recruits subjects worldwide. A total of 602 subjects completed the experiment.[44,45] We restricted our subjects to be English speakers. Overall 48.8% of the subjects were female, 47.7% were male, and 2.5% identified as neither of the two.[46] Data were collected during the month of March 2022. The experiment lasted between 10 and 15 minutes for 95% of the subjects. Subjects were paid an average of $4.00, corresponding to a rate of just over $19 per hour.[47]

8.2 Experimental Results

Below we analyze the data collected in the experiment. All the p-values reported for comparing data across the two rounds are the result of Wilcoxon matched-pairs signed-rank tests. We focus our analyses on the subjects whose choices in the risk elicitation task are consistent

[42]After the instructions in a given round, subjects were allowed to take the related quiz twice. If by the second attempt a subject failed to answer all three questions correctly they were removed from the experiment. The quiz questions serve both as a tool to exclude bots from our data and to ensure proper reading of the instructions. Overall, 76% of the subjects who started our experiment answered all questions in both quizzes correctly, among which about three quarters did so on the first attempt for each of the two quizzes.

[43]The risky gamble was identical in all rounds: subjects choosing the risky gamble would earn $1 with probability 1/2 and $2 with probability 1/2. The fixed amount varied from round to round, ranging from $1.25 to $1.75 in increments of 5 cents.

[44]For any particular ability level, we stopped collecting data once we had at least 50 observations.

[45]An additional 208 "participants" started the experiment but failed to advance because quiz questions were answered incorrectly. This proportion (25.7%) is not particularly noteworthy – aside from ensuring that participants have read and understood the instructions, the quizzes have a second purpose: ensuring that bots are unlikely to make it through to the main part of the study.

[46]The remaining 1% of the subjects preferred not to answer.

[47]The Prolific platform requires a minimum of $6.50 per hour. All subjects who completed the experiment were asked how much they earned per hour on average on Prolific. The average response was $9.65. The payments in this experiment were thus relatively high, almost triple the minimum required and double what subjects had earned in past studies.

with risk neutrality.[48] Additional results on non-risk-neutral subjects are in Appendix H. In reporting aggregate results we use the estimates from Section 7 to determine the appropriate weights for each ability level.

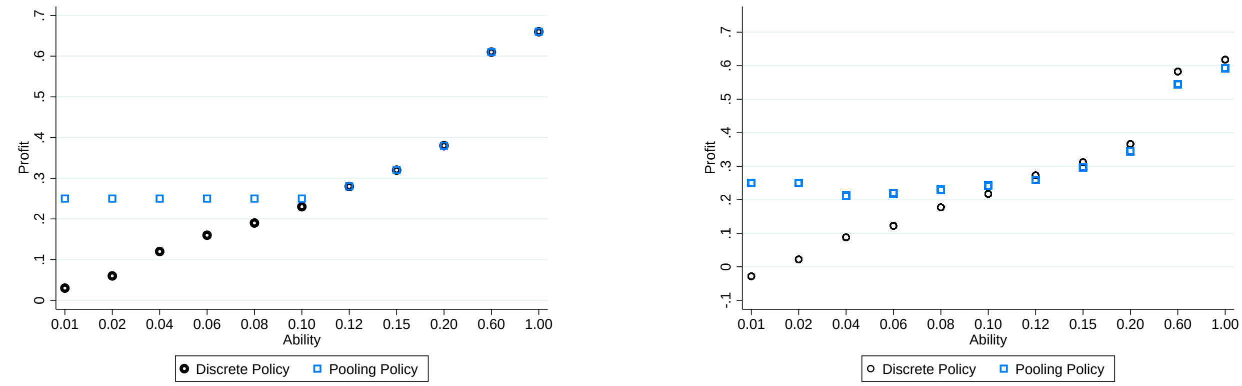

(a) Predictions *(b) Laboratory data*

Figure 7: Average profits under the discrete and pooling policies: predictions and experimental data.

Figure 7 shows the average net profits in the pooling policy and discrete policy rounds.[49] Overall, moving to the pooling policy led to a increases of 18.9% in aggregate profits, closely matching the theoretical prediction of 20.0%. Our analyses below focus on outcomes by ability level, allowing us to evaluate whether such a policy would be (weakly) helpful for each ability level.

We define "low-ability" subjects as those who are predicted to invest zero and choose the lottery in the pooling round, and therefore benefit from the pooling policy. With our parameters, these are subjects with ability of at most 0.10. Subjects with ability strictly greater than 0.10 are "high-ability" subjects. According to the theory, high ability subjects are indifferent between the pooling and discrete policies.

[48] Risk neutrality allows us to apply the model and its predictions with players' utilities equalling their potential earnings. Risk neutral subjects were identified using data from the Risk Elicitation task. Risk aversion can be determined by identifying at which point a subject switches away from the risky gamble to the fixed amount (84.6% of our subjects had a single cross-over point). Risk neutral subjects are identified as those who chose the risky gamble ($1 with probability 1/2 and $2 with probability 1/2) over the fixed amount when the fixed amount was strictly less than $1.50, but switched to the fixed amount when it was $1.50 or $1.55, indicating an indifference point "close to" the gamble's expected value of $1.50. Roughly 44% of the subjects with a single cross-over point behaved in this way. Over 85% of the remaining subjects who had a single crossover point from the lottery to the fixed amount were risk averse.

[49] These are profits net of the 100 token endowment. To make comparisons with the theory easier we re-scaled values to be between 0 and 1.

For low ability subjects, on average, moving to a pooling policy increased aggregate profits by 70.5%, more than the theoretical prediction of 65.7%.[50] Moreover, among low ability subjects, profits also increases for each ability level (the p-values in all pair-wise comparisons are all strictly lower than 0.001).

For high ability subjects, the theory predicts no change in aggregate profits across policies. In the experiment, moving to a pooling policy slightly decreases the aggregate profits (about 1.6%). This small drop is the result of the behavior of subjects whose ability is lowest among the high ability subjects, which in our setting corresponds to an ability of 0.12.[51] The behavior of these subjects was statistically different under the two policies, and the magnitude of the difference is large. Indeed, under the discrete policy, 89.5% of these subjects chose investment levels consistently with the theoretical predictions of the model. This fraction dropped to 36.8% under the pooling policy, with the remaining subjects opting for the lottery. This represents an aggregate profit loss of 5.1%.

Our experimental findings suggest that introducing the bottom pooling policy derived in Section 7 adapted to a college admissions setting with discretized ability and prize distributions should weakly or strictly improve the welfare of over 85% of the applicants.[52] But the pooling policy could make the remaining 15% of the applicants somewhat worse off, because the high ability subjects with the lowest ability did worse monetarily under pooling than under the discrete policy.

To evaluate the potential drawbacks of a pooling policy, we ask whether these slightly lower payoffs are the result of mistakes or a reflection of subjects' preferences. We argue that mistakes are an unlikely explanation. While subjects whose ability is closest to the threshold separating low and high ability are precisely those who one might think are most likely to make mistakes, we observed no such "mistakes" for low-ability subjects with the

[50]This is because under the discrete policy relatively fewer low ability subjects chose the profit-maximizing investment levels and their mistakes were larger in terms of forgone profits.

[51]All other high ability subjects behaved similarly across the two policies. The p-value comparing behavior across the two policies for subjects with ability 0.12 is $p = 0.006$, and the lowest p-value for all other comparisons is 0.125.

[52]In the experiment, subjects whose abilities are weakly less than .10 or weakly greater than 0.15 are in this category. Mapped to the Turkish student population, this corresponds to about 84% of individuals. This is a lower bound, since in the experiment ability jumps from 0.10 to 0.12 to 0.15, so we can only infer that the "turning point" is an ability between 0.12 and 0.15 but are unable to provide a more precise value.

highest ability, that is, subjects with an ability of 0.10.[53] Thus, deviations from predicted behavior are asymmetric around the threshold and confined to high-ability subjects with the lowest ability. This asymmetry, and the lack of such deviations at other ability levels, do not favor mistakes (or inattention) as a likely cause.[54,55]

One possible explanation for these deviations is that some subjects display a preference for randomization. This is the case for example in Dwenger, Kübler, and Weizsäcker (2016), who use both laboratory and non-laboratory data (from a clearinghouse for university admissions in Germany) to show that up to 50% of individuals choose lotteries between available allocations, indicating an explicit preference for randomization.[56] Agranov and Ortoleva (2017) show that when faced with "hard choices" a significant fraction of the population may prefer a lottery to making a deterministic choice. In our experiment, the difference between the expected net profit from the lottery and the theoretically predicted choice is the smallest among all subjects for precisely the subjects who deviate from the theoretical predictions, consistent with choosing a investment being "hard" for them.[57] We also point out that the optimal choice for subjects with ability 0.12 in the discrete round was still available in the pooling round, so for a wide range of preferences, subjects with ability 0.12 who switched to the lottery under the pooling policy round were likely made better off even if their monetary payoff decreased slightly.

Taken together, the results of the laboratory experiment and the empirical estimation of Section 7 suggest that Pareto improving policies of the kind we investigate likely exist in practice, can be identified and implemented, and have the potential to improve the welfare

[53]This is despite the fact that the difference in net profit between the two policies for subjects with ability 0.10 is even narrower than it is for subjects with ability 0.12.

[54]We also rule out order effects as there is no statistical difference between the groups who saw the pooling policy first and those who saw the discrete policy first ($p = 0.764$). This also rules out experimenter demand effects because these patterns exist also with subjects who saw the pooling policy first.

[55]Risk-seeking behavior is also an unlikely cause because we already restrict attention to players who appear to be risk-neutral. One caveat is that the coarseness of our measure of risk aversion may not identify mildly risk seeking subjects who would choose the lottery over the fixed amount. This explanation, however, would imply a relatively large fraction of subjects with such preferences, which is inconsistent with past work on risk aversion elicitation (see Holt and Laury (2002) for example).

[56]The authors discuss this in the context of responsibility aversion.

[57]Appendix H shows that if we remove the option to randomize by replacing the lottery with a fixed amount equal to the lottery's expected value, subjects of all ability levels make the profit maximizing choice under the pooling policy.

of millions of students who apply to colleges and universities every year.

9 Conclusion

This paper investigated how to improve college admissions based on centralized tests. Students engage in test-preparation activities to improve their ranking, but these activities are costly. Our main message is that coarse performance disclosure policies can benefit all students, regardless of their ability. These policies take a simple form and are easy to implement.

As a "proof of concept," we empirically estimated the key theoretical constructs, ability and prize distributions, using data on college admissions in Turkey. We used our theoretical results to simulate the equilibrium outcome of a college admissions contest based on these distributions, and demonstrated how to identify Pareto improving policies. We showed that a bottom pooling policy that pools together the majority of the students and randomly allocates them to the corresponding colleges would raise the welfare of these students without impacting the welfare of the other students. Finally, we conducted a laboratory experiment based on these empirical findings. The results of the experiment largely confirmed our theoretical predictions. Overall, our work suggests that Pareto improving performance disclosure policies of the kind we investigated often exist and have the potential to improve college admissions systems in practice.

References

[1] Ahn S-Y and Baek H-J (2013): "Academic achievement-oriented society and its relationship to the psychological well-being of Korean adolescents," In: Yi CC (eds) The psychological well-being of East Asian youth, quality of life in Asia, 265–279.

[2] Bodoh-Creed, A. and B. Hickman (2018): "College Assignment as a Large Contest," Journal of Economic Theory, 175: 88-126.

[3] Bodoh-Creed, A. and B. Hickman (2019): "Pre-College Human Capital Investment and Affirmative Action: A Structural Policy Analysis of US College Admissions," Mimeo.

[4] Boleslavsky, R. and C. Cotton (2015): "Grading Standards and Education Quality," American Economic Journal: Microeconomics, 7, 248–279.

[5] Chakravarty, S. and T. R. Kaplan (2013): "Optimal allocation without transfer payments," Games and Economic Behavior, 77(1), 1-20.

[6] Chao, H.-P, and R. Wilson (1987): "Priority Service: Pricing, Investment, and Market Organization," American Economic Review, 77(5), 899-916.

[7] Che, Y.-K., Condorelli, D., and J. Kim (2018): "Weak cartels and collusion-proof auctions," Journal of Economic Theory, 178, 398-435.

[8] Che, Y.-K. and I. Gale (1998): "Caps on Political Lobbying," American Economic Review, 88, 643-651.

[9] Condorelli, D. (2012): "What money can't buy: Efficient mechanism design with costly signals," Games and Economic Behavior, 75(2), 613-624.

[10] Damiano, E. and H. Li (2007): "Price discrimination and efficient matching", Economic Theory 30, 243-263.

[11] Dubey, P. and J. Geanakoplos (2010): "Grading exams: 100, 99, 98, . . . or A, B, C?" Games and Economic Behavior, 69, 72-94.

[12] Fang, D. and T. Noe (2018): "Lowering the bar and limiting the field: The effect of strategic risk-taking on selection contests," Mimeo.

[13] Fang, D., T. Noe, and P. Strack (2018): "A is the new B: Contests for grades under rational expectations," Mimeo.

[14] Fang, D., T. Noe, and P. Strack (2020): "Turning Up the Heat: The Demoralizing Effect of Competition in Contests," Journal of Political Economy, 128(5), 1940-1975.

[15] Frankel, A. and N. Kartik (2019): "Muddled Information," Journal of Political Economy, 127(4), 1739-1776.

[16] Gale, D. and L. Shapley (1962): "College Admissions and the Stability of Marriage," American Mathematical Monthly, 69(1), 9-15.

[17] Gottlieb, D. and K. Smetters (2014): "Grade Non-Disclosure," Mimeo.

[18] Hafalir, I., Hakimov, R., Kübler, D., and M. Kurino (2018): "College Admissions with Entrance Exams: Centralized versus Decentralized," Journal of Economic Theory, 176: 886-934.

[19] Harbaugh, R. and E. Rasmusen (2018): "Coarse grades: Informing the public by withholding information," American Economic Journal: Microeconomics, 10(1): 210-35.

[20] Holt, C.A., and S.K Laury (2002): "Risk aversion and incentive effects." American Economic Review 92 (5), 1644–1655.

[21] Hoppe, H., Moldovanu, B., and E. Ozdenoren (2011): "Coarse matching with incomplete information," Economic Theory, 47(1), 75-104.

