Pareto Improvements in the Contest for College Admissions*

Kala Krishna[†] Sergey Lychagin[‡] Wojciech Olszewski[§] Ron Siegel[¶]
Chloe Tergiman[‖]

December 2024

Abstract

Many countries base college admissions on a centrally-administered test. Students invest a great deal of resources to improve their performance on the test, and there is growing concern about the high costs associated with these activities. We consider modifying the test by introducing performance-disclosure policies that pool intervals of performance rankings. Pooling affects the equilibrium allocation of students to colleges,

*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, Ran Shorrer, 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, SITE market design conference, and the Southampton Winter Workshop in Economic Theory for very helpful comments and suggestions. Financial support from the NSF (grant SES-1325968) is gratefully acknowledged.

[†]Department of Economics, The Pennsylvania State University, University Park, PA 16801, USA, NBER, CES-IFO, and IGC. Email: kmk4@psu.edu.

[‡]WIFO Institute, Arsenal, Objekt 20, 1030 Vienna, Austria. Email: lychagin@gmail.com.

[§]Department of Economics, Northwestern University, Evanston, IL 60208, USA. Email: wo@northwestern.edu.

[¶]Department of Economics, The Pennsylvania State University, University Park, PA 16801, USA. Email: rus41@psu.edu. Corresponding author.

[‖]Smeal College of Business, The Pennsylvania State University, University Park, PA 16801, USA. Email: cjt16@psu.edu.

which hurts some students and benefits others, but also affects students' effort. We investigate how such policies can improve students' welfare in a Pareto sense, study the Pareto frontier of pooling policies, and 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. A laboratory experiment based on the estimated parameters generally supports our theoretical predictions.

1 Introduction

College and university admissions are often determined by students' performance on a centrally administered test. This is commonly the case in Brazil, China, Russia, South Korea, and Turkey. The students with the highest performance are admitted to the best colleges, those ranked below them are admitted to the next best colleges, etc. Other countries consider additional factors, but even then centralized tests typically play an important role.

Consequently, students invest a great deal of effort preparing for these tests. In China, Japan, South Korea, and Taiwan, students attend specialized "cram schools," which focus on improving students' performance on the tests. This consists of rote learning, solving many practice problems, and practicing test-taking strategies tailored to the specific test. Students also hire tutors, buy books, and take specialized courses to improve their test scores. But these costly activities are far less likely to generate substantial gains in students' productive human capital.

Reducing such activities is more difficult than it might initially appear. Passing laws to prohibit or limit them may be both difficult and ineffective.[1] 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 helped some students but hurt others would likely face significant resistance.

This paper investigates simple modifications to centralized tests that 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.[2] We consider performance-disclosure policies, which pool together intervals of performance and assign the same score to all performances in an

[1] 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 could not be seen from outside."

[2] Our analysis may also be applied to other large contest settings such as large corporate promotion contests (we thank a referee for suggesting this example) and large grant competitions.

interval. Students with the same score are randomly admitted to the corresponding fraction of colleges.[3] For example, a "bottom pooling" policy that pools some fraction of the lowest performing students assigns these students randomly to the same fraction of the lowest-ranked college seats. Performance-disclosure policies do not require changing the tests or introducing new components to the admissions process. They also respect the property that a higher score leads to a better expected college assignment than a lower score. This may help make such policies appealing to policy makers.

We study Pareto improving performance-disclosure policies, which benefit all students. Our notion of Pareto improvements is an interim one, once a student knows her score but before she learns her college assignment. Pareto improving policies often exist, because test preparation is costly. Relative to the baseline contest, introducing a performance-disclosure policy leads to some students being admitted to higher-ranked colleges with positive probability; this makes them better off even if they incur higher costs, provided the cost increase is not too large. Other students are admitted to lower-ranked colleges with positive probability; if they also incur lower costs they are made better off, provided the reduction in the costs is large enough.

We first characterize the Pareto improving policies that pool a single interval of performance. The characterization shows that such pooling 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 the uniform distribution. We then generalize this condition to policies with multiple pooling intervals and characterize the distributions of players' types and the distributions of college seats for which Pareto improving policies exist. We also consider mean-preserving contractions (MPCs) of the distribution of college seats, which corresponds to pooling groups of college seats.[4] We characterize the distributions of players' types and the distributions of college seats for which Pareto improving MPCs exist and the Pareto frontier of such MPCs.

[3] This can be viewed as making performance on the test noisier. Morgan et. al. (2022) suggest that other forms of noise can also be socially beneficial.

[4] We thank the referees for encouraging us to consider MPCs.

We then consider robust Pareto improving performance-disclosure 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. The characterization may be particularly useful for empirical work because it only requires obtaining an estimate of the students' ability distribution.

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) along with novel techniques 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. Among the many performance-disclosure policies on the Pareto frontier of Pareto improving policies, we focus on the one that maximizes the utility of the applicants with the lowest ability. This is a bottom pooling policy, which pools approximately 63 percent of the lowest test scores and increases applicants' estimated utility by approximately 32 percent. It also maximizes applicants' aggregate welfare among all bottom pooling policies.[5]

Finally, we conduct a laboratory experiment based on the calibrated distributions and the 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 aregue that a possible explanation for this behavior may be a preference for randomization, similarly to the findings of Dwenger, Kübler, and Weizsäcker (2016) in the context of school applications. 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, presents the equilibrium, and defines the notion of Pareto improvements. Section 3 investigates policies with a single pooling interval. Section 4 investigates policies with multiple pooling intervals. Section 5 derives the conditions for

[5]We additionally show that the robust Pareto improving policy is also a bottom pooling policy, which pools approximately 52 percent of the lowest test scores.

robust Pareto improvements. Section 6 discusses some limitations and extensions of the model. Section 7 describes the empirical exercise. Section 8 presents the experimental results. Section 9 concludes. The appendix contains proofs, considers the Pareto frontier of Pareto-improving policies that pool on multiple intervals, and examines peer effects. The online appendix contains additional details about the empirical exercise, estimation strategy, and counterfactuals, and provides additional details, results, and screenshots from our experiment.

1.1 Contribution to the literature

The work by Che et al. (2018) is the most closely related to the theoretical part of our paper. They study auction formats for a single object that are immune to collusion by bidders, and identify optimal cartels. Given a multi-bidder auction and an equilibrium of the auction, they model a cartel as a mechanism to which the bidders report their types and which bids on their behalf in the auction. The auction is immune to collusion if no such mechanism exists whose outcome is weakly preferred by all types of every bidder to the equilibrium of the auction, with a strict preference for some type of some bidder. Their Theorem 1 provides necessary and sufficient conditions for an auction to be immune to collusion. We study contests with many players and prizes. We use Olszewski and Siegel's (2016) large contest framework to show that any equilibrium is approximated by a single-agent mechanism that implements the assortative allocation of the prizes to agent types. Thus, the motivating questions are different, the settings are different (for example, they have multiple bidders whereas we have a single limiting agent), and the set of manipulations are different (all possible mechanisms in their setting versus pooling intervals of performance in our setting). Nevertheless, our characterization of contests for which no Pareto improvement exists (Theorem 1 and, for mean-preserving contractions of the prize distribution, Theorem 2) is similar to their Theorem 1, and relies on a condition very similar to their condition (PS). The intuition on their pages 411-412 for why the condition is necessary is similar in both results.[6] Our Theorem 3 also characterize the Pareto frontier of Pareto-improving mean-

[6]But the analysis, especially the proof of sufficiency, is different because of the different sets of manipulations. Perhaps most intuitive is the comparison between our setting and the "single-bidder version" of

preserving contractions, and this result is similar to Theorem 2 of Che et al. (2018) although its proof is different. However, the corresponding Pareto frontier of category rankings, which are the main objects of our analysis, differs from the Pareto frontier of mean-preserving contractions (see Example 2 in Appendix C), and our results on the Pareto frontier of category rankings have no counterparts in Che et al. (2018).

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.[7] In a two-sided matching model with ex-ante symmetric agents and costly signals, Hoppe et al. (2009) provide conditions 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. Other papers that compare contests and lotteries from the perspective of contestants' welfare include Taylor et al. (2003), Koh et al. (2006), Hoppe et al. (2011), and Chakravarty and Kaplan (2013). Condorelli (2012) characterizes the ex-ante efficient allocations of heterogeneous objects to heterogeneous agents with private valuations.[8] We are interested in interim Pareto improvements, and the set of ex-ante efficient allocations can be completely different from those that are interim Pareto improving. Olszewski and Siegel (2016) introduced the equilibrium approximation approach to large contests, which we use here.[9] Bodoh-Creed and Hickman (2018) use a similar approach to study quotas and affirmative action in college admissions. One of their findings is that a college assignment lottery would generate higher total student welfare than a college admissions contest. We show that a lottery may be improved upon for all students by partitioning the set of students

their setting in which the strategies of the other bidders are given. Then, if a subset of types in their setting deviates to bidding as some other type, the deviators obtain the same allocation as the other type. In our setting, because the set of prizes is fixed, the deviation allocation depends on the subset of deviating types, and this allocation usually differs from that of the other type.

[7]This happens when the ability distribution is sufficiently concave, in which case our results show that pooling is not Pareto improving. Dubey and Geanakoplos (2010) consider a game of status between students and show that coarse grading policies maximize effort.

[8]His main insights apply when all players' type distributions have monotone hazard rates. We do not require such a condition.

[9]Olszewski and Siegel (2020) use this approach to study performance-maximizing contests. Fang, Noe, and Strack (2020) study the effect of different prize structures on aggregate effort in symmetric all-pay auctions with complete information.

into several categories based on their performance and using a lottery in each category.[10]

Our empirical application is closely related to three strands of literature in empirical industrial organization. First, as we recover the distribution of placement values in the Turkish college admission system, we heavily rely on the method of estimating single-agent dynamic problems introduced in Hotz and Miller (1993); Krishna et al. (2018) adapt this method to the context of Turkey. Second, our estimates of the costs of effort are based on Berry (1994), who shows that mean values of choices can be backed out from the observed shares of agents making these choices. Finally, our approach to estimating the distribution of student ability is novel and thus does not have direct antecedents in the literature. To some degree, we draw our inspiration from Guerre et al. (2000), who infer the distribution of bidder private values from the distribution of observed bids in first-price auctions, which has parallels with the way we back out the distribution of ability from student test scores.

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 all the players agree that a prize with a higher value is better. 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. After privately observing her type, each player exerts costly effort to achieve her desired performance $t \geq 0$ on the test. 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). Thus, each player is admitted to the best college among those with available seats after all the players with a higher performance have been admitted. Ties are resolved by a fair

[10]The optimality of coarse partitions with random lotteries within elements of the partitions has been studied in less closely related papers, including Chao and Wilson (1987), Wilson (1989), and McAfee (2002) in the context of priority classes and Damiano and Li (2007) and Rayo (2013) in monopolistic settings. The effects of different grading policies has been studied by Ostrovsky and Schwarz (2010), Gottlieb and Smetters (2014), Boleslavsky and Cotton (2015), and Harbaugh and Rasmusen (2018). Frankel and Kartik (2019) show that test preparation not available to all students can diminish the signals contained in standardized tests.

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 $\lim_{t\to\infty} c(t) = \infty$.[11] 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 let $g_1(x)\,g_2(x) = x$. 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).[12] For convenience, throughout our theoretical analysis we will use utility (2), and in the empirical analysis of Section 7 we will use utility (3). All of our theoretical results hold without change for any utility (1). Section 6 discusses limitations of the model.

To solve for players' equilibrium behavior, we assume that players' types are drawn independently (but not necessarily identically) across players, and we apply the large contests results of Olszewski and Siegel (2016). These results show that players' equilibrium behavior is approximated by a particular single-agent mechanism, which assortatively allocates prizes to agent types and gives the lowest type a utility of 0. To describe this mechanism, we

[11] 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.

[12] Different functions g_1 have different implications for aggregate welfare, but this makes no difference for our theoretical analysis because we focus on policies that make all students better off (we provide a precise definition in Section 2.1).

denote by F the CDF of the average of players' type distributions and assume that F has a continuous, strictly positive density f. We denote by G the CDF that represents the empirical distribution of prizes, so the size of each atom of G corresponds to the number of seats in a particular college (or tier of colleges), with $G(0)$ representing the fraction of students in excess of the total number of college seats.[13] 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 (average) type distribution. 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

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

Olszewski and Siegel (2016) show that in any equilibrium of a contest with many players and prizes, for every player (except a small fraction of the players) the event that the player's type is some x, she chooses a performance close to $t^A(x)$, and obtains a prize close to $y^A(x)$, which gives her a utility close to $U(x)$, has probability close to 1.

The rest of the paper uses the approximating single-agent mechanism to investigate how different performance-disclosure policies affect students' welfare in a Pareto sense, which we define in the next subsection.

[13] If there are n players and k prizes have value y or lower (including "null prizes" with value 0), then $G(y) = k/n$.

2.1 The notion of Pareto improvements

We use the term "Pareto-improving" in reference to types' expected utility in an approximating mechanism. This utility corresponds to players' interim utility - after players' scores are realized but before they learn to which college they are admitted. A performance-disclosure policy is Pareto improving if in the corresonding approximating mechanism all types are better off, and a positive measure of types is strictly better off. Such an improvement implies that in a sufficiently large contest some players are strictly better of (in an interim sense) and no player is worse off by more than an arbitrarily small amount; moreover, the sum of these small amounts is arbitrarily small compared to the gains of the players who are strictly better off. We underscore that because 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.

3 Pooling on a single interval

We begin by considering pooling a single interval of performance. We denote the interval by $(q^*, q^{**}]$, where q^* and q^{**} are quantiles in the ranking of students' realized performance, with $0 \leq q^* < q^{**} \leq 1$.[14] Students whose performance ranking quantile lies in the interval are treated identically. As a group, these students still obtain the prizes they would in the baseline contest, that is, the prizes whose quantile ranking lies in the quantile interval $(q^*, q^{**}]$ of the prize distribution. But with pooling, these prizes are allocated uniformly at random to the students in the group. The resulting game is therefore different from the baseline contest. Nevertheless, we can describe the contest with pooling in an alternative, equivalent way that allows us to use the same contest framework to study the effect of pooling. We do this by considering an equivalent contest with no performance pooling in which the prizes that correspond to the pooled interval are replaced with the same mass $q^{**} - q^*$ of identical

[14] We use left-open intervals because the assortative allocation $y^A = G^{-1} \circ F$ is left continuous due to the type distribution F being continuous (by assumption) and the prize distribution G being right continuous (as a CDF). The same is true for other assignments of prizes to types.

prizes whose value is equal to the average value of the original prizes in the pooled interval.

Formally, let types x^* and x^{**} be the types at quantiles q^* and q^{**} in the type distribution, that is, $q^* = F(x^*)$ and $q^{**} = F(x^{**})$. We consider the prize distribution that results from replacing the prizes in quantile interval $(q^*, q^{**}]$ of the original prize distribution G with a mass $q^{**} - q^*$ of prize

$$y(q^*, q^{**}) = \frac{\int_{q^*}^{q^{**}} G^{-1}(z)\, dz}{q^{**} - q^*} = \frac{\int_{x^*}^{x^{**}} y^A(x) dF(x)}{F(x^{**}) - F(x^*)}. \tag{6}$$

The corresponding assortative allocation coincides with the original assortative allocation y^A for types lower than x^* and higher than x^{**}, and assigns prize $y(q^*, q^{**})$ to every type in $(x^*, x^{**}]$.

3.1 Welfare comparisons

Pooling affects both the equilibrium prize allocation and players' performance choices. To understand the overall welfare effects, our first result compares each type's utility in the approximating mechanisms with and without pooling.

Proposition 1. *Consider pooling on a quantile interval $(q^*, q^{**}]$. If all the prizes in the interval are identical, then pooling has no effect. If not all prizes in the interval are identical, then the following statements hold, where $q^* = F(x^*)$ and $q^{**} = F(x^{**})$.*

(a) Types lower than x^* are not affected.

(b) Type x in $(x^*, x^{**}]$ weakly benefits from pooling if and only if

$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x}) dF(\tilde{x})}{F(x^{**}) - F(x^*)} \geq \frac{\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}}{x - x^*}, \tag{7}$$

and strictly benefits if the inequality is strict.

(c) There is a type $\hat{x}$ in $(x^*, x^{**}]$ such that types in $(x^*, \hat{x})$ benefit from pooling and types in $(\hat{x}, x^{**})$ do not benefit from pooling.

(d) Pooling is Pareto improving if and only if it weakly benefits type x^{**}, that is,

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

and this holds if and only if pooling weakly benefits the highest type $x = 1$.

To understand the idea underlying Proposition 1, whose proof is in Appendix A, notice that players with types lower than x^* are unaffected because their performance and the prize they obtain do not change with pooling. Players with types in $(x^*, \hat{x})$ 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 pooling, which benefits them even though they choose a higher performance than without pooling. Players with types lower than but close to x^{**} obtain a prize lottery that is worse than the prize they obtain without pooling, so if pooling is Pareto improving, these players must choose a sufficiently lower performance with pooling that offsets the loss from the prize lottery.[15] To understand (7), which is the key condition in Proposition 1, multiply each side of (7) by $x - x^*$. The left-hand side of (7) is the prize lottery obtained by every type in $[x^*, x^{**}]$ in the contest with pooling, so the left-hand side of (7) multiplied by $x - x^*$ is the difference between the utilities of type $x > x^*$ and type x^* in the contest with pooling. The right-hand side of (7) multiplied by $x - x^*$ is

$$\int_{x^*}^{x} y^A\left(\tilde{x}\right)d\tilde{x} = \int_{0}^{x} y^A\left(\tilde{x}\right)d\tilde{x} - \int_{0}^{x^*} y^A\left(\tilde{x}\right)d\tilde{x},$$

which is the difference between the utilities of type $x > x^*$ and type x^* in the contest without pooling.[16] This difference in utilities is a linear function with pooling (since the prize $y(q^*, q^{**})$ is constant), and a convex function without pooling (because $y^A\left(\tilde{x}\right)$ increases in $\tilde{x}$), so type x^{**} is the most demanding: if type x^{**} benefits from pooling, all types in $[x^*, x^{**}]$

[15] If pooling is Pareto improving, then it reduces the aggregate cost of performance with utility (3) because the set of prizes is unchanged and all players are made (at least weakly) better off.

[16] Intuitively, type $\tilde{x} + d\tilde{x}$ can pretend to be type $\tilde{x}$, obtain prize $y^A\left(\tilde{x}\right)$, and enjoy a utility increase of $y^A\left(\tilde{x}\right)d\tilde{x}$ relative to type $\tilde{x}$.

benefit. This is illustrated in Figure 1.[17] The utility difference between types $x > x^{**}$ and type x^{**} is unaffected by pooling since their prize allocation does not change, which explains part (d) in the proposition.

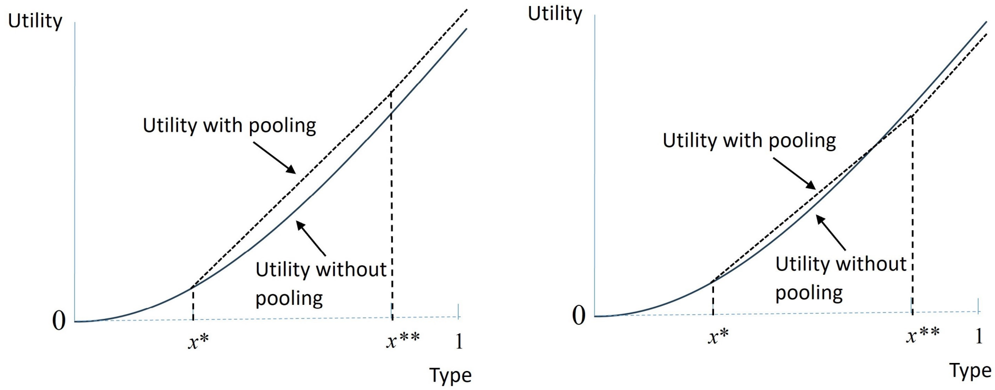

Figure 1: Pooling is Pareto improving (left) and pooling is not Pareto improving (right).

A particular class of pooling intervals are those with lower bound $q^* = 0$. We refer to such intervals as "bottom pooling" intervals, and identify each interval with its upper bound, q^{**}.[18] Such intervals will play an important role in the empirical and experimental parts of the paper.

4 Pooling on multiple intervals

We now consider more general performance disclosure policies, which pool on each of several intervals of performance ranking. We will use the term "category rankings" to describe such policies. One example is pooling above and below the median performance. Another example is pooling performances 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 K left-open intervals. The intervals are $J_k = (q_k^*, q_k^{**}]$ for $1 \leq k \leq K \leq n$,

[17] The same phenomenon is illustrated in Figure 4 of Che et al. (2018).

[18] Bottom pooling with $q^{**} = 1$ is a lottery over the entire set of prizes. This lottery is also a special case of "top pooling," which is the class of pooling intervals with upper bound $q^{**} = 1$.

where $0 \leq q_1^* < q_1^{**} \leq \cdots \leq q_K^* < q_K^{**} \leq 1$. The interpretation is that the fraction $q_k^{**} - q_k^*$ 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. Similarly to the case of a single interval, this contest is equivalent to a contest with no performance pooling in which the prizes that correspond to each pooled interval J_k are replaced with a mass $q_k^{**} - q_k^*$ of prize

$$y(J_k) = \frac{\int_{q_k^*}^{q_k^{**}} G^{-1}(z)\, dz}{q_k^{**} - q_k^*} = \frac{\int_{x_k^*}^{x_k^{**}} y^A(x) dF(x)}{F(x_k^{**}) - F(x_k^*)}, \tag{8}$$

where $q_k^* = F(x_k^*)$ and $q_k^{**} = F(x_k^{**})$. Thus, the category ranking $\mathcal{J}$ induces a partition $\mathcal{I}$ of the set of types $X = [0,1]$ into singletons and the K intervals $I_k = (x_k^*, x_k^{**}]$, such that in the approximating mechanism of the category ranking all types in interval I_k choose the same performance and obtain the same prize $y(J_k)$, and singleton types obtain the prize they did in the original approximating mechanism. In what follows, it will be convenient to consider such partitions of the set of types into singletons and left-open intervals 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.

4.1 Welfare comparisons

For category rankings that include more than one interval, a generalization of the conditions in Proposition 1 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, I_1 say, may increase the utility of types in a higher interval, I_2 say, to such a degree that the net effect on all types is positive even if pooling on I_2 in isolation lowers the utility of some types in I_2 relative to the baselines contest. Relatedly, a category ranking that benefits the highest type $x = 1$ is no longer necessarily Pareto improving.[19]

[19]Both of these phenomena arise because with more than one pooling interval the equivalent of part (a) in Proposition 1 no longer holds: for an interval I_k that is not the lowest one in the category ranking (so

To obtain the sufficient conditions, consider a category ranking $\mathcal{I}$ that consists of 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 1 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 result as a corollary of Proposition 1.

Proposition 2. *(a) Type x in $I_k = (x_k^*, x_k^{**}]$ weakly benefits from the category ranking $\mathcal{I} = \{I_1, ..., I_K\}$ if*

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

and strictly benefits if any of the inequalities is strict.

(b) The category ranking is Pareto improving if

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

The next subsection characterizes when Pareto improving category rankings exist and discusses the Pareto frontier of category rankings.

4.2 Existence of Pareto improving category rankings and their Pareto frontier

We now provide a condition that characterizes contests for which Pareto improvements exist. For the condition, denote by $\widehat{F}$ the concavification (concave closure) of F, that is, the lowest-

$k > 1$), the utility of types lower than x_k^* may be lower or higher than in the baseline contest, depending on the effect of pooling on lower intervals.

valued concave function that is (pointwise) weakly higher than F.

Theorem 1. *There does not exist a Pareto improving category ranking if and only if for any interval of types on which $\widehat{F}$ is linear, in the assortative allocation all the types in the interval obtain the same prize, that is, $y^A(\cdot)$ is constant on the interval.*

The proof of Theorem 1 is in Appendix A. For some intuition, recall that by part (d) of Proposition 1, pooling on an interval is Pareto improving if and only if it benefits the highest type in the interval. When not all the types in an interval obtain the same prize, without pooling the highest (and best-performing) type obtains a prize that is (interim) better than the prize lottery induced by pooling. But pooling also eliminates the competition between the types in the interval, since they all obtain the same prize lottery. This competition is more intense, leading to more costly performance, if there are relatively more high types in the interval, since they are willing to pay relatively more for the higher prizes in the interval. The relative frequency of the types in the interval, which determines this trade-off between a worse (interim) prize and a less costly performance, is precisely captured by the comparison between F and $\widehat{F}$.

The condition in Theorem 1 is essentially condition (PS) of Che et al. (2018). The statement of Theorem 1 is very similar to what would be a single-bidder version of their Theorem 1.[20] The proof of the necessity of the condition is similar in both results, and the proof of the sufficiency of the condition in our result is much simpler because the settings are completely different.

For the distribution F that we estimate in Section 7 there is a single (maximal) interval of types on which the concavification $\widehat{F}$ is linear and not all the types in the interval obtain the same prizes. The lower bound of this interval is 0, and we investigate Pareto improving bottom pooling policies that include this interval. The distribution and its concavification are illustrated in Figure 6.

Turning to the Pareto frontier of Pareto improving category rankings, in Appendix C we provide a method for checking whether a Pareto-improving category ranking can be

[20] But their setting stipulates multiple bidders. In a setting like theirs with a single bidder, or with multiple bidders when fixing the equilibrium strategies of all but one of the bidders, the remaining bidder has no profitable manipulation (no "Pareto improvement"), regardless of his type distribution.

further Pareto improved. Using this method, we provide the following sufficient condition for a Pareto-improving category ranking to belong to the Pareto frontier. The proof is in Appendix C.

Proposition 3. *A Pareto-improving category ranking belongs to the Pareto frontier of Pareto-improving category rankings if on any interval of types on which $\widehat{F}$ is linear, all types in the interval obtain the same prize (or prize lottery) in the allocation induced by the category ranking.*

The condition in Proposition 3 is only sufficient. Example 2 in Appendix C shows that the Pareto frontier may also contain category rankings that do not satisfy the condition. As we discuss in the next subsection, however, the condition becomes sufficient and necessary if one considers mean-preserving contractions of the prize distribution instead of category rankings.

