Tournaments for Talent: Competitive Selection and Right-Tail Human Capital

Ian Calaway Delong Meng Farrell Wu

July 7, 2026

Abstract

We develop a class of developmental tournament models in which a competition both selects talent and produces it. Participants of varying talent exert costly effort to build skill, an organizer awards winner status on the basis of a noisy assessment, and winning carries two benefits: a direct prize and a productive augmentation that scales with the winner's own skill. Because effort builds skill, the organizer's rules shape not only who is recognized but the entire post-competition skill distribution. We show that effort is hump-shaped in talent, so that middle-ability participants work hardest while the most talented coast and the least talented give up. We also show that assessment noise, though costly for selection, can raise aggregate skill production by keeping more participants in contention. We then test the model's central predictions using two decades of novel data from Texas University Interscholastic League academic competitions, exploiting cross-event differences in assessment precision and two institutional reforms to identify how organizers' design choices reshape the distribution of skill.

We thank Zhen Fang, Jingsai Ni, and Sitta Tantikul for excellent research assistance. The authors have no finanical disclosure or conflict of interest.

1 Introduction

Formally organized competitions are among humanity's oldest institutions, spanning ancient athletic games to modern STEM Olympiads. Although the skills at the center of these competitions vary widely, these institutions often share a common structure: an organizer decides on a set of rules, prizes, and evaluation criteria, while participants of varying talent decide how much effort to exert in skill development. By setting up a competition and evaluating and rewarding skill, an organizer can shape the skill distribution, influencing the post-competition abilities of both winners and losers. Optimal competition design, therefore, depends on the utility an organizer derives from different parts of the skill distribution.

In this paper we develop a class of *developmental tournament models* that captures the interplay between organizers and competitors. Participants differ in latent talent, choose costly effort, and produce private skill that is the product of talent and effort. The organizer awards winner status on the basis of a noisy assessment of that skill, and winning results in two potential prizes: a direct prize that does not depend on skill, and a productive augmentation whose value scales with the winner's own skill. A developmental tournament is therefore both a selection device and a skill-formation device: it changes who is recognized, how much effort is exerted, how much skill is produced, and where along the talent distribution that excess skill is realized.

The framework makes the organizer's design problem explicit through three levers: the *advancement threshold* (how high measured skill must be to win, and thus how many are recognized), the *award structure* (the size and type of prize, fixed or skill-augmenting), and the *sharpness* of the assessment (how precisely the contest translates true skill into observed performance).

Our first result concerns how effort is distributed across the talent spectrum. Although the skill a participant produces is everywhere increasing in talent, effort is not: for an intermediate level of sharpness, effort first rises and then falls in ability, generating a hump-shaped profile in which middle-ability participants exert the most effort. The most talented win so easily that they coast, the least talented have too little chance of winning to do more than baseline effort, and the participant near the threshold must work hardest of all to secure a win. In the deterministic limit, the contest collapses to a cutoff rule and effort concentrates at the threshold.

Our second result concerns the organizer's levers, and shows that the relationship between assessment accuracy and skill production is not what a pure selection logic would suggest. From a selection standpoint, noise is unambiguously costly: a noisier contest more often recognizes the wrong participant. But noise also keeps a wider range of participants in contention, and participants who remain in contention continue to exert effort and produce skill. An organizer who cares about the skill of the whole field, rather than only about identifying the top, therefore faces a genuine tradeoff in choosing how sharply to measure.

To test the model predictions, we use novel data from more than two decades of academic competitions organized by the University Interscholastic League (UIL) in Texas. These data include 628,110 event-year level records for 299,157 students across 1,630 schools and eleven events, from 2004 to 2026. The setting maps cleanly onto the model. Because a fixed number of participants

advance from each district, the advancement threshold is set endogenously by the local distribution of talent, and districts vary widely in how strong that distribution is. This generates notable variation in the threshold that participants face. The UIL also administers many distinct competitions under similar advancement rules but with very different measurement instruments, generating wide variation in sharpness while holding tournament structure fixed. We measure event sharpness in two ways. The first is the same-year, within-student correlation between advancing students' district and regional performances. The second is the between-year, within-student correlation between returning students' two appearances in the competition. These produce nearly identical event rankings (a correlation of 0.97).

Two institutional changes provide exogenous variation in the model's key parameters. In 2015 the UIL increased the number of conferences, reshuffling which schools competed against one another and thereby raising the threshold for schools moved into stronger pools and lowering it for others. In 2023 the UIL raised the number of entrants a school could field per event from four to six, so that for school-level events—where the team score aggregates individual entries—drawing from more students reduces the variance of measured performance and sharpens the contest. We plan to leverage these sources of variation to test whether school and student skill responses align with our model's predictions.

Our work contributes to the literature on talent selection. Rosen (1981) argues that even a small difference in talent could result in a large difference in income, because the reward function is convex in quality. In contrast our model examines how effort and contest sharpness shape the final skill distribution. In Holmstrom (1999) the agent produces an outcome that depends on her intrinsic ability, her effort, and a noise. A key difference with our model is that in Holmstrom (1999) the ability is fixed, whereas in our setting the effort increases the agent's ability (e.g. working on math problems makes the student better at math). Lazear and Rosen (1981) show that a rank-order tournament could induce an efficient outcome even if the outcome is noisy. In their setting the principal motivates effort through a reward (a transfer), whereas in our setting the reward also includes skill augmentation. Nalebuff and Stiglitz (1983) argue that in the presence of a common shock, the principal should pay the workers based on relative performance instead of their absolute performance. In our model the contest sharpness serves as a common shock, and it is part of the contest designer's problem.

Our paper also relates to the growing literature on contest design. Moldovanu and Sela (2001) investigate how a contest designer should allocate a fixed budget (e.g. first prize, second prize, and so on). Our model only has one prize tier (winners vs losers), but winners get different rewards because skill augmentation depends on the initial talent. Moldovanu and Sela (2006) study the trade-off between having a single grand contest vs dividing the participants into sub-contests. Our model has only one contest, but the designer chooses the contest sharpness. Olszewski and Siegel (2016) analyze a large contest model with heterogeneous agents and prizes, and they reduce the problem to a single-agent mechanism with assortative prize allocation. In a follow-up paper Olszewski and Siegel (2020) solve for the performance maximizing large contest and show that

the prize size is gradually decreasing. In our setting the contest has a single cutoff score, but the rewards still differ due to skill augmentation.