[22] Hoppe, H., Moldovanu, B., and A. Sela (2009): "The Theory of Assortative Matching Based on Costly Signals," Review of Economic Studies, 76, 253-281.

[23] Hsin, A. and Y. Xie (2014): "Explaining Asian Americans' academic advantage over whites," Proceedings of the National Academy of Sciences, 111(23), 8416–8421.

[24] Koh, W. T., Z. Yang and L. Zhu (2006): "Lottery rather than waiting-line auction," Social Choice and Welfare, 27, 289-310.

[25] Krishna, K., Lychagin, S., and V. Frisancho (2018): "Retaking in High Stakes Exams: Is Less More?," International Economic Review, 59, 449-477.

[26] Lien, J., J. Zheng, and X. Zhong (2017): "Ex-ante fairness in the Boston and serial dictatorship mechanisms under pre-exam and post-exam preference submission," Games and Economic Behavior, 100, 98-120.

[27] McAfee, P. (2002): "Coarse matching," Econometrica, 70(5), 2025-2034.

[28] McIntyre, S. and J. Munson (2008): "Exploring Cramming: Student Behaviors, Beliefs, and Learning Retention in the Principles of Marketing Course," Journal of Marketing Education, 30(3), 226-243.

[29] Moldovanu, B., Sela, A., and X. Shi (2007): "Contests for status," Journal of Political Economy, 115(2), 338-363.

[30] Morgan J., Tumlinson J., and F. Vardy (2022): "The Limits of Meritocracy," Journal of Economic Theory 201, Article 105414.

[31] Olszewski, W. and R. Siegel (2016): "Large Contests," Econometrica 84, 835-854.

[32] Olszewski, W. and R. Siegel (2020): "Performance-Maximizing Large Contests," Theoretical Economics, 15(1), 57-88.

[33] Ostrovsky, M. and M. Schwarz (2010): "Information Disclosure and Unraveling in Matching Markets," American Economic Journal: Microeconomics, 2(2), 34-63.

[34] Rayo, L. (2013): "Monopolistic signal provision," B.E. Journal in Theoretical Economics, 13(1), 27-58.

[35] Spence, M. (1973). "Job Market Signaling." Quarterly Journal of Economics, 87, 355-374.

[36] OSYM (2002a). Handbook for the ÖSS Exam Applicants. https://www.osym.gov.tr/TR,2193/oss-kilavuzu.html

[37] OSYM (2002b). ÖSS Exam Statistics.

[38] OSYM (2002c). Survey of the Applicants to the University Entrance Examination (ÖSS) and Their Performance.

[39] Taylor, G. A., K. K. Tsui, and L. Zhu (2003): "Lottery or waiting-line auction?" Journal of Public Economics, 87, 1313-1334.

[40] Wilson, R. (1989): "Efficient and competitive rationing," Econometrica, 57(1), 1-40.

Appendices

A Pareto frontier of category rankings

Using Proposition 2, we provide a method for checking whether a category ranking belongs to the Pareto frontier of category rankings. For this, we will need another concept. Let $\mathcal{I}$ be a category ranking, and let $x^* < x^{**}$ be an arbitrary pair of types that belong to two different elements $I \neq I'$ (intervals or singletons) of $\mathcal{I}$, so $x^* \in I \in \mathcal{I}$ and $x^{**} \in I' \in \mathcal{I}$. We define a new category ranking $\mathcal{I}(x^*, x^{**})$ that groups all types between x^* and x^{**} into one category as follows: (i) if $I = (a, b]$ and $I' = (a', b']$, replace I, I', and all elements of $\mathcal{I}$ between I and I' with $(a, x^*]$, $(x^*, x^{**}]$, and $(x^{**}, b']$; (ii) if $I = \{x^*\}$ and $I' = (a', b']$, replace I' and all elements of $\mathcal{I}$ between I and I' with $(x^*, x^{**}]$, and $(x^{**}, b']$; (iii) if $I = (a, b]$ and $I' = \{x^{**}\}$, replace I, I', and all elements of $\mathcal{I}$ between I and I' with $(a, x^*]$ and $(x^*, x^{**}]$; (iv) if $I = \{x^*\}$ and $I' = \{x^{**}\}$, replace I' and all elements of $\mathcal{I}$ between I and I' with $(x^*, x^{**}]$.

Proposition 5. *A category ranking $\mathcal{I}$ belongs to the Pareto frontier of category rankings if* *and only if there is no pair of types $x^* < x^{**}$ such that*

$$x^* = a \text{ for some } I = (a, b] \in \mathcal{I} \text{ or } x^* = d \text{ for some } I = \{d\} \in \mathcal{I} \text{ and } x^{**} \in I' \neq I \in \mathcal{I},$$

and type x^{**} weakly prefers ranking $\mathcal{I}(x^*, x^{**})$ to ranking $\mathcal{I}$.

The proof of Proposition 5 establishes and uses a somewhat involved formula for the utility of various types given a general category ranking. To see the main idea, restrict attention to two-interval rankings, and consider the category ranking $\mathcal{I}$ that consists of $(0, b]$ and $(b, 1]$. Any two-interval category ranking $\mathcal{J}$ that dominates $\mathcal{I}$ must consist of $(0, c]$ and $(c, 1]$, with $c > b$; indeed, if $c < b$, then types close to 0 prefer $\mathcal{I}$ to $\mathcal{J}$. If $\mathcal{J}$ dominates $\mathcal{I}$, then $\mathcal{J} = \mathcal{I}(x^*, x^{**})$ for $x^* = a = 0$ and $x^{**} = c$, so the necessary and sufficient condition in Proposition 5 for $\mathcal{I}$ to belong to the Pareto frontier fails.

Conversely, if $x^{**} = c$ weakly prefers $\mathcal{J} = \mathcal{I}(x^*, x^{**})$ to $\mathcal{I}$, then types smaller than x^{**} prefer $\mathcal{J}$ to $\mathcal{I}$ by the argument that underlies part (d) of Proposition 1 or part (b) of Proposition 2. Types greater than x^{**} prefer $\mathcal{J}$ to $\mathcal{I}$ because the difference between the utility of type $x > x^{**}$ and the utility of type x^{**} under $\mathcal{J}$ increases faster than the same difference under $\mathcal{I}$. This last observation follows, because type x^{**} is indifferent between the two intervals of $\mathcal{J}$; so, if type x^{**} exerted under $\mathcal{J}$ the effort that yields a prize from $(c, 1]$, the efforts of types x and x^{**} in each of the two rankings would be equal, and their prizes would be better, under $\mathcal{J}$ than under $\mathcal{I}$.

The necessary and sufficient condition from Proposition 5 looks cumbersome. However, it helps to characterize the Pareto frontier by substantially reducing the set of category rankings to which any given ranking must be compared, as the following example demonstrates.

Example 5. *Revisit Example 3. It is easy to verify that any interval that satisfies condition (15) must be contained in $(0, 3/4]$ or $(3/4, 1]$. Thus, any candidate for a Pareto-improving category ranking consists of an interval partition of $(0, 3/4]$ and an interval partition of $(3/4, 1]$. But by the general payoff formula (20) in the proof of Proposition 5, if a partition of $(0, 3/4]$ (or a partition of $(3/4, 1]$) includes more than one element, then $(x^*, x^{**}] = (0, 3/4]$ ($(x^*, x^{**}] = (3/4, 1]$, respectively) violates the condition from Proposition 5. Indeed, the* *payoffs for x^{**} are equal under $\mathcal{I}$ and under $\mathcal{I}(x^*, x^{**})$. Thus, the Pareto frontier has only one element, the category ranking* $\{(0, 3/4], (3/4, 1]\}$.

B Proofs

Proof of Proposition 1. Part (a) follows because with top pooling types $x < x^*$ choose effort $t^A(x)$ and obtain prize $y^A(x)$. For part (b), note that the utility of type x^* is the same in the approximating mechanisms of the original contest and in the one with top pooling. Consider first the utility of a type $x > x^*$ in the approximating mechanism of the original contest. By (5), this utility exceeds that of type x^* by

$$\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}.$$

In the approximating mechanism with top pooling, the utility of type x exceeds that of type x^* by

$$(x - x^*)\frac{\int_{x^*}^{1} y^A(\tilde{x})\, dF(\tilde{x})}{1 - F(x^*)},$$

since both types' performance is M, and both types' prize is chosen randomly from the mass $1 - F(x^*)$ of the highest prizes. Thus, top pooling increases the utility of type x if and only if (10) holds.

For part (c), note that the derivative with respect to x of the utility gain of type x is

$$\frac{\int_{x^*}^{1} y^A(\tilde{x})\, dF(\tilde{x})}{1 - F(x^*)} - y^A(x)\,. \tag{19}$$

The fraction in (19) is a weighted average of $y^A(\tilde{x})$ over types in $[x^*, 1]$, so (19) is positive for types x close to x^*, monotonically decreases as x increases, and becomes negative for types x close to 1. Thus, the utility gain resulting from top pooling for types $x > x^*$ first increases and then decreases in the type. In particular, the utility of all types $x < 1$ strictly increases if the utility of type 1 weakly increases, which gives part (d).

Proof of Proposition 2. Let $r^* = F(x^*)$ and $r^{**} = F(x^{**})$. Then, any type $x < x^*$

provides effort $t(x) = t^A(x)$ and obtains prize $y(x) = y^A(x)$. Types in $(x^*, x^{**}]$ provide a certain effort t and obtain a fair lottery over prizes $y \in (G^{-1}(r^*), G^{-1}(r^{**})]$. Since type x^* is indifferent between the two options, we have that

$$x^* y^A(x^*) - c(t^A(x^*)) = \int_0^{x^*} y^A(\tilde{x})\, d\tilde{x} = x^* \frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\, dF(\tilde{x})}{F(x^{**}) - F(x^*)} - c(t).$$

By (5) we have that

$$U^{\mathcal{I}}(x) - U(x) = (x - x^*) \frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\, dF(\tilde{x})}{F(x^{**}) - F(x^*)} - \int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}$$

$$= (x - x^*) \left[\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\, dF(\tilde{x})}{F(x^{**}) - F(x^*)} - \frac{\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}}{x - x^*} \right].$$

Thus, (15) is a necessary and sufficient condition for types x in $(x^*, x^{**}]$ to be better off. To show that (15) is a necessary and sufficient condition for Pareto improvement, observe that types $x \leq x^*$ are indifferent. Any type $x > x^{**}$ obtains prize $y^A(x)$ and has payoff

$$U^{\mathcal{I}}(x) = \int_0^{x^*} y^A(\tilde{x})\, d\tilde{x} + (x^{**} - x^*) \frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\, dF(\tilde{x})}{F(x^{**}) - F(x^*)} + \int_{x^{**}}^{x} y^A(\tilde{x})\, d\tilde{x},$$

which is no lower than $U(x)$ if and only if condition (15) is satisfied.

The last equality is obtained directly from (5) by noticing that the contest under our category ranking $\mathcal{I}$ is equivalent to a baseline contest in which prizes $y^A(\tilde{x})$, for $\tilde{x}$ in $(x^*, x^{**}]$, are replaced with the certainty equivalents of the lottery faced by types $\tilde{x}$ in $(x^*, x^{**}]$ under our category ranking $\mathcal{I}$.

Proof of Proposition 5. It will be helpful to provide first a general formula for the utility of type $x \in [0, 1]$ under category ranking $\mathcal{I}$. This utility exceeds $U(x)$ given by (5) by the expression

$$\sum_{(a,b] \in \mathcal{I}, a < b < x} \left[(b - a) \frac{\int_a^b y^A(\tilde{x})\, dF(\tilde{x})}{F(b) - F(a)} - \int_a^b y^A(\tilde{x})\, d\tilde{x} \right] + \tag{20}$$

$$(x-a)\frac{\int_a^b y^A(\tilde{x})\,dF(\tilde{x})}{F(b)-F(a)} - \int_a^x y^A(\tilde{x})\,d\tilde{x} \text{ for } x \in (a,b] \in \mathcal{I}.$$

This formula follows directly from the fact that types $\tilde{x} \in (a,b] \in \mathcal{I}$ obtain a fair lottery over prizes $y^A(\tilde{x}')$ for $\tilde{x}' \in (a,b]$.

We will first show that when a pair $x^* < x^{**}$ satisfies the condition in Proposition 5, the category ranking $\mathcal{J} = \mathcal{I}(x^*, x^{**})$ Pareto improves over $\mathcal{I}$. Types $x \in [0, x^*]$ are obviously indifferent between the two category rankings. By assumption, the utility of type x^{**} is no lower under $\mathcal{J}$ than under $\mathcal{I}$. We will now show that the utility of types $x \in (x^*, x^{**})$ is strictly higher under $\mathcal{J}$ than under $\mathcal{I}$. Indeed, the derivative on $(x^*, x^{**}]$ of type x's utility under $\mathcal{J}$, $U^{\mathcal{J}}(x)$, is constant and equal to

$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\,dF(\tilde{x})}{F(x^{**})-F(x^*)}.$$

In turn, the derivative on $(x^*, x^{**}]$ of type x's utility under $\mathcal{I}$, $U^{\mathcal{I}}(x)$, is equal to $y^A(x)$ if x does not belong to any non-degenerate interval $(a,b] \in \mathcal{I}$, and is equal to

$$\frac{\int_a^b y^A(\tilde{x})\,dF(\tilde{x})}{F(b)-F(a)}$$

if $x \in (a,b] \in \mathcal{I}$. This means that the derivative increases in x, and increases strictly except on intervals $(a,b] \in \mathcal{I}$. So, $U^{\mathcal{I}}(x)$ is a convex non-linear function. Since $U^{\mathcal{J}}(x)$ is linear on $(x^*, x^{**}]$, $U^{\mathcal{I}}(x^*) = U^{\mathcal{J}}(x^*)$, and $U^{\mathcal{I}}(x^{**}) \leq U^{\mathcal{J}}(x^{**})$, we obtain that $U^{\mathcal{I}}(x) \leq U^{\mathcal{J}}(x)$ for all $x \in (x^*, x^{**}]$, and the inequality is strict for all types $x \in (x^*, x^{**}]$. Similarly, the derivative of $U^{\mathcal{J}}(x)$ on $(x^{**}, b']$ exceeds that of $U^{\mathcal{I}}(x)$ if $a' < x^{**} < b'$ for some $(a', b'] \in \mathcal{I}$, and the two derivatives are equal for $x > b'$, which completes the proof that $\mathcal{J}$ Pareto improves over $\mathcal{I}$.