4.3 Mean-preserving contractions

Category rankings pool intervals of performance. We investigated the effect of these policies by using the fact that they are equivalent to pooling intervals of prizes. We now consider mean-preserving contractions of the prize distribution G, which generalize category rankings when the latter are viewed as pooling intervals of prizes instead of intervals of performance. Formally, distribution H is a *mean-preserving contraction* (henceforth: MPC) of G (defined on the same domain as G) if H second-order stochastically dominates G. In particular, category rankings can be viewed as MPCs of the original prize distribution. An MPC is an inverse of a mean-preserving spread, the latter of which is obtained by adding to each outcome of the original distribution a random variable with mean zero.[21] Adding a random variable can be interpreted as distributing the "mass" assigned to each outcome over possibly other outcomes. Therefore, an MPC of the prize distribution can be interpreted as dividing the prizes into into groups, and replacing the mass of each group with its expected value (or, equivalently, with a lottery over the prizes in the group). Moreover, any grouping of this kind results in an MPC of the original prizes. Thus, all MPCs of the the prize distribution,

[21] See Mas-Colell et al. (1995), Example 6.D.2 and Proposition 6.D.2.

and not only category rankings, can be implemented with an appropriate grouping policy. But implementing some MPCs involves policies that seem less realistic, or at least more controversial, than category rankings. The following example demonstrates this.

Example 1. *Suppose there are equal masses of colleges of quality 1, 3, 4, and 6. Consider an MPC that groups colleges 1 and 4 together and colleges 3 and 6 together. The average college quality is 2.5 in the first group and 4.5 in the second group. Applicants would be classified as "high performance" or "low performance" depending on whether their performance exceeds the median. Low performance applicants would be assigned to a college in the first group, and high performance applicants would be assigned to a college in the second group. Thus, some high performance applicants would be assigned to college 3 while some low performance applicants would be assigned to college 4. This is likely to be controversial.*

Even though some MPCs may be unrealistic, MPCs may be of theoretical interest. We therefore provide two results, whose proof is in Appendix B. The first result characterizes when Pareto-improving MPCs exist, generalizing Theorem 1.

Theorem 2. *There does not exist a Pareto-improving MPC if and only if for every interval of types on which $\widehat{F}$ is linear, in the assortative allocation all the types in the interval obtain the same prize, that is, $y^A(\cdot)$ is constant on the interval.*

The second result characterizes the Pareto frontier of Pareto-improving MPCs.[22]

Theorem 3. *An MPC belongs to the Pareto frontier of Pareto-improving MPCs if and only if on any interval of types on which $\widehat{F}$ is linear, all types in the interval obtain the same prize (or prize lottery) in the allocation induced by the MPC.*

5 Robust Pareto improvements

The results in the previous sections involve both the type distribution F and the prize distribution G (via the assortative allocation $y^A = G^{-1} \circ F$). We now present more robust

[22]Recall that Proposition 3 only provided a sufficient condition for a Pareto-improving category ranking to belong to the Pareto frontier.

results that involve only the type distribution F. These results may be useful for empirical work because their lack of reliance on G frees the analyst from making assumptions about how various college attributes that students may value are aggregated into a unidimensional prize value.[23] We point out, however, that the results still assume that students agree on the ranking of colleges. 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."[24] Our main robustness results characterize robust Pareto improving single-interval policies and category rankings, as well as their corresponding Pareto frontiers.

We begin with a characterization of robust Pareto improving single-interval policies.

Proposition 4. *Pooling on a quantile interval $(q^*, q^{**}]$ is robust Pareto improving if and only if*

$$\frac{F(x) - F(x^*)}{F(x^{**}) - F(x^*)} \leq \frac{x - x^*}{x^{**} - x^*} \tag{9}$$

for every x in $(x^*, x^{**}]$, where $q^* = F(x^*)$ and $q^{**} = F(x^{**})$.

Proposition 4 follows from part (d) of Proposition 1 because $y^A(\tilde{x}) = G^{-1}(F(x))$ can be an arbitrary increasing function with values in $[0, 1]$ for an appropriate G, and (9) states that distribution F restricted to the interval $[x^*, x^{**}]$ first-order stochastically dominates (FOSD) the uniform distribution on $[x^*, x^{**}]$.

The following result characterizes robust Pareto improving category rankings.

Proposition 5. *Category ranking $\mathcal{I} = \{I_1, ..., I_K\}$ is robust Pareto improving if and only if for every interval $I_k = (x_k^*, x_k^{**}]$ in $\mathcal{I}$ we have that*

$$\frac{F(x) - F(x_k^*)}{F(x_k^{**}) - F(x_k^*)} \leq \frac{x - x_k^*}{x_k^{**} - x_k^*} \tag{10}$$

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

Sufficiency of the condition in Proposition 5 follows from part (b) of Proposition 2 because (10) states that distribution F restricted to each interval $[x_k^*, x_k^{**}]$ first-order stochastically

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

[24] This is different from robustness to the underlying information structure studied in the mechanism design literature.

dominates the uniform distribution on $[x_k^*, x_k^{**}]$. Necessity follows from Proposition 4 by observing that for every interval $I_k = (x_k^*, x_k^{**}]$ there are prize distributions G such that $y^A(\tilde{x}) = G^{-1}(F(x))$ is an arbitrary increasing function on $(x_k^*, x_k^{**}]$ with values in $[0,1]$ and is constant below x_k^* and above x_k^{**}.

Proposition 4, Proposition 5, and Theorem 1 imply the following characterizations of the Pareto frontiers of robust Pareto improving single-interval pooling policies and category rankings. The proofs are in Appendix A.

Theorem 4. *A single-interval pooling policy is on the Pareto frontier of robust Pareto improving single-interval pooling policies if and only if the interval is a maximal interval on which $\widehat{F}$ is linear.*

Theorem 4 shows that to obtain an unimprovable Pareto improving single-interval pooling policy that is robust to the prize distribution, one should pool on a maximal interval on which $\widehat{F}$ is linear. If, however, one is able to estimate the prize distribution, then other single-interval pooling policies may be on the Pareto frontier. Section 7 shows this in our empirical exercise.

Theorem 5. *If the number of maximal intervals on which $\widehat{F}$ is linear is finite, then the Pareto frontier of robust Pareto improving category rankings consists of the single category ranking that pools on every maximal interval on which $\widehat{F}$ is linear. Otherwise, the Pareto frontier is empty.*

The distinction in Theorem 5 between a finite and infinite number of maximal intervals on which $\widehat{F}$ is linear arises because, by definition, a category ranking consists of a finite number of pooling intervals.[25]

[25]This definition suffices for empirical applications and leads to technically simple proofs. Modifying the definition to allow for a countably infinite number of intervals would lead to the Pareto frontier consisting of the single category ranking that pools on every maximal interval on which $\widehat{F}$ is linear even when the number of such intervals is infinite.

6 Discussion: limitations and extensions

Our analysis relies on several assumptions. First, we stipulate a common ordinal ranking of college quality across students.[26] In reality, students often vary in how they rank colleges. Second, we assume that students can choose their performance precisely, without any "noise." This is obviously not the case in practice. In reality, students' test performance is often noisy, and the noisier the relationship between students' preparation efforts and their performance on the test, the less applicable our analysis of Pareto improvements. We make these restrictive assumptions because they are required for our use of Olszewski and Siegel's (2016) large contest framework.[27] Third, we assume that test preparation is costly, as in Spence (1973). This should be interpreted as net of any direct benefit from the preparation activities. This is most appropriate for activities geared specifically toward improving students' performance on the test, as discussed in the introduction. The model is less suitable for countries in which high-school performance or other activities that have significant direct benefits at moderate levels of investment play an important role in college admissions.[28]

In our model, a student's valuation for being admitted to a college does not depend on which other students are admitted to the same college. In Appendix D we consider a setting in which each type x exerts a peer effect $p(x)$, and each student in a college experiences the average effects of the other students in the college. 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. This generates a new prize distribution, and the rest of the analysis is unchanged.

Our focus on pooling intervals of performance on the test has some practical advantages.

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

[27] With heterogeneous ordinal rankings, a higher type would not necessarily choose a higher performance. With noisy performance, every effort choice would map to an endogenous distribution over prizes. In both cases, the limiting allocation of prizes to types would no longer be assortative, which is required for our techniques.

[28] But even there the costs may exceed the benefits past a certain point, as Bodoh-Creed and Hickman (2024) demonstrate in the context of college admissions in the United States. 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.

First, pooling intervals of performance does not require changing any of the prizes (college seats); it only requires that the prizes corresponding to a set of students with pooled performances be allocated uniformly at random to these students. This makes such a policy potentially easy to implement. Second, pooling intervals of performance maintains the property that a higher performance is better: a higher performance leads to a lottery over prizes that are all better than the prizes in any different lottery associated with a lower performance. As we discuss in Section 4.3, MPCs of the prize distribution, which generalize pooling intervals of prizes, may violate this property.

7 Empirical illustration with Turkish data

We provide an empirical "proof of concept" by applying our theory to obtain Pareto improvements in a college admissions setting. We first estimate the primitives of our theory: a type distribution F, a prize distribution G, and a cost function c. Our estimation uses data on college applications of Turkish high school students.[29] These students invest in tutoring and obtaining admission to selective schools, and take the college entrance exam at the end of high school. They can retake the exam every year, and only the last attempt is considered. We extend the structural model of Krishna et al. (2018) to leverage the richness of the data, including students' choice to retake the exam, and estimate the primitives of the theoretical model.[30] We then use the estimated primitives to simulate a contest based on our theoretical model and derive Pareto improvements for this contest. We first apply our results on robust Pareto improvements and show that the Pareto optimal robust Pareto improving category ranking is a bottom pooling policy. We then consider the Pareto frontier of (non-robust) Pareto improving category rankings and show that the category ranking on the frontier that maximizes the utility of the students with the lowest ability is also a bottom pooling policy, which pools an even larger fraction of the students.

In the structural model, students participate in a large contest in a stationary overlapping

[29]Details regarding the data and the university entrance exam system in Turkey can be found in Krishna et al. (2018).

[30]We distinguish between the *theoretical model*, laid out in Section 2, and the *structural model*, which we estimate below.

generations environment. Students are heterogeneous in ability, and each student solves a single-agent infinite-horizon dynamic problem. In each period, the student faces uncertainty over future exam scores. The student makes an initial effort choice, which corresponds to schooling and tuition and determines her expected competitive standing when she takes the college entrance exam for the first time. Then the student learns her score, which is equal to her expected score plus noise, and decides whether to accept the corresponding placement or retake the exam. If the student retakes the exam, she incurs a retaking cost, which can depend on the number of times she has retaken the exam but is otherwise homogeneous across students. Students who retake the exam draw a new score without new effort choices, and so on.[31] In the steady state of this problem, the mass of students taking the exam at any point in time (fresh high school graduates and retakers) is constant.

We use the structural model and students' test retaking decisions to estimate, for every realized score, the value of obtaining that score in a particular retaking attempt.[32] This value includes the value of placement at that realized score and the option value of retaking the exam. From these estimates, we compute function $W(t)$, which is the value of reaching the *expected* score t in the first exam attempt. We also use the estimates to recover the prize distribution G because these estimates contain information on the value of placement (the prize). This part of the estimation process borrows from Krishna et al. (2018).

We then use the structural model to infer students' mean cost of effort, $C(t)$, at each score t. This is distinct from the cost $c(t)$ in the theoretical model, as we clarify below. To recover $C(t)$, we use students' private tutoring and high school choices (before the exam).[33] These choices are costly, but generate benefits captured by $W(t)$ (estimated above), which depend only on the expected score in the first exam attempt. The costs are the unknown parameters to be estimated. The investment choices are discrete, so we employ a discrete choice setting (mixed logit). We estimate the logit model, and from the observed choice shares, we back out the net values of investments, which are the values of placement minus the investment costs. We remove these net values from the gross benefits of investments reflected in $W(t)$

[31] Students who retake also obtain a learning shock to their score that is estimated to have a positive mean and decreases with the number of retaking attempts.

[32] This value may change across the first several attempts.

[33] There is no data on the cost of effort, so we have to estimate this cost indirectly from students' behavior.

and obtain the cost $C(t)$ of score t by averaging them across all the students with score t.

After estimating functions $W(t)$ and $C(t)$, we use the theoretical model to obtain its remaining primitives: the type distribution $F(x)$ and cost function $c(t)$. This part of the estimation has no counterpart in Krishna et al. (2018) and is a contribution in itself. It can be used independently of the procedure used to estimate $W(t)$ and $C(t)$ in the previous step. To obtain the type distribution $F(x)$ and cost function $c(t)$, we suppose that the students compete in a large contest that corresponds to the theoretical model with prizes given by the estimated function $W(t)$, and derive the type distribution $F(x)$ such that when each type chooses an optimal score, the score distribution matches the one in the data, and the cost of each score t is given by the estimated cost $C(t)$. The assumed optimality of students' choices gives us the function $x(t)$, which maps each score t to the type $x(t)$ that chooses it. This, together with $C(t)$, allows us to back out the cost function $c(t)$.

Having estimated prize function $G(t)$, cost function $c(t)$, and type distribution $F(x)$, we simulate a contest that corresponds to the theoretical model with these primitives and derive the effort, prize allocation, and utility of each type. We then use our theoretical results to investigate Pareto improvements for this contest. In what follows, we elaborate on these steps, focusing on the intuition and leaving the details to Online Appendices 1 and 2.

7.1 Estimation

7.1.1 Estimating G and W using the structural model

Each prize y corresponds to the value of a seat in the college system.[34] To back out the distribution G of y and function $W(t)$, which describes the value of obtaining an expected score of t in the first exam attempt, we first estimate a value function that maps every realized rank (realized score quantile) in the exam in a particular retaking attempt to a value that includes the value of placing with this rank, the option value of retaking the exam, and the cost of retaking following that attempt. Given a student's realized score rank on the exam, the student is more likely to retake the exam if the value function increases more sharply

[34] All values are in utility terms relative to the value of the best available seat, which is normalized to 1. The value of the worst available seat is normalized to zero, and all the costs and value functions are measured in these units.

at this rank. The placement value of a student with a rank r is $G^{-1}(r)$. Once we know the value function, we obtain $G^{-1}(r)$ and the retaking costs that rationalize this function by using Bellman's equation that describes the student's decision whether to retake the exam. Thus, the value function, retaking costs, and $G(y)$ are pinned down by the observed exam retaking rates. Since the cost of retaking is assumed homogeneous, it is pinned down by the average retaking rates. The variation in retaking rates by student rank, in turn, pins down the curvature of $G(y)$.

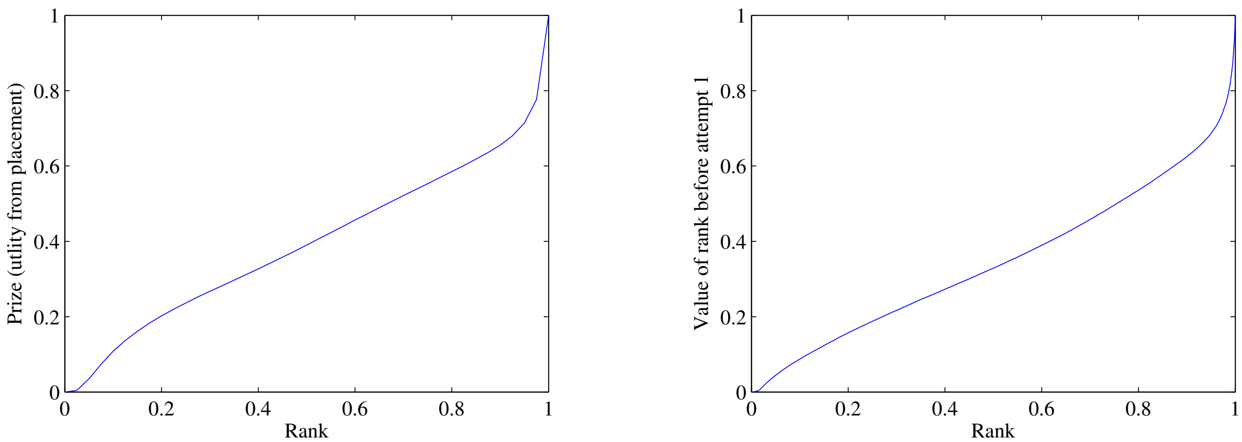

(a) Inverse CDF of prize distribution, $G^{-1}(r)$, mapping one's rank to prize.

(b) Value function before the first exam attempt, $W(t(r))$.

Figure 2: The value of placement and the value of rank in the first attempt

Figure 2a depicts the estimate of $G^{-1}(r)$, the utility of placement, obtained after normalizing the range of prizes to $[0, 1]$ and approximating $G^{-1}(r)$ in a flexible non-parametric way. This function increases sharply at the very top. Students with close-to-perfect scores retake, which can only be rationalized by a sharply increasing utility of placement.

Taking an expectation of the value function over score rank realizations, we obtain the value of a given expected rank in the first attempt, $W(t(r))$, where $t(r)$ is the expected score needed to obtain rank r ($W(t(r))$ is formally defined in equation (28) in Online Appendix 1). This value includes the utility of placement and the option value of retaking and its cost. It is depicted in Figure 2b.[35]

[35] The rank in Figure 2a is restricted to the set of students who accept placement, while that in Figure 2b includes all students taking the exam.

7.1.2 Estimating C using the structural model

To back out the remaining two primitives of the theoretical model, the cost function and the distribution of types, we rely on the observable data on pre-test effort. There are two ways in which students can increase their test scores: by going to private or selective public schools, and by taking extra preparatory courses after school. Both ways are costly: selective schools require entrance exams of their own and involve much effort, and private schools and preparatory courses charge tuition. We refer to selective public schools as 'exam schools.'

There is a clear relationship between pre-test investment and test outcomes. Figure 3 plots the shares of students taking preparatory courses and/or being enrolled in selective and private schools conditional on the exam score. Students at the bottom of the score distribution are predominantly educated in public schools; only about a third of them take preparatory courses before the exam. Students at the top of the score distribution are predominantly educated in exam schools, and nearly all of them take extra preparatory courses. Private schools are the middle ground between public schools and exam schools.

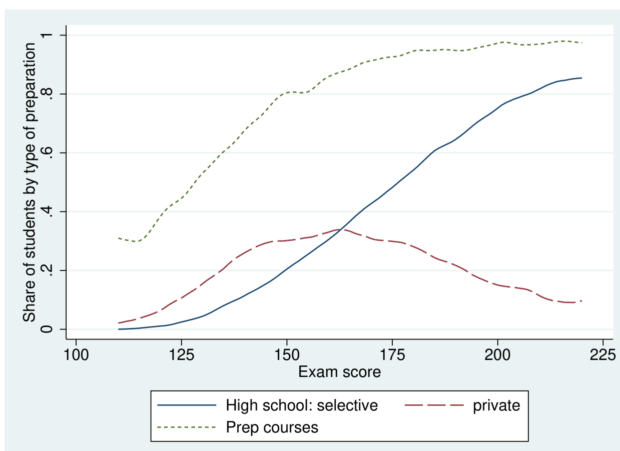

Figure 3: Schooling investments before the first exam attempt

To determine the costs of the three types of pre-test investment, we use a mixed logit discrete choice model with data on individual choices of students and the value of expected rank derived from the utility of placement estimated above. Each middle-school student chooses between public, private, and exam high schools, with or without extra preparatory courses, which results in six effort options in total. Each effort option is associated with an expected gain in score estimated from a regression with the score on the left-hand side, controlling for demographics, middle-school GPA, and other relevant observables. The utility gain from effort is the increase in the value of the expected rank, which comes from Figure 2b. The student maximizes the net gain, which is the difference between the above utility gain and the associated cost of investment.[36]

The net utilities in the logit model are directly related to the shares in the data of the options chosen by the students, that is, the net utilities are obtained by inverting the shares. Having net utilities and gross gains allows us to infer the costs incurred as their difference. These inferred costs vary across students both because of variations in student background and the random shocks to scores and costs of effort in the structural model. Having obtained the cost of each of the six effort levels and the parameters of the cost shocks, we use these parameters and the data on each student's effort level and realized score to compute, for each score, the average effort cost incurred by the students who obtained this score. These inferred costs, $C(t)$, are depicted in Figure 4.

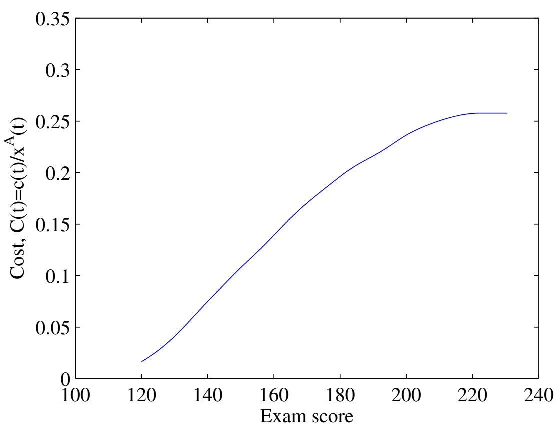

Figure 4: The estimated cost of effort, by score, in the first attempt

[36]Enrolling into an exam school may not be feasible for students who cannot pass the selection exams, especially if they choose the type of high school shortly before these exams. However, since these choices are often made well in advance, students can (and do) prepare for the selection exams. Thus, in principle, exam schools are available to everyone given sufficient effort. The cost of effort and the returns can vary across students, and our structural model allows for this. In particular, the mean cost of effort $\gamma_{g_{i0}}(hs, pt)$ in our model depends on students' middle-school GPA, g_{i0}, which captures the fact that to pass the school selection tests may be harder for students with lower levels of achievement.

7.1.3 Estimating F and c using the theoretical model

We have estimated $G(y)$, $W(t)$, and $C(t)$ using the structural model. To fully calibrate the theoretical model, it remains to estimate $F(x)$ and $c(t)$. To do so, we now think of the Turkish students as competing in a one-shot contest given by our theoretical model with prizes given by $W(t)$, which includes in reduced form the value of any future test retaking.[37] We suppose that the data corresponds to an equilibrium of this model, so that, in particular, $c(t)/x^A(t) = C(t)$.[38] We use this equation to obtain $F(x)$ and $c(t)$ as explained below.

From equation (3) with $W(t)$ instead of y, we obtain that the utility of type x is $W(t) - c(t)/x$. Differentiating this with respect to t, we obtain the optimality condition for type $x = x^A(t)$,

$$\frac{c'(t)}{x^A(t)} = W'(t). \tag{11}$$

Differentiating $C(t) = c(t)/x^A(t)$ and substituting for $\frac{c'(t)}{x^A(t)}$ using (11) gives

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

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

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

At this point, we know the ratio $\frac{W'(t)-C'(t)}{C(t)}$. We also know that the highest-ability student chooses the highest score. This allows us to integrate the equation numerically to obtain $x^A(t)$. Once we have $x^A(t)$, we invert this function to obtain the distribution of ability since we have the distribution of scores in the data.[39] We also obtain $c(t)$ as $c(t) = C(t)x^A(t)$. These remaining primitives of the model are presented in Figure 5. Panels 5a and 5b plot the

[37] In particular, we suppose that the students can choose a deterministic non-negative score instead of choosing one of the six levels of effort that translate into a noisy score.

[38] Recall that the theoretical model predicts that students with higher ability choose higher scores; $x^A(t)$ associates each score t with the type $x^A(t)$ that chooses it in equilibrium.

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

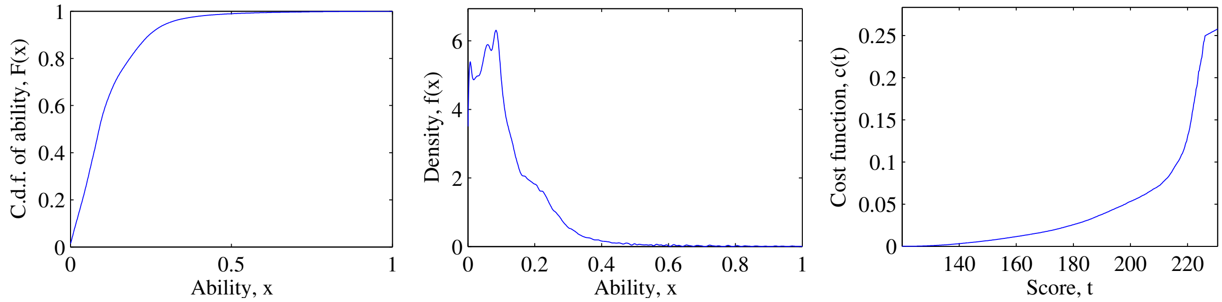

(a) Distribution function, $F(x)$. *(b) Density, $f(x)$.* *(c) Cost function $c(t)$*

Figure 5: Estimated elements of the model

ability distribution $F(x)$ and its density $f(x)$. The estimate for $c(t)$ is shown in Figure 5c.

7.2 Simulating Pareto improvements in the theoretical model

Having estimated $F(x)$, $G(y)$, and $c(t)$, we simulate a contest based on our theoretical model with these primitives and utility $y - c(t)/x$ (that is, (3)).[40] We then investigate Pareto improvements for this contest.