Finally, our paper expands on empirical studies related to childhood competitions, both athletic and academic. Ashby (2023) finds improvement in high school swimmers after assignment to more competitive environments, which mirrors our skill threshold analysis. In contrast, Bottan and Bernhardt (2024) show that when swimmers compete against higher age students (as opposed to those of the same age) their performance suffers. Our study offers an explanation that reconciles these seemingly contradictory results. On incentives, our model is consistent with results from a field experiment that gave 5th-8th graders in Utah a lump-sum incentive to study for the AMC 8 Cotton et al. (2022), which yielded broad performance gains. Ellison and Swanson (2010) and Ellison and Swanson (2016) find interesting demographic and geographic patterns in the participation and success of students and schools in the AMC 12. In future versions of this work, we hope to better understand the roles through which competition structure and parameters might influence gender gaps and the prevalence of "star" schools found in these studies. Our data and empirical strategy also enable us to speak on the impacts specific to right-tail talent, which has been shown to be a key input in scientific advancement and innovation by Agarwal and Gaulé (2020) and Calaway (2024).

2 Model

This section sets up a tournament model with heterogeneous talent, costly effort, multiplicative skill production, and noisy recognition. A tournament is summarized by a skill cutoff s_c, a sharpness parameter ρ, and a vector of winner-benefit components $\boldsymbol{\beta} = (\beta_0, \beta_1)$. These components allow winner status to generate payoffs that are fixed or scale with the winner's own skill.

2.1 Talent, Effort, and Skill Production

We model talent as following a **Pareto Type-II distribution** with location parameter $\mu = 0$ and scale normalized to one, i.e. a Lomax distribution:

$$F(\theta) = 1 - (1+\theta)^{-\gamma}, \qquad f(\theta) = \gamma(1+\theta)^{-(\gamma+1)}, \qquad \theta \geq 0.$$

The shape parameter $\gamma > 2$ governs the thickness of the right tail. A lower γ implies a fatter tail and greater mass of high-ability participants, while a higher γ concentrates the population closer to zero. We impose $\gamma > 2$ so total skill is integrable even in the no-tournament case.

A continuum of participants is characterized by a scalar *talent* parameter $\theta \in \mathbb{R}_+$. Each participant chooses effort $e \geq 0$ at cost $C(e)$. We assume C is twice differentiable and satisfies

$$C : \mathbb{R}_+ \to \mathbb{R}_+, \qquad C(0) = 0, \qquad C'(e) > 0 \text{ for } e > 0, \qquad C''(e) \geq 0 \text{ for all } e \geq 0.$$

Thus effort is costly, marginal cost is positive away from zero, and marginal cost is weakly increasing.

In the absence of any tournament benefit, a type-θ participant who exerts effort e produces *skill* $s(\theta, e) = \theta \cdot e$. Skill is thus the product of innate talent and chosen effort. This captures the complementarity between ability and investment: the same unit of effort generates more skill for a higher-talent participant. Skill is rewarded in the labor market at wage rate $r_w \geq 0$, so baseline earnings are $r_w\, s(\theta, e)$.

For the analysis below, we focus on the quadratic-cost case $C(e) = e^2/2$. In the no-tournament benchmark, a type-θ participant solves $\max_{e \geq 0} \{r_w \theta e - e^2/2\}$, so effort is $e^0(\theta) = r_w \theta$ and skill is $s^0(\theta) = r_w \theta^2$.

2.2 Tournament Structure and Noise Specification

For this paper we consider a tournament with one recognition level: a participant either receives winner status or does not. The organizer sets a skill cutoff $s_c > 0$ and a contest sharpness parameter $\rho > 0$. A participant with skill s wins with probability

$$q(s; s_c, \rho) = \pi_\rho \left(\frac{s}{s_c} \right),$$

where $\pi_\rho(\cdot)$ is increasing and differentiable. The ratio specification reflects the multiplicative nature of skill: proportional skill differences, rather than additive differences, determine distance from the cutoff. This also pairs naturally with the right-skewed type distribution introduced above.

A leading example is the **binary Tullock tournament**, in which

$$q(s; s_c, \rho) = \frac{s^\rho}{s^\rho + s_c^\rho} = \frac{(s/s_c)^\rho}{1 + (s/s_c)^\rho}.$$

The cutoff s_c is the point at which the probability of winning equals one half. The sharpness parameter ρ governs how sensitively winning probabilities respond to skill differences. A higher ρ makes the tournament more discriminating: participants above the cutoff are more likely to win and participants below it are less likely to win. A lower ρ compresses winning probabilities and represents a noisier evaluation process. A selection process based on one test rather than multiple independent assessments can therefore be represented as having a lower ρ.

Remark 1 (Log-skill interpretation). For $s > 0$, the binary Tullock form is equivalent to a noisy cutoff rule in log skill. Suppose the organizer observes $\log s + \varepsilon$ and awards winner status when this value exceeds $\log s_c$. If ε has a logistic distribution with scale $1/\rho$, then

$$\Pr(\log s + \varepsilon \geq \log s_c) = \frac{1}{1 + \exp\{\rho(\log s_c - \log s)\}} = \frac{s^\rho}{s^\rho + s_c^\rho}.$$

Thus ρ is an inverse noise parameter in log skill.

Under the binary Tullock success function the marginal winning probability is

$$\frac{\partial q(s, s_c; \rho)}{\partial s} = \frac{\rho}{s} \cdot q(1 - q).$$

(We use q as a shorthand for $q(s, s_c; \rho)$). The term $q(1-q)$ plays a central role in the tournament's incentive effects. It is largest when $q = 1/2$, which occurs when $s = s_c$. Thus, the tournament creates the strongest marginal incentives near the winning cutoff. Far below the cutoff, the probability of winning is small, so effort has little effect on expected prize receipt. Far above the cutoff, the probability of winning is already close to one, so additional effort again has little effect.