Suppose now that another category ranking $\mathcal{I}'$ Pareto improves over $\mathcal{I}$. Recall that $\mathcal{I}$ consists of singletons and a finite number of intervals $[x_1, x_1'], (x_2, x_2'], ..., (x_k, x_k']$, with $x_i' < x_{i+1}$. Denote by x' the highest type such that $\mathcal{I}$ and $\mathcal{I}'$ coincide up to x', and suppose that x' is the lower endpoint of an interval $(x_l, x_l']$ in $\mathcal{I}$. (A similar argument to the one that follows applies if x' is a singleton.)

Now, x' must be the endpoint of a non-trivial interval in $\mathcal{I}'$, which we denote by $(x^*, x^{**}]$, where $x' = x^* < x^{**}$. Otherwise, for types x slightly higher than x_l the utility of these types under $\mathcal{I}$ would exceed their utility under $\mathcal{I}'$ by (20). It also cannot be that $x^{**} < x_l'$, since it would then follow from (20) that x^{**} strictly prefers $\mathcal{I}$ to $\mathcal{I}'$.

Thus $x_l' < x^{**}$, and since $\mathcal{I}'$ Pareto improves over $\mathcal{I}$, type x^{**} weakly prefers $\mathcal{I}'$ to $\mathcal{I}$. And since (by (20)) the payoff of type x^{**} under any ranking depends only on the intervals up to the one that contains x^{**}, type x^{**} is indifferent between ranking $\mathcal{I}'$ and ranking $\mathcal{J} = \mathcal{I}(x^*, x^{**})$, and therefore prefers ranking $\mathcal{J}$ to ranking $\mathcal{I}$.

Proof of Lemma 1. We have to show for every $x \in I_1 \cup I_2$ that

$$\frac{F(x) - F(x_1)}{F(x_4) - F(x_1)} \leq \frac{x - x_1}{x_4 - x_1}. \tag{21}$$

Consider first $x \in I_1$. By the definition of FOSD on I_1, we have

$$\frac{F(x) - F(x_1)}{F(x_3) - F(x_1)} \leq \frac{x - x_1}{x_3 - x_1},$$

so for (21) it suffices to show that

$$\frac{F(x_3) - F(x_1)}{F(x_4) - F(x_1)} \leq \frac{x_3 - x_1}{x_4 - x_1}.$$

Suppose instead that

$$\frac{F(x_3) - F(x_1)}{F(x_4) - F(x_1)} > \frac{x_3 - x_1}{x_4 - x_1}. \tag{22}$$

This implies that

$$\frac{F(x_4) - F(x_3)}{F(x_4) - F(x_1)} < \frac{x_4 - x_3}{x_4 - x_1} \tag{23}$$

In addition,

$$\frac{F(x_3) - F(x_2)}{F(x_3) - F(x_1)} = 1 - \frac{F(x_2) - F(x_1)}{F(x_3) - F(x_1)} \geq 1 - \frac{x_2 - x_1}{x_3 - x_1} = \frac{x_3 - x_2}{x_3 - x_1}, \tag{24}$$

which multiplied by (22) gives

$$\frac{F(x_3)-F(x_2)}{F(x_4)-F(x_1)} > \frac{x_3-x_2}{x_4-x_1}. \tag{25}$$

Dividing (23) by (25) we obtain

$$\frac{F(x_4)-F(x_3)}{F(x_3)-F(x_2)} < \frac{x_4-x_3}{x_3-x_2} \Rightarrow \frac{F(x_3)-F(x_2)}{F(x_4)-F(x_3)} + 1 > \frac{x_3-x_2}{x_4-x_3} + 1 \Rightarrow$$

$$\frac{F(x_4)-F(x_2)}{F(x_4)-F(x_3)} > \frac{x_4-x_2}{x_4-x_3} \Rightarrow \frac{F(x_4)-F(x_3)}{F(x_4)-F(x_2)} < \frac{x_4-x_3}{x_4-x_2}.$$

This last inequality is a contradiction, since FOSD on I_2 implies the opposite weak inequality, similarly to (24). Therefore, (21) holds for $x \in I_1$.

Now consider $x \in [x_3, x_4]$. Instead of (21) we will show the equivalent inequality

$$\frac{F(x_4)-F(x)}{F(x_4)-F(x_1)} \geq \frac{x_4-x}{x_4-x_1}.$$

From FOSD on I_2 we have

$$\frac{F(x_4)-F(x)}{F(x_4)-F(x_2)} \geq \frac{x_4-x}{x_4-x_2}.$$

Thus, it suffices to show that

$$\frac{F(x_4)-F(x_2)}{F(x_4)-F(x_1)} \geq \frac{x_4-x_2}{x_4-x_1}.$$

This inequality holds, because otherwise we would have

$$\frac{F(x_2)-F(x_1)}{F(x_4)-F(x_1)} > \frac{x_2-x_1}{x_4-x_1},$$

which would violate (21) for $x = x_2 \in I_1$.

C Peer effects

We can model peer effects in a way that does not change any of our results and requires only a transformation of the prize distribution. The idea is that each student exerts a type-dependent effect on all his peers (those attending the same college), and the effects are additive. We will show that such peer effects fit easily into our framework, as does the change in the endogenous set of peers brought about by pooling. Of course, other ways of modeling peer effects would lead to different impacts of pooling because of the change in peers that pooling induces.

We will consider a limit prize distribution that consists of atoms, where each atom represents a mass of seats in a particular college. Students who attend a particular college experience peer effects from other students attending the same college.

To model this, take an equilibrium of a finite contest and denote by $I(y)$ the set of players admitted to university y (for a particular realization of types and bids). The utility of a player of type x admitted to university y by bidding t is

$$xy + x\frac{\sum_{i\in I(y)} p(x_i)}{|I(y)|} - c(t) = x\underbrace{\left(y + \frac{\sum_{i\in I(y)} p(x_i)}{|I(y)|}\right)}_{\tilde{y}} - c(t),$$

where $p(x)$ captures the peer effect exerted by a player of type x. We refer to $\tilde{y}$ as the effective prize for player i, which is the sum of the value of the college and the average peers effects of the other students attending the college. Note that the effective prize depends on the equilibrium because it is determined by the equilibrium allocation of prizes.

The limit approximating mechanism still implements the assortative allocation. To see this, consider two universities $y < y'$ and the corresponding limit effective prizes $\tilde{y}$ and $\tilde{y}'$. If the allocation of students to universities is such that $\tilde{y} \geq \tilde{y}'$, then all types would make the lower bid to get into $\tilde{y}$. But then ties will be broken randomly, which would generate the same peer effects for both prizes, so $y < y'$ would give us $\tilde{y} < \tilde{y}'$, a contradiction. If $\tilde{y} < \tilde{y}'$, then standard mechanism design results tell us that higher types get higher prizes, so it will be the highest types that are assigned to y', and the lower types will be assigned to y. Thus, in the limit we get the assortative allocation.

This means that in the limit approximating mechanism, for each prize y in the support of the limit prize distribution G we have that the effective prize is

$$\tilde{y} = y + \frac{\int_{x_L^y}^{x_H^y} p(\tilde{x})\, dF(\tilde{x})}{F(x_H^y) - F(x_L^y)} = \frac{\int_{x_L^y}^{x_H^y} (y + p(\tilde{x}))\, dF(\tilde{x})}{F(x_H^y) - F(x_L^y)}, \tag{26}$$

where (x_L^y, x_H^y) is the interval of types that are allocated prize y in the assortative allocation (so $x_L^y = F^{-1}(\lim_{y' \uparrow y} G(y'))$ and $x_H^y = F^{-1}(G(y))$). Now, replace the limit prize distribution G with distribution $\tilde{G}$ in which every prize y is replaced with the effective prize $\tilde{y}$. The assortative allocation y^A is replaced with $\tilde{y}^A$, so $\tilde{y}^A(x)$ is the effective prize for type x under the assortative allocation. Then, all our results on the characterization of Pareto improvements continue to hold.

To see this, it is enough to consider two consecutive prizes and determine the effect of pooling all the types that are allocated these prizes. Denote by $y < y'$ two consecutive prizes in the support of the limit prize distribution G, so $y = y^A(x)$ for x in $(x_L^y, x_H^y]$ and $y' = y^A(x)$ for x in $(x_L^{y'}, x_H^{y'}]$ (with $x_H^y = x_L^{y'}$). By pooling types on interval $(x_L^y, x_H^{y'}]$, the two prizes y and y' are combined to create an average prize y''. The corresponding effective prize is

$$\begin{aligned}
\tilde{y}'' &= \frac{\int_{x_L^y}^{x_H^y} y\, dF(\tilde{x}) + \int_{x_L^{y'}}^{x_H^{y'}} y'\, dF(\tilde{x}) + \int_{x_L^y}^{x_H^{y'}} p(\tilde{x})\, dF(\tilde{x})}{F(x_H^{y'}) - F(x_L^y)} \\
&= \frac{(F(x_H^y) - F(x_L^y))\, \tilde{y} + (F(x_H^y) - F(x_L^y))\, \tilde{y}'}{F(x_H^{y'}) - F(x_L^y)} \\
&= \frac{\int_{x_L^y}^{x_H^y} \tilde{y}^A(\tilde{x})\, dF(\tilde{x}) + \int_{x_L^{y'}}^{x_H^{y'}} \tilde{y}^A(\tilde{x})\, dF(\tilde{x})}{F(x_H^{y'}) - F(x_L^y)} \\
&= \frac{\int_{x_L^y}^{x_H^{y'}} \tilde{y}^A(\tilde{x})\, dF(\tilde{x})}{F(x_H^{y'}) - F(x_L^y)},
\end{aligned}$$

where the first equality follows from (26). As in the proof of Proposition 2, pooling is a

Pareto improvement if and only if

$$\tilde{y}'' - \frac{\int_{x_L^y}^{x_H^{y'}} \tilde{y}^A(x)\,dx}{x_H^{y'} - x_L^y} \geq 0 \iff \frac{\int_{x_L^y}^{x_H^{y'}} \tilde{y}^A(x)\,dF(x)}{F(x_H^{y'}) - F(x_L^y)} \geq \frac{\int_{x_L^y}^{x_H^{y'}} \tilde{y}^A(x)\,dx}{x_H^{y'} - x_L^y}.$$

D More general utility functions

We will now show how our results for top pooling can be extended to more general student utility functions.

Consider the separable utility functions $h(x,y) - c(t)$ and $h(y) - c(x,t)$, where $h(x,0) = c(x,0) = 0$ for all x, c is strictly increasing in t when $x > 0$ and decreasing in x when $t > 0$, and h is strictly increasing in y when $x > 0$ and strictly increasing in x when $y > 0$. These utilities generalize utilities (2) and (3), respectively. We extend the results of Section 4 to these utility functions.[58] Part (a) of Proposition 1 follows from the results of Olszewski and Siegel (2016), as these results hold for the more general utility functions. (Of course, the formula defining M will be different for the more general utility functions.)

The effort $t^A(x)$ is determined by

$$t^A(x) = c^{-1}\left(h(x, y^A(x)) - \int_0^x h_1\left(z, y^A(z)\right) dz\right), \tag{27}$$

for the utility $h(x,y) - c(t)$, and satisfies

$$h(y^A(x)) - c(x, t^A(x)) = -\int_0^x c_1(z, t^A(z))dz$$

for the utility $h(y) - c(x,t)$. By differentiating this equation, assuming that F and G are differentiable, we obtain the following differential equation

$$h'(y^A(x))(y^A)'(x) = c_2(x, t^A(x))(t^A)'(x).$$

[58]The results of Section 4 can be extended to a more general separable utility function $h(x,y) - c(x,t)$. However, the results for this more general function would be a combination the results for the two more specific functions, and the analysis would be less transparent.

This equation, together with the initial condition $t^A(0) = 0$, uniquely determines $t^A(x)$, assuming that the involved functions satisfy the Lipschitz condition.

For the utility $h(x,y) - c(t)$ we have that

$$U(x) = h(x, y^A(x)) - c(t^A(x)) = \int_0^x h_1\left(z, y^A(z)\right) dz$$

and

$$\begin{aligned} U^{x^*}(x) &= \frac{\int_{x^*}^1 h(x, y^A(z))dF(z)}{1 - F(x^*)} - c(M) \\ &= \frac{\int_{x^*}^1 h(x, y^A(z))dF(z)}{1 - F(x^*)} - \frac{\int_{x^*}^1 h(x^*, y^A(z))dF(z)}{1 - F(x^*)} + \int_0^{x^*} h_1\left(z, y^A(z)\right) dz. \end{aligned}$$

This yields

$$U^{x^*}(x) - U(x) = \frac{\int_{x^*}^1 h(x, y^A(z))dF(z)}{1 - F(x^*)} - \frac{\int_{x^*}^1 h(x^*, y^A(z))dF(z)}{1 - F(x^*)} - \int_{x^*}^x h_1\left(z, y^A(z)\right) dz.$$

Comparing this expression to 0, we obtain an analogue of condition (10) from part (b) of Proposition 1.

The derivative of $U^{x^*}(x) - U(x)$ is equal to

$$\frac{\int_{x^*}^1 h_1(x, y^A(z))dF(z)}{1 - F(x^*)} - h_1\left(x, y^A(x)\right).$$

If we assume that $h_1(x, y)$ increases in y, then this derivative is positive at $x = x^*$ and negative at $x = 1$. Assuming that the derivative changes its sign only once, we obtain part (c) of Proposition 1; in addition, this yields the following analogue of the condition from part (d):

$$\frac{\int_{x^*}^1 h(1, y^A(z))dF(z)}{1 - F(x^*)} - \frac{\int_{x^*}^1 h(x^*, y^A(z))dF(z)}{1 - F(x^*)} - \int_{x^*}^1 h_1\left(z, y^A(z)\right) dz \geq 0.$$

For the utility $h(y) - c(x,t)$ we have that

$$U(x) = h(y^A(x)) - c(x, t^A(x)) = -\int_0^x c_1(z, t^A(z))dz$$

and

$$U^{x^*}(x) = \frac{\int_{x^*}^1 h(y^A(z))dF(z)}{1-F(x^*)} - c(x,M) = c(x^*,M) - c(x,M) - \int_0^{x^*} c_1(z,t^A(z))dz,$$

which yields

$$U^{x^*}(x) - U(x) = c(x^*,M) - c(x,M) + \int_{x^*}^{x} c_1(z,t^A(z))dz.$$

Comparing this expression to 0, we obtain an analogue of condition (10) from part (b) of Proposition 1.