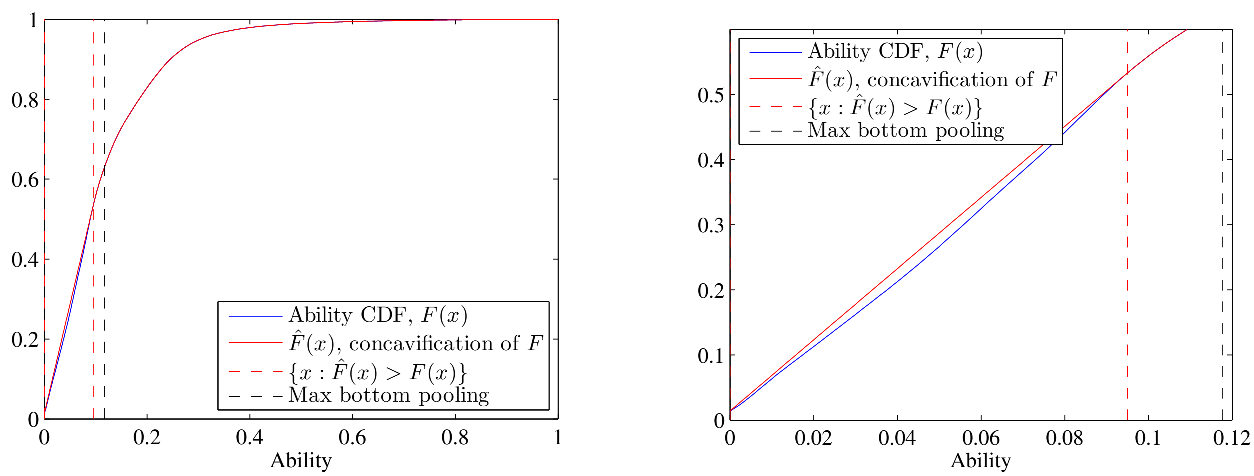

(a) C.d.f. of ability and its concavification *(b) C.d.f. of ability and its concavification, zoomed in*

Figure 6: Concavification of the ability distribution

Our theoretical results on Pareto-improving pooling intervals and category rankings refer to the concavification $\widehat{F}(x)$ of $F(x)$. Figure 6a plots $F(x)$ in blue and its concavification

[40]This contest has no exam retaking and uses prize distribution $G(y)$, which does not include the value of retaking, so that players face the rank-to-prize mapping in Figure 2a. One may consider using $W(t)$ instead of $G(t)$ as the prize distribution, but although $W(t)$ correctly reflects the value of placement and retaking in the baseline contest equilibrium, it is not a primitive of the theoretical model. A counterfactual pooling policy would affect retaking incentives, thereby altering the prize distribution relative to the baseline contest.

$\widehat{F}(x)$ in red. The two functions look very similar on $[0,1]$, but the concavification is in fact significantly above F approximately on the interval $[0, 0.09]$.[41] This interval, which pools approximately 52 percent of the students, is highlighted in Figure 6b and its upper bound is depicted by the dashed red line, which clearly shows that $\widehat{F}(x)$ is linear on this interval and is not linear elsewhere. We refer to this interval as the minimal bottom pooling interval.

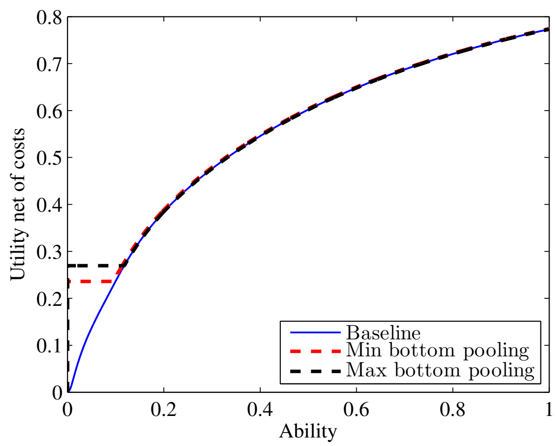

(a) Welfare gains for all types

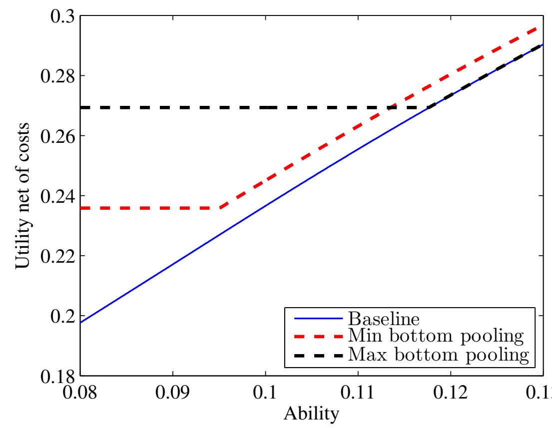

(b) Welfare gains at the top of the pooling interval

Figure 7: Welfare gains under minimal and maximal bottom pooling

We begin with robust Pareto improvements, which depend only on the type distribution. By Theorem 5, the Pareto frontier of robust Pareto improving category rankings (which may involve pooling on multiple intervals), consists of the single policy that pools every interval on which $\widehat{F}(x)$ is linear. Thus, because $\widehat{F}(x)$ is linear only the minimal bottom pooling interval, the bottom pooling policy that pools on this interval is the unique Pareto optimal robust Pareto improving policy.

We now turn to Pareto-improving policies that depend on the prize distribution and are not necessarily robust. Because the assortative allocation for the estimated type and prize distributions does not assign all types in the minimal bottom pooling interval the same prizes (see Figure 2a), by Theorem 1 Pareto improving policies exist. Moreover, by Proposition 3, the minimal bottom pooling policy (which is the Pareto optimal robust Pareto improving policy) is on the Pareto frontier of (non-robust) Pareto-improving category rankings. This

[41] Everywhere else the distance between $F(x)$ and its concavification, if it is positive, is more than a hundred times smaller. Such regions arise mostly due to rounding and smoothing errors for extreme values of x as the data get sparse in these regions.

Pareto frontier also includes other policies, which benefit different types to different degrees. Among them, we focus on the policy that benefits the lowest types the most.

To identify this policy, notice that the lowest types only benefit from policies that include a bottom-pooling interval.[42] Moreover, the longer the bottom-pooling interval, the greater the benefit of the low types because any bottom pooling interval is associated with a score of zero, and the higher the upper bound of the interval, the better the distribution of prizes associated with the interval. Thus, to benefit the lowest types the most, we first identify the largest Pareto-improving bottom-pooling interval. As we clarify below, this interval strictly contains the minimal bottom pooling interval, and the utility of the type at the top of the interval is the same as in the baseline contest (otherwise the interval can be further increased), and therefore so are the utilities of all higher types. Since $\widehat{F}(x)$ is not linear on any interval above the minimal bottom pooling one, Theorem 1 applied to the prize distribution corresponding to the largest Pareto-improving bottom pooling policy shows that no additional pooling interval can be added while remaining Pareto improving. That is, the longest Pareto-improving bottom pooling interval is the category ranking on the Pareto frontier of Pareto-improving category rankings that benefits the lowest types the most.

To identify the longest bottom pooling interval, we first pool on the minimal bottom pooling interval.[43] We simulate the baseline equilibrium payoffs, as well as those under minimal bottom pooling. Figure 7a shows these payoffs as a function of student type. The former and the latter payoffs are shown in solid blue and dashed red respectively. Figure 7b highlights these differences by zooming in on lower abilities.

Note that minimal bottom pooling makes everyone better off. It creates slack for those not pooled in terms of their utility relative to the baseline contest. This slack lets us increase the upper bound of the pooled region while keeping the non-pooled above their baseline payoffs. As raising the upper bound raises the expected prize for the pooled types, these types gain,

[42]Any category ranking without bottom pooling leaves unchanged the chosen score, prize allocation, and utility of the types in $[0, x]$ for some $x > 0$.

[43]Recall that there are other intervals on which $\widehat{F}(x)$ is slightly above $F(x)$, which we ignore. To verify that we do not lose anything by restricting ourselves to the lowest interval only, we compare the minimal bottom pooling to the policy that, in addition, pools on all these additional intervals. The payoffs under these two policies are depicted in Figure 3.1 in Online Appendix 3, which shows that they are essentially identical.

while the non-pooled ones lose. At some point, as we continue to raise the upper bound, the slack is exhausted, so that the type at the top of the pooling interval has the same payoff as in the status quo. At this point, we reach the maximal bottom pooling policy, under which the lowest types obtain their maximal payoff across all Pareto-improving policies on the Pareto frontier. The payoffs under this policy are depicted in Figure 7 in dashed black. This policy pools together approximately 63 percent of the students.[44]

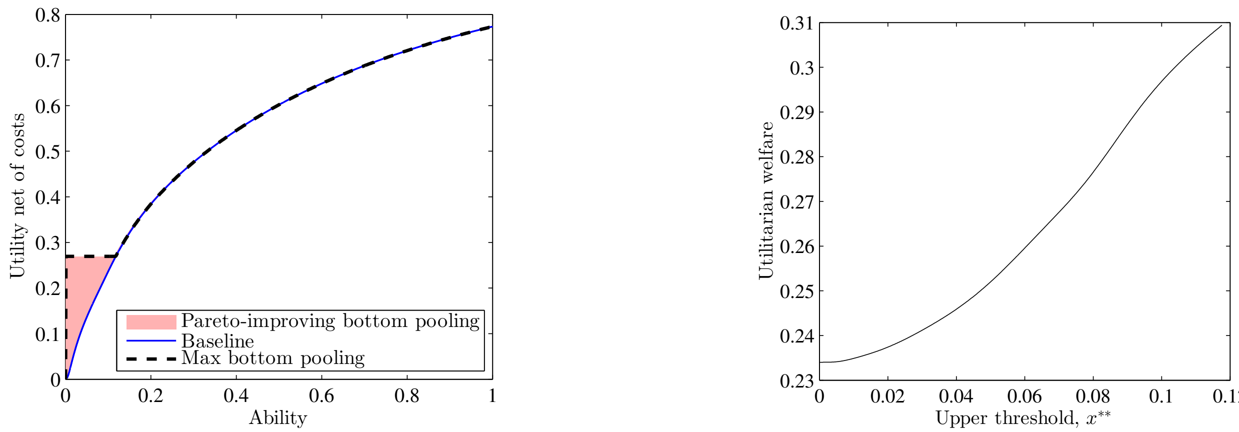

(a) Range of net utilities under Pareto-improving bottom pooling vs. maximum bottom pooling

(b) Aggregate utility for all Pareto-improving bottom pooling policies as a function of x^{**}

Figure 8: All Pareto-improving bottom pooling policies

In addition to maximizing the utility of the lowest types among all the policies on the Pareto frontier of Pareto-improving category rankings, in our setting the maximal bottom pooling policy also maximizes the aggregate utility across all bottom pooling policies. In Figure 8a, the pink area gives the payoffs obtained across all Pareto-improving bottom pooling policies, while the dashed black line depicts payoffs under maximal bottom pooling (as in Figure 7). Figure 8b depicts the aggregate utility for each upper bound of the Pareto-improving bottom-pooling policies. The figure shows that the aggregate utility increases in the upper bound, and is therefore maximized by maximal bottom-pooling.

To better understand the source of gains under maximal bottom pooling, consider Fig-

[44]As we mention in Section 6, our theoretical model does not allow for noise in scores or test retaking. To check if our main results still hold in a more complex environment, we simulate the maximal bottom pooling policy with the *structural* model from Section 7 instead of the theoretical model from Section 2 and present the results in Online Appendix 3. The structural model explicitly allows for shocks to scores (noise), multiple dimensions in student ability, and the option to retake the university entrance exam. We show that bottom pooling still achieves a Pareto improvement: irrespective of student's initial standing before high school, maximal bottom pooling raises student utility.

ure 9. Figure 9a shows types' utilities, $y(x) - c(t(x))/x$, under maximal bottom pooling and in the baseline contest. Figure 9b shows the equilibrium scores under the two policies. Figures 9c and 9d decompose the payoffs into what is explained by placement, y, and what is explained by the cost of effort, $c(t)/x$.

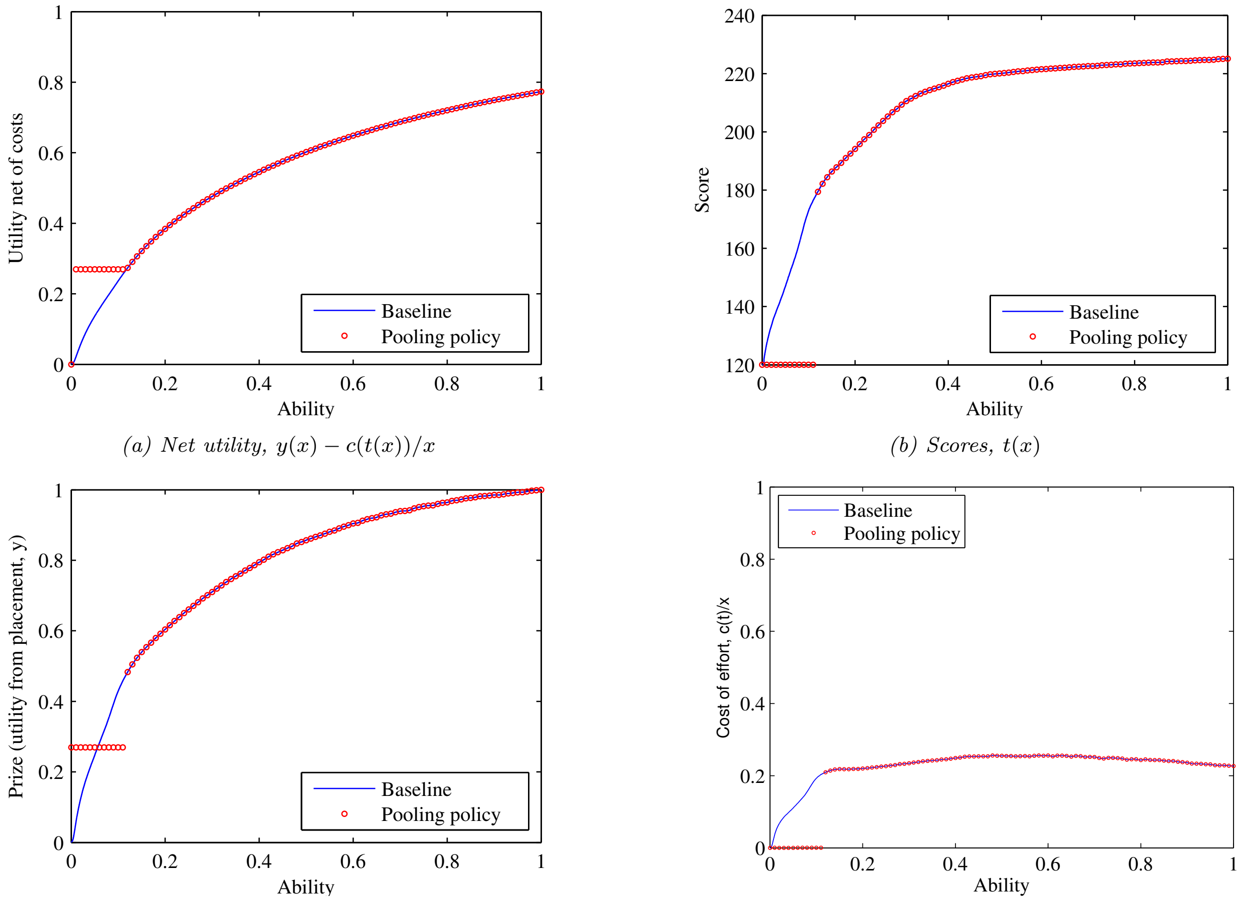

(a) Net utility, $y(x) - c(t(x))/x$

(b) Scores, $t(x)$

(c) Prize in equilibrium, $y(x)$

(d) Cost of effort in equilibrium, $c(t(x))/x$

Figure 9: Equilibrium payoffs and effort in the baseline contest and under bottom pooling.

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. Pooling the prizes at the bottom raises the placement payoff for the lower-end types within the pooled interval and reduces the placement payoff for the higher-end types (Figure 9c). However, since the costs of effort are zero for all the pooled types (Figure 9d), everyone in the pooled interval gains, with the lower-end types gaining more than the higher-end types (as evident from Figure 9a). Higher types, those above the pooled interval, are not affected because neither their placements nor their effort is affected by the change in policy from the baseline contest to the maximal bottom pooling one. Overall, the mean utility increases by 32 percent, while the pooled types gain 83 percent on average compared to the baseline contest.

8 Experimental exercise

We conducted a laboratory experiment based on a discretized version of the calibration exercise and the maximal Pareto improving (non-robust) bottom pooling policy identified in Section 7. While there are substantial differences in stakes between the experiment and real-world decisions that influence college admissions, the experiment may help us identify unanticipated behaviors policy-makers may observe when implementing a pooling policy, and what the resulting welfare implications would be.

To conduct the experiment, we transformed the game into an individual decision-making problem in a discretized setting without strategic uncertainly. The costs and benefits, exogenous from the subjects' point of view, correspond to a situation in which the admission criteria are known in advance. This is often the case in college admissions settings that involve a large number of applicants and therefore entail little uncertainty.

In the experiment, each college, or tier of pooled colleges, had an admission threshold. Subjects in the experiment chose how much to invest in costly "virtual study materials" to reach their desired threshold under a discrete policy and under a pooling policy. Additional details are in Online Appendices 4 and 5. Online Appendix 6 contains all the experimental materials subjects faced.

8.1 Main experimental results

Figure 10 shows subjects' average utility by ability level under the discretized theoretical pooling policy (panel a) and in the experiment (panel b). Overall, adopting the pooling policy in the experiment increased aggregate welfare by 18.9 percent, closely matching the theoretical prediction of 20 percent[45].

[45] This theoretical prediction is computed for the discretized economy, which has a mix of the eleven ability types shown in Figure 10. Since this economy is not identical to the continuous-type economy in Section 7,

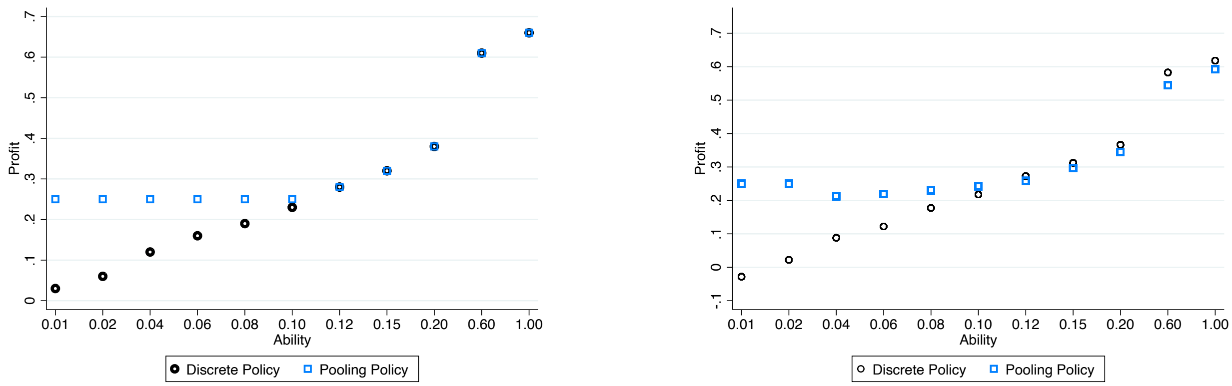

(a) Predictions *(b) Laboratory data*

Figure 10: Average profits under the discrete and pooling policies: predictions and experimental data.

Considering the impact of the pooling policy on each ability level separately, we find strong agreement with the theoretical predictions for low ability subjects: their welfare increases by over 70 percent, close to the theoretical prediction of 65 percent in the discretized model. The theoretical prediction for high ability subjects is that their utility should be the same across the two policies. Instead, their welfare in the experiment decreased by 1.6 percent under the pooling policy. This was due to the behavior of subjects with ability levels of 0.12, the lowest level among the high ability subjects (whose utility was predicted not to change). These subjects opted to invest nothing and be assigned to the pooled set of colleges, so they faced a lottery, instead of investing and obtaining a better and deterministic college seat (further details are in Online Appendix 4).

Mapped to the Turkish student population, the experimental results imply that at least 85 percent of applicants should see their utility strictly or weakly increase, and at most 15 percent of the population may see their utility slightly decrease.

We explore possible reasons for the apparently suboptimal behavior of subjects with ability 0.12. We argue in Online Appendix 4 that this behavior is unlikely due to mistakes, experimental procedures, or risk-seeking behavior, and instead is consistent with preferences for randomization. Dwenger, Kübler, and Weizsäcker (2016) explored such preferences in the context of school applications and showed that up to 50 percent of individuals choose lotteries between available allocations, indicating an explicit preference for randomization.[46]

the welfare gains reported here are not identical to those in Section 7.

[46] We also highlight that in our experiment the optimal choice for subjects with ability 0.12 without pooling

9 Conclusion

This paper investigates how to improve college admissions based on centralized tests. Our main message is that coarse performance disclosure policies can benefit all students, regardless of their ability, when test preparation is costly. These policies take a simple form and are easy to implement. As a "proof of concept," we empirically estimated the key theoretical constructs using Turkish college admissions data. We used our theoretical results to simulate the equilibrium outcome of a college admissions contest based on these estimates, and demonstrated how to identify Pareto improving policies. We showed that a policy that pools together the majority of the lowest-performing students would benefit the lower-ability students the most, raising the welfare of the pooled students without impacting the welfare of the other students. We also conducted a laboratory experiment based on these empirical findings, which 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.

was still available in the round with pooling. Thus, these subjects could obtain the same utility in both rounds by maintaining the same behavior. By revealed preference, those subjects with ability 0.12 who switched to the lottery in the round with pooling were likely made better off, even though their monetary payoff decreased slightly, which is consistent with a wide range of preferences.

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] Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," The RAND Journal of Economics, 25(2), 242–262.

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

[4] Bodoh-Creed, A. and B. Hickman (2024): "Pre-College Human Capital Investment and Affirmative Action: A Structural Policy Analysis of US College Admissions," Mimeo.

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

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

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

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

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

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

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

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

[13] Dwenger, N., Kubluer, D. and G. Weizsacker (2016): "Flipping a coin: Evidence from university applications", Journal of Public Economics 167, 240-250.

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

[15] 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.

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

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

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

[19] Guerre, E.,, Perrigne, I., and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," Econometrica, 68(3), 525–574.

[20] 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.

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

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

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

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

[25] Hotz, V. J., and R. A. Miller (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” Review of Economic Studies, 60, 497–529.

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

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

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

[29] 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.

[30] Magnac T., and D. Thesmar (2002): “Identifying Dynamic Discrete Decision Processes,” Econometrica, 70 (2), 801-816.

[31] Mas-Colell A., M. Whinston, and J. Green (1995): “Microeconomic Theory,” Oxford University Press.

[32] McAfee, P. (2002): “Coarse matching,” Econometrica, 70(5), 2025-2034.

[33] 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.

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

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

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

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

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

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

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

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

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

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

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

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

Appendices

A Proofs of Proposition 1, Theorem 1, Theorem 4, and Theorem 5

Proof of Proposition 1. Part (a) follows because with pooling types $x \leq x^*$ 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 with and without pooling. By (5), in the approximating mechanism of the original contest, the utility of a type x in $(x^*, x^{**}]$ exceeds that of type x^* by $\int_{x^*}^{x} y^A(\tilde{x})\, d\tilde{x}$. In the approximating mechanism with pooling, the utility of type x exceeds that of type x^* by

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

Thus, pooling increases the utility of type x if and only if (7) holds.

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

$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})\, dF(\tilde{x})}{F(x^{**}) - F(x^*)} - y^A(x)\,. \tag{12}$$

The fraction in (12) is a weighted average of $y^A(\tilde{x})$ over types in $(x^*, x^{**}]$, so (12) is positive for types x close to x^*, monotonically decreases as x increases, and becomes negative for types x close to x^{**}. Thus, the utility gain from pooling for types x in $(x^*, x^{**}]$ first increases and then decreases in the type. In particular, if not all prizes in the pooled interval are identical, the utility of all types x in (x^*, x^{**}) strictly increases if the utility of type x^{**} weakly increases. And the difference between the utilities of type $x > x^{**}$ and type x^{**} is $\int_{x^{**}}^{x} y^A(\tilde{x})\, d\tilde{x}$ in both approximating mechanisms, so pooling weakly benefits type x^{**} if and only if it weakly benefits type $x = 1$ if and only if it is Pareto improving, which gives part (d).

Proof of Theorem 1. Suppose first that there is a type interval $(x^*, x^{**}]$ on which $\widehat{F}$ is linear but not all types in the interval obtain the same prize, and without loss of generality suppose that $[x^*, x^{**}]$ is maximal, that is, $\widehat{F}$ is not linear on any interval that strictly contains $[x^*, x^{**}]$. By definition of $\widehat{F}$, we have that $\widehat{F}(x^*) = F(x^*)$, $\widehat{F}(x^{**}) = F(x^{**})$, and $\widehat{F}(x) \geq F(x)$ for every x in $[x^*, x^{**}]$. Thus, restricted to $[x^*, x^{**}]$, F first-order stochastically dominates $\widehat{F}$. Therefore,

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

where the first equality follows because $\widehat{F}$ is linear on $[x^*, x^{**}]$, so part (d) of Proposition 1 holds and pooling on $(x^*, x^{**}]$ is Pareto improving.

For the other direction, suppose that on every interval on which $\widehat{F}$ is linear all types obtain the same prize, and assume for contradiction that there exists a Pareto improving category ranking $\mathcal{I}$. Denote by $(x^*, x^{**}]$ the lowest type interval in the category ranking $\mathcal{I}$ that includes some types that strictly benefit from the category ranking.[47] Since $\mathcal{I}$ is Pareto improving, all types not higher than x^* are indifferent between the baseline contest and the category ranking. This implies that pooling on the single interval $(x^*, x^{**}]$ is Pareto improving. In particular, not all types in the interval obtain the same prize, otherwise pooling on $(x^*, x^{**}]$ has no effect. Let quantiles q^* and q^{**} and types x' and x'' be such that $q^* = \widehat{F}(x') = F(x^*)$ and $q^{**} = \widehat{F}(x'') = F(x^{**})$. Because $F \leq \widehat{F}$, we have that $x' \leq x^*$ and $x'' \leq x^{**}$. In addition, $G^{-1} \circ F = G^{-1} \circ \widehat{F}$. Indeed, $G^{-1}(F(x)) = G^{-1}(\widehat{F}(x))$ whenever $F(x) = \widehat{F}(x)$. And for a type x with $F(x) < \widehat{F}(x)$, by definition $\widehat{F}$ is linear on an interval that includes x. Consider the maximal interval on which $\widehat{F}$ is linear that includes x. All the types in the interval obtain the same prize, and at the endpoints of the interval F and $\widehat{F}$ coincide (by definition of $\widehat{F}$), so on this interval $G^{-1} \circ F$ and $G^{-1} \circ \widehat{F}$ coincide. These observations imply that

$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x}) dF(\tilde{x})}{F(x^{**}) - F(x^*)} = \frac{\int_{x'}^{x''} G^{-1}(\widehat{F}(x)) d\widehat{F}(\tilde{x})}{\widehat{F}(x'') - \widehat{F}(x')} \leq \frac{\int_{x'}^{x''} G^{-1}(\widehat{F}(x)) d\tilde{x}}{x'' - x'} \leq \frac{\int_{x^*}^{x^{**}} y^A(\tilde{x}) d\tilde{x}}{x^{**} - x^*}, \tag{13}$$

where the equality follows because both expressions are equal to the expected prize $\int_{q^*}^{q^{**}} G^{-1}(z)\, dz / (q^{**} - q^*)$ in quantile interval $[q^*, q^{**}]$, the first inequality follows because $\widehat{F}$ is concave so is first-order stochastically dominated by the uniform distribution when both are restricted to the interval $[x', x'']$, and the second inequality follows because $x' \leq x^*$, $x'' \leq x^{**}$, and $G^{-1} \circ F = G^{-1} \circ \widehat{F}$. Moreover, the second inequality is strict if $x' < x^*$ or $x'' < x^{**}$ because not all types in the interval $(x^*, x^{**}]$ obtain the same prize. And the first inequality is strict if $x' = x^*$ and $x'' = x^{**}$ because then not all types in $(x', x'']$ obtain the same prize and $\widehat{F}$ is strictly concave on $[x', x'']$.