2.3 Tournament Benefits and Participant Payoffs

Winner status can provide benefits through two payoff forms:

$$B_W(s; \boldsymbol{\beta}) = \underbrace{\beta_0}_{\text{direct prize}} + \underbrace{\beta_1 r_w s}_{\text{skill augmentation}}, \qquad \boldsymbol{\beta} \in \mathbb{R}^2_+.$$

This decomposition is categorical: winner rewards may be fixed or scale with the winner's own skill. Because these forms enter the effort problem differently, they imply different effort choices and therefore different resulting skill.

First, β_0 is a *direct prize*: a lump-sum reward that does not scale with skill. This category includes cash prizes, scholarships, in-kind awards, or fixed human-capital benefits such as a single course enrollment.

Second, β_1 is an *augmenting prize*: an augmentation or access benefit whose value scales with the participant's own skill. This includes uncapped mentorship, advanced coursework, internships at prominent labs or firms, and access to higher-return markets such as a higher-paying labor market. The last example can be interpreted through standard capital-skill complementarity.

The expected value of a type-θ participant choosing effort e is

$$\begin{aligned} v(\theta, e; s_c, \boldsymbol{\beta}, \rho) &= r_w\, s(\theta, e) + q(s(\theta, e); s_c, \rho) B_W(s(\theta, e); \boldsymbol{\beta}) - C(e) \\ &= \underbrace{r_w\, s}_{\substack{\text{baseline} \\ \text{skill value}}} + q\Big[\underbrace{\beta_0}_{\substack{\text{direct} \\ \text{prize value}}} + \underbrace{\beta_1 r_w\, s}_{\substack{\text{augmenting} \\ \text{prize value}}} \Big] - C(e), \end{aligned}$$

where $s = s(\theta, e)$, and the shorthand q in the second line means $q(s(\theta, e); s_c, \rho)$.

The next section uses this notation to study the participant's effort choice and comparative statics.

3 The Participant's Problem and Comparative Statics

In this section we assume that $C(e) = \frac{e^2}{2}$ and $\rho > 1$. The first assumption simplifies the marginal cost function, and the second assumption makes the marginal benefit function bounded.

The Choice of Skill We analyze the student's choice of skill. Let $s(\theta)$ denote the chosen skill for the student with talent θ. We have

$$s(\theta) = \arg\max_{s} \; r_w s + q \cdot [\beta_0 + \beta_1 r_w s] - \frac{s^2}{2\theta^2}.$$

If the student is indifferent between multiple skills, she chooses the largest one.

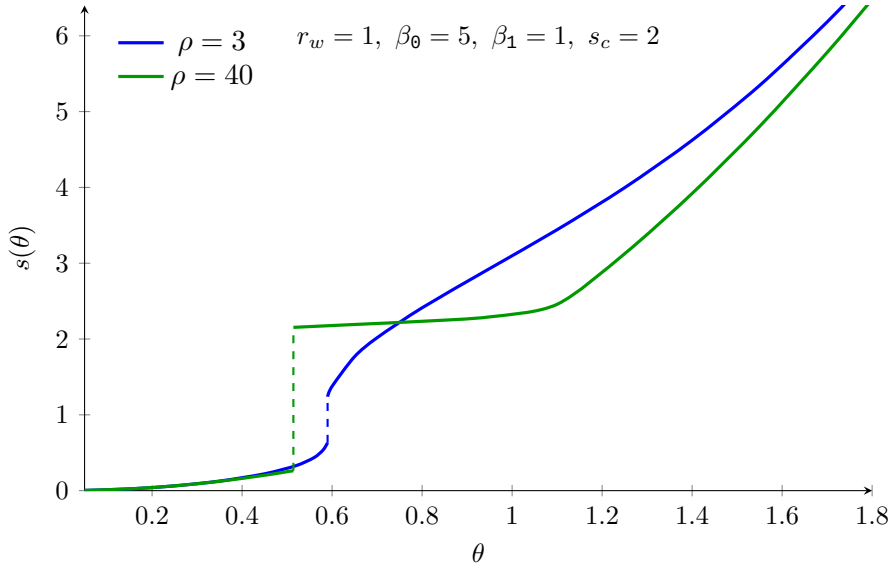


Figure 1: The Student's Choice of Skill

Figure 1 plots two sample skill functions. These two graphs offer three observations. First students with higher talent choose higher skill. Second the two graphs seem to converge as talent goes to 0 or infinity. Finally both graphs have exactly one jump. The next result formalizes these observations.

Theorem 1. *The skill function has the following properties.*

1. *We have $s(\theta)$ is non-decreasing in θ.*
2. *As $\theta \to 0$ we have $s(\theta) \approx r_w\theta^2$, and as $\theta \to 1$ we have $s(\theta) \approx (1+\beta_1)r_w\theta^2$.*
3. *The function $s(\theta)$ has at most one point of discontinuity, and it has exactly one jump if ρ is sufficiently large.*

Intuitively if the student's talent θ is close to 0, she makes the baseline effort $r_w\theta$, so her skill becomes $r_w\theta^2$. On the other hand if the student's talent θ is sufficiently large, she is almost assured of winning the competition, so she makes effort $r_w(1+\beta_1)\theta$ (the optimal effort when the probability of winning is equal to 1).

For intermediate talent levels there is potentially a jump in the effort, because the student either chooses a skill close to the cutoff s_c (fighting hard to beat the cutoff), or simply chooses the

baseline effort (giving up because the chance of winning is low). This jump occurs especially when the contest is sharp (ρ is large). In this case the marginal benefit of skill increases dramatically as the skill gets close to the cutoff s_c. In fact as Figure 1 demonstrates when ρ is large, the skill choice after the jump is nearly flat; these students choose the skill close to s_c, and they are almost guaranteed to win.

It is worth noting that although the skill choice $s(\theta)$ is monotonic in talent θ, the effort is not monotonic. In particular when ρ is large, Figure 1 shows that $s(\theta)$ is flat after the jump, so the effort actually decreases within this interval. Indeed for students in this interval the choice of skill is nearly constant; they all choose skill levels close to the cutoff s_c, so students with higher talent actually exert less effort. In other words students right on the cutoff must work extra hard to win the competition.

Contest Sharpness We now analyze the comparative statics with respect to ρ. How does the function $s(\theta)$ change as we increase ρ? Figure 1 shows that when ρ increases, the function $s(\theta)$ decreases when θ is close to 0 and when θ is sufficiently large. In other words students with low and high talents exert less effort as the contest gets more deterministic. The next result formalizes this observation.