The derivative of this expression is $c_1(x,t^A(x)) - c_1(x,M)$. Assume that $c_1(x,t)$ decreases with t. Then, $c_1(x,t^A(x)) - c_1(x,M) \geq 0$ when $t^A(x) \leq M$, and $c_1(x,t^A(x)) - c_1(x,M) \leq 0$ when $t^A(x) \geq M$. Since $U^M(x) - U(x) = 0$ when $x = x^*$, part (c) of Proposition 1 must hold; in addition, this yields the following analogue of the condition from part (d):

$$c(x^*,M) - c(1,M) + \int_{x^*}^{1} c_1(z,t^A(z))dz \geq 0.$$

Corollary 1 also generalizes to utilities $h(x,y) - c(t)$ and $h(y) - c(x,t)$, and its proof requires only minor changes. Consider first the utility $h(x,y) - c(t)$. We have

$$U^{x^*}(x) = \frac{\int_{x^*}^1 h(x,y^A(z))dF(z)}{1-F(x^*)} - \frac{\int_{x^*}^1 h(x^*,y^A(z))dF(z)}{1-F(x^*)} + \int_0^{x^*} h_1\left(z,y^A(z)\right)dz,$$

so, $U^{x^*}(x)$ can be represented as $\phi(x,x^*) + \psi(x^*)$, where

$$\phi(x,x^*) = \frac{\int_{x^*}^1 h(x,y^A(z))dF(z)}{1-F(x^*)}$$

and

$$\psi(x^*) = \int_0^{x^*} h_1\left(z,y^A(z)\right)dz - \frac{\int_{x^*}^1 h(x^*,y^A(z))dF(z)}{1-F(x^*)}.$$

Assume that for all $x'' > x'$, the difference $h(x'',y) - h(x',y)$ strictly increases in y. This implies that $\phi(x,x_1^*) - \phi(1,x_1^*) > \phi(x,x_2^*) - \phi(1,x_2^*)$. Together with $\phi(1,x_1^*) + \psi(x_1^*) \geq \phi(1,x_2^*) + \psi(x_2^*)$, this yields part (a).

To show part (b), suppose that type 1 strictly prefers x_2^* to x_1^*. By the assumption that

$h(x'', y) - h(x', y)$ strictly increases in y, it follows that $U^{x_1^*}(x) - U^{x_2^*}(x)$ strictly decreases in x. Thus, there exists an x^{**} such that types $x < x^{**}$ strictly prefer x_1^* and types $x > x^{**}$ strictly prefer x_2^*. And since all players with types close to 1 strictly prefer x_2^* to x_1^*, we have that $x^{**} < 1$. Finally, it can be readily checked that the derivative of ψ with respect to x^* is negative. Hence, $\psi(x_1^*) > \psi(x_2^*)$, and so $x^{**} > 0$.

For the utility $h(y) - c(x, t)$, we have that

$$U^{x^*}(x) = c(x^*, M) - c(x, M) - \int_0^{x^*} c_1(z, t^A(z))dz.$$

So, $U^{x^*}(x)$ can be represented as $\phi(x, x^*) + \psi(x^*)$, where

$$\phi(x, x^*) = -c(x, M)$$

and

$$\psi(x^*) = c(x^*, M) - \int_0^{x^*} c_1(z, t^A(z))dz.$$

Assume that for all $x'' > x'$, the difference $c(x', M) - c(x'', M)$ strictly increases in M. This implies that $\phi(x, x_1^*) - \phi(1, x_1^*) > \phi(x, x_2^*) - \phi(1, x_2^*)$. Together with $\phi(1, x_1^*) + \psi(x_1^*) \geq \phi(1, x_2^*) + \psi(x_2^*)$, this yields part (a).

To show part (b), suppose that a player of type 1 strictly prefers x_2^* to x_1^*. By the assumption that $c(x', M) - c(x'', M)$ strictly increases in M, it follows that $U^{x_1^*}(x) - U^{x_2^*}(x)$ strictly decreases in x. Thus, there exists an x^{**} such that types $x < x^{**}$ strictly prefer x_1^* and types $x > x^{**}$ strictly prefer x_2^*. And since all players with types close to 1 strictly prefer x_2^* to x_1^*, we have that $x^{**} < 1$. Finally, it can be readily checked that the derivative of ψ with respect to x^* is negative. Hence, $\psi(x_1^*) > \psi(x_2^*)$, so $x^{**} > 0$.

E Relationship to Krishna et al. (2018)

The model in Krishna et al. (2018) serves as a bridge between our theory and the data. Most variables of our model, such as agent ability, costs of effort and prize values, are unlikely to be observed in a typical administrative dataset. The model of Krishna et al. (2018), on the other hand, can be estimated using existing data and produce close analogs of the above variables, via simulation or directly.[59] Our model is then calibrated to these quantities. In what follows, we explain how the elements of our model are connected to their counterparts in Krishna et al. (2018).

The model of Krishna et al. (2018) describes a single-agent problem faced by students in the Turkish centralized college placement system. The timing is as follows:

1. Each student chooses between three types of high school (public, private or selective school) and two levels of private tutoring. These choices affect subsequent exam scores.

2. The agent studies in high school and gets a signal about his innate abilities.

3. The agent takes the exam and learns his placement score.

4. The agent decides whether to be placed with the current score or retake the exam next year. Retaking is costly, but may improve the score. Placement is the terminal state. Each retaker draws a permanent and a transitory shock to the scores and decides whether to retake again, and so on.

Students are heterogeneous in many dimensions. Each student i is characterized by a vector of socioeconomic characteristics and academic performance prior to high school, X_i. Each of the six levels of pre-exam investment of effort, e, is associated with a cost $C_i(e) = \gamma_e + w_{ie}$. The idiosyncratic part of the cost, w_{ie}, is modeled as a mixed logit shock and is observed by the agent. Choice-specific parameters γ_e capture the cost component common to all agents. Effort affects the exam score in the first attempt, t_{i1}:

$$t_{i1} = X_i'\beta + \rho_e + \varepsilon_i,$$

The idiosyncratic shock to the score consists of two components: $\varepsilon_i = \varepsilon_{i0} + \varepsilon_{i1}$. The first one is persistent and is revealed at stage 2. It captures residual variation in innate ability not

[59]We estimate the version of the structural model from Krishna et al. (2018) that allows for pre-test investment into exam preparation. In contrast to the main specification in the above paper, we restrict all income categories to have the same utility and cost structure. Unlike Krishna et al. (2018), we do not focus on income-related asymmetries in placement outcomes here, and removing this extra dimension of heterogeneity greatly simplifies the calibration exercise.

explained by X_i. The second one, ε_{i1}, is transitory and accounts for random shocks during the exam. This shock is revealed to the agent at stage 3. In what follows, it is convenient to separate the persistent part of the score, $\bar{t}_{i1} = X_i'\beta + \rho_e + \varepsilon_{i0}$, from the transitory shock.

The dynamic decision problem in stage 3 is described by Bellman's equation. From the agent's perspective, the state of the world at the exam attempt $a = 1, 2, \dots$ in this dynamic problem is described by the pair $(\bar{t}_{ia}, \varepsilon_{ia})$, where $\bar{t}_{ia}$ is the persistent component of i's score and ε_{ia} is the transitory shock affecting the score in attempt a:

$$W_a(\bar{t}_{ia}, \varepsilon_{ia}) = \max\left\{U(r(\bar{t}_{ia} + \varepsilon_{ia})), VC_a(\bar{t}_{ia})\right\},$$
$$VC_a(\bar{t}_{ia}) = \delta E_{\lambda_{ia+1}, \varepsilon_{ia+1}}\left[W_a(\bar{t}_{ia} + \lambda_{ia+1}, \varepsilon_{ia+1})\right] - \psi_a$$

Parameter ψ_a captures the cost of retaking after attempt a. The persistent part of the score stochastically evolves over time: $\bar{t}_{ia+1} = \bar{t}_{ia} + \lambda_{ia+1}$. The shock λ_{ia+1} captures learning between the attempts. The function $U(r)$ describes the utility of being placed with the rank r, while $r(t)$ is the rank that the score t can buy in the equilibrium. We refer the reader to Krishna et al. (2018) for more details on this part.

Let $V_0(\bar{t}_{i1})$ be the function that captures the value of reaching the expected score $\bar{t}_{i1} = X_i'\beta + \rho_e + \varepsilon_{i0}$ in stage 2, when the agent is just about to take the exam for the first time:

$$V_0(\bar{t}_{i1}) = E_{\varepsilon_{i1}}[V_1(\bar{t}_{i1}, \varepsilon_{i1})] \tag{28}$$

This function's value accounts for the option to retake and the payoff of eventual placement.

As explained in Section 7, the distribution of prizes, $G(y)$, is directly related to the utility function $U(r)$. Consider an agent whose placement rank is r. In the notation of Krishna et al. (2018), the agent's placement payoff is $U(r)$. Since prizes are allocated according to the ranking, he gets $y = G^{-1}(r)$. Thus, $U(r) = G^{-1}(r)$.

Another object that we use to calibrate our model is $C(t) = c(t)/x^A(t)$, the score-to-cost mapping in the Turkish equilibrium. In our model, students vary in a single dimension, ability, which leads to one-to-one mapping between scores, rank and ability. In Krishna et al. (2018), agents are heterogeneous in multiple dimensions: idiosyncratic shocks w_{ie}

and ε_{i1} affect their costs and scores, while the observable characteristics X_i explain pre-existing differences in scores related to i's socioeconomic background. As a result, there is no one-to-one mapping between the cost of effort and the exam score. To preserve the general relationship between the costs and the scores, while abstracting away from the extra dimensions of student heterogeneity, we simulate the costs $C_i(e_i^*)$ of the optimal effort level e_i^* and the scores t_{i1} for each agent using the estimated model of Krishna et al. (2018) and fit a non-parametric regression of $C_i(e_i^*)$ on t_{i1} to get the average cost at every score and use this as the cost for the student who obtains this score.[60] That is, we use the regression estimate of $E[C_i(e_i^*)|t_{i1} = t]$ in place of $C(t)$.

Finally, in order to back out the distribution of ability in Section 7, we need to know the marginal payoff from improving the score in the first attempt, $V'(t)$. We use the value function V_0 as defined in (28) in place of $V(t)$.

F Robust Pareto Improvement

In this section, we find the robust policy characterized in Proposition 4. The policy is found using grid search:

1. We set $x_1^* = 0$, $x_1^{**} = 1$ and gradually reduce x_1^{**} until the truncated distribution of types on $[x_1^*, x_1^{**}]$ dominates the uniform distribution in the same interval:

$$\frac{x - x_1^*}{x_1^{**} - x_1^*} \geq \frac{F(x) - F(x_1^*)}{F(x_1^{**}) - F(x_1^*)} \text{ for all } x \in [x_1^*, x_1^{**}]$$

2. If such upper boundary x_1^{**} exists and the mass of agents in the interval, $F(x_1^{**}) - F(x_1^*)$, is above 1 percent, we save $[x_1^*, x_1^{**}]$ as the first pooling interval, set x_2^* at x_1^{**} plus a small increment, x_2^{**} at unity, and start searching for the second pooling interval.

3. If the conditions in the previous step do not hold for any x_i^{**}, we add a small increment to x_i^* and start over.

[60] It is also worth noting that the mean cost parameters γ_e are identified up to an additive constant. We address this issue by normalizing the cost of attending a public school with no private tutoring to zero.

The optimal policy is a category ranking consisting of 16 intervals, depicted in Figure F1 as shaded areas and documented in detail in Table F1.

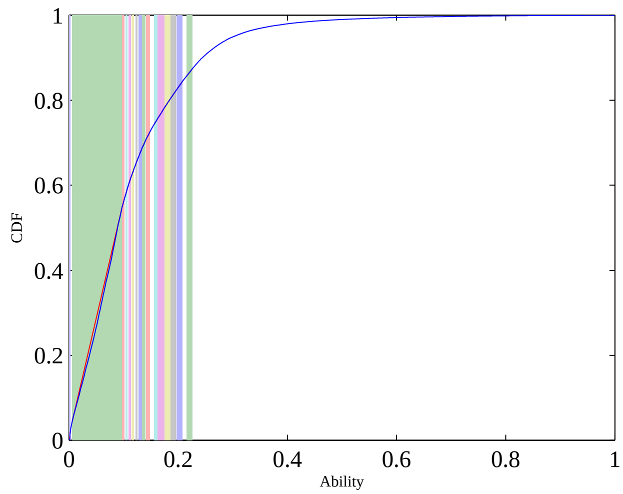

Figure F1: Robust Pareto-improving pooling policy.

Table F1: Robust Pareto-improving pooling policy.