[47]Such an interval exists otherwise all types weakly prefer the baseline contest to the category ranking, so the category ranking is not Pareto improving.

We conclude that
$$\frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})dF(\tilde{x})}{F(x^{**}) - F(x^*)} < \frac{\int_{x^*}^{x^{**}} y^A(\tilde{x})d\tilde{x}}{x^{**} - x^*},$$
so type x^{**} is strictly harmed by pooling on type interval $(x^*, x^{**}]$ (as are types slightly lower than x^{**}), contradicting that pooling on $(x^*, x^{**}]$ is Pareto improving.

Proof of Theorem 4. We first observe that any robust Pareto improving pooling interval is contained in an interval on which $\widehat{F}$ is linear. To see this, take an interval $(x^*, x^{**}]$ that is not so contained. Consider a distribution G of prizes that gives the same prizes to types in any interval on which $\widehat{F}$ is linear, and gives different prizes to all other types.[48] Then the proof of the "if" direction of Theorem 1 shows that pooling on $(x^*, x^{**}]$ hurts type x^{**} (and nearby types) because not all types in $(x^*, x^{**}]$ obtain the same prize (since $(x^*, x^{**}]$ is not contained in an interval on which $\widehat{F}$ is linear). Thus, pooling on $(x^*, x^{**}]$ is not robust Pareto improving. Now, pooling on a maximal interval on which $\widehat{F}$ is linear is robust Pareto improving by definition of $\widehat{F}$ and Proposition 4.

For the Pareto frontier, we show that if $(x', x'']$ and $(x^*, x^{**}]$ are robust Pareto improving pooling intervals and $(x', x''] \subseteq (x^*, x^{**}]$, then $(x^*, x^{**}]$ is robust Pareto preferred to $(x', x'']$. This is because for any prize distribution G, pooling on interval $(x', x'']$ is equivalent to using a baseline contest in which types $(x', x'']$ are allocated identical prizes that are the average of the prizes they are allocated under G. Starting from this modified prize distribution and pooling on the interval $(x^*, x^{**}]$ is equivalent to pooling on the interval $(x^*, x^{**}]$ directly, which is robust Pareto improving. Finally, consider any two distinct maximal intervals on which $\widehat{F}$ is linear. For each interval consider a prize distribution that gives all types below the interval the same prize, all types above the interval the same prize, and different prizes to the types in the interval. With this prize distribution pooling on the interval is Pareto improving, but pooling on the other interval has no effect. Thus, neither of the two robust Pareto improving intervals is robust Pareto preferred to the other.

Proof of Theorem 5. Similarly to the proof of Theorem 4, any interval that is part of a robust Pareto improving category ranking is contained in an interval on which $\widehat{F}$ is linear, pooling on maximal intervals on which $\widehat{F}$ is linear is robust Pareto improving by definition of

[48] For example, start with assigning prize $y = x$ for every type x, and then for any maximal interval $(x', x'']$ on which $\widehat{F}$ is linear set the prize of every type in the interval to be the average of $y = x'$, and $y = x''$.

$\widehat{F}$ and Proposition 5, if every interval in one category ranking is contained in some interval of another category ranking then the second is Pareto preferred to the first, and if one category ranking contains a maximal interval on which $\widehat{F}$ is linear that is not contained in some other category ranking, then the other category ranking is not robust Pareto preferred to the first. Finally, taking a category ranking that consists of one or more maximal intervals on which $\widehat{F}$ is linear and adding to it another maximal interval on which $\widehat{F}$ is linear generates a category ranking that is robust Pareto preferred to the original category ranking. This is because the effect of the original category ranking is identical to the effect of using a baseline contest in which the prizes allocated to each pooled interval of types are replaced with the same mass of the average of these prizes, and then pooling on the additional maximal interval on which $\widehat{F}$ is linear. This proves the result.

B Proofs of Theorems 2 and 3 (mean-preserving contractions)

To prove Theorem 2, we will need the following two lemmas. The first lemma seems to belong to "statistics folklore". We give its proof for completeness; for simplicity, we will restrict attention to continuous, strictly increasing G and H.

Lemma 1. *Suppose that H and G are two CDFs of distributions with the same domain, say $[0,1]$. Then, H second-order stochastically dominates G if and only if*

$$\int_0^z [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})]d\widetilde{z} \leq 0 \tag{14}$$

for all z, with equality for $z = 1$.

Proof. We will first prove necessity. Observe first that (14) holds whenever $G^{-1}(z) = H^{-1}(z)$. Indeed, in the system of coordinates with z on the vertical axis and with $G^{-1}(z)$ and $H^{-1}(z)$ on the horizontal axis, $\int_0^z G^{-1}(\widetilde{z})d\widetilde{z}$ is the area between the graph of G^{-1}, the vertical axis, and the horizontal line at the level of z. This area is equal to the area of the rectangle $[0, G^{-1}(z)] \times [0, z]$ minus the area between the graph of G, the horizontal axis, and the vertical line at the level of $G^{-1}(z)$. Similarly, $\int_0^z H^{-1}(\tilde{z})d\tilde{z}$ is equal to the area of the rectangle $[0, H^{-1}(z)]\times[0, z]$ minus the area between the graph of H, the horizontal axis, and the vertical line at the level of $H^{-1}(z)$. Thus, (14) follows from the fact that

$$\int_0^x G(\tilde{x})d\tilde{x} \geq \int_0^x H(\tilde{x})d\tilde{x}$$

for all x (in particular $x = G^{-1}(z) = H^{-1}(z)$) when H second-order stochastically dominates G.

Note here that the last inequality holds with equality for $x = 1$, by second-order dominance, which yields (14) with equality for $z = 1$.

Suppose now that $G^{-1}(z) > H^{-1}(z)$. Let $\overline{z} = \min\{\tilde{z} \geq z : G^{-1}(\tilde{z}) = H^{-1}(\tilde{z})\}$. Then

$$\int_0^{\overline{z}} [G^{-1}(\tilde{z}) - H^{-1}(\tilde{z})]d\tilde{z} \geq \int_0^z [G^{-1}(\tilde{z}) - H^{-1}(\tilde{z})]d\tilde{z},$$

because $G^{-1}(\tilde{z}) \geq H^{-1}(\tilde{z})$ for $\tilde{z} \in [z, \overline{z}]$. Thus, since (14) holds for $\overline{z}$, it holds for z.

Suppose finally that $G^{-1}(z) < H^{-1}(z)$. Let $\underline{z} = \max\{\tilde{z} \leq z : G^{-1}(\tilde{z}) = H^{-1}(\tilde{z})\}$. Then

$$\int_0^z [G^{-1}(\tilde{z}) - H^{-1}(\tilde{z})]d\tilde{z} \leq \int_0^{\underline{z}} [G^{-1}(\tilde{z}) - H^{-1}(\tilde{z})]d\tilde{z},$$

because $G^{-1}(\tilde{z}) \leq H^{-1}(\tilde{z})$ for $\tilde{z} \in [\underline{z}, z]$. Thus, since (14) holds for $\underline{z}$, it holds for z. This completes the proof of necessity.

The proof of sufficiency obtains by applying the proof of necessity to the functions $\widetilde{G} = H^{-1}$ and $\widetilde{H} = G^{-1}$. □

Lemma 2. *Let $f, g : [0,1] \to \mathbb{R}$ be bounded Lebesgue measurable functions. Suppose that f is weakly increasing and g has the property that $\int_x^1 g(\tilde{x})d\tilde{x} \geq 0$ for every $x \in (0,1]$ and $\int_0^1 g(\tilde{x})d\tilde{x} = 0$. Then $\int_0^1 f(\tilde{x})g(\tilde{x})d\tilde{x} \geq 0$. Moreover, $\int_0^1 f(\tilde{x})g(\tilde{x})d\tilde{x} > 0$ if one of the following conditions is satisfied: (a) $\int_{\overline{x}}^1 g(\tilde{x})d\tilde{x} > 0$ for some $\overline{x}$, and f is not constant on any interval $(\underline{x}, \overline{x}]$; (b) $\int_{\underline{x}}^1 g(\tilde{x})d\tilde{x} > 0$ for some $\underline{x}$, and f is not constant on any interval $[\underline{x}, \overline{x})$.*

Proof. Assume w.l.o.g. that f takes values in $(0,1)$. Otherwise, consider an affine transformation $cf + d$ of f with a positive slope $c > 0$, which takes values in $(0,1)$. The lemma for $cf + d$ implies the lemma for f. Represent f as the pointwise limit of functions $f_n = \sum_{i=1}^{n}(1/n)\chi_{A_i^n}$, where $A_i^n = \{x \in [0,1] : f(x) > i/n\}$ and $\chi_{A_i^n}$ takes value 1 on A_i^n and value 0 on $[0,1]\backslash A_i^n$. Since f is weakly increasing, $A_i^n = (x_i^n, 1]$ or $A_i^n = [x_i^n, 1]$ for some x_i^n.

Since f, f_n, and g are bounded,

$$\begin{aligned}\int_0^1 f(\widetilde{x})g(\widetilde{x})d\widetilde{x} &= \lim_n \int_0^1 g(\widetilde{x})\left[\sum_{i=1}^{n}(1/n)\chi_{A_i^n}(\widetilde{x})\right]d\widetilde{x} \\ &= \lim_n \sum_{i=1}^{n}(1/n)\int_0^1 g(\widetilde{x})\chi_{A_i^n}(\widetilde{x})d\widetilde{x} \\ &= \lim_n \sum_{i=1}^{n}(1/n)\int_{x_i^n}^1 g(\widetilde{x})d\widetilde{x},\end{aligned} \tag{15}$$

This yields the first part of the lemma, that is, $\int_0^1 f(\widetilde{x})g(\widetilde{x})d\widetilde{x} \geq 0$, because$\int_{x_i^n}^1 g(\widetilde{x})d\widetilde{x} \geq 0$ for all i.

For part (a) of the second part of the lemma, notice that if $\int_{\overline{x}}^1 g(\widetilde{x})d\widetilde{x} > 0$ for some $\overline{x}$, then there is a constant $c > 0$ such that $\int_x^1 g(\widetilde{x})d\widetilde{x} > c$ for all x from an interval $(\underline{x}, \overline{x}]$. If f is not constant on this interval, then for any large enough n there is a fraction of i's that is bounded away from zero such that $x_i^n \in (\underline{x}, \overline{x}]$. This implies that the last sum in (15) is bounded away from zero, uniformly across all large enough n's. The argument for part (b) is analogous. □

We can now prove Theorem 2.

Proof of Theorem 2. The "only if" follows from the "only if" of Theorem 1. For the other direction, suppose that for every interval of types on which $\widehat{F}$ is linear, in the assortative allocation all the types in the interval obtain the same prize but there exists an MPC H of G that is Pareto-improving. We can assume w.l.o.g. that the allocation under H is also constant on every interval $[x', x'']$ on which $\widehat{F}$ is linear; otherwise, it can be replaced by its Pareto-improving contraction that pools each maximal such interval $[x', x'']$. This pooling composed with H is a Pareto-improving MPC of G as a composition of two mappings with these two properties.

The utility of type $x = 1$ when the prizes are distributed according to G and according to H are

$$\int_0^1 G^{-1} \circ F(x)\, dx \text{ and } \int_0^1 H^{-1} \circ F(x)\, dx,$$

respectively. Since in the assortative allocation all the types of any interval on which $\widehat{F}$ is linear obtain the same prize, none of the two values will be affected when we replace F with $\widehat{F}$.

By substituting $z = \widehat{F}(x)$, we obtain that the utility of type $x = 1$ is

$$\int_0^1 G^{-1}(z) \frac{1}{\widehat{f}(\widehat{F}^{-1}(z))} dz \text{ and } \int_0^1 H^{-1}(z) \frac{1}{\widehat{f}(\widehat{F}^{-1}(z))} dz,$$

respectively. The difference between the two values is

$$\int_0^1 [G^{-1}(z) - H^{-1}(z)] \frac{1}{\widehat{f}(\widehat{F}^{-1}(z))} dz. \tag{16}$$

Since H second-order stochastically dominates G, we have that $\int_0^z [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} \leq 0$ for all $z < 1$, and $\int_0^1 [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} = 0$. So, $\int_z^1 [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} \geq 0$ for all $z < 1$ and $\int_0^1 [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} = 0$. Since $\widehat{F}$ is concave, $1/\widehat{f}(\widehat{F}^{-1}(z))$ is a weakly increasing function of z. Thus, (16) is nonnegative by Lemma 2.

However, to show that the contraction H strictly decreases the utility of type $x = 1$ and obtain a contradiction, we must show that (16) is strictly positive. To show this, we apply the second part of Lemma 2. To be Pareto improving, H must be a nontrivial contraction, that is, it must be that $\int_z^1 [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} > 0$ for some $z \in (0, 1)$. Moreover, we can assume w.l.o.g. that $z = F(x')$ or $z = F(x'')$ for a maximal interval $[x', x'']$ on which $\widehat{F}$ is linear. Indeed, since both G^{-1} and H^{-1} are constant on $[F(x'), F(x'')]$ for each each such interval $[x', x'']$, if $\int_z^1 [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} > 0$ for some $z \in [F(x'), F(x'')]$, then $\int_z^1 [G^{-1}(\widetilde{z}) - H^{-1}(\widetilde{z})] d\widetilde{z} > 0$ for $z = F(x')$ or $F(x'')$. If $\overline{z} := z = F(x')$, then $\widehat{f}(\widehat{F}^{-1}(z))$ cannot be constant on any interval $(\underline{z}, \overline{z}]$ because $[x', x'']$ is maximal. If $\underline{z} := z = F(x')$, then $\widehat{f}(\widehat{F}^{-1}(z))$ cannot be constant on any interval $[\underline{z}, \overline{z})$ because $[x', x'']$ is maximal.

To prove Theorem 3 we need the following lemma.

Lemma 3. *Let $f : [0, x] \to \mathbb{R}_+$ and $h : [0, x] \to \mathbb{R}$ be bounded Lebesgue measurable functions. Suppose that f is weakly decreasing and h has the property that $\int_0^y h(\widetilde{y}) d\widetilde{y} \geq 0$ for every*

$y \in (0, x]$. *Then* $\int_0^x f(\widetilde{x})g(\widetilde{x})d\widetilde{x} \geq 0$.

Proof. Assume w.l.o.g. that f takes values in $[0, 1)$. Otherwise, consider a linear transformation cf of f with a positive slope $c > 0$, which takes values in $(0, 1)$. The lemma for cf implies the lemma for f. Represent f as the pointwise limit of functions $f_n = \sum_{i=1}^{n}(1/n)\chi_{A_i^n}$, where $A_i^n = \{y \in [0, x] : f(y) > i/n\}$ and $\chi_{A_i^n}$ takes value 1 on A_i^n and value 0 on $[0, x]\backslash A_i^n$. Since f is weakly decreasing, $A_i^n = [0, x_i^n]$ or $A_i^n = [0, x_i^n)$ for some x_i^n.

Since f, f_n, and g are bounded,

$$\begin{aligned}
\int_0^x f(\widetilde{x})h(\widetilde{x})d\widetilde{x} &= \lim_n \int_0^x h(\widetilde{x})\left[\sum_{i=1}^{n}(1/n)\chi_{A_i^n}(\widetilde{x})\right]d\widetilde{x} \\
&= \lim_n \sum_{i=1}^{n}(1/n)\int_0^x g(\widetilde{x})\chi_{A_i^n}(\widetilde{x})d\widetilde{x} \\
&= \lim_n \sum_{i=1}^{n}(1/n)\int_0^{x_i^n} g(\widetilde{x})d\widetilde{x},
\end{aligned} \tag{17}$$

This completes the proof, because $\int_0^{x_i^n} g(\widetilde{x})d\widetilde{x} \geq 0$ for all n and i. □

We can now prove Theorem 3.

Proof of Theorem 3. The "only if" direction follows from part (d) of Proposition 1, similarly to the proof of the "only if" direction of Theorem 1. For the other direction, consider an MPC H of G, and suppose that for every interval of types on which $\widehat{F}$ is linear, in the allocation induced by H all the types in the interval obtain the same prize, but there exists an MPC $\widetilde{H}$ of G that Pareto improves on H. We can assume w.l.o.g. that the allocation induced by $\widetilde{H}$ is also constant on every interval $[x', x'']$ on which $\widehat{F}$ is linear; otherwise, it can be replaced by a Pareto-improving MPC that pools each maximal such interval $[x', x'']$.[49]

We will show that $\widetilde{H}$ must be an MPC of H. Since both H and $\widetilde{H}$ are contractions of G,

$$\int_0^1 (\widetilde{H})^{-1}(\widetilde{x})d\widetilde{x} = \int_0^1 G^{-1}(\widetilde{x})\, d\widetilde{x} = \int_0^1 H^{-1}(\widetilde{x})\, d\widetilde{x}.$$

[49]This pooling is a Pareto-improving MPC of G as a Pareto-improving MPC of the Pareto-improving MPC $\widetilde{H}$ of G.

So, by Lemma 1, it remains to show that

$$\int_0^x [(\widetilde{H})^{-1}(\widetilde{x}) - H^{-1}(\widetilde{x})]d\widetilde{x} \geq 0 \tag{18}$$

for all $x \in (0,1)$.

Since $\widetilde{H}$ Pareto improves H, by (5) it must be that

$$\int_0^x (\widetilde{H})^{-1} \circ F(\widetilde{x})\, d\widetilde{x} \geq \int_0^x H^{-1} \circ F(\widetilde{x})\, d\widetilde{x} \tag{19}$$

for all $x \in [0,1]$. And since both H and $\widetilde{H}$ are constant on every interval $[x', x'']$ on which $\widehat{F}$ is linear, we can replace F with $\widehat{F}$ in (19), that is,

$$\int_0^x (\widetilde{H})^{-1} \circ \widehat{F}(\widetilde{x})\, d\widetilde{x} \geq \int_0^x H^{-1} \circ \widehat{F}(\widetilde{x})\, d\widetilde{x}, \tag{20}$$

for all $x \in [0,1]$. By substituting $\widetilde{y} = \widehat{F}(\widetilde{x})$, (20) is equivalent to

$$\int_0^y [(\widetilde{H})^{-1}(\widetilde{y}) - H^{-1}(\widetilde{y})]\frac{1}{\widehat{f}(\widehat{F}^{-1}(\widetilde{y}))}d\widetilde{y} \geq 0 \tag{21}$$

for all $y \in [0,1]$.

Now, apply Lemma 3 to $h(y) = [(\widetilde{H})^{-1}(y) - H^{-1}(y)]/\widehat{f}(\widehat{F}^{-1}(y))$ and $f(y) = \widehat{f}(\widehat{F}^{-1}(y))$ (which is weakly decreasing in y since $\widehat{F}$ is concave) to obtain (18).

Thus, $\widetilde{H}$ is an MPC of H. This is a contradiction to Theorem 2 applied to prize distribution H, because for every interval $[x', x'']$ on which $\widehat{F}$ is linear all types in $[x', x'']$ obtain the same prize in the allocation induced by H.

C Pareto frontier of category rankings

First, using Proposition 2, we will provide a method for checking whether a category ranking belongs to the Pareto frontier of category rankings. Next, we use this method to show that there is no Pareto-improving, category ranking of a category ranking that is constant on each interval $(x', x'']$ on which $\widehat{F}$ is linear. Finally, we give an example showing that there may exist Pareto-frontier category rankings that are not constant on some intervals $(x', x'']$ on which $\widehat{F}$ is linear.

Let $\mathcal{I}$ be a category ranking, and let $x^* < x^{**}$ be a pair of types such that $x^* = a$ for an interval $I = (a, b] \in \mathcal{I}$ or $x^* = d$ for $\{d\} \in \mathcal{I}$, and $x^{**} \in I' \in \mathcal{I}$ with $I' \neq I$. 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 $x^* = a$ for an interval $I = (a, b]$, and $I' = (a', b']$, then replace I, I', and all the elements of $\mathcal{I}$ between I and I' with $(x^*, x^{**}]$ and $(x^{**}, b']$; (ii) if $x^* = d$ for $\{d\} \in \mathcal{I}$, and $I' = (a', b']$, then replace I' and all the elements of $\mathcal{I}$ between $\{d\}$ and I' with $(x^*, x^{**}]$ and $(x^{**}, b']$; if $x^* = a$ for an interval $I = (a, b]$, and $I' = \{x^{**}\}$, then replace I, I', and all the elements of $\mathcal{I}$ between I and I' with $(x^*, x^{**}]$; if $x^* = d$ for $\{d\} \in \mathcal{I}$, and $I' = \{x^{**}\}$, then replace I' and all the elements of $\mathcal{I}$ between $\{d\}$ and I' with $(x^*, x^{**}]$.

Proposition 6. *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*

$$\begin{gathered}x^* = a \text{ for some } I = (a, b] \in \mathcal{I} \text{ or } x^* = d \text{ for some } I = \{d\} \in \mathcal{I} \text{ and } x^{**} \in I' \neq I \in \mathcal{I}, \\ \text{and type } x^{**} \text{ weakly prefers ranking } \mathcal{I}(x^*, x^{**}) \text{ to ranking } \mathcal{I}.\end{gathered}$$

Proof of Proposition 6. 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

$$\sum_{(\tilde{a}, \tilde{b}] \in \mathcal{I}, \tilde{a} < \tilde{b} < x} \left[(\tilde{b} - \tilde{a}) \frac{\int_{\tilde{a}}^{\tilde{b}} y^A(\tilde{x})\, dF(\tilde{x})}{F(\tilde{b}) - F(\tilde{a})} - \int_{\tilde{a}}^{\tilde{b}} y^A(\tilde{x})\, d\tilde{x} \right] + \tag{22}$$

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

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 6, the category ranking $\mathcal{J} = \mathcal{I}(x^*, x^{**})$ Pareto improves over $\mathcal{I}$. Types $x \in [0, x^*]$ are indifferent between the two category rankings, because their allocation and performance are the same in both cases. 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.)

Then x' must be the lower endpoint of a non-trivial interval in $\mathcal{I}'$. Denote this interval 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 (22). It also cannot be that $x^{**} < x_l'$, since it would then follow from (22) 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 (22)) 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 $\mathcal{I}$.

Proposition 6 provides a method for checking whether a category ranking belongs to the Pareto frontier of category rankings. Proposition 6 and the arguments developed for the proof of Theorem 2 enable us to derive the simpler to pursue sufficient condition in Proposition 3 for a category ranking to be an element of the Pareto frontier of category rankings.

Proof of Proposition 3. Theorem 2 establishes that there is no Pareto-improving, mean-preserving contraction of the category ranking $\mathcal{I}$ that is constant on each interval $(x', x'']$ on which $\widehat{F}$ is linear. In particular, there is no such category ranking. However, this does not yet mean that $\mathcal{I}$ is on the Pareto frontier. It could be dominated by a category ranking that is not a contraction of $\mathcal{I}$. The rest of our proof is devoted to showing that this is impossible.

Suppose that the allocation induced by $\mathcal{I}$ is constant on all intervals on which $\widehat{F}$ is linear, but $\mathcal{I}$ is not on the Pareto frontier. Then there exist a pair of types $x^* < x^{**}$ described in Proposition 6. In particular, type x^{**} weakly prefers ranking $\mathcal{I}(x^*, x^{**})$ to ranking $\mathcal{I}$. Define another ranking $\mathcal{I}'$ that is obtained from ranking $\mathcal{I}$ by changing the interval $I' = (a', b']$ (from Proposition 6). If $x^{**} \in (a', b')$, then $I' = (a', b']$ is replaced with two intervals: $(a', x^{**}]$ and $(x^{**}, b']$. Otherwise, that is, if $I' = \{x^{**}\}$ or $x^{**} = b'$, then $\mathcal{I}' = \mathcal{I}$.

We will now show that types $x \leq x^{**}$ weakly prefer ranking $\mathcal{I}(x^*, x^{**})$ to ranking $\mathcal{I}'$. Note first that type x^{**} weakly prefers $\mathcal{I}$ to $\mathcal{I}'$. This is so, because type x^* is indifferent between the two rankings, and the difference between the payoff of $x \in [x^*, x^{**}]$ and the payoff of x^* increases faster under $\mathcal{I}$ than under $\mathcal{I}'$. Compare now ranking $\mathcal{I}(x^*, x^{**})$ to ranking $\mathcal{I}'$. Types $x \leq a$ (where $I = (a, b]$ in Proposition 6) are indifferent. Further, as observed in the proof of Proposition 6, the payoff of type x is a convex function of x. Since this payoff function is linear on $[a, x^{**}]$ for $\mathcal{I}(x^*, x^{**})$, and types a and x^{**} weakly prefer $\mathcal{I}(x^*, x^{**})$, so must do all types $x \in [a, x^{**}]$. (Type x^{**} prefers $\mathcal{I}(x^*, x^{**})$ to $\mathcal{I}$ by Proposition 6, and we have noticed earlier that x^{**} prefers $\mathcal{I}$ to $\mathcal{I}'$.)