Proposition 1. *As ρ increases we have $s(\theta)$ decreases for sufficiently small or sufficiently large θ.*

Intuitively as ρ increases the outcome of the contest is less random. In the extreme case of $\rho = \infty$ the contest is deterministic; students with skill $s > s_c$ are guaranteed to win. As a result for large ρ the marginal benefit of skill acquisition is concentrated at the cutoff s_c. In particular the marginal benefit decreases for skills away from the cutoff, while the marginal benefit increases near the cutoff. As the contest outcome gets more deterministic students with high talent (large θ) put in less effort because they are assured of winning. On the other hand students with low talent (small θ) also put in less effort because they are assured of losing.

Skill Cutoff We also study the comparative statics with respect to the skill cutoff s_c. If we increase s_c, what happens to the skill function $s(\theta)$?

In Figure 2 as we increase the cutoff s_c the chosen skill $s(\theta)$ decreases for small θ and increases for large θ. In other words students with low talent decrease their effort, while students with high talent increase their effort. The next result formalizes this observation.

Proposition 2. *As s_c increases we have $s(\theta)$ decreases for sufficiently small θ and increases for sufficiently large θ.*

Intuitively when the cutoff skill increases, the marginal benefit shifts to the right. In other words for students with low talent the marginal benefit of acquiring skill decreases, because their chance of winning the contest is smaller. However for students with high talent the marginal benefit increases, because they already have a high chance of winning, and the extra effort helps them secure a win. Therefore increasing the cutoff motivates high talent students to work harder, but discourages low talent students from exerting effort.

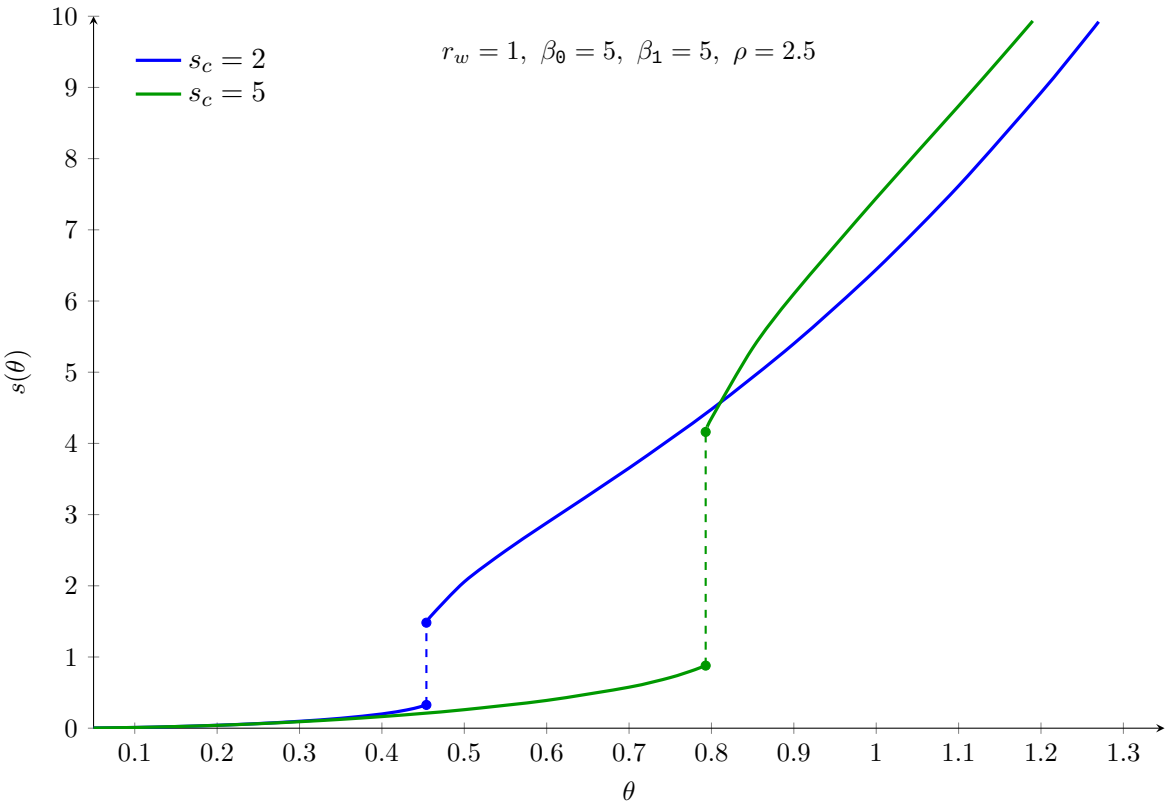


Figure 2: Changing the Cutoff Skill

4 Empirical Evidence

4.1 University Interscholastic League Competitions

The University Interscholastic League (UIL) organizes academic, athletic, and fine arts competitions for Texas public school students in grades K–12. The academic competitions span a diverse range of disciplines, including mathematics, science, social studies, language arts, and computer science. For the purposes of this paper, we focus on the subset of high school UIL academic events that produce continuous individual scores: Calculator Applications (CAL), Mathematics (MTH), Number Sense (NUM), Science (SCI), Computer Science (CSC), Accounting (ACC), Current Issues and Events (CIE), Computer Applications (COM), Literary Criticism (LIT), Social Studies (SOC), and Spelling and Vocabulary (SPV). Together, these events cover a wide range of skill domains, from rapid numerical computation to knowledge of history, and share a common tournament structure that maps cleanly onto the model. A description of each of these competitions is provided in Table 1.

UIL competitions are organized hierarchically. Students first compete at the district level, where the top performers advance to regional competition, with further advancement possible to the state level. The number of students who advance from a given district is fixed, so the winner share α and advancement threshold s_c are determined endogenously by the distribution of scores within each district. We focus primarily on the district competition, which represents each student's first

encounter with the tournament mechanism and where the model's predictions about effort and skill production are most directly testable. This level also offers the most complete coverage: every participating student competes at the district level, whereas the regional and state samples are selected by construction.

Our data consist of individual student-level results covering all Texas districts and UIL conferences from 2004 to 2026.[12] The dataset contains 628,110 individual competition records across 299,157 unique students, 1,630 schools, and 6,336 district-event-year cells. We supplement this individual-level data with school-level team competition records, which aggregate individual scores to the school level and are used in Section 4.4.