Interval, i	$[x_i^*, x_i^{**}]$	$[F(x_i^*), F(x_i^{**})]$	Mass
1	$[0.000, 0.002]$	$[0.00, 0.02]$	0.020
2	$[0.006, 0.097]$	$[0.04, 0.55]$	0.506
3	$[0.097, 0.101]$	$[0.55, 0.57]$	0.019
4	$[0.103, 0.107]$	$[0.58, 0.59]$	0.016
5	$[0.110, 0.114]$	$[0.60, 0.62]$	0.015
6	$[0.114, 0.119]$	$[0.62, 0.64]$	0.016
7	$[0.122, 0.125]$	$[0.65, 0.66]$	0.012
8	$[0.128, 0.134]$	$[0.67, 0.69]$	0.018
9	$[0.134, 0.140]$	$[0.69, 0.70]$	0.017
10	$[0.142, 0.149]$	$[0.71, 0.73]$	0.018
11	$[0.156, 0.163]$	$[0.74, 0.76]$	0.014
12	$[0.163, 0.175]$	$[0.76, 0.78]$	0.025
13	$[0.175, 0.186]$	$[0.78, 0.80]$	0.022
14	$[0.186, 0.196]$	$[0.81, 0.82]$	0.017
15	$[0.198, 0.208]$	$[0.83, 0.84]$	0.018
16	$[0.215, 0.226]$	$[0.86, 0.87]$	0.017

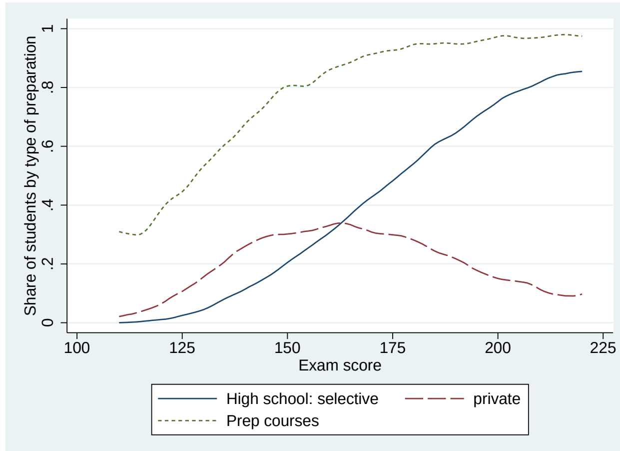

Figure F2: Types of pre-test investment conditional on the final score

G Cost-to-School Mapping

G.1 Construction of the cost-to-school mapping

We adapt the techniques of Section 7 to a discretized economy in which there is a unit mass of students and 10 schools. Each school has a capacity for 0.1 of the student population.

The cost of getting the score t is $c(t)/x$, where x is student's ability. Using the same cost function and the distribution of student abilities as in the estimated continuous economy, we simulate equilibrium cutoffs for the 10 schools in the discrete economy. Net payoffs $(y - c(t)/x)$ in the discretized economy resemble those in the continuous one, and the placement payoffs u are approximating those in the continuous economy. Figure G1 shows how the ability cutoffs (G1a) and net payoffs (G1b) in the discretized economy relate to those in the continuous economy.

The bottom pooling policy we found previously for the continuous economy is necessarily Pareto-improving in the discrete economy. Further, the two policies are qualitatively similar: roughly the same set of students get seats in the pooled school under both policies, and their placement payoffs are almost identical. Figure G2 illustrates how the pooling policy is Pareto improving in the discretized economy (to see how this is qualitatively similar to the continuous economy, compare this figure to Figure 6 in the main text). Student net payoffs in the discretized economy under the discrete policy and the pooling policy and displayed

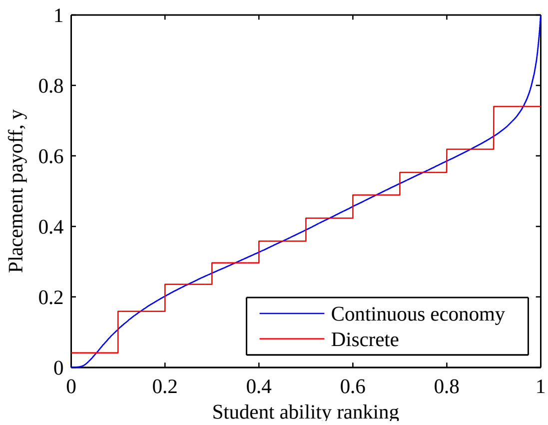

(a) Mapping from student's ability ranking to prize.

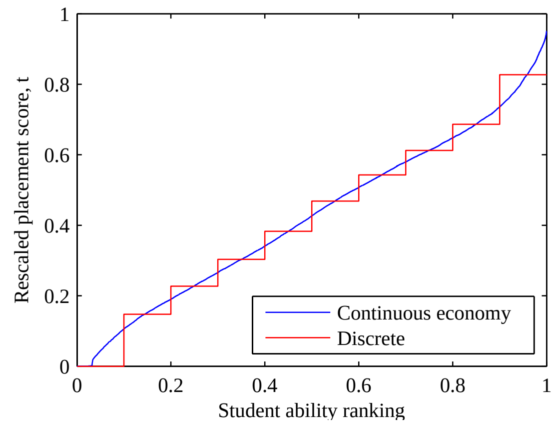

(b) Student net payoffs as a function of ability ranking in the no-pooling scenario.

Figure G1: Discretized vs. continuous economy.

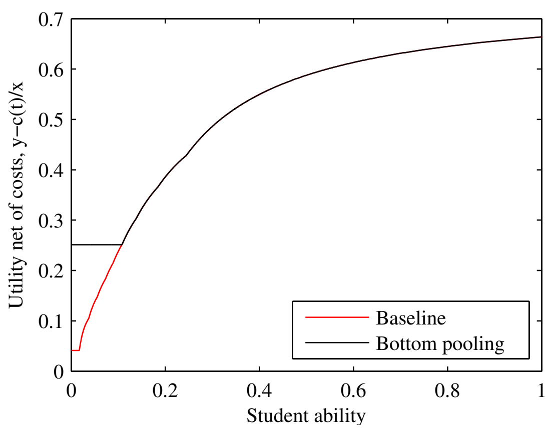

(a) Net payoffs by ability

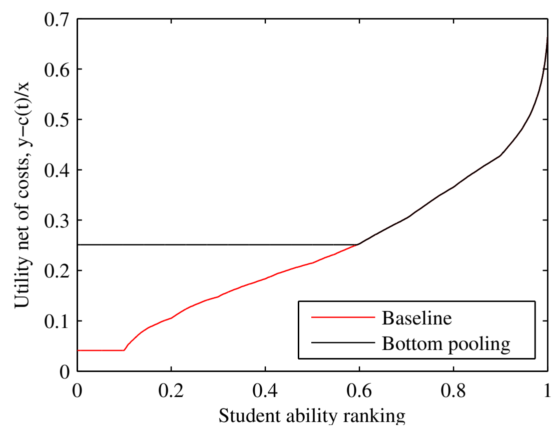

(b) Net payoffs by ability ranking

Figure G2: Student net payoffs as a function of ability in the discretized economy

as a function of their ability (G2a), and a function of their ability ranking (G2b).

G.2 Parameter selection of the cost-to-school mapping

Here we present the particular cost-to-school mapping for the online experiment. The experiment uses 6 ability levels that are to the left of the pooling threshold and 5 that are to the right of it: subject in the experiment was randomly assigned to one of the following ability levels: 0.01, 0.02, 0.04, 0.06, 0.08, 0.10, 0.12, 0.15, 0.20, 0.60, 1.00. These were chosen so that the majority of subjects had abilities below the pooling threshold (as is the case in the Turkish data – see Section 7).

Table G1 shows the mapping from investment levels to cost under the discrete policy for each ability level in our experiment, using the results of the empirical estimation for the discretized economy described above. Table G2 shows the parameters that were used under the pooling policy. The parameters in this table are those the subjects saw on their screens.[61] Highlighted in yellow are the profit maximizing choices.

As can be seen, subjects whose abilities fall below the pooling threshold (i.e. subjects with ability less than or equal to 0.10) are predicted to opt for the pooling lottery under the pooling policy, while "high ability" subjects instead are predicted to behave no differently across the discrete and pooling policies.

		Ability										
School	Enrollment Bonus	0.01	0.02	0.04	0.06	0.08	0.1	0.12	0.15	0.2	0.6	1
J	3	0	0	0	0	0	0	0	0	0	0	0
I	16	20	9	7	3	3	2	2	1	1	0	0
H	24	49	24	12	10	8	5	4	3	2	1	1
G	30	83	41	21	14	13	10	8	6	4	1	1
F	35	127	64	32	21	16	14	11	9	6	2	1
E	42	185	93	46	31	25	19	17	12	10	3	2
D	49	257	128	64	43	34	28	21	19	16	4	3
C	55	346	173	86	57	43	35	30	23	21	6	4
B	62	466	233	117	78	57	47	39	33	24	8	5
A	74	762	381	191	127	95	76	64	51	43	13	8

Notes: For example, a subject assigned to ability level 0.06 would have to invest 14 Tokens out of the 100 Token endowment in order to meet the requirement for School G and obtain a 30-Token enrollment bonus. If that subject had instead invested 21 Tokens, they would enroll in School F and obtain 35 Tokens.

Table G1: Discrete policy: cost of targeting particular schools by ability.

.

[61] The parameters are scaled up by a factor of 100 relative to the empirical estimation.

School	Enrollment Bonus	Ability										
		0.01	0.02	0.04	0.06	0.08	0.1	0.12	0.15	0.2	0.6	1
J												
I	LOTTERY											
H	Equal Chances of											
G	3, 16, 24,						0					
F	30, 35, 42											
E												
D	49	257	128	64	43	34	28	21	19	16	4	3
C	55	346	173	86	57	43	35	3	23	21	6	4
B	62	466	233	117	78	57	47	39	33	24	8	5
A	74	762	381	191	127	95	76	64	51	43	13	8

Notes: For example, a subject assigned to ability level 0.06 who invested 0 Tokens out of the 100 Token endowment would enter an enrollment lottery and have equal chances of being admitted to school E, F, G, H, I, J and earn the bonus associated with the school they ultimately enrolled in. If that subject had instead invested 57 Tokens, they would enroll in School C and obtain 55 Tokens.

Table G2: Cost of targeting particular schools by ability under the pooling policy.

H Experiment: Additional Results

H.1 Results on non-risk-neutral subjects

In this section we provide additional statistics on policy evaluation regarding subjects who are not risk neutral. We find that risk aversion and behavior in our setting follow what would be expected. The p-values reported below correspond to the exact p-values of a Wilcoxon rank-sum test.

High Ability Players. First, risk averse subjects whose abilities are above the pooling threshold should behave similarly to risk neutral subjects in the Lottery treatment when faced with the pooling policy. This is exactly what we find (the lowest p-value when comparing behavior for each ability level across the risk averse and risk neutral players is $p = 0.1031$). However, subjects who are risk-seeking should either be more likely or equally likely to choose the lottery over the certain outcome compared with risk-neutral, players but never less likely. We find that subjects with the three highest ability levels are more likely to choose the lottery (the highest p-value is $p = 0.0689$), and subjects with the lower two ability levels are equally likely (the lowest p-value is $p = 0.3285$).

Low Ability Players. Risk seeking subjects whose abilities are below the pooling threshold should behave similarly to risk neutral subjects in the Lottery treatment when faced with the pooling policy. This is exactly what we find (the lowest p-value is $p = 0.3517$). For risk averse players, they should either be more likely to opt out of the lottery, or, depending on how risk averse they are, behave as risk neutral players do. We find no difference in behavior between risk averse and risk neutral players who are of low ability (the lowest p-value is $p = 0.3748$).

H.2 Additional data collection

We ran some sessions in which under the pooling policy the lottery was replaced with a fixed amount of 0.25, which corresponded to the expected value of the lottery under the pooling policy. Below we present the results from these sessions. Figure H1 shows the average profits under the discrete and pooling policies. Just as in the data presented in the main text, subjects here on average do better under the pooling policy.

In fact, statistically speaking, results are largely similar with those from the main text. The only noteworthy difference is that here subjects who are high ability but whose ability are just above the pooling threshold behave as predicted: removing the lottery removes the ability to randomize and subjects make the profit maximizing choice.

I Screen Shots

I.1 Discrete Policy First

Below we present the screen shots that subjects saw on their screens. This particular sequence corresponds to the subjects who saw the discrete policy first.

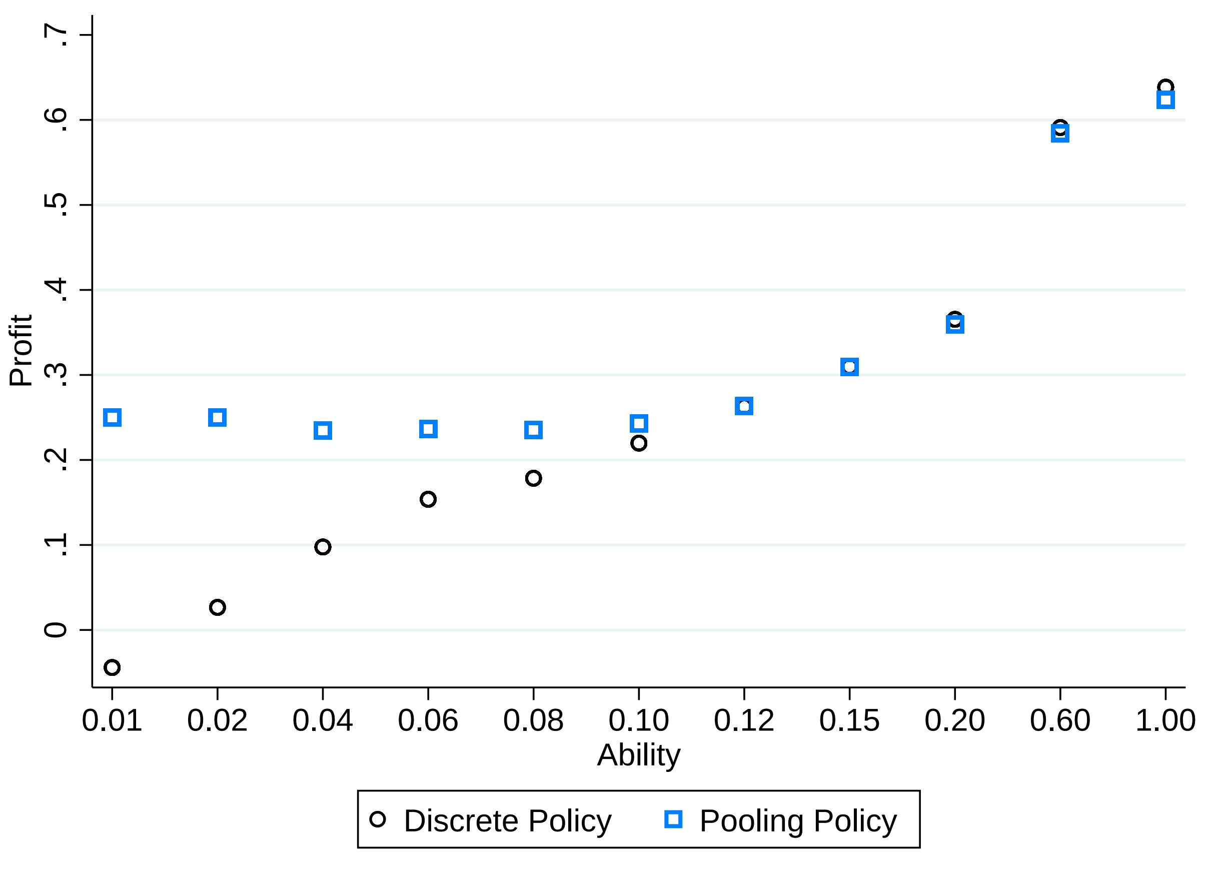

Figure H1: Average profits under the discrete and pooling policies.