Suppose first that x^{**} does not belong to the interior of any interval on which $\widehat{F}$ is linear. Then, ranking $\mathcal{I}(x^*, x^{**})$ restricted to interval $[0, x^{**}]$ Pareto dominates ranking $\mathcal{I}'$ restricted to this interval. However, this contradicts Theorem 4 from Appendix D, because $\mathcal{I}(x^*, x^{**})$ is a mean-preserving contraction of $\mathcal{I}'$ on $[0, x^{**}]$, and the allocation induced by $\mathcal{I}'$ is constant on any interval on which $\widehat{F}$ is linear.

So, suppose that x^{**} belongs to the interior of an interval on which $\widehat{F}$ is linear. Take the longest interval $(x', x'']$ that contains x^{**} on which $\widehat{F}$ is linear. It exists because the union of intervals that contain x^{**} on which $\widehat{F}$ is linear also has the two properties. This longest interval must be contained in an interval of ranking $\mathcal{I}$, because the allocation induced by $\mathcal{I}$ is constant on any interval on which $\widehat{F}$ is linear. So, $(x', x^{**}]$ is contained in an interval of ranking $\mathcal{I}'$ restricted to $[0, x^{**}]$. (More precisely, it is contained in $I' \cap [0, x^{**}]$.) This implies that the allocation induced by $\mathcal{I}'$ on $[0, x^{**}]$ is constant on any interval on which $\widehat{F}$ is linear. Since ranking $\mathcal{I}(x^*, x^{**})$ restricted to $[0, x^{**}]$ is a mean-preserving contraction of ranking $\mathcal{I}'$ restricted to $[0, x^{**}]$, $\mathcal{I}(x^*, x^{**})$ cannot Pareto dominate $\mathcal{I}'$ on $[0, x^{**}]$.

The following example shows that the Pareto frontier of Pareto-improving category rankings may include category rankings that do not satisfy the condition in Proposition 3.

Example 2. *Let G be the CDF of two prizes: 0 and 1, with mass $1/2$ on each. Let F be a strictly increasing CDF that satisfies the following conditions: (a) $F(1/2) = 1/2$; (b) there is an $x^* > 3/4$ such that $1 = (1/2)/F(1/2) > x^*/F(x^*)$ and $x/F(x) > x^*/F(x^*)$ for all $x > x^*$. We will not describe any parametric CDF with these properties, but it is easy to show that such CDFs exist just by drawing them.*

Consider a category ranking $\mathcal{I}$ that consists of two intervals: $[0, x^*]$ and $(x^*, 1]$. We claim that this category ranking is on the Pareto frontier of category rankings. To see why, notice first that any category ranking $\mathcal{J}$ that Pareto dominated $\mathcal{I}$ would have to have an interval $[0, x]$ for some $x \geq x^*$. Indeed, types sufficiently close to 0 would be worse off otherwise. (By condition (a), they would obtain a prize of 1 with a lower probability.) Notice next that pooling the types from $(x, 1]$ would have no welfare effects. Therefore, if there is a category ranking that Pareto dominates $\mathcal{I}$, then a category ranking $\mathcal{J}$ that consists of two intervals: $[0, x]$ and $(x, 1]$, where $x > x^*$, also Pareto dominates $\mathcal{I}$.

We will show that type 1 is worse off under $\mathcal{J}$ than under $\mathcal{I}$, and in this manner we will obtain a contradiction. Under both category rankings, type 1 obtains prize 1. The *performance required to obtain this prize in $\mathcal{J}$ and in $\mathcal{I}$ is determined by the indifference condition of types x and x^*, respectively. The first condition is*

$$x\frac{F(x)-F(1/2)}{F(x)} = x - t,$$

and the second condition obtains by replacing x with x^*. Thus,

$$t = x\frac{F(1/2)}{F(x)} \text{ and } t^* = x^*\frac{F(1/2)}{F(x^*)}.$$

This completes the argument because $t > t^*$ by the second part of condition (b).

It remains to show that $\mathcal{I}$ Pareto improves on the assortative allocation. Indeed, types from $[1/2, 1]$ are better off because of the first part of condition (b). And types from $[0, 1/2]$ are better off because $\mathcal{I}$ gives them a chance to obtain a prize of 1 at zero effort.

D 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-dependant effect on all her peers (those attending the same college), and the effects are additive. We will show that such peer effects fit 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 a finite number 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, 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 $p(x_i)$ captures the peer effect exerted by a player of type x_i. 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 allocation of prizes to students.

For the approximation, consider a mechanism that implements the assortative allocation of prizes to types. Then, for each prize y in the support of the limit prize distribution G we have that the effective prize is

$$\tilde{y} = y + \frac{\int_{x_L^y}^{x_H^y} p(\tilde{x})\, dF(\tilde{x})}{F(x_H^y) - F(x_L^y)} = \frac{\int_{x_L^y}^{x_H^y} (y + p(\tilde{x}))\, dF(\tilde{x})}{F(x_H^y) - F(x_L^y)}, \tag{23}$$

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. While Olszewski and Siegel's (2016) large contest framework does not formally accommodate prizes whose values depend on their allocation, it is easy to show that the behavior specified by the mechanism forms an ε-equilibrium for sufficiently large contests. And all our results on the characterization of Pareto improvements continue to hold for this mechanism, as we now show.

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

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

where the first equality follows from (23). As in the proof of Proposition **1**, pooling is Pareto improving if and only if

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

For online publication only

1. Online Appendix 1: Modeling the Turkish system
2. Online Appendix 2: Estimation strategy
3. Online Appendix 3: Counterfactuals
4. Online Appendix 4: Experiment: additional results
5. Online Appendix 5: Experiment: procedures and cost-to-school mapping
6. Online Appendix 6: Experiment: screen shots

1 Modeling the Turkish system

The students play a game in a stationary overlapping generations environment with a unit mass of college-bound students graduating from high school each year. Each student is solving a single-agent infinite-horizon dynamic problem. The time periods in this problem correspond to the last year in the middle school, first attempt at the college entrance exam and all the subsequent years.

Figure 1.1 summarizes the timing of the model and the shocks affecting students. Each student chooses the type of high school and tutoring in high school, which affects the college placement score at graduation. Then, the student decides whether to retake the exam or accept placement. Students who retake draw a new set of shocks to the score, and decide whether to retake again. The decisions on high school choice, tutoring and retaking constitute the game strategy. In each period, the student faces uncertainty over future exam scores. In this section, we outline student's decision problem and characterize the steady state Markov perfect equilibrium of this game.

1.1 High school period

Before enrolling into high school, the student chooses one of the three broad types of high school (public, private and Anatolian/science) and whether to get private tutoring. The respective choices of student i are denoted as $hs_i \in HS = \{PUB, PRIV, ANAT\_SC\}$ and $pt_i \in \{0, 1\}$. In total, there are six combinations to choose from: public school with no tutoring, public school with tutoring, private school with no tutoring, and so on.

1.1.1 Costs of schooling

Before making the choice (hs_i, pt_i), student i observes X_{0i}, a vector of her own background characteristics; this includes parental education, internet access, how the student is planning to fund her education, and the population of the city of the high school attended. The student also knows g_{i0}, her middle-school GPA that takes one of three values $(A, B$ or $C)$.

Schooling choices are associated with costs, which capture the fees and effort of keeping up with curriculum. Choosing high school category hs and private tutoring pt entails a cost of

$$c_i(hs,pt) = \gamma_{g_{i0}}(hs,pt) + u_{i,hs} + v_{i,pt} + w_{i,hs,pt},$$
$$hs \in \{PUB, PRIV, ANAT\_SC\},\ pt \in \{0,1\}, \quad (24)$$

This allows costs to vary with middle school GPA, which serves an early signal of ability for the student. The coefficients γ vary across choices (hs,pt). Schooling costs also depend on unobservables u_i, v_i and w_i. The first two reflect costs specific to formal schooling and private tutoring respectively. Idiosyncratic shocks $w_{i,hs,pt}$ are independently drawn from the standard Gumbel distribution for each (i,hs,pt). Public school with no tutoring is set as the baseline option so that $\gamma_{PUB,0,g_0} = u_{i,PUB} = v_{i,0} = 0$. We assume that the remaining shocks $[u_{i,PRIV}, u_{i,ANAT\_SC}, v_{i,1}]$ are jointly normal with a zero mean and a covariance matrix, Σ_c.

1.1.2 Schooling choices and payoffs

Choice of schooling (hs,pt) affects expected college placement scores in the subsequent periods. Student i's placement score after high school is given by

$$s_{i1}(hs,pt) = \underbrace{\rho_{g_{i0}}(hs,pt) + X_{i0}\chi_{g_{i0}}}_{t_{i0}(hs,pt)} + \lambda_{i1} + \varepsilon_{i1} \quad (25)$$

Parameter $\rho_{g_{i0}}(hs,pt)$ captures the effect of schooling on the expected score, while χ controls for demographics, X_{i0}. The subscript g_{i0} indicates that the parameters can vary with middle school GPA. The composite error term $(\lambda_{i1} + \varepsilon_{i1})$ represents the shock to the score, with λ_{i1} capturing its persistent part ("learning") and ε_{i1} being transitory noise. It is convenient to combine the first two terms and label them as t_{i0}, i's expected score conditional on the observables, X_{i0} and g_{i0}, and effort, (hs_i, pt_i).

In the theoretical model, there is no uncertainty, and we have the variable t, which is the score of the student. In our structural model, there is randomness, which is why the expected score t_{i0} is the counterpart of t in the theoretical model.

The student's objective is to maximize the expected net payoff

$$\max_{(hs,pt)} \left[bW(t_{i0}(hs,pt)) - c_i(hs,pt) \right] \tag{26}$$

conditional on the observables and the cost shocks. The function $W(t_{i0})$, captures the net present value of aiming for the expected placement score t_{i0} at graduation from high school, and b is the coefficient on it. We flesh out the definition of W in the following section. The choices in high school and the their consequences are depicted in Figure 1.1. The arrows show the direction of influence.

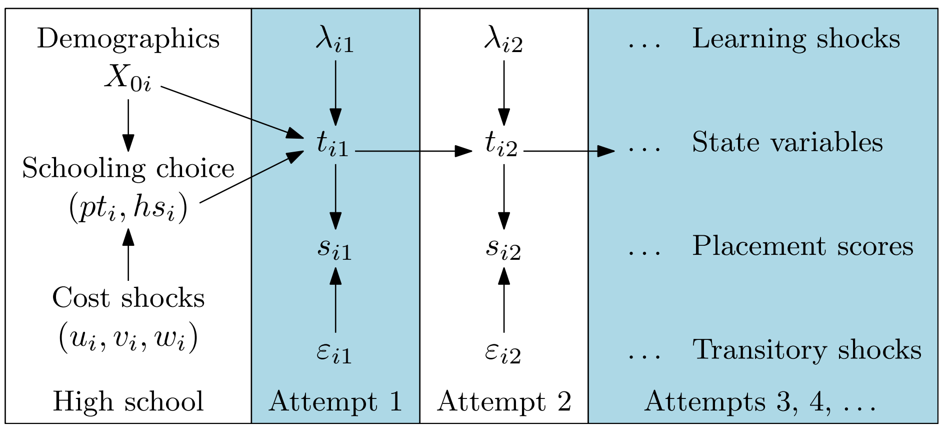

Figure 1.1: Shocks affecting students and their timing

We assume that students do not have private information about λ_{i1} when they make schooling choices; their expectations of the future scores are fully formed by the observables contained in X_{i0} and the GPA g_{i0}. The learning shock λ_{i1} and the transitory noise ε_{i1} do not affect this choice, which rules out endogenous selection into schooling choices. Figure 1.1 illustrates this by the absence of arrows coming from λ_{i1} and ε_{i1} to schooling choices.

1.2 College placement periods and retaking choices

In periods $\tau = 1, 2, \ldots$, students take the college entrance exam and, after each attempt, decide whether to be placed or keep taking the exam. Student's placement score in period τ is given by

$$s_{i\tau} = t_{i\tau} + \varepsilon_{i\tau}, \quad \text{where } t_{i\tau} = t_{i\tau-1} + \lambda_{i\tau}.$$

The pair $(t_{i\tau}, \varepsilon_{i\tau})$ defines student's state in period τ. Both shocks, $\lambda_{i\tau}$ and $\varepsilon_{i\tau}$, are drawn independently from the normal distribution and are only revealed to the student at the beginning of period τ as illustrated in Figure 1.1.

Once the student learns $\lambda_{i\tau}$ and $\varepsilon_{i\tau}$, she can retake the exam or get placed with the score $s_{i\tau}$ determining her placement ranking. Only the latest score matters for placement. Once the student is placed, she cannot take the exam again.

To formally characterize this decision problem, let V_τ be the value function for attempt τ and VR_τ be the expected payoff from retaking:[50]

$$\begin{aligned} V_\tau(t_{i\tau}, \varepsilon_{i\tau}) &= \max\left\{U\left(r^*\left(t_{i\tau} + \varepsilon_{i\tau}\right)\right), VR_\tau\left(t_{i\tau}\right)\right\} \\ VR_\tau(t_{i\tau}) &= \delta E_{\lambda_{i\tau+1}, \varepsilon_{i\tau+1}}\left[V_{\tau+1}\left(t_{i\tau} + \lambda_{i\tau+1}, \varepsilon_{i\tau+1}\right) \middle| t_{i\tau}\right] - \psi_\tau. \end{aligned} \tag{27}$$

where $U(\cdot)$ is a non-decreasing function mapping one's ranking in college placement to utility, $r^*(s)$ maps placement score to ranking, ψ_τ is the cost of retaking and δ is the discount rate. The student's state is captured by the persistent part of the score, $t_{i\tau}$ and the transitory shock, $\varepsilon_{i\tau}$. The expectation in the second line of equation (27) is taken over the shocks drawn in attempt $\tau + 1$, $\lambda_{i\tau+1}$ and $\varepsilon_{i\tau+1}$.

The utility function $U(r)$ is a structural parameter of the problem and is one of the key objects that map the empirical component to the theory.[51] We assume, as in the theoretical model, that all students have the same preferences so that being placed in a college seat ranked r delivers the same payoff $U(r)$ to any student. The ranking function $r^*(\cdot)$ is an equilibrium object, so it will change with the demand and supply of college seats.

Since t_{i0} is affected by investments made in high school, the value of effort is related to $W(t_{i0}(hs, pt))$, the expectation of V_1 at the time before the shocks λ_{i1} and ε_{i1} are realized, but after the schooling choices (hs_i, pt_i) have already been made:

$$W(t_{i0}(hs, pt)) = E_{\lambda_{i1}, \varepsilon_{i1}}\left[V_1(t_{i0}(hs, pt) + \lambda_{i1}, \varepsilon_{i1}) | t_{i0}(hs, pt)\right] \tag{28}$$

[50]Note that $t_{i\tau}$ carries over to the next period, while $\varepsilon_{i\tau}$ does not. This is why VR is a function of $t_{i\tau}$, but not of $\varepsilon_{i\tau}$.

[51]Recall that in the theoretical model, the distribution of prizes, y, is $G(y)$. This function gives the ranking needed to get the prize y. Therefore, $r = G(y)$, or $y = G^{-1}(r)$. Since the prize y is the counterpart of utility $U(\cdot)$, $U(r) = G^{-1}(r)$.

Thus, the valuations of the schooling choices in problem (26) are directly derived from the impact these choices have on subsequent placement outcomes in periods $\tau = 1, 2, \ldots$.

1.3 Equilibrium

In equilibrium, the strategies of all students and the supply of college seats determine the mapping $r^*(s)$, and the mapping determines the strategies so that the actions and the expectations are aligned. Student i's strategy is the best response to r^*. This strategy prescribes schooling choices in period $\tau = 0$ and retaking decisions in $\tau = 1, 2, \ldots$ for each feasible state $(t_{i\tau}, \varepsilon_{i\tau})$.

We restrict ourselves to stationary Markov-perfect equilibria in this game. In this equilibrium, a new cohort of students enters high school every period. Each student is infinitesimally small; there is a unit mass of students in the cohort. The distributions of demographics, middle-school GPA and all the shocks are the same in all cohorts.

There is a unit mass of college seats. As the equilibrium is stationary, all seats are filled each period so that the number of students entering the game matches the number of those who leave. As a result, the mass and the composition of retakers will also be constant in the stationary equilibrium. Consequently, the ranking function $r^*(\cdot)$ is not changing over time.

In equilibrium, given the ranking function $r^*(s)$ students play their optimal strategies: the choice of schooling maximizes the payoff in (26), while the retaking strategy (29) is a solution to the Bellman's equation (27). Each student ignores her own impact on r^* as she is infinitesimally small.

The highest ranking is $r^* = 1$, while the lowest one is $r^* = 0$.[52] The rank of a student with score s, $r^*(s)$, is one minus the mass of students from all the active cohorts (which include the current cohort and retakers) who are: (a) not yet placed, (b) whose optimal strategies call for placement in the current period, and (c) whose scores are greater or equal to s.

The optimal retaking strategy can be characterized by a threshold rule. As the placement utility, $U(r)$, is non-decreasing in r, and the ranking function, $r^*(s)$, is non-decreasing in s,

[52]Off-equilibrium, $r^* < 0$ is possible; this corresponds to not being placed anywhere in the current period.

the payoff from being placed, $U(r^*(t_{i\tau} + \varepsilon_{i\tau}))$ is non-decreasing in $\varepsilon_{i\tau}$. Since the value of retaking, VR_τ, does not depend on $\varepsilon_{i\tau}$, while $U\left(r^*\left(t_{i\tau} + \varepsilon_{i\tau}\right)\right)$ does, student's decision should follow a simple threshold rule: retake in attempt τ if

$$\varepsilon_{i\tau} < e_\tau(t_{i\tau}). \tag{29}$$

where $e_\tau(t_{i\tau})$ denotes the retaking threshold for i.

2 Estimation strategy

Our goal is to estimate the model's structural parameters. We proceed in five steps:

1. We estimate the variances of the score shocks in the first attempt, $\sigma_{\lambda 1}$ and $\sigma_{\varepsilon 1}$.

2. Then, we turn to the dynamic problem (27). Our objective here is to find the parameters of the learning shock distributions, $\mu_{\lambda\tau}$ and $\sigma_{\lambda\tau}$, $\tau = 2, 3, \ldots$, which determine how student states evolve over attempts. We also non-parametrically estimate the cutoff function $e_\tau(t_{i\tau})$, which describes the equilibrium retaking strategy (29).

3. We use the above estimates to find payoff-related parameters: the costs of retaking, ψ_τ, and the placement payoff function, $U(r)$. The latter is identified non-parametrically for all values of $r \in [0, 1]$ using a flexible approximation.

4. Next, we estimate the improvement in score associated with each level of pre-test effort and the coefficients for predicting the score based on demographics, $\rho_{g_{i0}}(hs, pt)$ and $\chi_{g_{i0}}$ in (25).

5. Lastly, we recover cost and benefit parameters $\gamma_{hs,pt}(g_{i0})$, b and Σ_c that rationalize schooling choices in problem (26).

2.1 Step 1: distribution of λ_{i1} and ε_{i1}.

Our main objective in this step is to obtain the residual in the score equation (25) for the first attempt that cannot be explained by observable demographics or the choice of schooling and attribute its variation to two parts: λ_{i1}, which represents immutable ability, and ε_{i1}, which stands for pure randomness.

The exam consists of four subject sections: mathematics, natural sciences, social studies and Turkish language. The unexplained gains in score during high school that persist between the exam attempts are also likely to affect all exam sections in the same attempt. For example, having a good fit between the student and the school would raise GPA and performance in all subjects and attempts. Conversely, shocks that reflect pure luck are unlikely to affect scores across all subjects or survive between attempts. Following this intuition, we separate the unexplained variation in score that persists across subjects from the subject-specific variation, and attribute the former to λ_{i1}, and the latter to ε_{i1}.

More formally, let the subject exam scores in the first attempt and the high school GPA be specified as follows:

$$\begin{aligned} s_{ij1} &= X_i'\beta_j + \theta_{iv}\alpha_{jv} + \theta_{iq}\alpha_{jq} + \varepsilon_{ij1}, \quad j \in \{Math, NatSci, SocStudies, Turkish\} \qquad (30) \\ g_i &= X_i'\beta_g + \theta_{iv}\alpha_{gv} + \theta_{iq}\alpha_{gq} + \varepsilon_{ig} \end{aligned}$$

The vector X_i includes observable student characteristics, pre-test schooling choices and tutoring. Gains in student ability during high school unexplained by these observables are captured by student-specific common factors $\theta_i = (\theta_{iq}, \theta_{iv})$, which drive persistent performance in quantitative and verbal tasks, while $(\varepsilon_{ij1}, \varepsilon_{ig})$ are the idiosyncratic shocks. The weights on θ_i reflect the relevance of each skill for each subject.[53] Abilities captured by $(\theta_{iq}, \theta_{iv})$ will therefore affect performance across the board, in multiple subjects and between the attempts. In contrast, idiosyncratic gains captured by $(\varepsilon_{ij1}, \varepsilon_{ig})$ are attempt- and subject-specific.

We start step 1 by estimating the coefficients (β_j, β_g) in equations (30) using OLS on the sample of first-time takers. The residuals represent $\theta_i'\alpha_j + \varepsilon_{ij1}$ and $\theta_i'\alpha_g + \varepsilon_{ig}$. To recover the distribution of $(\theta_i, \varepsilon_{ij1}, \varepsilon_{ig})$, we assume that all these shocks are jointly normal and independent with the exception of θ_{iq} and θ_{iv}, which can be correlated as a good match of

[53] The loadings in the math and Turkish equations are normalized to $\alpha_M = [1, 0]'$ and $\alpha_T = [0, 1]'$ respectively: in other words, quantitative ability is a common factor that affects the math score but not the Turkish score and vice versa for verbal ability.

the school and the student would likely affect both verbal and quantitative skills. Let R_{ig} and R_{ij} be the residuals from the GPA and the subject score equations. The covariance matrix for the vector $R_i = [R_{ig}, R_{iM}, R_{iT}, R_{iSc}, R_{iSS}]'$ is related to the factor loadings, the variances of the idiosyncratic shocks and the covariance matrix of θ_i as follows:

$$Var[R_i] = \alpha' \Sigma_\theta \alpha + \Sigma_\varepsilon, \tag{31}$$

where α denotes the matrix of factor loadings and Σ_ε is a matrix with the variances of ε_{ig} and ε_{ij1}'s on the main diagonal. Parameters α, Σ_θ and Σ_ε, which capture the distribution of unobservables, are estimated via GMM using equations in (31) as the identifying moment conditions.

Note that the aggregate exam score, which is used to determine the placement rank, is composed of the subject scores and the GPA with the weights w[54]. This connects the factor model in (30) to the earlier notation in equation (25):

$$s_{i1} = \underbrace{X_i \left(\sum_j \beta_j w_j + w_g \beta_g \right)}_{t_{i0}} + \underbrace{\theta_i' \left(\sum_j \alpha_j w_j + \alpha_g w_g \right)}_{\lambda_{i1}} + \underbrace{\left(\sum_j \varepsilon_{ij1} w_j + \varepsilon_{ig} w_g \right)}_{\varepsilon_{i1}}$$

2.2 Step 2: State transitions and students' policy function

This subsection and the ones that follow are almost exactly the same as sections 3 and 4 in Krishna et al (2018). Needed changes in notation have been made.

Estimating gains in student scores between the exam attempts is confronted by two obstacles. First, retaking is endogenous and can be partly driven by student's known ability. Second, as we only observe one cross section of exam takers, we can only identify learning gains by comparing cohorts rather than by tracking individual students and their scores between exam attempts.

Our estimation strategy is inspired by the conditional choice probability approach.[55] We impose the exclusion restriction that high school GPA is not directly related to retaking.

[54] These weights are given by the central exam authority.

[55] For example, see Hotz and Miller (1993).

In other words, conditional on the noise-free score, $t_{i\tau}$, we assume that just having a high GPA does not make one more or less likely to retake. This makes the GPA distribution of repeat takers differ from that of first time takers only because of selection. The distribution of exam scores of repeat takers, in contrast, is affected by both selection and learning. Thus, by comparing the distributions of scores and GPAs across attempts, we are able to distinguish learning from selection. We assume steady state so that second time takers in a given year can be thought of as identical to retakers from today's cohort of first time takers and so on.

Below we depict how selection and learning operate in a very simple example that highlights the essential intuition. We use first and second time takers for concreteness. Assume there are two types, labeled high and low, with half of the agents being of each type. GPA and the placement scores in the two attempts are

$$\begin{aligned} g_i &= \theta_i + \varepsilon_{ig}, \\ s_{i1} &= \theta_i + \varepsilon_{i1}, \\ s_{i2} &= \theta_i + \lambda_{i2} + \varepsilon_{i2}, \end{aligned}$$

where $\theta_i = 2$ for i if she is a high type and $\theta_i = 1$ if she is a low type. The shocks ε_{ig}, ε_{i1}, ε_{i2} are drawn from a uniform distribution over $[-1, 1]$. Learning shocks are deterministic and the same for both types, so $\lambda_{i2} = \bar{\lambda}$.[56]

The joint distribution of g_i and s_{i1} for the universe of agents is depicted in Figure 2.1, where we have the high school GPA and the placement score on the two axes. The shaded lower left square corresponds to the low type, while the shaded upper right one depicts the high type. The retaking rule is to retake if $\varepsilon_{i1} < e_1(t_{i1})$ and hence if $s_{i1} < t_{i1} + e_1(t_{i1})$. Thus, the low type retakes if her score s_{i1} is below $1 + e_1(1)$, and the high type does so when $s_{i1} < 2 + e_1(2)$.[57] In this simple model, t_{i1} is just θ_i so that the two types have different cutoffs; for example they could be as depicted by the two vertical dashed lines in Figure 2.1.