Raw scores are not comparable across events or years due to differences in how performance is assessed and year-to-year competition difficulty. We therefore standardize scores within each event-year cell: $z_{it} = \dfrac{\text{score}_{it} - \mu_t}{\sigma_t}$

where μ_t and σ_t denote the mean and standard deviation of scores in competition cell c in year t. This normalization places all students on a common scale and allows direct comparison of the skill distributions and advancement thresholds across events and over time. Figure 5 displays the distribution of standardized scores for each event at the district level, pooling across all years and conferences.

4.2 Measuring Parameters of Interest

The model features two key design parameters: the skill threshold s_c and the contest sharpness σ. Both are recoverable from the data.

Skill Threshold. For each district-level competition cell, we define the empirical skill threshold as the standardized score of the lowest-performing student who advances to the regional competition:

$$\hat{s}_c = \min_{i2A_c} z_{ict}, \tag{1}$$

where A_c denotes the set of advancers in cell c. This construction reflects the model directly: s_c is the score an organizer sets as the minimum to receive the prize, and empirically the threshold is revealed by the marginal advancer. Because the number of advancement slots is fixed by UIL rules, districts with greater participation mechanically face a lower winner share α — a feature of the data that motivates the analysis in Section 4.3. Figure 6 displays the distribution of the implied $\hat{s}_c$ across district-year-event cells. The substantial variation in $\hat{s}_c$ reflects differences in both district competitiveness and participation rates, and provides the identifying variation we exploit in Section 4.3.

[1] UIL organizes schools into conferences based on enrollment, ranging from 1A (smallest) to 6A (largest). This enrollment-based classification is relevant to our analysis of the 2015 conference restructuring in Section 4.3.

[2] This excludes 2020 due to the cancellation of UIL events during the COVID-19 pandemic.

Sharpness. Sharpness ρ captures the precision with which the tournament translates true skill into observed performance. A higher ρ implies a sharper contest in which scores more faithfully reflect underlying ability. We measure ρ using variation generated by UIL's multi-level structure: students who advance from districts to regionals compete twice within the same season, providing two noisy observations of the same underlying skill level. Because the time between district and regional competitions is short, skill itself cannot reasonably change between the two competitions. Any deviation in a student's relative performance across the two levels therefore reflects measurement noise rather than genuine skill growth.

Formally, for each student i competing in event e in year t, we compute the district-to-regional correlation of standardized scores among advancers:

$$\rho_a = \text{Corr}(z_{i,D,t},\, z_{i,R,t}), \tag{2}$$

where D and R index the district and regional competitions, respectively. A higher ρ therefore corresponds to a sharper contest. A limitation of this strategy is that it only leverages data from those students who qualified for the regional round and are necessarily top performers.

As a complementary measure, we also compute the year-over-year test-retest correlation for individual students who compete in the same event at the district level in consecutive years:

$$\rho_b = \text{Corr}(z_{i,e,t},\, z_{i,e,t-1}), \tag{3}$$

where the sample is restricted to students observed in consecutive years. This measure avoids the restriction to top performers at the cost of allowing for a longer time over which true skill may change.

The two measures produce a highly correlated cross-event ranking ($r = 0.97$), consistent with both capturing the same underlying signal-to-noise ratio. The ranking aligns with prior intuitions about the precision of each event's evaluation mechanism. Events with long, high-item-count formats and well-defined content — Accounting (ACC, 80–100 questions), the various math competitions (CAL/MTH/NUM, 70/60/80 questions), and Spelling & Vocabulary (SPV, 60 questions) — exhibit the highest sharpness. Literary Criticism (LIT), Science (SCI), and Social Studies (SOC) exhibit relatively lower sharpness; with their much broader content base, noise may be introduced by whether the content studied by a student actually appears on the test. Of these three, Social Studies (SOC) has the lowest sharpness, which may reflect the presence of an essay portion. The Current Issues and Events (CIE) contest also requires an essay and has a very dynamic content base, so it being the noisiest is unsurprising. Figure 8 displays both measures.

4.3 Shifting the Skill Threshold

In 2015, UIL reorganized its conference structure by introducing a sixth conference (6A) for the largest Texas schools. This restructuring changed the composition of schools competing against one another at the district level, generating exogenous variation in the skill threshold s_c faced

by individual schools. Some schools came to compete against a stronger pool of opponents, and therefore faced a higher s_c, while others transitioned into weaker pools and faced a lower s_c. Because s_c is a key lever governing equilibrium effort and skill production in the model, this restructuring offers a natural source of identifying variation.

To measure the change in competitive environment facing each school, we construct a counterfactual threshold. For each school i, we compute the average s_c it faced based on the schools it competed against in the pre-reform period (2004–2014). We then compute the average s_c it *would have faced* during the same pre-reform period had it been competing against the schools it is matched with in the post-reform period (2015–2022). The difference between these two quantities is our measure of the exogenous shift in the skill threshold:

$$\Delta \hat{s}_{c,i} = \bar{s}_{c,i}^{\mathsf{post}} - \bar{s}_{c,i}^{\mathsf{pre}}, \tag{4}$$

where both averages are computed using pre-reform skill distributions, so that the variation in $\Delta \hat{s}_{c,i}$ reflects only the change in competitive groupings rather than any endogenous response by students or schools.

We propose to estimate the effect of the threshold shift on student performance using the specification

$$z_{iet} = \beta \cdot \Delta \hat{s}_{c,i} \cdot \mathbf{1}[t \geq 2015] + \gamma_i + \delta_{et} + \varepsilon_{iet}, \tag{5}$$

where γ_i denotes school fixed effects, δ_{et} denotes event-by-year fixed effects, and β captures the effect of a one-unit increase in the threshold on average standardized performance. We further plan to examine heterogeneity by a school's pre-reform proximity to the threshold, asking whether schools whose talent was already close to s_c respond more strongly to a change in the competitive environment. The model predicts precisely this pattern: schools near the threshold face the steepest part of the effort-incentive curve and should therefore exhibit the largest response to a marginal change in s_c.