Figure G1: Consent screens

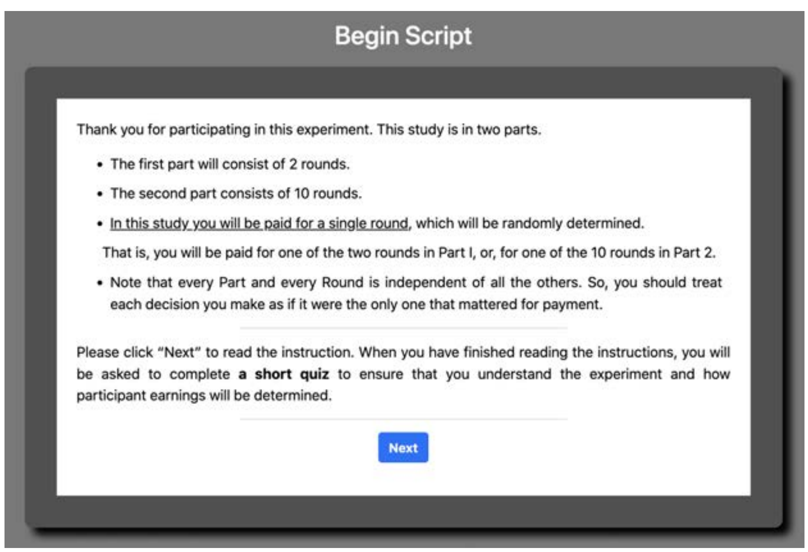

Figure G2: Preamble

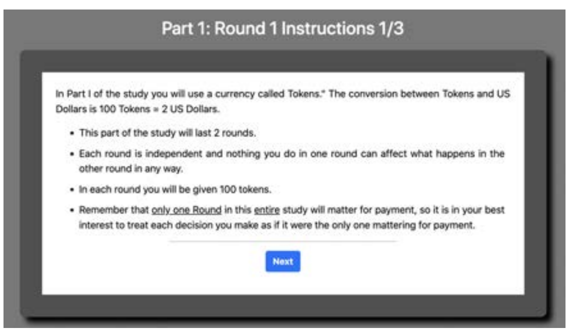

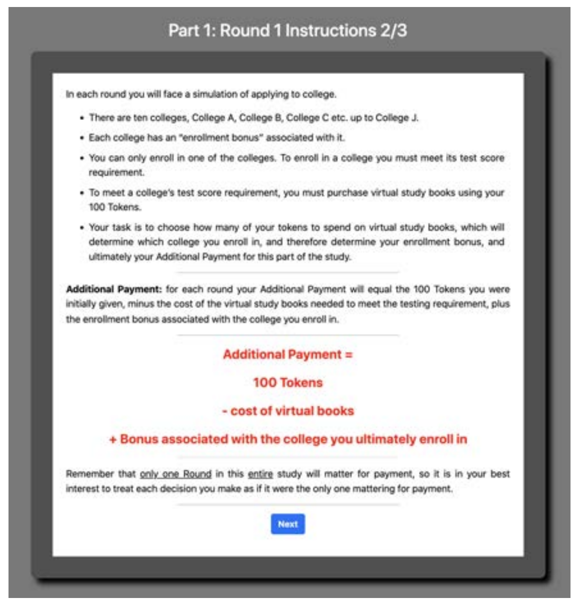

Figure G3: Part 1 General instructions

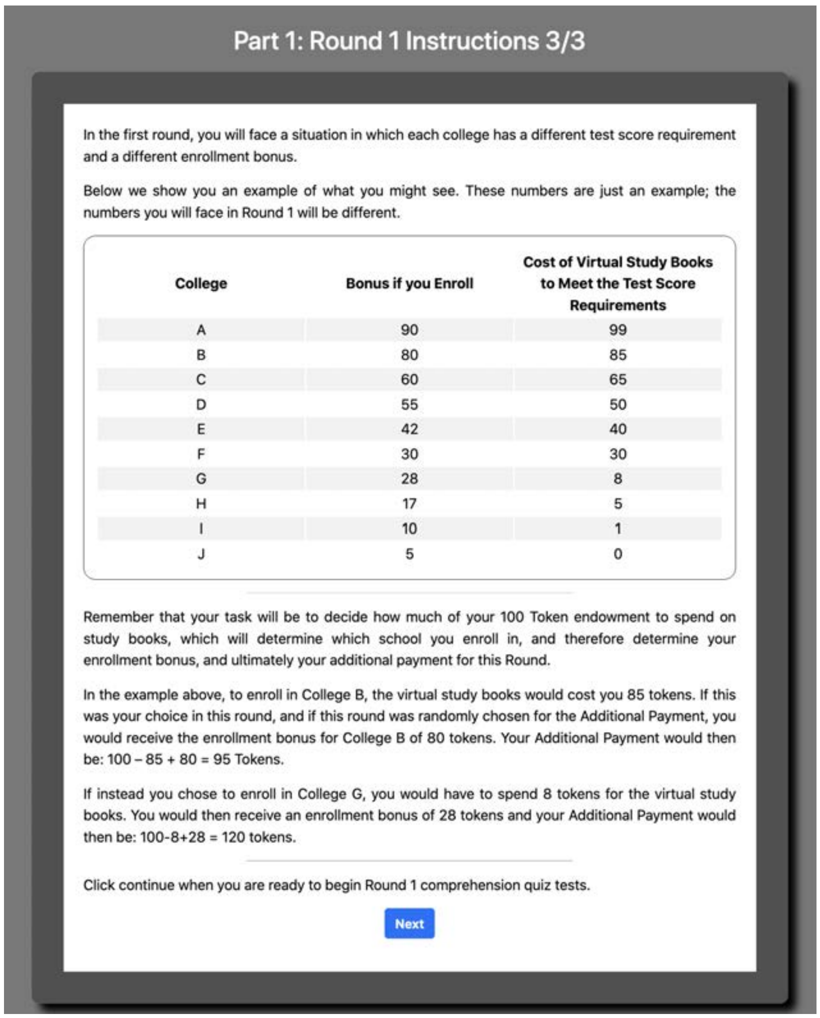

Part 1: Round 1 Instructions 3/3

In the first round, you will face a situation in which each college has a different test score requirement and a different enrollment bonus.

Below we show you an example of what you might see. These numbers are just an example; the numbers you will face in Round 1 will be different.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	28	8
H	17	5
I	10	1
J	5	0

Remember that your task will be to decide how much of your 100 Token endowment to spend on study books, which will determine which school you enroll in, and therefore determine your enrollment bonus, and ultimately your additional payment for this Round.

In the example above, to enroll in College B, the virtual study books would cost you 85 tokens. If this was your choice in this round, and if this round was randomly chosen for the Additional Payment, you would receive the enrollment bonus for College B of 80 tokens. Your Additional Payment would then be: 100 – 85 + 80 = 95 Tokens.

If instead you chose to enroll in College G, you would have to spend 8 tokens for the virtual study books. You would then receive an enrollment bonus of 28 tokens and your Additional Payment would then be: 100-8+28 = 120 tokens.

Click continue when you are ready to begin Round 1 comprehension quiz tests.

Next

Figure G4: Instructions for Round 1 if discrete policy is seen first.

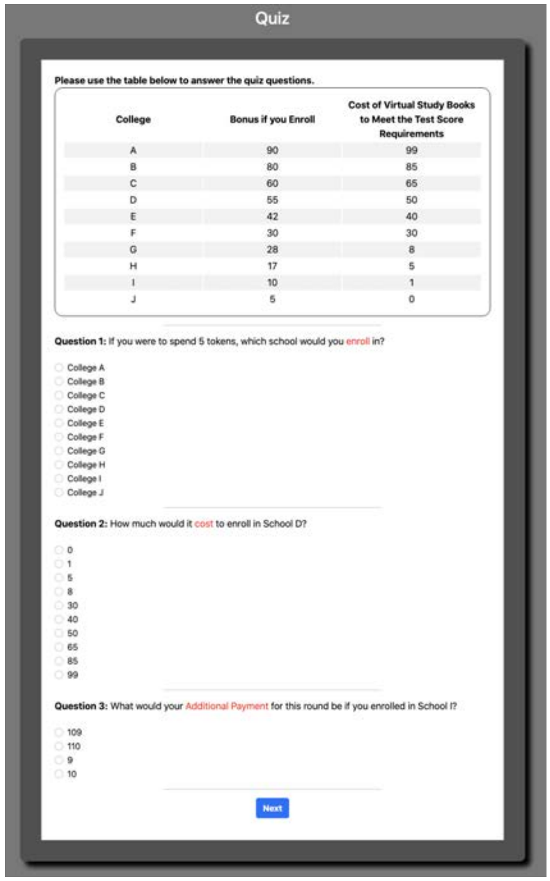

Quiz

Please use the table below to answer the quiz questions.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	28	8
H	17	5
I	10	1
J	5	0

Question 1: If you were to spend 5 tokens, which school would you enroll in?

- College A
- College B
- College C
- College D
- College E
- College F
- College G
- College H
- College I
- College J

Question 2: How much would it cost to enroll in School D?

- 0
- 1
- 5
- 8
- 30
- 40
- 50
- 65
- 85
- 99

Question 3: What would your Additional Payment for this round be if you enrolled in School I?

- 109
- 110
- 9
- 10

Next

Figure G5: Quiz after discrete round instructions.

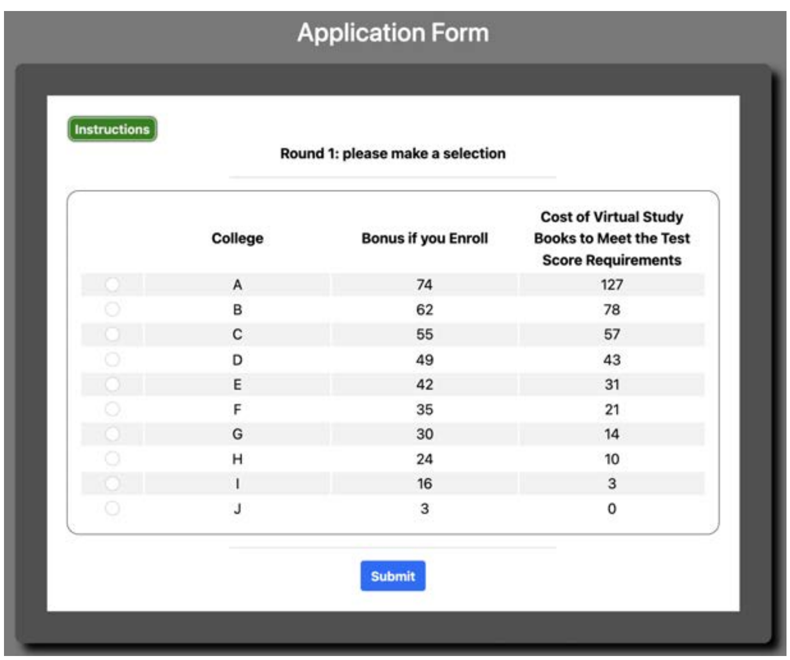

Application Form

Instructions

Round 1: please make a selection

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	74	127
B	62	78
C	55	57
D	49	43
E	42	31
F	35	21
G	30	14
H	24	10
I	16	3
J	3	0

Submit

Figure G6: Example of selection screen for discrete policy.

You have now finished Part 1, Round 1. Please click the button below to continue with the study.

Next

Figure G7: Transition to Round 2.

Part 1: Round 2 Instructions

You are now going to complete the second round of this part of the study.

Just like in the previous round, your task in this round will be to decide how much of your 100 Token endowment to spend on virtual study books, which will determine the college you enroll in, and therefore your Enrollment Bonus and Additional Payment.

However, in this round, there has been a policy change and now some colleges are pooled together. The colleges that are pooled together have the same test score requirement and the same cost of meeting this requirement. If you choose to spend your tokens on meeting the test score requirement for those colleges that are pooled together, you will enter an enrollment lottery and be randomly enrolled in one of those colleges (with an equal chance of enrolling in each of those pooled colleges). The enrollment bonus you receive will then equal the enrollment bonus of the college you are randomly assigned to enroll in.

Below is an example of a situation in which some colleges have been pooled together. Note that these numbers are all just examples and the numbers you will actually face in Round 2 will be different.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	**Enrollment lottery** Equal chances of enrolling in: College G with a bonus of 28 tokens College H with a bonus of 17 tokens College I with a bonus of 10 tokens College J with a bonus of 5 tokens	0
H		0
I		0
J		0

Remember that your task will be to decide how much of your 100 Token endowment to spend on study books, which will determine which school you enroll in, and therefore determine your enrollment bonus, and ultimately your additional payment for this Round.

In this example, Colleges G, H, I and J have been pooled together. If you choose to spend 0 tokens on study books you will enter an enrollment lottery where you will have an equal chance of ultimately enrolling in any one of the colleges that are pooled together. Your enrollment bonus will then equal the enrollment bonus of the college you were randomly assigned to enroll in. That is, if you spend 0 tokens, you will receive an enrollment bonus of 28 or 17 or 10 or 5 tokens, each with an equal chance. Your Additional Payment would then be: 100 - 0 + enrollment bonus.

If instead you wanted to enroll in college D you would have to spend 50 tokens on virtual study books. You would then receive a 55 token enrollment bonus and your Additional Payment would be: 100 - 50 + 55 = 105 tokens.

Click continue when you are ready to begin Round 2 comprehension quiz tests.

Next

Figure G8: Instructions for Round 2 if discrete policy is seen first.