[56]Note that this example differs from the estimated model as it assumes placement scores do not include GPA as a component (i.e., $w_{ig} = 0$). As a result, retaking cutoffs vary only by ability type. In contrast, in our estimated model, retaking cutoffs depend on the permanent part of the score (i.e., type), which includes the GPA among other components of the score.

[57]This rule depends on the shape of the utility function.

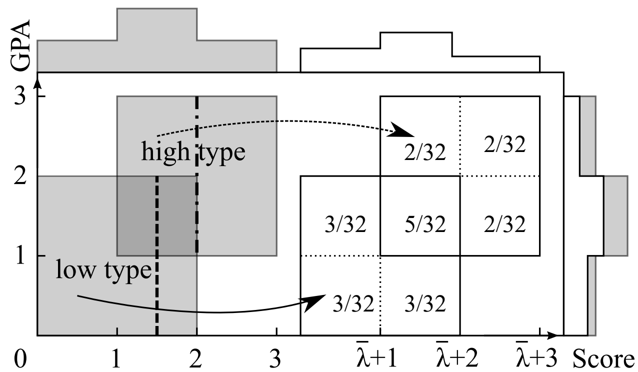

Figure 2.1: Identifying learning and selection using scores and high school GPA

Notes: The boxes inside the diagram depict the joint distributions of scores and GPAs in attempts 1 and 2. The histograms at the top are marginal distributions of scores (inflated by the number of students). The histograms on the right hand side of the diagram show marginal distributions of GPA. Distributions related to the first attempt are shaded, while distribution of second time takers are in white. The dashed lines depict the retaking thresholds, $t_{i\tau} + e_1(t_{i\tau})$, for each student type.

We have low types being more likely to retake in this figure. The shaded histograms on the right and the top of the figure denote the marginal distribution of GPA and scores among first time takers, respectively.

The unshaded histograms depict the distributions of GPA and scores, joint and marginal, of second time takers. Half the high type and three quarters of the low type of agents retake. This results in the mass in the upper tail of the GPA distribution of second time takers being lower than that in the lower tail. The learning shocks are depicted by the arrows. Given our assumptions on learning shocks, the distribution of scores moves to the right by $\bar{\lambda}$. The number in each unshaded box is equal to the mass of second-time takers there. The total mass equals the share of first time takers who choose to retake.

We identify the cutoff scores for the two types by comparing the distributions of GPAs of first and second time takers. The mass at any point in the upper tail of the GPA distribution of second time takers gives us the cutoff for the high type, while the mass at any point in the lower tail pins down the cutoff for the low type. The shift in the score distribution between the first and second attempt pins down the learning shock. If learning shocks are stochastic, in addition to being shifted to the right as depicted in Figure 2.1, the distribution of scores gains more variance due to variation in λ_{i2}. This extra variance allows us to identify $\sigma^2_{\lambda 2}$, the variance of learning shocks in attempt 2. A similar argument applies to third versus second time takers, and so on.

2.2.1 Estimation algorithm for attempts 1–3.

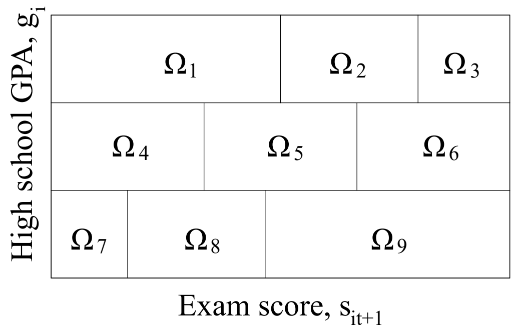

Figure 2.2: Sets of GPAs and scores used to construct the GMM estimator

The retaking threshold function and the parameters of learning shocks, $e_\tau(t)$, $\mu_{\lambda\tau+1}$ and $\sigma_{\lambda\tau+1}$, are estimated sequentially using a separate GMM routine for each attempt $\tau = 1, 2, 3$. This GMM estimator is designed to match predicted numbers of exam takers whose GPAs and scores fall into sets $\Omega_1, \ldots, \Omega_9$, depicted in Figure 2.2. The horizontal lines in this figure are the 1st and the 2nd tercile of the GPA distribution, while the vertical lines are exam score terciles conditional on GPA being between the respective horizontal lines.

Let $T_{i\tau} = [t_{i1} \ldots t_{i\tau}]$ denote the trajectory of noise-free scores that student i would face if she takes the exam at least τ times, and $\mathcal{F}(T_{i\tau}|s_{i\tau}, g_i, X_i)$ be the probability measure in the space of such trajectories conditional on the observables in attempt τ. Let $a_{i\tau}$ be a dummy that equals one if i is not placed after τ attempts. The set of moment conditions that are used to estimate $e_\tau(\cdot)$, $\mu_{\lambda\tau}$ and $\sigma_{\lambda\tau}$ take the following form:

$$\begin{aligned}\Pr[(g_i, s_{i\tau+1}) \in \Omega_k, a_{i\tau+1} = 1|a_{i\tau} = 1] = &\int \left[\int \Pr\{a_{i\tau} = 1|T_{i\tau}\}\Phi\left(\frac{e_\tau(t_{i\tau})}{\sigma_\varepsilon}\right)\right.\\ &\left.\times \Pr\{(g_i, s_{i\tau+1}) \in \Omega_k|g_i, t_{i\tau}; \mu_{\lambda\tau+1}, \sigma_{\lambda\tau+1}\}d\mathcal{F}(T_{i\tau}|s_{i\tau}, g_i, X_i)\right] \frac{d\mathcal{G}(s_{i\tau}, g_i, X_i|a_{i\tau} = 1)}{\Pr\{a_{i\tau} = 1|s_{i\tau}, g_i, X_i\}}\\ &k = 1, \ldots, 9. \quad (32)\end{aligned}$$

where $\mathcal{G}$ denotes the distribution of the observables among τ-time takers.

The expression on the right hand side of (32) predicts the share of applicants who retakes after attempt τ and whose score-GPA combination in attempt τ+1 ends up in the cell Ω_k. We take an expectation over all observables among τ-time takers, $s_{i\tau}$, g_i and X_i, and trajectories $T_{i\tau}$ that the student's noise-free scores can follow conditional on the observables. The first term inside the inner integral is the probability that the trajectory is not interrupted by a placement decision before attempt τ. The second one is the probability that the student retakes at τ. Finally, the third term gives the probability that the trajectory of scores lands into Ω_k in attempt $\tau+1$. The inner integral is divided by the probability of surviving τ attempts conditional on observables in attempt τ:

$$\Pr\{a_{i\tau}=1|s_{i\tau},g_i,X_i\}=\int \Pr\{a_{i\tau}=1|T_{i\tau}\}d\mathcal{F}(T_{i\tau}|s_{i\tau},g_i,X_i) \tag{33}$$

In order to use conditions (32) in a GMM estimator, one has to compute the integrands in (32) and (33) and approximate the expected values with the finite-sample analogs.[58]

First, the probability of surviving τ attempts along the trajectory $\bar{S}_{i\tau}$ is found from (29). The student keeps retaking if her ε shocks received in all prior attempts stay below the retaking thresholds. Therefore, the estimate for the above probability is

$$\widehat{\Pr}\{a_{i\tau}=1|T_{i\tau}\}=\prod_{l=1}^{\tau-1}\Phi\left(\frac{\widehat{e}_l(t_{il})}{\hat{\sigma}_\varepsilon}\right) \tag{34}$$

The probability that the combined shock $\lambda_{i\tau+1}+\varepsilon_{i\tau+1}$ takes the student from $t_{i\tau}$ into Ω_k is estimated by

$$\widehat{\Pr}\{(g_i,s_{i\tau+1})\in\Omega_k|g_i,t_{i\tau};\mu_{\lambda\tau+1},\sigma_{\lambda\tau+1}\}$$
$$=\begin{cases}0, \text{if } (g_i,s)\notin\Omega_k\ \forall s,\\ \Phi\left(\frac{S_{ku}-t_{i\tau}-\mu_{\lambda\tau+1}}{\sqrt{\hat{\sigma}^2_\varepsilon+\sigma^2_{\lambda\tau+1}}}\right)-\Phi\left(\frac{S_{kl}-t_{i\tau}-\mu_{\lambda\tau+1}}{\sqrt{\hat{\sigma}^2_\varepsilon+\sigma^2_{\lambda\tau+1}}}\right), \text{ otherwise,}\end{cases}$$

where S_{ku} and S_{kl} are the upper and the lower boundaries of scores in Ω_k.

Integrating over $T_{i\tau}$ in (32) and (33) requires the knowledge of joint density of $t_{i1},\dots,t_{i\tau}$ conditional on $s_{i\tau}$, g_i and X_i. Note that this density describes the whole population, not just the retakers who survive τ attempts. The latter fact allows us to use normality and independence of all shocks in the model. This assumption implies that the variables $t_{i1},\dots,t_{i\tau},s_{i\tau}$, g_i are jointly normal conditional on X_i in the original cohort of students. The mean and the

[58] For more details and the formal proofs, see the online appendix to Krishna et al (2018).

covariance matrix of this distribution depend on $\{\mu_{\lambda l}, \sigma_{\lambda l}\}_{l=2}^{\tau}$ and the parameters estimated in step 1. This implies that the sequence $t_{i1}, \ldots, t_{i\tau}$ is jointly normal, too, conditional on $s_{i\tau}$, g_i and X_i. The mean and the covariance matrix of this distribution can be easily derived from the above parameters.

In order to obtain finite-sample analogs for the moment conditions in (32), we approximate the left hand side by the number of $(\tau+1)$-time takers in the set Ω_k divided by the total number of τ-time takers in the data. The outer integral on the right hand side is approximated by the average over $(s_{i\tau}, g_i, X_i)$ of τ-time takers in the data. The inner integral is computed numerically using Gauss-Hermite quadratures.

We approximate the unknown retaking thresholds $e_\tau(t)$ by piecewise-linear functions defined on three grid points, $t^*_{1,\tau}, t^*_{2,\tau}, t^*_{3,\tau}$:

$$e_\tau(t) = \begin{cases} e^*_{1,\tau}, & t \le t^*_{1,\tau} \\ e^*_{1,\tau} + (e^*_{2,\tau} - e^*_{1,\tau})\frac{t - t^*_{1,\tau}}{t^*_{2,\tau} - t^*_{1,\tau}}, & t^*_{1,\tau} \le t \le t^*_{2,\tau} \\ e^*_{2,\tau} + (e^*_{3,\tau} - e^*_{2,\tau})\frac{t - t^*_{2,\tau}}{t^*_{3,\tau} - t^*_{2,\tau}}, & t^*_{2,\tau} \le t \le t^*_{3,\tau} \\ e^*_{3,\tau}, & t^*_{3,\tau} \le t \end{cases}$$

The grid points $t^*_{1,\tau}$ and $t^*_{3,\tau}$ are located at the 20^{th} and the 80^{th} percentiles of $s_{i\tau}$ among τ-time takers, while $t^*_{2,\tau} = (t^*_{1,\tau} + t^*_{3,\tau})/2$. In total, we have nine equations in (32) to identify five parameters, $e^*_{1,\tau}$, $e^*_{2,\tau}$, $e^*_{3,\tau}$, $\mu_{\lambda\tau+1}$ and $\sigma_{\lambda\tau+1}$, for cohorts $\tau = 1, 2, 3$.

2.2.2 Estimation for attempts greater than 3.

In our data, the number of attempts is censored at five. Thus, we cannot use moment conditions (32) for $\tau \ge 4$. Thus, we assume that the students stop improving their scores and the costs of retaking stop changing after the fourth attempt.[59] We rely on one important implication of this assumption: student's future stream of payoffs does not depend on τ after $\tau = 4$ as the dynamic problem in (27) becomes stationary. Thus, $e_\tau(t) = e_4(t)$ for any $\tau \ge 4$. In order to completely describe student's behavior after fourth attempt, one has to pin down a single threshold function $e_4(t)$.

[59] Taking the exam four times gives the student at least 3 extra years to prepare, and it is natural to assume that the investing preparation time has diminishing returns.

The set of moment conditions for $e_4(\cdot)$ is obtained by summing both sides of equation (32) for attempts $\tau \geq 4$ and expressing the right hand side in terms of variables observed in attempt 4:

$$\sum_{\tau\geq 4} \Pr[(g_i, s_{i\tau+1}) \in \Omega_k, a_{i\tau+1} = 1 | a_{i4} = 1] = \int \left[\int \Pr\{a_{i4} = 1 | T_{i4}\} \frac{\Phi\left(\frac{e_4(t_{i4})}{\sigma_\varepsilon}\right)}{1 - \Phi\left(\frac{e_4(t_{i4})}{\sigma_\varepsilon}\right)} \right.$$
$$\left. \times \Pr\{(g_i, s_{i\tau+1}) \in \Omega_k | g_i, t_{i4}\} d\mathcal{F}(T_{i4} | s_{i4}, g_i, X_i) \right] \frac{d\mathcal{G}(s_{i4}, g_i, X_i | a_{i4} = 1)}{\Pr\{a_{i4} = 1 | s_{i4}, g_i, X_i\}} \qquad (35)$$
$$k = 1, \ldots, 9.$$

Our GMM estimation procedure is organized sequentially. First, we estimate parameters associated with the first retaking decision: $e_1(\cdot)$, $\mu_{\lambda 2}$ and $\sigma_{\lambda 2}$. Then, we set up and run the GMM estimator for $e_2(\cdot)$, $\mu_{\lambda 3}$ and $\sigma_{\lambda 3}$, using $\widehat{e}_1(\cdot)$, $\widehat{\mu}_{\lambda 2}$ and $\widehat{\sigma}_{\lambda 2}$ to compute the probability of survival in (34) and the distribution of noise-free scores $\mathcal{F}$. Then, we obtain the estimates for $\tau = 3$ in a similar way; we use the GMM estimator based on (35) for $\tau = 4$.[60]

2.3 Step 3: payoff-related parameters

In the final step of our estimation procedure, we find the fundamental components of the model related to payoffs: the costs of retaking, ψ_τ, and the utility function, $U(r)$. We do so along the lines of step 2 in Hotz-Miller's CCP algorithm.

First, note that the continuation value function $VR_\tau(t_{i\tau})$ from Bellman's equation (27) can be found by integrating the net present value of future payoffs over all trajectories $t_{i\tau+1}, t_{i\tau+2}, \ldots$ and all future student actions. Since we have estimated the parameters of learning shocks in step 2, we know the law that generates state transitions: $(t_{i\tau+1} - t_{i\tau}) \sim N[\mu_{\lambda\tau+1}, \sigma_{\lambda\tau+1}]$, $\varepsilon_{i\tau} \sim N[0, \sigma_\varepsilon]$. We also know the decision rule that students use in equilibrium: retake in attempt τ if $\varepsilon_{i\tau} < e_\tau(t_{i\tau})$. Thus, given candidate values of ψ_t and $U(r)$, we can compute the continuation value $VR_\tau(t_\tau)$ on a grid of t_τ by simulating shocks that hit

[60]Estimating parameters sequentially is less efficient that using one GMM routine with optimal weights. However, a joint GMM routine requires computing the probability in (33), which is very demanding computationally. If we were estimating all parameters in one run, we would have to recompute this probability on every iteration of the GMM algorithm.

students, student reactions to these shocks and the associated payoffs. Knowing who chooses placement and what their scores are, we can also construct the inverse cutoff function, $r^*(s)$.

Once the continuation values are known for every income group, attempt and grid point, one can impose the assumption that students maximize their utility. That is, a rational student in the state $(t_{i\tau}, \varepsilon_{i\tau})$ retakes if the utility of placement is lower than the continuation value:

$$U(r^*(t_{i\tau} + \varepsilon_{i\tau})) < VR_\tau(t_{i\tau}).$$

We use this inequality to find the threshold for ε_τ, $\tilde{e}_\tau(t_{i\tau})$, below which a rational student chooses to retake.

Finally, we plug the threshold functions $\tilde{e}_\tau(t_{i\tau})$ into moment conditions (32) and (35) in place of $e_\tau(t_{i\tau})$. We find the estimates for ψ_τ and $U(r)$ by minimizing the objective function for unweighted GMM employing all moment equations in (32) and (35). In contrast to step 2, we use all moment conditions simultaneously since one of the parameters being estimated (the utility of placement) is common to all attempts.

The utility function is parameterized in a flexible manner as:

$$U(r^*(s)) = \sum_{j=1}^{10} \gamma_j \Phi\left(\frac{s - s_j}{h}\right), \quad \gamma_j \geq 0, \quad \sum_{j=1}^{10} \gamma_j = 1.$$

The coefficients γ are allowed to differ by income group. The normalization of $\sum_{j=1}^{10} \gamma_j = 1$ ensures that the utility at $s = \infty$ is unity. As $\Phi(.)$ is increasing in s, constraining $\gamma_j \geq 0$ ensures that the utility function is non-decreasing. The larger is h, the smoother is the function; we set $h = 15$.

Note that γ is not a structural parameter; the function approximated by γ depends on $r^*(\cdot)$, an equilibrium outcome, which changes in response to policy interventions. After obtaining the estimates of γ, we use the simulated $r^*(\cdot)$ to find $U(r)$ for the values of $r = 0, 0.01, \ldots, 1$. In total, we have four cost and ten utility parameters: $\psi_1 \ldots \psi_4$, $\gamma_1 \ldots \gamma_{10}$. They are identified using 36 moment conditions in equations (32) and (35): nine conditions (one for each cell Ω_k) for each attempt $\tau = 1, 2, 3$ in (32) and nine more in (35).

The economics behind identification of the payoff parameters is as follows: if the marginal utility of a higher score increases sharply at s, then students close to s will be risk loving and hence tend to retake the exam more than students with s where utility of the score is less convex. Thus, the observed local retaking rates pin down the curvature of the utility function. The retaking costs are pinned down by the overall retaking rates in the given attempt. We do not attempt to estimate δ as it is well known that discount factors are hard to identify in such settings (see Magnac and Thesmar (2002)); we set δ at 0.9 for all students.

2.4 Step 4: returns to pre-test effort

In order to estimate the effect of pre-test effort on the exam scores in the first attempt, we estimate equation (25) using OLS. The key parameter of interest is $\rho_{g_{i0}}(hs, pt)$ for every combination of middle school GPA, $g_{i0} \in \{A, B, C\}$, high school type, $hs \in \{PUB, PRIV, ANAT\_SC\}$, and the dummy for getting private tutoring, $pt \in \{0, 1\}$.

Estimating returns to schooling raises usual concerns over endogenous selection into treatment. For instance, students with higher ability and more educated parents may be more likely to choose better schools and obtain higher scores at the entrance exam. Unless one controls for ability, the true effect of schooling could be confounded by this relationship.

To address these concerns, we include three sets of controls: middle school GPA, parental occupation and parental level of education, separately for both parents. Other controls include access to the Internet, the expected sources of funds to cover education, the population in the home city, the number of siblings and gender.

2.5 Step 5: costs and benefits of pre-test effort

In this step, we need to estimate the parameters responsible for the costs and benefits of schooling choices before the first exam attempt. For students with each level of middle school GPA ($g_{i0} = A$, B or C), we estimate the set of coefficients $\gamma_{g_{i0}}(hs, pt)$ for six schooling levels: public, private or Anatolian high school with or without extra tutoring. These costs are allowed to depend on middle school GPA as many good schools select students based on performance. Since we use public school with no tutoring as the baseline, we are left with 5 parameters to estimate. We also need to estimate the covariance matrix Σ_c of student-specific cost shocks $[u_{i,PRIV}, u_{i,ANAT\_SC}, v_{i,1}]$, which would allow us to capture relationships in individual cost shocks driven by unobservables that can potentially cause bias. Finally, we need to estimate the coefficient b that determines the tradeoff between benefits and costs in the payoff function (26).

We start by computing the value of effort $W(t_{i0}(hs, pt))$ for each first-time exam taking student. We use the estimates from step 4 to find the expected score $t_{i0}(hs, pt) = \widehat{\rho}_{g_{i0}}(hs, pt) + X_{i0}\widehat{\chi}_{g_{i0}}$ for all six possible schooling options (hs, pt) conditional on the observables and middle-school GPA. Then, we recursively solve Bellman's equation (27) using the parameter estimates obtained in step 3, which describe the payoff structure and the distribution of shocks. In particular, we obtain the value function $V_1(t_{i1}, \varepsilon_{i1})$. We numerically integrate this function over the values of the learning shock λ_{i1} using the Gauss-Hermite quadrature formula and compute the payoff to each schooling option $W(t_{i0}(hs, pt))$ as prescribed by equation (28).

Once we know the value of effort associated with each choice, the choice of schooling in (26) just becomes the standard mixed logit discrete choice problem. In equations (24) and (26), the payoffs depend on the unknown parameters (b, $\gamma_{g_{i0}}(hs, pt)$ and Σ_c). The idiosyncratic shock w is drawn from the standard Gumbel distribution, the error components u and v can be viewed as random coefficients on the school type dummies. We estimate the parameters b, $\gamma_{g_{i0}}(hs, pt)$ and Σ_c in this model using maximum likelihood.

The estimates of b and $\gamma_{g_{i0}}(hs, pt)$ are important for calibrating our theoretical model. Intuitively, if b is low, i.e. performing better has little value, then the cost shocks are more important in driving choices rather than the benefits of scoring high. Consequently, we would see little relationship between the payoffs from costly schooling and the share of students opting for such schools. The γ's are identified by comparing gains from the various options to the shares of students opting for them. If, for example, public schools are chosen by many students despite low gains in scores, this means the costs of attending public schools are low.

3 Counterfactuals

3.1 Robust Pareto frontier and bottom pooling

While the robust Pareto frontier policy should pool agents in all the intervals where the distribution of types, $F(x)$ lies below its concavification, $\widehat{F}(x)$, all the upper intervals combined include just a small fraction of the population. We therefore compare the robust Pareto frontier policy to the minimal bottom policy that only pools agents in the lowest interval (robust bottom pooling). Simulated gains under these policies are indistinguishable, as is evident from Figure 3.1. For this reason, we restrict our attention to single-interval policies (which need not be robust) including bottom pooling, top pooling, and any single-interval policy in between.

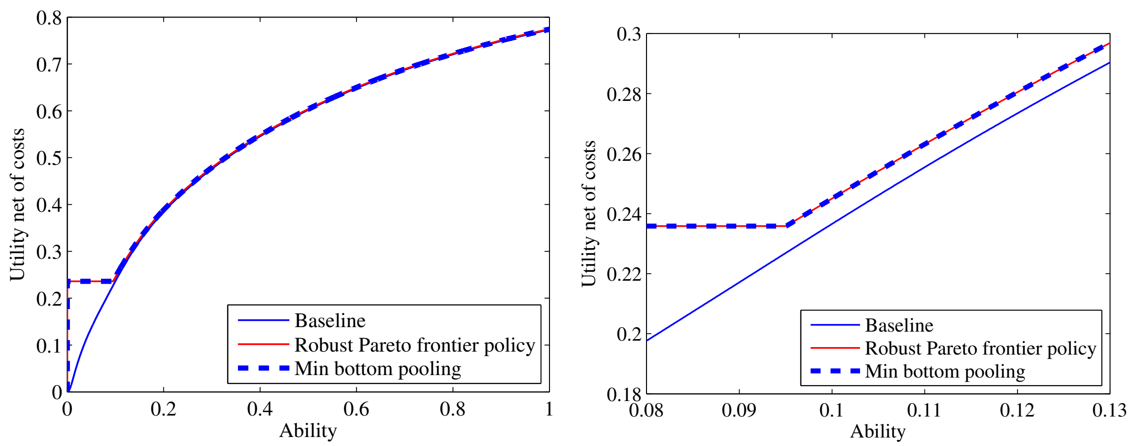

(a) Welfare gains for all types

(b) Welfare gains, zoomed at the upper bound of the pooling interval

Figure 3.1: Net utility $(y - c(t)/x)$: robust Pareto frontier policy vs minimal bottom pooling

3.2 Bottom pooling in the structural model

While the theoretical model is useful in producing tractable intuition, one may wonder if the results survive in a more realistic setting. In particular, two features of the model may raise readers' concern: (a) in the model, exam scores are noise-free, which means that students can perfectly target their desired ranking, (b) the model does not allow for exam retaking, while retaking is prevalent in the data.

In this subsection, we use numerical simulations to explore the consequences of adding noise to the exam scores and allowing retaking. We simulate the outcomes of the maximal bottom pooling policy found in Section 7.2, but we do so using the structural model outlined in Section 7.1 rather than the theoretical model from Section 2. The structural model allows for multiple sources of uncertainty in scores and explicitly accounts for retaking.