4.4 Increasing Sharpness

In 2023, UIL amended its competition rules to allow schools to enter up to six students per event at the district level, up from the previous limit of four. For events with a team component, where the school's aggregate score sums the individual performances of its entries, this change increases the number of draws contributing to the team score. By the law of large numbers, averaging over more independent draws reduces the variance of the team score around its true mean, lowering σ and making the contest sharper. This rule change therefore provides a direct, institution-generated increase in sharpness that we can use to study the model's predictions about the consequences of changing σ.

We propose to proceed in two steps. First, we plan to verify that the rule change produced the anticipated increase in sharpness using our district-to-regional reliability measure: for affected events, the correlation between district and regional team scores should rise in the post-2023 period,

while events without a team component serve as a placebo and should show no corresponding change.

Second, we plan to examine how the increase in sharpness affected the distribution of skill, focusing on schools near the threshold. The model predicts that a reduction in σ sharpens the incentive for near-threshold schools to invest in skill, producing a concentration of performance near s_c. We propose to estimate

$$z_{iet} = \sum_k \beta_k \cdot \mathbf{1}[\text{bin}_k(\hat{s}_{c,i} - z_{i,\text{pre}})] \cdot \mathbf{1}[t \geq 2023] + \gamma_i + \delta_{et} + \varepsilon_{iet}, \tag{6}$$

where the bins k partition schools by their pre-reform distance to the threshold, $\hat{s}_{c,i} - z_{i,\text{pre}}$. The model predicts that β_k is largest for schools in the bin closest to zero, whose pre-reform performance placed them just below the advancement cutoff, and smaller for schools farther from the threshold.

5 Conclusion

In this paper we proposed a model in which students invest effort into a contest in return for a direct reward and skill augmentation. We also conducted comparative statics on the participants' behavior as we vary the contest sharpness and cutoff score.

In the next step we plan to investigate the contest organizer's problem. In particular the organizer derives utility from post-competition skill outcomes, but must also bear the costs of prizes and augmentation awarded to winners. The organizer's preferences over individual skill outcomes are captured by $u(S)$, which should be thought of as flexible: a linear $u(S) = S$ indicates the organizer values skill proportionally across the distribution; a convex $u(S)$ indicates the organizer derives disproportionately more utility from top performers, capturing a preference for cultivating exceptional rather than broadly distributed talent. The organizer's net utility across the full talent distribution is

$$U = \int_0^1 (1 - q(s(\theta, e); s_c, \rho))\, u(s(\theta, e)) + q(s(\theta, e); s_c, \rho) \left[u((1 + \beta_1)s(\theta, e)) - (\beta_0 + C^A(\beta_1))\right] dF(\theta),$$

where the integrand groups outcomes by type: losers contribute skill utility $u(s(\theta, e))$ and winners contribute skill utility $u((1 + \beta_1)s(\theta, e)) - (\beta_0 + C^A(\beta_1))$, where costs β_0 and $C^A(\beta_1)$ are incurred only for winners, which is reflected in their appearance inside the winner term.

The designer chooses the contest sharpness ρ and the cutoff skill s_c to maximize her objective over the final skill distribution. Naively one might think that the designer always prefers a sharper contest (higher ρ), because the contest better identifies talent. However as we showed in the comparative statics a less sharp contest might actually induce more effort for both high ability and low ability students, which makes the designer's problem much more subtle. Similarly the designer also faces a trade-off in the cutoff skill: a higher cutoff skill induces more effort among high ability students, but reduces the effort from low ability students, so the optimal cutoff depends on the convexity of the designer's payoff function.

References

Agarwal, Ruchir and Patrick Gaulé, "Invisible Geniuses: Could the Knowledge Frontier Advance Faster?," *American Economic Review: Insights*, 2020, *2* (4), 409–424. Working paper dated January 2019.

Ashby, Nathan J., "An Examination of Peer Effects Using High School Competition Realignments," *Journal of Economic Behavior & Organization*, 2023, *206*, 122–135.

Bottan, Nicolas L. and Dan Bernhardt, "Small Fish in a Big Pool: The Discouraging Effects of Relative Assessment," Warwick Economics Research Paper 1511, University of Warwick, Department of Economics 2024.

Calaway, Ian, "Early mentors for exceptional students," *Available at SSRN 6677338*, 2024.

Cotton, Christopher, Brent Hickman, and Joseph Price, "Affirmative Action and Human Capital Investment: Evidence from a Randomized Field Experiment," *Journal of Labor Economics*, 2022, *40* (1), 157–185.

Ellison, Glenn and Ashley Swanson, "The Gender Gap in Secondary School Mathematics at High Achievement Levels: Evidence from the American Mathematics Competitions," *Journal of Economic Perspectives*, 2010, *24* (2), 109–128.

_ and _, "Do Schools Matter for High Math Achievement? Evidence from the American Mathematics Competititions," *American Economic Review*, 2016, *106* (6), 1244–1277.

Holmstrom, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 1999, *66* (1), 169–182.

Lazear, Edward P. and Sherwin Rosen, "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 1981, *89* (5), 841–864.

Moldovanu, Benny and Aner Sela, "The Optimal Allocation of Prizes in Contests," *American Economic Review*, 2001, *91* (3), 542–558.

_ and _, "Contest Architecture," *Journal of Economic Theory*, 2006, *126* (1), 70–96.

Nalebuff, Barry J. and Joseph E. Stiglitz, "Prizes and Incentives: Towards a General Theory of Compensation and Competition," *The Bell Journal of Economics*, 1983, *14* (1), 21–43.

Olszewski, Wojciech and Ron Siegel, "Large Contests," *Econometrica*, 2016, *84* (2), 835–854.

_ and _, "Performance-Maximizing Large Contests," *Theoretical Economics*, 2020, *15* (1), 57–88.

Rosen, Sherwin, "The Economics of Superstars," *American Economic Review*, 1981, *71* (5), 845–858.

A Omitted Proofs

A.1 Proof of Theorem 1

The Monotonicity of the Skill Function $s(\theta)$ We have $s(0) = 0$, so we only need to prove the monotonicity for strictly positive θ.

For the ease of notation let $G(s) = r_w s + \frac{s^\rho}{s^\rho + s_c^\rho} \cdot [\beta_0 + \beta_1 r_w s]$.