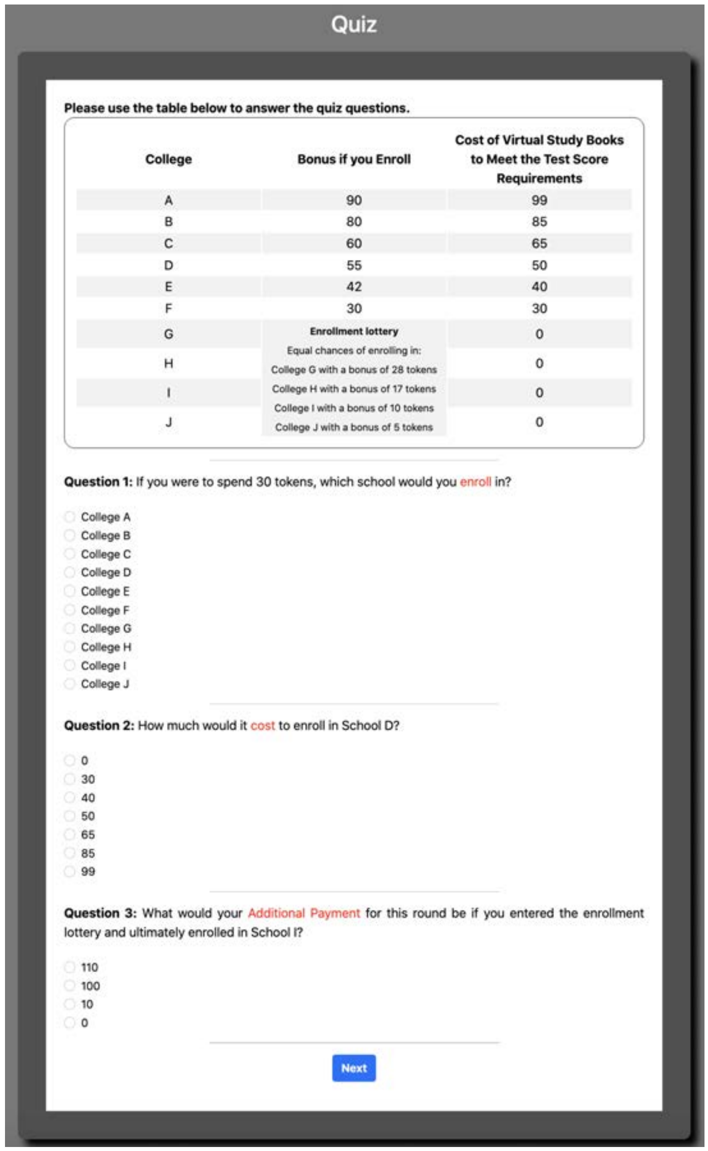

Quiz

Please use the table below to answer the quiz questions.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	**Enrollment lottery** Equal chances of enrolling in: College G with a bonus of 28 tokens College H with a bonus of 17 tokens College I with a bonus of 10 tokens College J with a bonus of 5 tokens	0
H		0
I		0
J		0

Question 1: If you were to spend 30 tokens, which school would you enroll in?

- College A
- College B
- College C
- College D
- College E
- College F
- College G
- College H
- College I
- College J

Question 2: How much would it cost to enroll in School D?

- 0
- 30
- 40
- 50
- 65
- 85
- 99

Question 3: What would your Additional Payment for this round be if you entered the enrollment lottery and ultimately enrolled in School I?

- 110
- 100
- 10
- 0

Next

Figure G9: Quiz after pooling round instructions.

Application Form

Instructions

Round 1: please make a selection

Colleges E, F, G, H, I and J have been pooled together. If you choose to spend 0 tokens on study books you will enter an enrollment lottery where you will have an equal chance of ultimately enrolling in any one of the colleges that are pooled together. Your enrollment bonus will then equal the enrollment bonus of the college you were randomly assigned to enroll in. That is, if you spend 0 tokens, you will receive an enrollment bonus of 42 or 35 or 30 or 24 or 16 or 3 tokens, each with an equal chance (for your information, on average this equals 25 tokens).

	College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
	A	74	127
	B	62	78
	C	55	57
	D	49	43
	E	**Enrollment lottery** Equal chances of enrolling in: College E with a bonus of 42 tokens College F with a bonus of 35 tokens College G with a bonus of 30 tokens College H with a bonus of 24 tokens College I with a bonus of 16 tokens College J with a bonus of 3 tokens	0
	F		0
	G		0
	H		0
	I		0
	J		0

Submit

Figure G10: Example of selection screen for pooling policy.

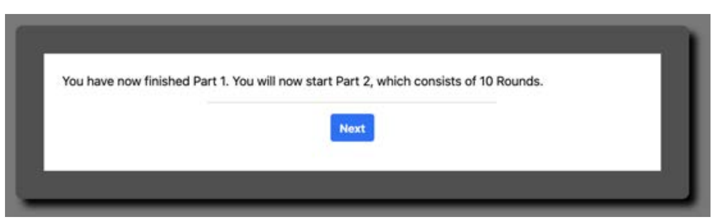

You have now finished Part 1. You will now start Part 2, which consists of 10 Rounds.

Next

Figure G11: Transition from Part 1 to Part 2.

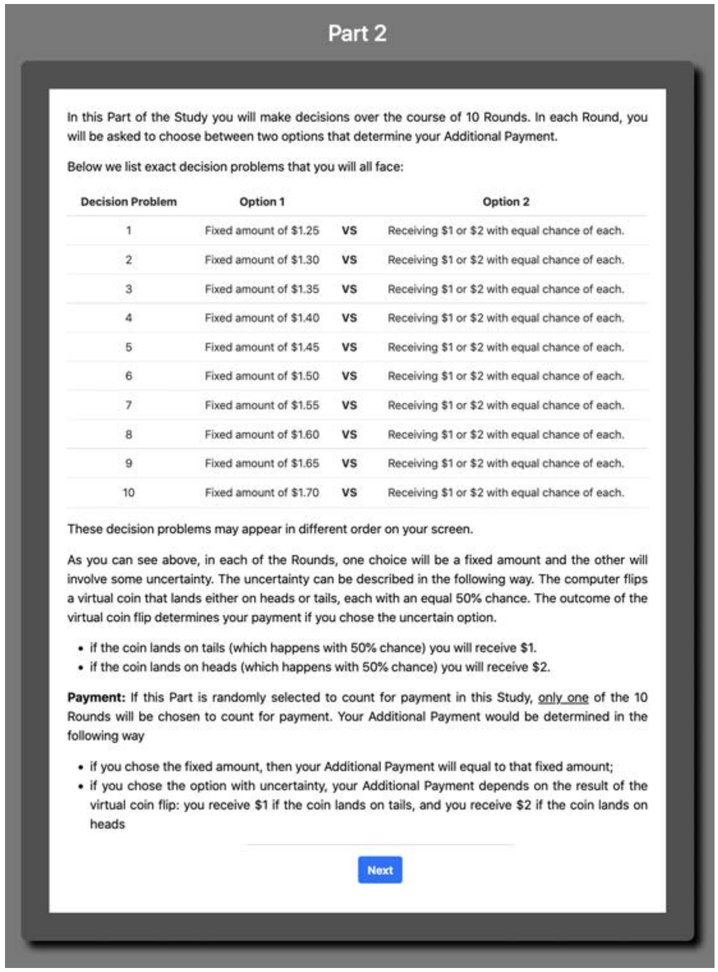

Part 2

In this Part of the Study you will make decisions over the course of 10 Rounds. In each Round, you will be asked to choose between two options that determine your Additional Payment.

Below we list exact decision problems that you will all face:

Decision Problem	Option 1		Option 2
1	Fixed amount of $1.25	VS	Receiving $1 or $2 with equal chance of each.
2	Fixed amount of $1.30	VS	Receiving $1 or $2 with equal chance of each.
3	Fixed amount of $1.35	VS	Receiving $1 or $2 with equal chance of each.
4	Fixed amount of $1.40	VS	Receiving $1 or $2 with equal chance of each.
5	Fixed amount of $1.45	VS	Receiving $1 or $2 with equal chance of each.
6	Fixed amount of $1.50	VS	Receiving $1 or $2 with equal chance of each.
7	Fixed amount of $1.55	VS	Receiving $1 or $2 with equal chance of each.
8	Fixed amount of $1.60	VS	Receiving $1 or $2 with equal chance of each.
9	Fixed amount of $1.65	VS	Receiving $1 or $2 with equal chance of each.
10	Fixed amount of $1.70	VS	Receiving $1 or $2 with equal chance of each.

These decision problems may appear in different order on your screen.

As you can see above, in each of the Rounds, one choice will be a fixed amount and the other will involve some uncertainty. The uncertainty can be described in the following way. The computer flips a virtual coin that lands either on heads or tails, each with an equal 50% chance. The outcome of the virtual coin flip determines your payment if you chose the uncertain option.

- if the coin lands on tails (which happens with 50% chance) you will receive $1.
- if the coin lands on heads (which happens with 50% chance) you will receive $2.

Payment: If this Part is randomly selected to count for payment in this Study, only one of the 10 Rounds will be chosen to count for payment. Your Additional Payment would be determined in the following way

- if you chose the fixed amount, then your Additional Payment will equal to that fixed amount;
- if you chose the option with uncertainty, your Additional Payment depends on the result of the virtual coin flip: you receive $1 if the coin lands on tails, and you receive $2 if the coin lands on heads

Next

Figure G12: Instructions for Part 2 (risk elicitation).

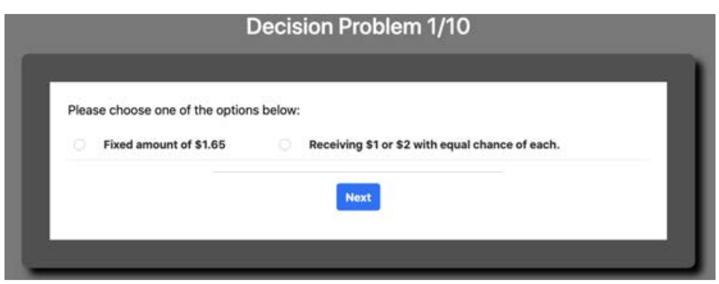

Decision Problem 1/10

Please choose one of the options below:

Fixed amount of $1.65 **Receiving $1 or $2 with equal chance of each.**

Next

Figure G13: Example of decision screen for Part 2.

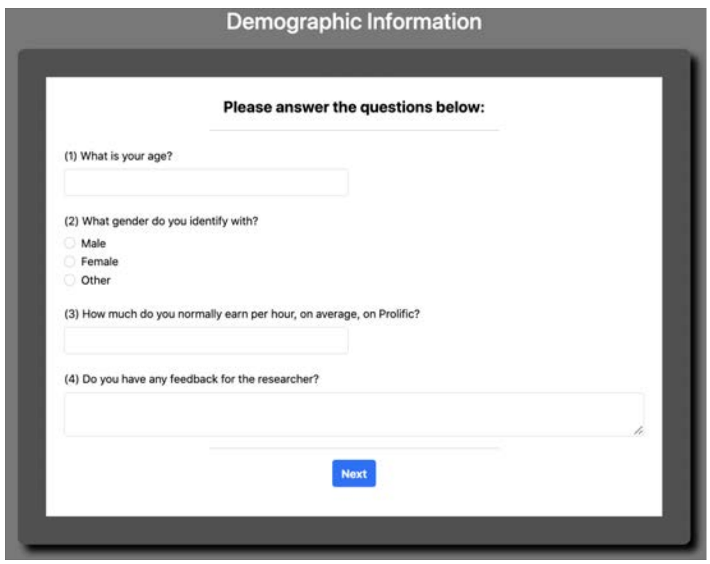

Figure G14: Questionnaire.

I.2 Pooling Policy First

If a subject saw the pooling policy first the sequence of screens was slightly modified from the Baseline in which subjects saw the discrete round first. Figures G4, G5, G6, G8, G9 and G10 are replaced, in order, with the figures below (Figures G15, G16, G17, G18, G19 and G20, respectively).

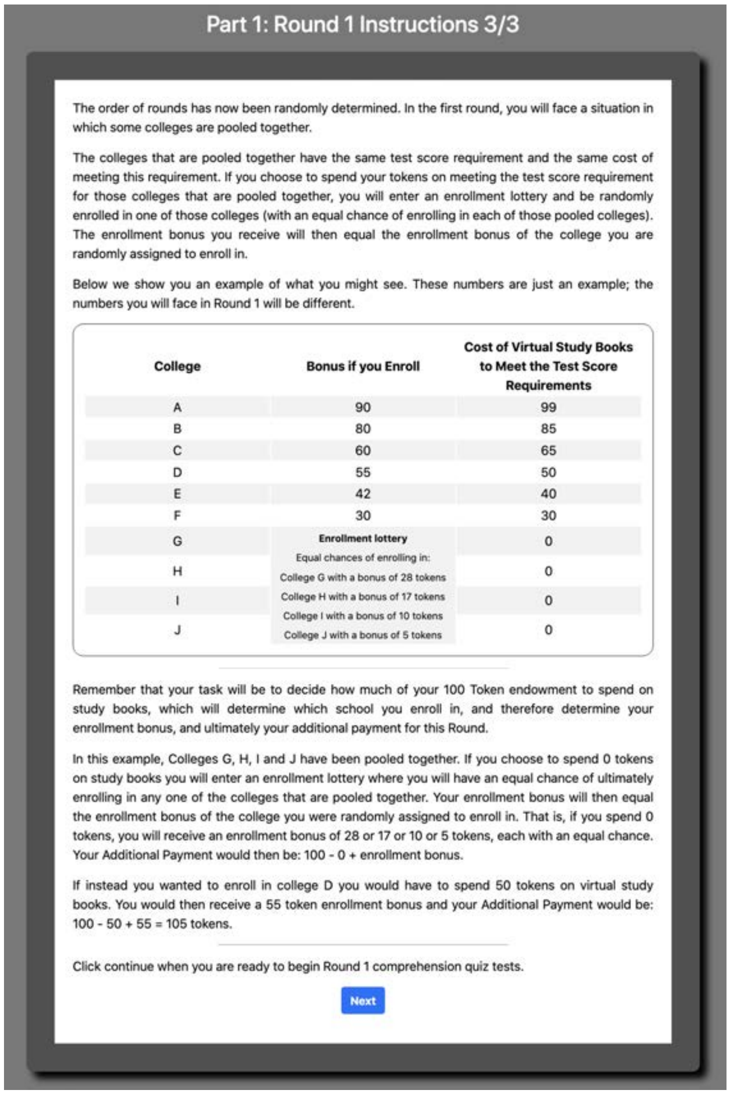

Part 1: Round 1 Instructions 3/3

The order of rounds has now been randomly determined. In the first round, you will face a situation in which some colleges are pooled together.