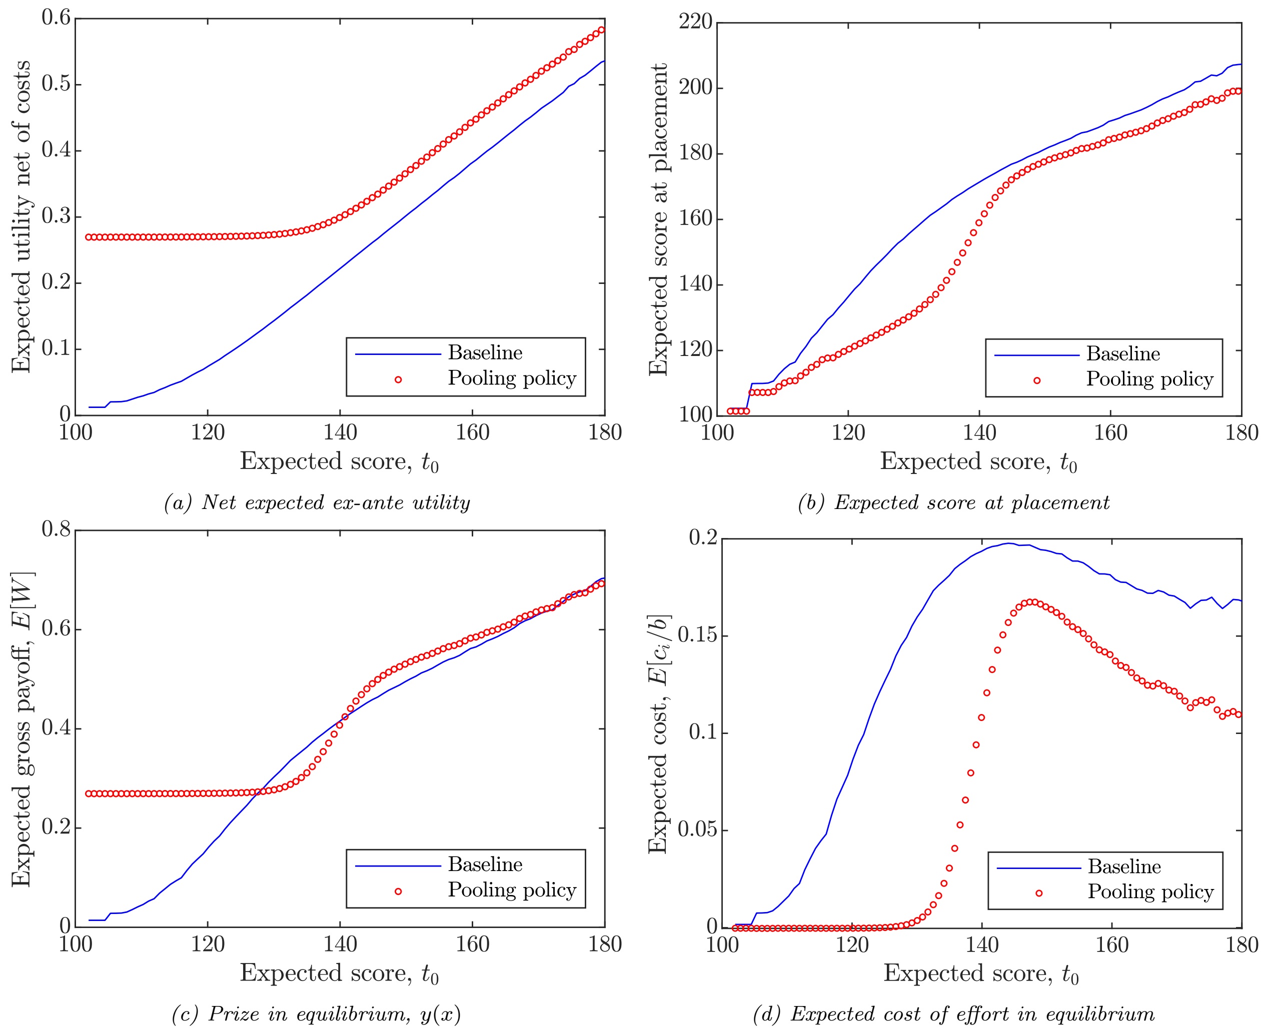

(a) Net expected ex-ante utility

(b) Expected score at placement

(c) Prize in equilibrium, $y(x)$

(d) Expected cost of effort in equilibrium

Figure 3.2: Equilibrium payoffs and effort in the structural model without retaking under assortative matching and bottom pooling.

We start by simulating the structural model without the option to retake. Given the estimated structural parameters, we simulate two equilibria: a status-quo equilibrium taking the estimated distribution of prizes as given, and a counterfactual one imposing the maximal bottom pooling policy found in Section 7. To facilitate comparison with the main results, we focus on variables similar to those depicted in Figure 9. The expected score $t_{i0}(PUB, 0)$ (see equation 25) conditional on choosing public school and no private tutoring is a natural counterpart for ability. The expected placement payoff $E[W(t_{i0}(hs, pt))|(hs, pt)$ solves (26)] is the counterpart of the prize y. The expected cost of schooling $E[c_i(t_{i0}(hs, pt))/b|(hs, pt)$ solves (26)] is the analog of the cost of effort.

Figure 3.2 summarizes the impact of bottom pooling in the structural model without retaking. Qualitatively, the key results from the theoretical model hold. Bottom pooling results in the reduction of effort, especially at the lower end of the ability distribution. Lower effort leads to a reduction in placement scores, which puts most low-ability students in the pooled bottom interval. If one defines utility as the expected payoff with all the unobservable shocks to scores and costs of schooling integrated out, the maximal bottom pooling policy is Pareto-improving as evident from Figure 3.2a. The most notable difference between this set of results and those from the theoretical model in Figure 9 is that high-ability students are not indifferent between the pooling policy and the status quo here.

To make the model even more realistic and reflect the prevalence of retaking in Turkey, we allow for the option to retake the exam in the model and repeat the simulation exercise again. We find the baseline equilibrium and the equilibrium under the maximal bottom pooling policy. Figure 3.3 summarizes the results. All the main conclusions stay the same. Most importantly, we find again that the maximal bottom pooling policy is Pareto-improving.

4 Experiment: additional results

4.1 Main experimental results

Our analyses in the main text focus on the subjects whose choices in the risk elicitation task are consistent with risk neutrality. Risk neutrality allows us to apply the model and its predictions with players' utilities equalling their potential earnings. Reasoning for this choice as well as additional results on non-risk-neutral subjects are in the next subsection of this appendix. The tests and corresponding p-values for comparing data across the discrete and pooling rounds are the result of Wilcoxon matched-pairs signed-rank tests. In reporting aggregate results we use the estimates from Section 7 to determine the appropriate weights

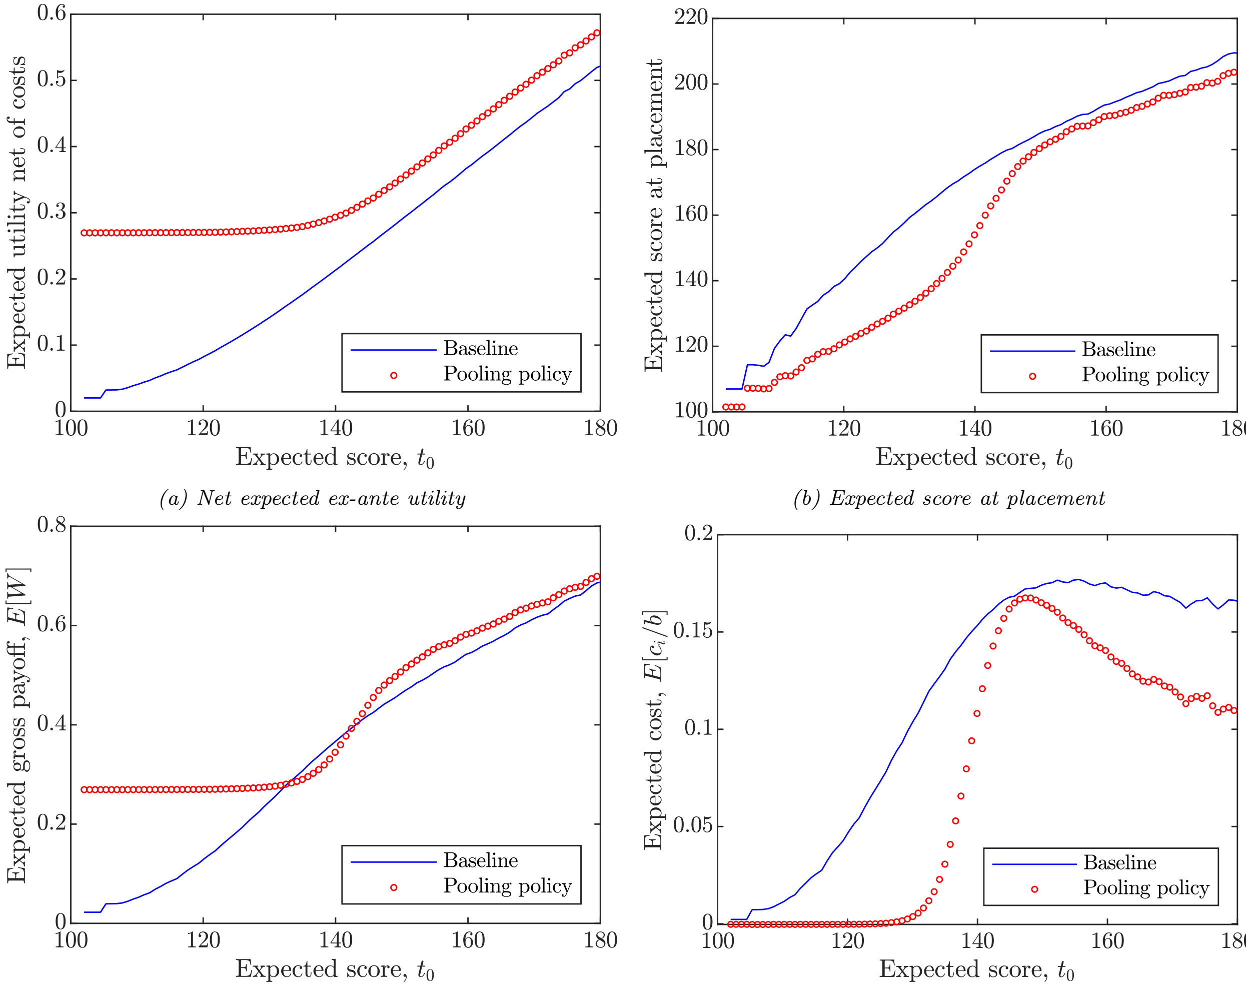

(a) Net expected ex-ante utility

(b) Expected score at placement

(c) Prize in equilibrium, $y(x)$

(d) Expected cost of effort in equilibrium

Figure 3.3: Equilibrium payoffs and effort in the structural model with retaking under assortative matching and bottom pooling.

for each ability level.

4.2 Results on non-risk-neutral subjects

Risk aversion can be determined by identifying at which point a subject switches away from the risky gamble to the fixed amount (84.6 percent 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 percent of the subjects with a single cross-over point behaved in this way. Over 85 percent of the remaining subjects who had a single crossover point from the lottery to the fixed amount were risk averse.

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 $p = 0.3748$).

4.3 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 4.1 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.

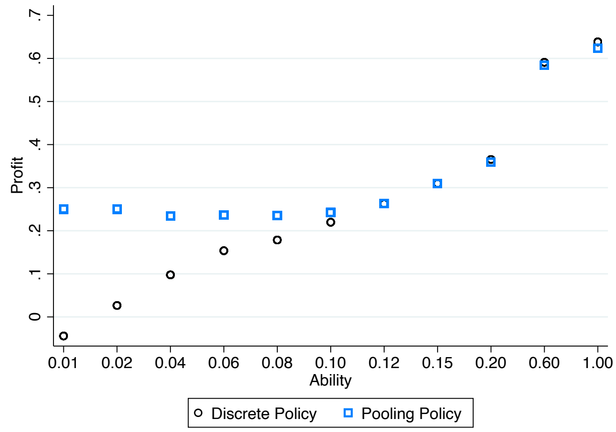

Figure 4.1: Average profits under the discrete and pooling policies.

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.

4.4 Behavior or low and high ability subjects

Theoretically, the low-ability subjects (in our experiment those with ability levels of at most 0.10) are the ones who benefit from the pooling policy, while the high-ability (ability levels of 0.12 and above) ones should be indifferent between the pooling and discreet policies.

Empirically, we indeed find that among low ability subjects, profits increases for each ability level under pooling compared to under the discrete policy (the p-values in all pair-wise comparisons are all strictly lower than 0.001). At the aggregate moving to a pooling policy increased aggregate profits by 70.5 percent, more than the discretized theoretical prediction of 65.7 percent.[61]

For high ability subjects, we instead find that moving to a pooling policy slightly decreases the aggregate profits (about 1.6 percent). We identify that this small drop is the result of the behavior of subjects whose ability is lowest among the high ability subjects (ability of 0.12).

[61] 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.

The behavior of these subjects was statistically different under the two policies ($p = 0.006$), and the magnitude of the difference is large. Indeed, under the discrete policy, 89.5 percent of these subjects chose investment levels consistently with the theoretical predictions of the model. This fraction dropped to 36.8 percent under the pooling policy, with the remaining subjects opting for the lottery. This represents an aggregate profit loss of 5.1 percent. All other high ability subjects behaved similarly across the two policies (the lowest p-values for pair-wise comparisons between the discrete and pooling policies is 0.125).

4.5 Deviations from predictions

Our experimental results show that subjects whose ability levels are the lowest among those who should exert positive effort instead on average exert none and choose the lottery. In this section we explore differetnt possible explanations.

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 highest ability, that is, subjects with an ability of 0.10. 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. 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.

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.

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).

In the main text we argue instead that preferences for randomization may influence these subjects' choices. Preferences for randomization have been explored by 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 percent of individuals choose lotteries between available allocations, indicating an explicit preference for randomization. The authors discuss this in the context of responsibility aversion. 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. Further, we showed above 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.

5 Experiment: procedures and cost-to-school mapping

5.1 Procedures – additional details

The experiment had two parts, each with multiple rounds. Online Appendix 6 contains all the experimental materials subjects faced.

Part 1: the college admissions task. Part 1 consisted of two rounds. Each round corresponded to a college admissions setting, one round with a bottom pooling policy (the "pooling policy" round) and one round without any pooling (the "discrete policy" round). 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 payoffs, costs, and abilities are derived from a discrete version of the calibration and equilibrium results of Section 7.

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 to determine the college in which the subject enrolled, with the associated payment.[62] Figure 10a in the main text presents the expected profits predicted by the theory in both rounds for each ability level in the discretized model. The overall weighted profits, using the appropriate weights for each ability level, are predicted to be 20 percent higher under the pooling policy than under the discrete policy. The predicted increase for low ability subjects is 65.7 percent.

Part 2: risk elicitation. In Part 2 subjects participated in a series of ten rounds in which they had to choose between a fixed amount and a risky gamble (for more details see 6). This task allows us to identify subjects whose choices were consistent with risk neutrality, an important element for the analysis of the pooling-induced lottery.

Procedures. A total of 602 subjects completed the experiment, recruited from the Prolific online platform.

Subject selection. For any particular ability level, we stopped collecting data once we had at least 50 observations. An additional 208 “participants” started the experiment but failed to advance because quiz questions were answered incorrectly. This proportion (25.7 percent) 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. We restricted our subjects to

[62] To isolate the possible effects of the lottery per se, 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. The results from these sessions are in Online Appendix 4.

be English speakers. Overall 48.8 percent of the subjects were female, 47.7 percent were male, and 2.5 percent identified as neither of the two. The remaining 1 percent of the subjects preferred not to answer. Data were collected during the month of March 2022. The experiment lasted between 10 and 15 minutes for 95 percent of the subjects. Subjects were paid an average of \$4.00, corresponding to a rate of just over \$19 per hour.[63]

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. 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 percent 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.

5.2 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 5.1 shows how the ability cutoffs (5.1a) and net payoffs (5.1b) 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:

[63]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.

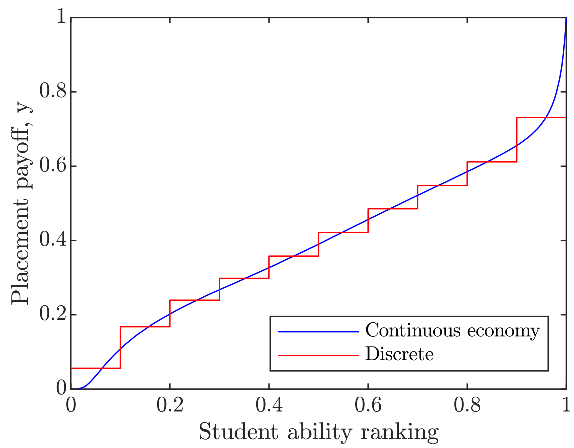

(a) Mapping from student's ability ranking to prize.

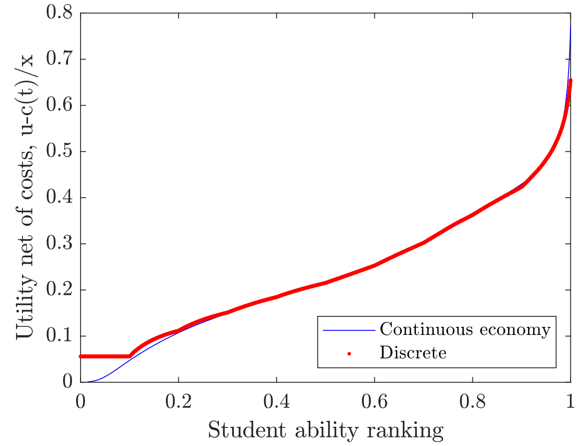

(b) Student net payoffs as a function of ability ranking in the no-pooling scenario.

Figure 5.1: Discretized vs. continuous economy.

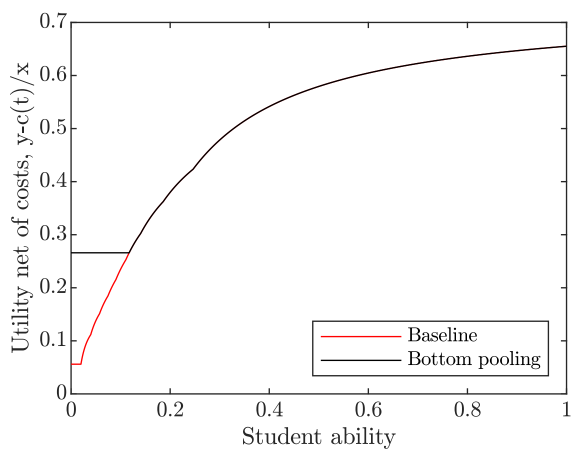

(a) Net payoffs by ability

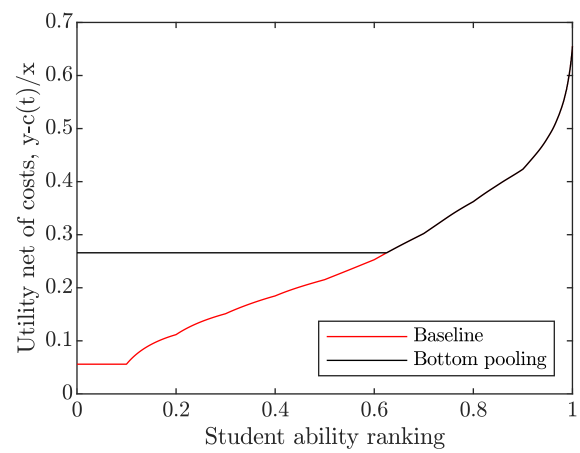

(b) Net payoffs by ability ranking

Figure 5.2: Student net payoffs as a function of ability in the discretized economy

roughly the same set of students get seats in the pooled school under both policies, and their placement payoffs are almost identical. Figure 5.2 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 9 in the main text). Student net payoffs in the discretized economy under the discrete policy and the pooling policy and displayed as a function of their ability (5.2a), and a function of their ability ranking (5.2b).

5.3 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).

		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

Table 5.1: Discrete policy: cost of targeting particular schools by ability.

Table 5.1 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. 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 5.2 shows the parameters that were used under the pooling policy. The parameters in this table are those the subjects saw on their screens.[64] Highlighted in yellow are the profit maximizing choices. For example, a subject assigned to ability level 0.06 who invested 0 Tokens out of the

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

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.

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.

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 H G F E	LOTTERY Equal Chances of 3, 16, 24, 30, 35, 42						0					
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

Table 5.2: Cost of targeting particular schools by ability under the pooling policy.

6 Experiment: screen shots

6.1 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 6.1: Consent screens

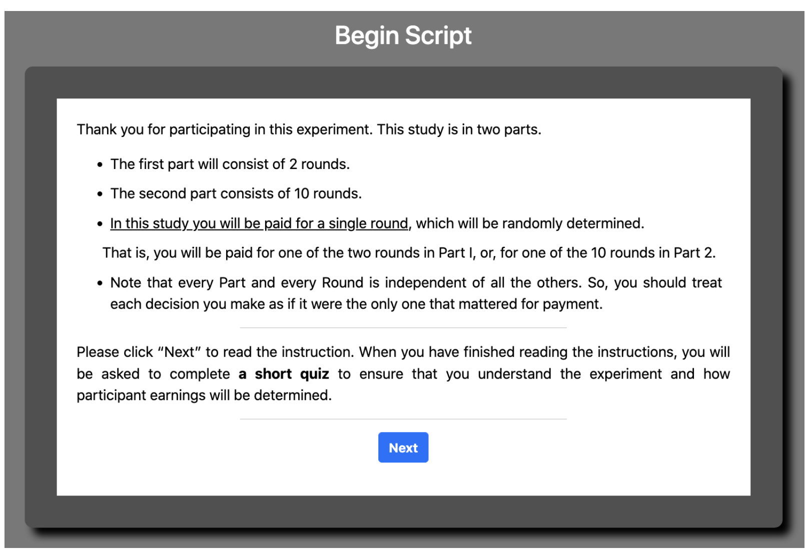

Figure 6.2: Preamble

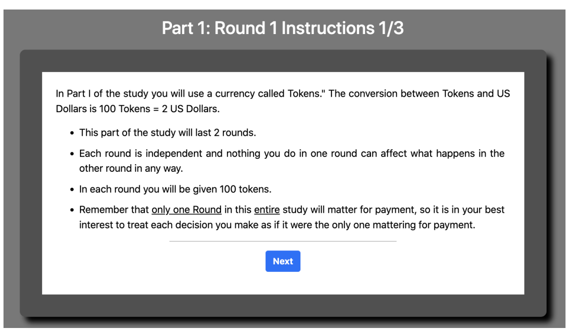

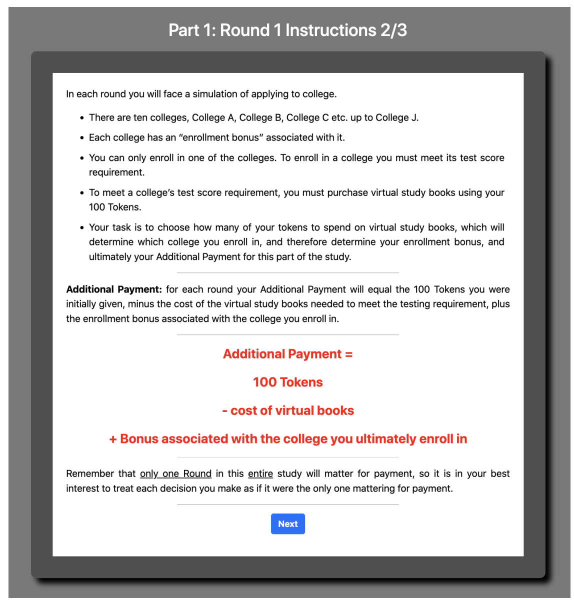

Figure 6.3: Part 1 General instructions

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

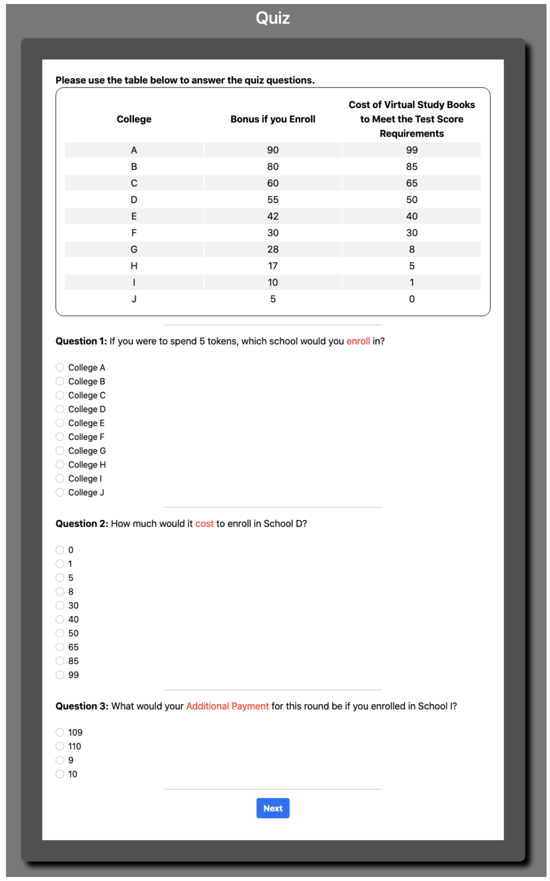

Quiz

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 6.5: Quiz after discrete round instructions.

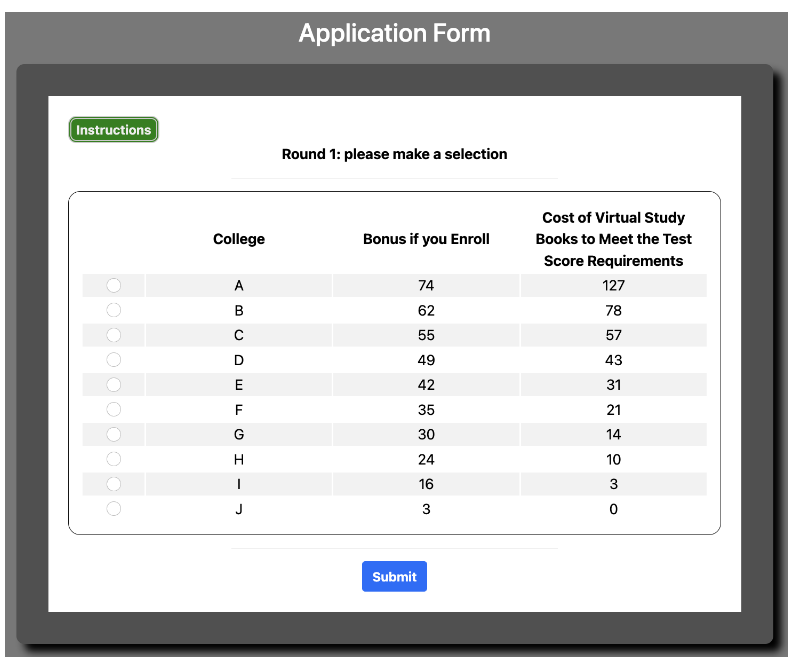

Application Form

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 6.6: Example of selection screen for discrete policy.