We have $s(\theta)$ is the maximizer of

$$G(s) - \frac{s^2}{2\theta^2}.$$

Suppose $0 < \theta_1 < \theta_2$. We have

$$G(s(\theta_1)) - \frac{s(\theta_1)^2}{2\theta_1^2} \geq G(s(\theta_2)) - \frac{s(\theta_2)^2}{2\theta_1^2},$$

and

$$G(s(\theta_2)) - \frac{s(\theta_2)^2}{2\theta_2^2} \geq G(s(\theta_1)) - \frac{s(\theta_1)^2}{2\theta_2^2},$$

Indeed the first inequality says type θ_1 student prefers skill $s(\theta_1)$ to $s(\theta_2)$, and the second inequality says type θ_2 student prefers skill $s(\theta_2)$ to $s(\theta_1)$.

We deduce that

$$\frac{s(\theta_2)^2 - s(\theta_1)^2}{2\theta_1^2} \geq \frac{s(\theta_2)^2 - s(\theta_1)^2}{2\theta_2^2}.$$

Since $0 < \theta_1 < \theta_2$, we have $s(\theta_2)^2 - s(\theta_1)^2 \geq 0$, which means $s(\theta_2) \geq s(\theta_1)$. Therefore the skill is non-decreasing in θ.

The Limits of the Skill Function $s(\theta)$ The marginal cost for type θ student is linear in s:

$$MC(s, \theta) = \frac{s}{\theta^2}.$$

The marginal benefit is

$$MB(s) = r_w + \beta_0 \cdot \frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} + \beta_1 r_w \cdot \left(\frac{s^\rho}{s^\rho + s_c^\rho} + \frac{\rho s_c^\rho s^\rho}{(s^\rho + s_c^\rho)^2} \right).$$

In particular the marginal benefit is independent of θ.

We now discuss the limits of $s(\theta)$.

- Suppose $\theta \to 0$. Then the marginal cost $\frac{s}{\theta^2}$ (as a function of s) is a line passing through the origin with slope near infinity.

  Since the marginal benefit is independent of θ, as $\theta \to 0$ the marginal cost intersects the marginal benefit near 0, which means $s(\theta) \to 0$.

  As $s \to 0$ the marginal benefit $MB(s)$ converges to r_w, so at any intersect of MB and MC we have $\frac{s}{\theta^2} \approx r_w$. We deduce that as $\theta \to 0$ we have $s(\theta) \approx r_w \theta^2$.

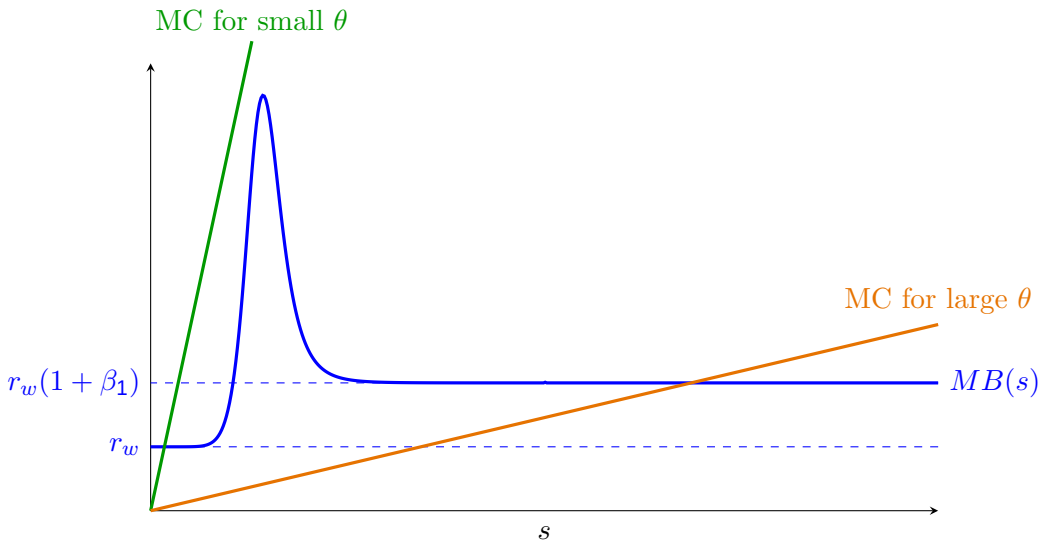


Figure 3: Changing θ: MB is fixed; the slope of MC changes.

- Suppose $\theta \to 1$. Then the marginal cost $\frac{s}{\theta^2}$ (as a function of s) is a line passing through the origin with slope near zero.

  Since the marginal benefit is independent of θ and is bounded below by r_w, as $\theta \to 1$ the intersect of the marginal cost and the marginal benefit goes to infinity, which means $s(\theta) \to 1$.

  As $s \to 1$ we have $MB(s)$ converges to $r_w(1+\beta_1)$, so at any intersect of MB and MC we have $\frac{s}{\theta^2} \approx r_w(1+\beta_1)$. We deduce that as $\theta \to 1$ we have $s(\theta) \approx r_w(1+\beta_1)\theta^2$.

At Most One Jump in the Skill Function $s(\theta)$ We need to prove that the marginal benefit function has the following properties.

- Either $\frac{MB(s)}{s}$ is decreasing in s, or $\frac{MB(s)}{s}$ is decreasing, increasing, and decreasing.
- If ρ is sufficiently large, then $\frac{MB(s)}{s}$ is decreasing, increasing, and decreasing.

If $\frac{MB(s)}{s}$ is decreasing in s, then the choice of skill $s(\theta)$ is continuous in θ. If $\frac{MB(s)}{s}$ is decreasing, increasing, and decreasing, then the choice of skill $s(\theta)$ has exactly one jump.

The connection between $\frac{MB(s)}{s}$ and the continuity of $s(\theta)$ is as follows. Figure 3 shows that the marginal benefit $MB(s)$ is independent of θ, while the marginal cost $\frac{s}{\theta^2}$ is linear in s. Therefore as we increase θ the marginal benefit is fixed, and the slope of the marginal cost decreases. If $\frac{MB(s)}{s}$ is decreasing in s, then the MB and MC have a unique intersection, so as we increase θ the intersect $s(\theta)$ changes continuously. On the other hand if $\frac{MB(s)}{s}$ is decreasing, increasing, and decreasing, then for some values of θ the MB and MC have three intersections as Figure 4 shows. In this case we have two possible candidates for the optimal skill $s(\theta)$. For small θ the left intersect is optimal, but for large θ the right intersect becomes optimal, so the optimal skill $s(\theta)$ has a jump (the point of discontinuity occurs for the θ where both intersects are optimal).