The colleges that are pooled together have the same test score requirement and the same cost of meeting this requirement. If you choose to spend your tokens on meeting the test score requirement for those colleges that are pooled together, you will enter an enrollment lottery and be randomly enrolled in one of those colleges (with an equal chance of enrolling in each of those pooled colleges). The enrollment bonus you receive will then equal the enrollment bonus of the college you are randomly assigned to enroll in.

Below we show you an example of what you might see. These numbers are just an example; the numbers you will face in Round 1 will be different.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	**Enrollment lottery** Equal chances of enrolling in: College G with a bonus of 28 tokens College H with a bonus of 17 tokens College I with a bonus of 10 tokens College J with a bonus of 5 tokens	0
H		0
I		0
J		0

Remember that your task will be to decide how much of your 100 Token endowment to spend on study books, which will determine which school you enroll in, and therefore determine your enrollment bonus, and ultimately your additional payment for this Round.

In this example, Colleges G, H, I and J have been pooled together. If you choose to spend 0 tokens on study books you will enter an enrollment lottery where you will have an equal chance of ultimately enrolling in any one of the colleges that are pooled together. Your enrollment bonus will then equal the enrollment bonus of the college you were randomly assigned to enroll in. That is, if you spend 0 tokens, you will receive an enrollment bonus of 28 or 17 or 10 or 5 tokens, each with an equal chance. Your Additional Payment would then be: 100 - 0 + enrollment bonus.

If instead you wanted to enroll in college D you would have to spend 50 tokens on virtual study books. You would then receive a 55 token enrollment bonus and your Additional Payment would be: 100 - 50 + 55 = 105 tokens.

Click continue when you are ready to begin Round 1 comprehension quiz tests.

Next

Figure G15: Instructions for Round 1 if pooling policy is seen first.

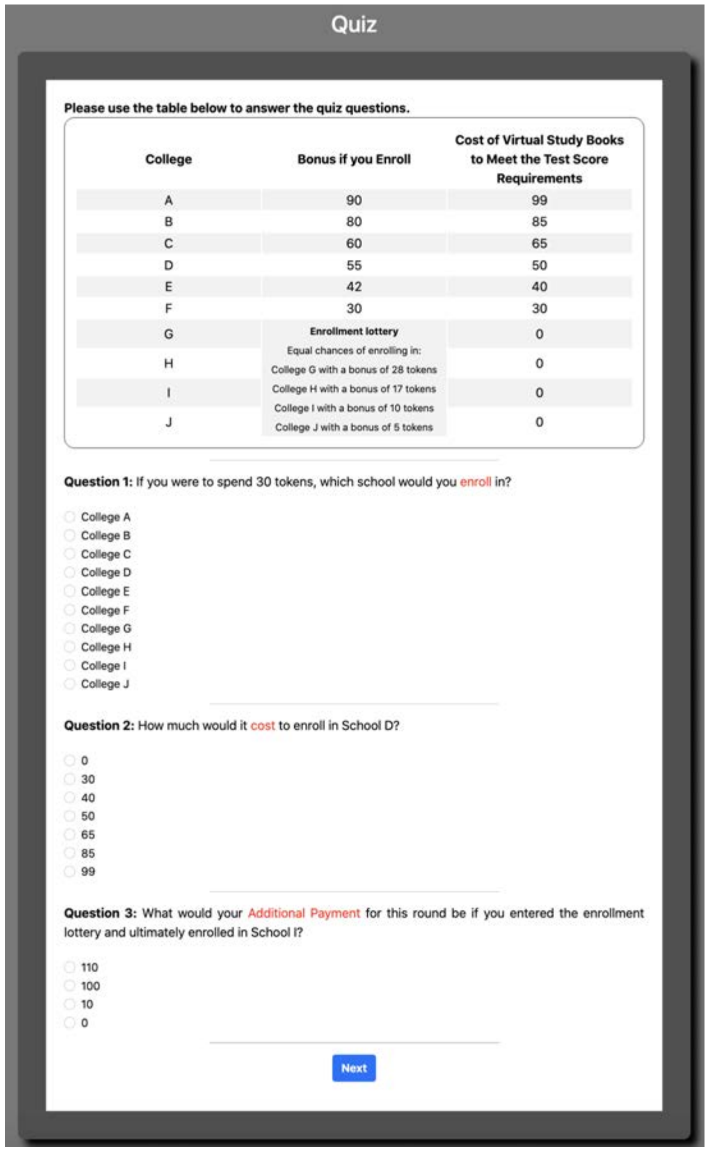

Quiz

Please use the table below to answer the quiz questions.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	**Enrollment lottery** Equal chances of enrolling in: College G with a bonus of 28 tokens College H with a bonus of 17 tokens College I with a bonus of 10 tokens College J with a bonus of 5 tokens	0
H		0
I		0
J		0

Question 1: If you were to spend 30 tokens, which school would you enroll in?

- College A
- College B
- College C
- College D
- College E
- College F
- College G
- College H
- College I
- College J

Question 2: How much would it cost to enroll in School D?

- 0
- 30
- 40
- 50
- 65
- 85
- 99

Question 3: What would your Additional Payment for this round be if you entered the enrollment lottery and ultimately enrolled in School I?

- 110
- 100
- 10
- 0

Next

Figure G16: Quiz after pooling round instructions.

Application Form

Instructions

Round 1: please make a selection

Colleges E, F, G, H, I and J have been pooled together. If you choose to spend 0 tokens on study books you will enter an enrollment lottery where you will have an equal chance of ultimately enrolling in any one of the colleges that are pooled together. Your enrollment bonus will then equal the enrollment bonus of the college you were randomly assigned to enroll in. That is, if you spend 0 tokens, you will receive an enrollment bonus of 42 or 35 or 30 or 24 or 16 or 3 tokens, each with an equal chance (for your information, on average this equals 25 tokens).

	College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
	A	74	127
	B	62	78
	C	55	57
	D	49	43
	E	**Enrollment lottery** Equal chances of enrolling in:	0
	F	College E with a bonus of 42 tokens	0
	G	College F with a bonus of 35 tokens	0
	H	College G with a bonus of 30 tokens	0
	I	College H with a bonus of 24 tokens	0
	J	College I with a bonus of 16 tokens College J with a bonus of 3 tokens	0

Submit

Figure G17: Example of selection screen for pooling policy.

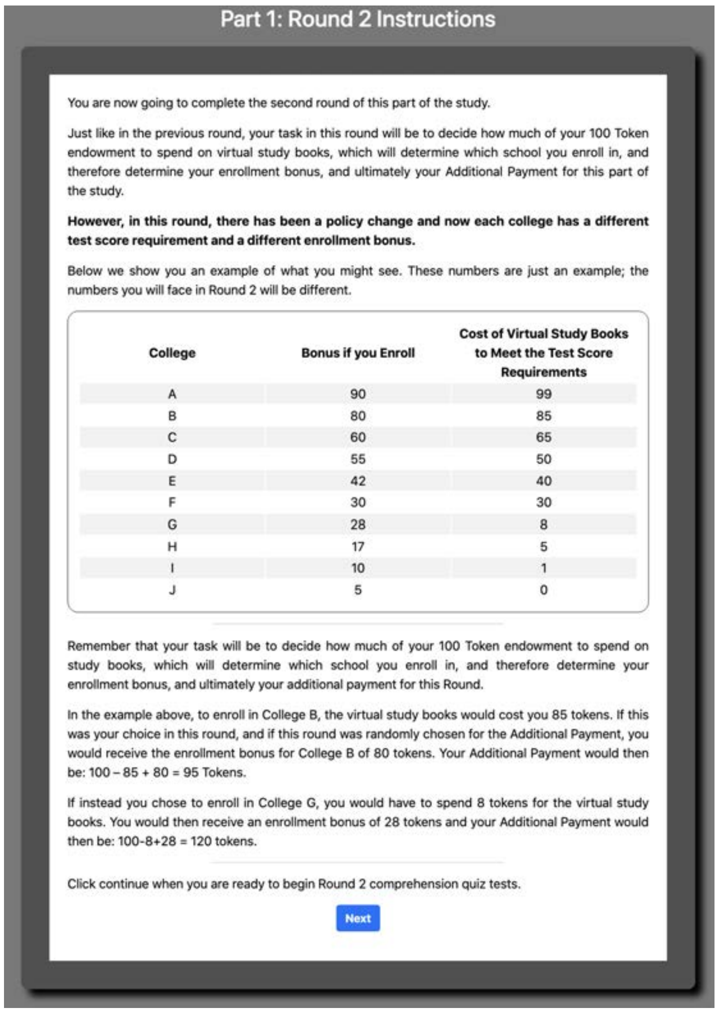

Part 1: Round 2 Instructions

You are now going to complete the second round of this part of the study.

Just like in the previous round, your task in this round will be to decide how much of your 100 Token endowment to spend on virtual study books, which will determine which school you enroll in, and therefore determine your enrollment bonus, and ultimately your Additional Payment for this part of the study.

However, in this round, there has been a policy change and now each college has a different test score requirement and a different enrollment bonus.

Below we show you an example of what you might see. These numbers are just an example; the numbers you will face in Round 2 will be different.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	28	8
H	17	5
I	10	1
J	5	0

Remember that your task will be to decide how much of your 100 Token endowment to spend on study books, which will determine which school you enroll in, and therefore determine your enrollment bonus, and ultimately your additional payment for this Round.

In the example above, to enroll in College B, the virtual study books would cost you 85 tokens. If this was your choice in this round, and if this round was randomly chosen for the Additional Payment, you would receive the enrollment bonus for College B of 80 tokens. Your Additional Payment would then be: 100 – 85 + 80 = 95 Tokens.

If instead you chose to enroll in College G, you would have to spend 8 tokens for the virtual study books. You would then receive an enrollment bonus of 28 tokens and your Additional Payment would then be: 100-8+28 = 120 tokens.

Click continue when you are ready to begin Round 2 comprehension quiz tests.

Next

Figure G18: Instructions for the discrete policy.

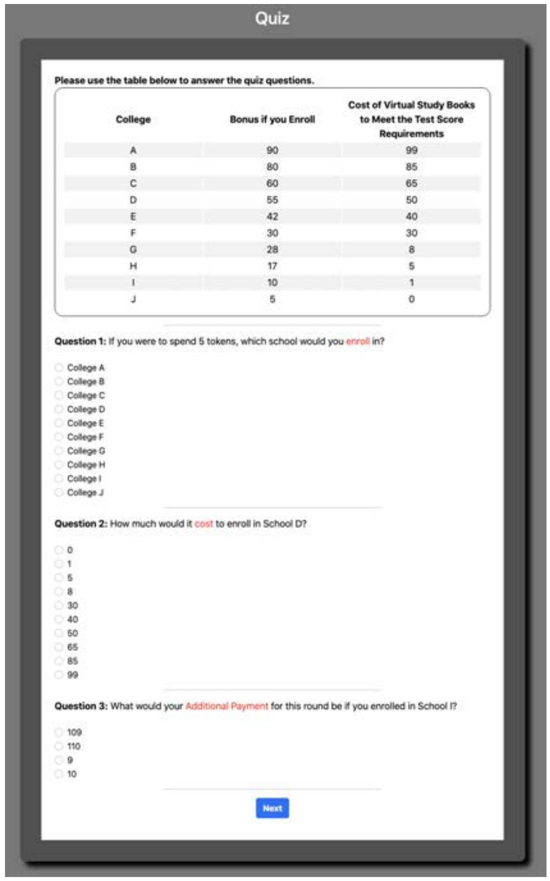

Quiz

Please use the table below to answer the quiz questions.

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	90	99
B	80	85
C	60	65
D	55	50
E	42	40
F	30	30
G	28	8
H	17	5
I	10	1
J	5	0

Question 1: If you were to spend 5 tokens, which school would you enroll in?

- College A
- College B
- College C
- College D
- College E
- College F
- College G
- College H
- College I
- College J

Question 2: How much would it cost to enroll in School D?

- 0
- 1
- 5
- 8
- 30
- 40
- 50
- 65
- 85
- 99

Question 3: What would your Additional Payment for this round be if you enrolled in School I?

- 109
- 110
- 9
- 10

Next

Figure G19: Quiz after discrete round instructions.

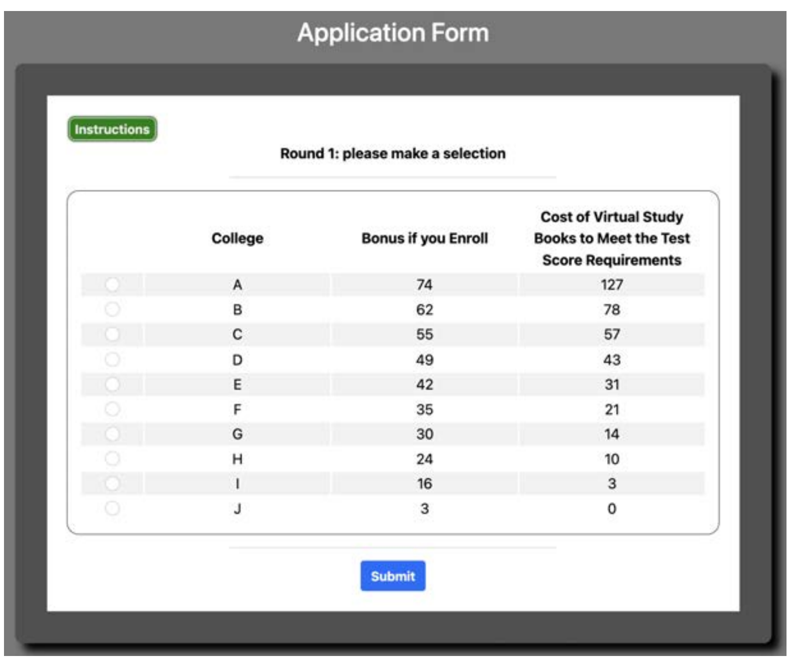

Application Form

Instructions

Round 1: please make a selection

College	Bonus if you Enroll	Cost of Virtual Study Books to Meet the Test Score Requirements
A	74	127
B	62	78
C	55	57
D	49	43
E	42	31
F	35	21
G	30	14
H	24	10
I	16	3
J	3	0

Submit

Figure G20: Example of selection screen for discrete policy.