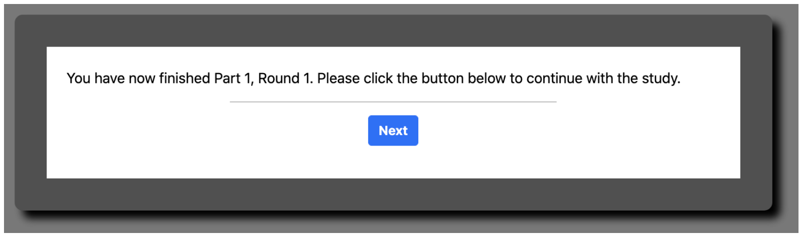

You have now finished Part 1, Round 1. Please click the button below to continue with the study.

Next

Figure 6.7: Transition to Round 2.

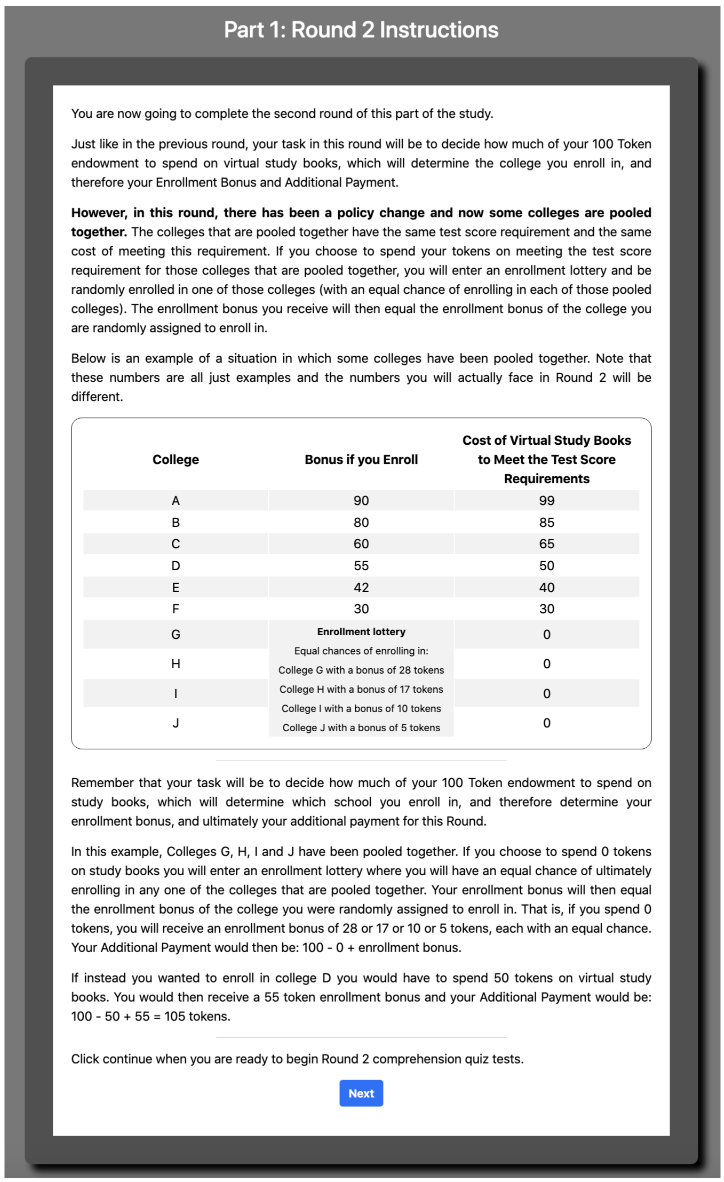

Part 1: Round 2 Instructions

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

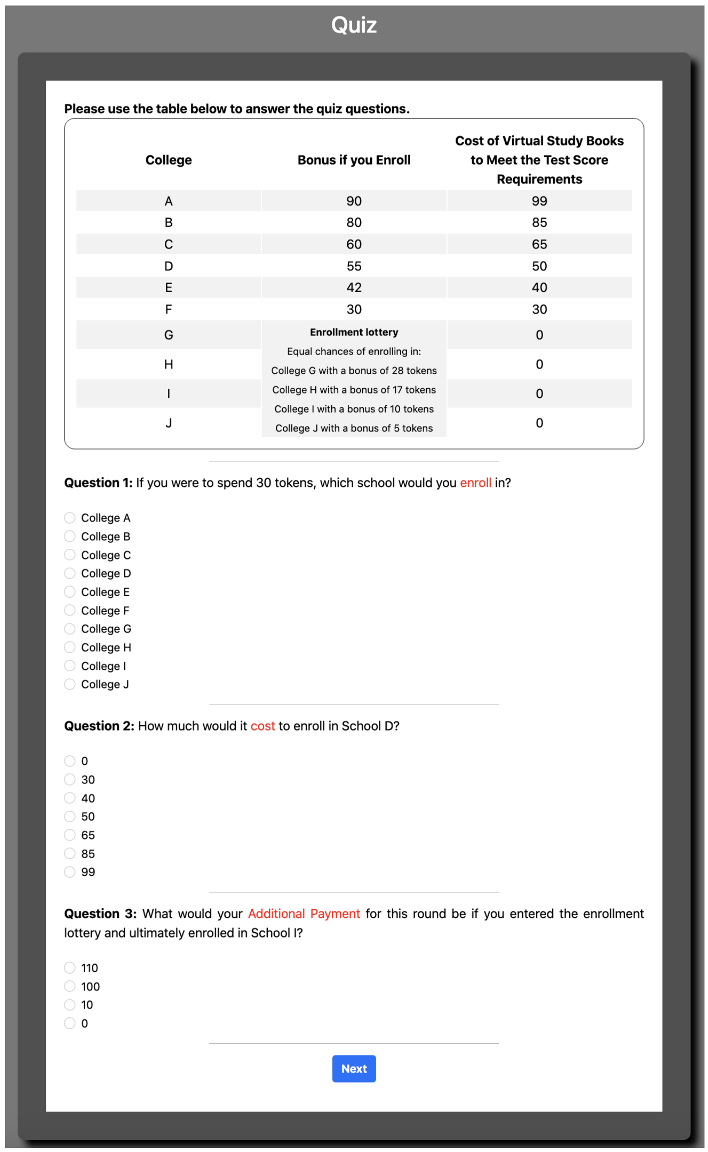

Quiz

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 6.9: Quiz after pooling round instructions.

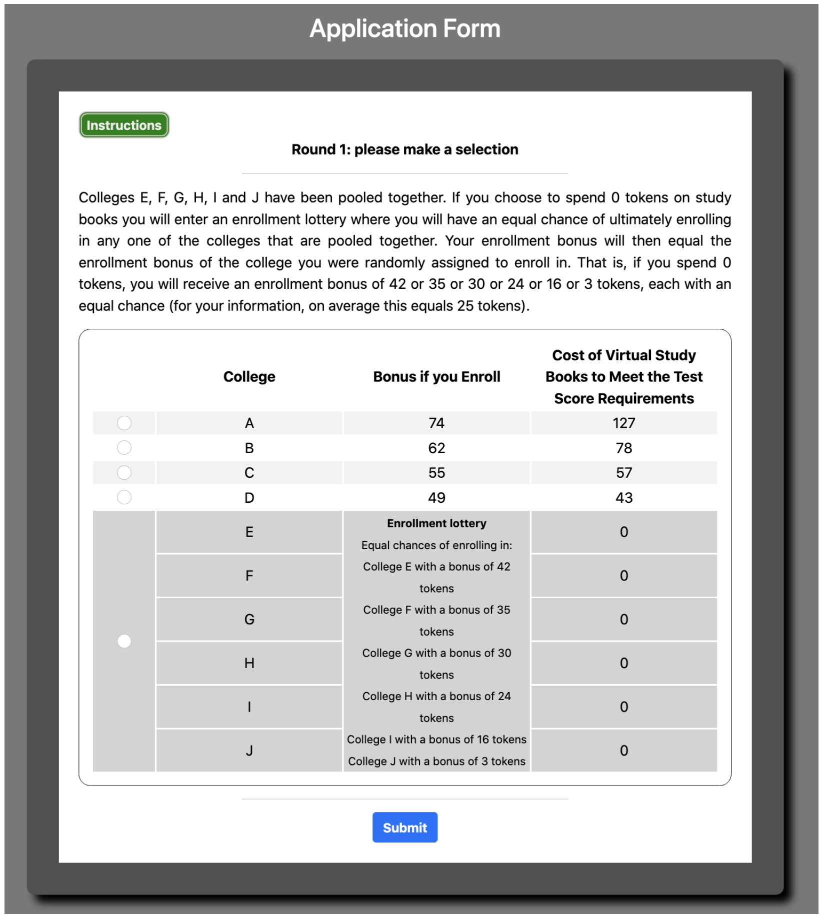

Application Form

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).

<table>
<tr><th></th><th>College</th><th>Bonus if you Enroll</th><th>Cost of Virtual Study Books to Meet the Test Score Requirements</th></tr>
<tr><td></td><td>A</td><td>74</td><td>127</td></tr>
<tr><td></td><td>B</td><td>62</td><td>78</td></tr>
<tr><td></td><td>C</td><td>55</td><td>57</td></tr>
<tr><td></td><td>D</td><td>49</td><td>43</td></tr>
<tr><td rowspan="6"></td><td>E</td><td rowspan="6">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</td><td>0</td></tr>
<tr><td>F</td><td>0</td></tr>
<tr><td>G</td><td>0</td></tr>
<tr><td>H</td><td>0</td></tr>
<tr><td>I</td><td>0</td></tr>
<tr><td>J</td><td>0</td></tr>
</table>

Submit

Figure 6.10: Example of selection screen for pooling policy.

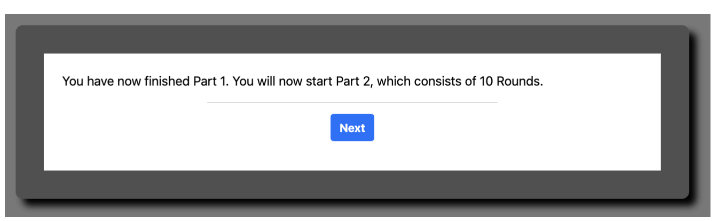

You have now finished Part 1. You will now start Part 2, which consists of 10 Rounds.

Next

Figure 6.11: Transition from Part 1 to Part 2.

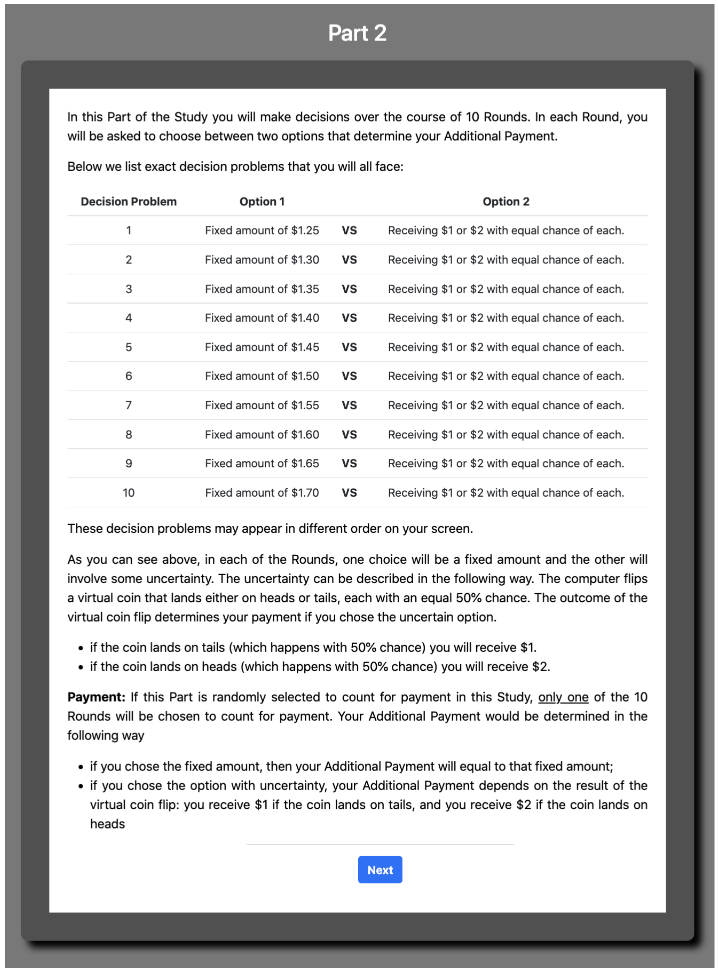

Part 2

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 6.12: Instructions for Part 2 (risk elicitation).

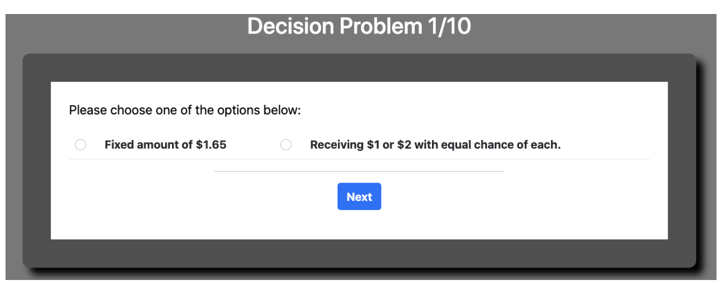

Decision Problem 1/10

Please choose one of the options below:

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

Next

Figure 6.13: Example of decision screen for Part 2.

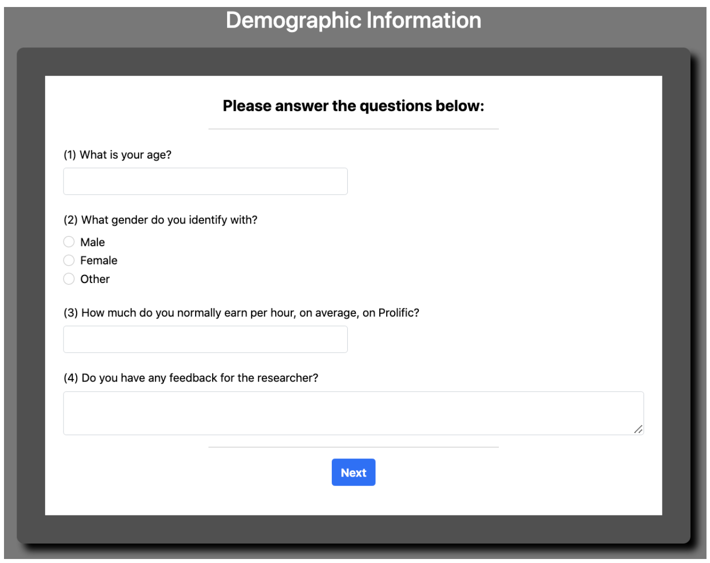

Figure 6.14: Questionnaire.

6.2 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 6.4, 6.5, 6.6, 6.8, 6.9 and 6.10 are replaced, in order, with the figures below (Figures 6.15, 6.16, 6.17, 6.18, 6.19 and 6.20, respectively).

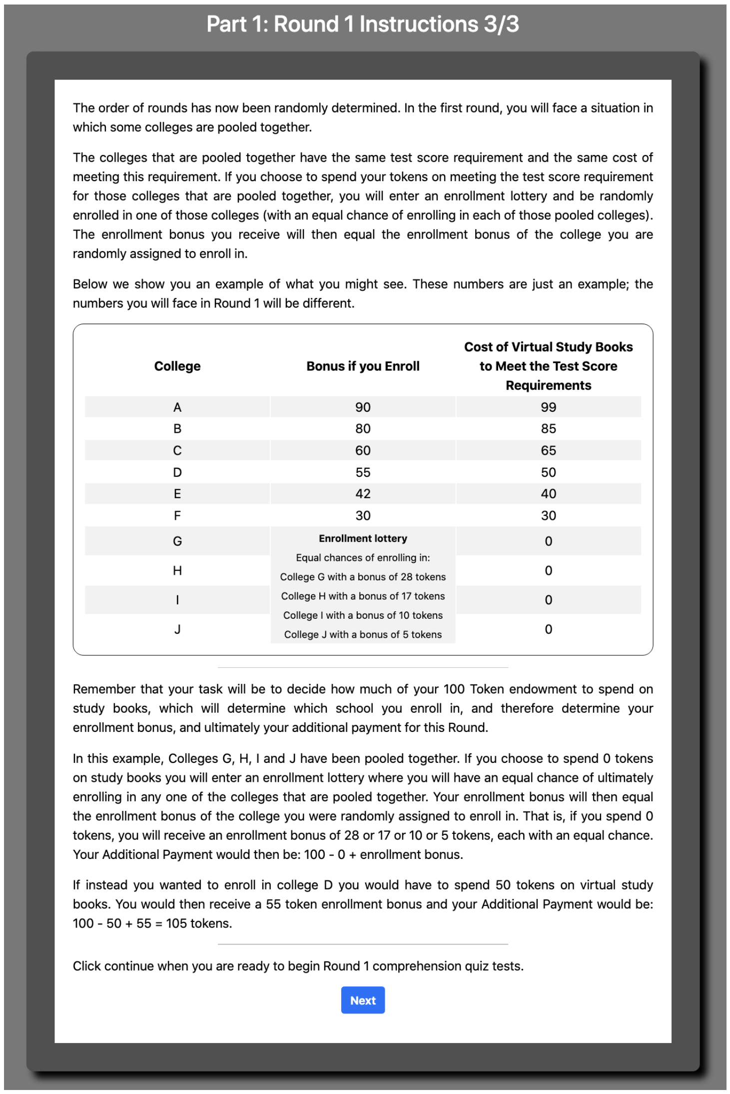

Part 1: Round 1 Instructions 3/3

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:	0
H	College G with a bonus of 28 tokens	0
I	College H with a bonus of 17 tokens College I with a bonus of 10 tokens	0
J	College J with a bonus of 5 tokens	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 tokens 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 6.15: Instructions for Round 1 if pooling policy is seen first.

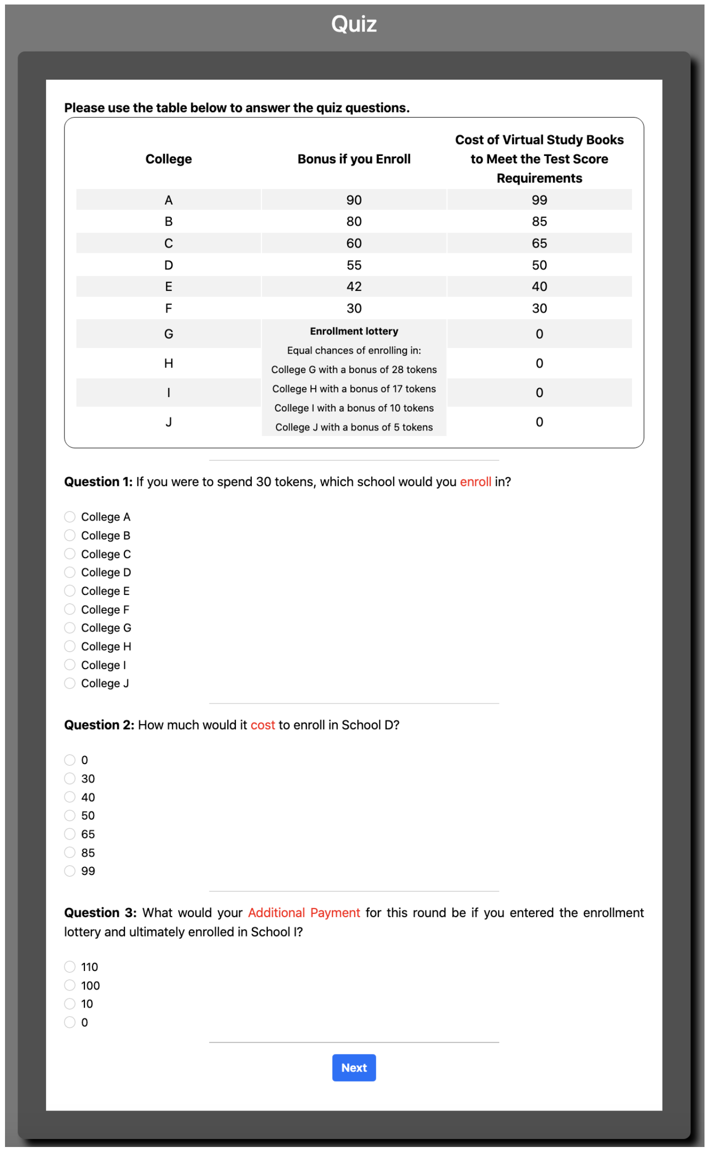

Figure 6.16: Quiz after pooling round instructions.

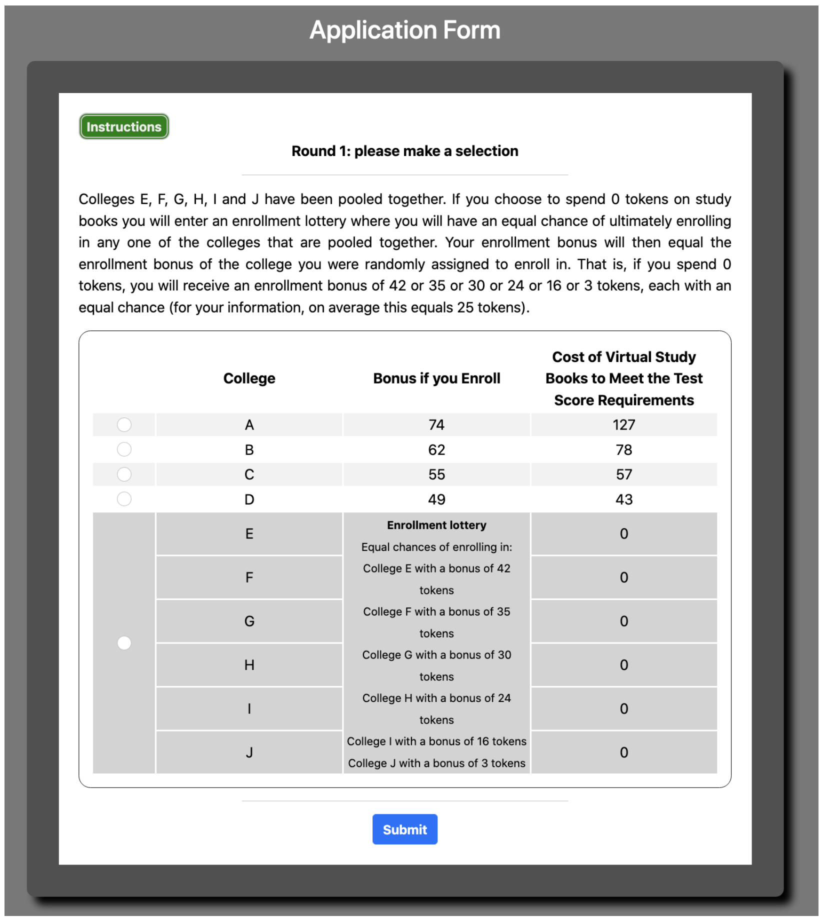

Application Form

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 6.17: 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 6.18: Instructions for the discrete policy.

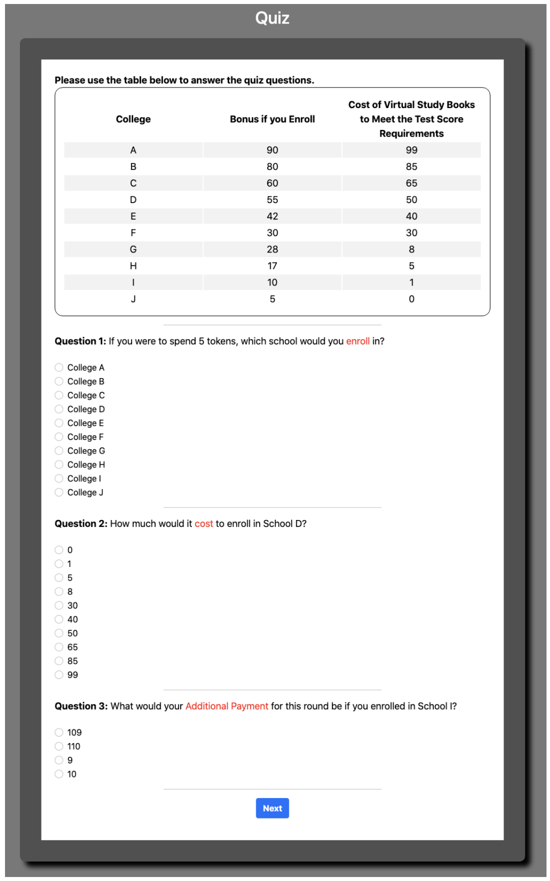
Quiz

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 6.19: Quiz after discrete round instructions.

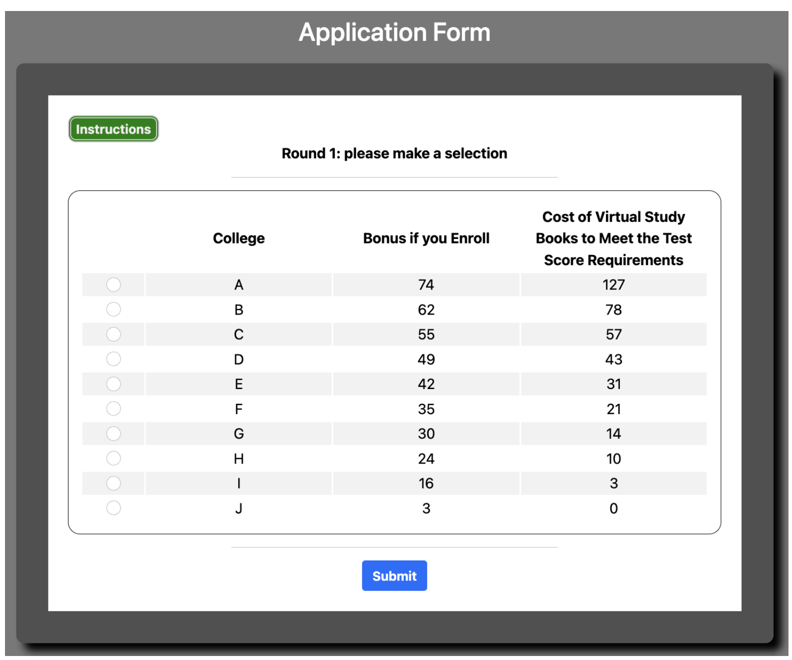
Application Form

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 6.20: Example of selection screen for discrete policy.