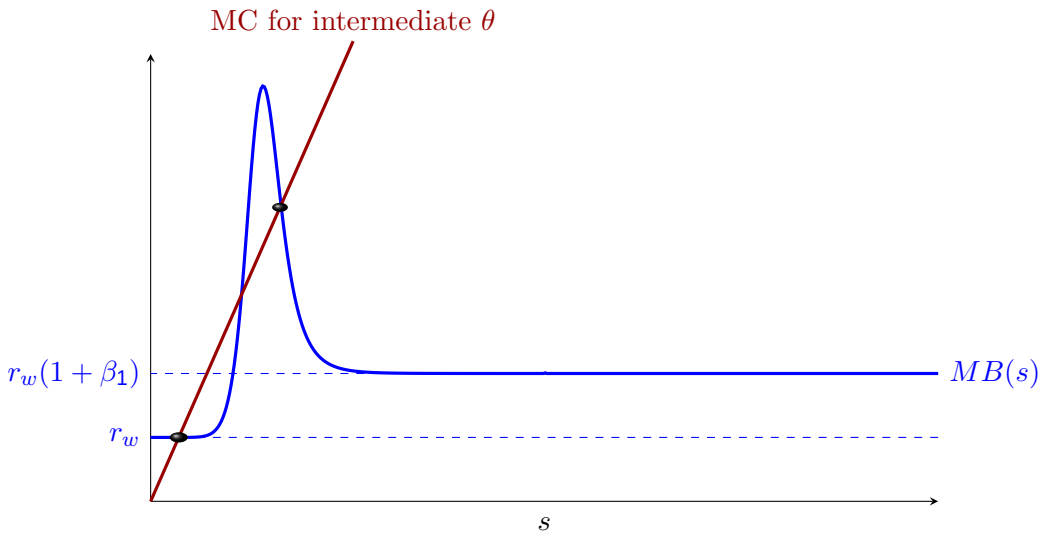


Figure 4: Two Possible Optimal Skill Choices

We next focus on the shape of $\frac{MB(s)}{s}$. We need to compute its derivative, and for the ease of notation we define

$$x = \frac{s}{s_c}, \qquad a = \frac{\beta_0}{r_w s_c}, \qquad b = \beta_1, \qquad Q(x) = \frac{x^\rho}{1+x^\rho}.$$

We have

$$MB(s) = r_w \left[1 + aQ^{\emptyset}(x) + b\big(Q(x) + xQ^{\emptyset}(x)\big)\right].$$

Consequently we have

$$\frac{MB(s)}{s} = \frac{r_w}{s_c} \cdot \frac{1 + aQ^{\emptyset}(x) + b\big(Q(x) + xQ^{\emptyset}(x)\big)}{x}.$$

A direct calculation gives

$$\frac{d}{ds}\left(\frac{MB(s)}{s}\right) = -\frac{r_w}{s_c} \cdot \frac{P_\rho(x)}{x^2(1+x^\rho)^3},$$

where

$$\begin{aligned} P_\rho(x) = 1 &+ a\rho(2-\rho)x^{\rho-1} + \left[3 + b(1-\rho^2)\right]x^\rho \\ &+ a\rho(\rho+2)x^{2\rho-1} + \left[3 + b(2+\rho^2)\right]x^{2\rho} + (1+b)x^{3\rho}. \end{aligned}$$

We deduce that the sign of $\frac{d}{ds}\left(\frac{MB(s)}{s}\right)$ is the opposite of the sign of $P_\rho(x)$.

Now P_ρ is a generalized polynomial with exponents

$$0, \quad \rho-1, \quad \rho, \quad 2\rho-1, \quad 2\rho, \quad 3\rho.$$

Because $\rho > 1$, these exponents are strictly increasing:

$$0 < \rho - 1 < \rho < 2\rho - 1 < 2\rho < 3\rho.$$

The corresponding coefficient sequence is

$$1, \quad a\rho(2-\rho), \quad 3 + b(1-\rho^2), \quad a\rho(\rho+2), \quad 3 + b(2+\rho^2), \quad 1 + b.$$

The first coefficient is positive, and the last three coefficients are nonnegative, so this sequence has at most two sign changes. By the generalized Descartes rule, P_ρ has at most two positive roots. Moreover it is easy to see that P_ρ is positive when x is near zero or as x goes to infinity. Therefore either P_ρ is always positive, or it is positive, negative, and positive. As a result either $\frac{MB(s)}{s}$ is always decreasing, or it decreases, increases, then decreases.

Finally we show that if ρ is sufficiently large, then $\frac{MB(s)}{s}$ is increasing for some interval of s. Notice that for all $s < s_c$, we have $MB(s) \to r_w$ as $\rho \to \infty$. On the other hand $MB(s_c) \to \infty$ as $\rho \to \infty$. As a result for any fixed s below s_c, if ρ is sufficiently large, we have $\frac{MB(s)}{s} < \frac{MB(s_c)}{s_c}$, because the left hand side is bounded, while the right hand side goes to infinity. We conclude that $s(\theta)$ has a jump when ρ is sufficiently large.

A.2 Proof of Proposition 1

Decrease in Skill for Sufficiently Small θ Suppose $\rho_1 < \rho_2$. Let $s_1(\theta)$ and $s_2(\theta)$ denote the skill function under ρ_1 and ρ_2 respectively. We need to show that for sufficiently small θ we have $s_1(\theta) > s_2(\theta)$.

We only need to show that $MB_1(s) > MB_2(s)$ for sufficiently small s. (We use MB_1 and MB_2 to denote marginal benefits under ρ_1 and ρ_2.) Indeed for sufficiently small θ the slope of the marginal cost $\frac{s}{\theta^2}$ is large, and the MC intersects MB at exactly one point. Since the MC is a line through the origin, decreasing the MB also decreases the intersect.

If s is sufficiently small, we have

$$\frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} \approx \frac{\rho s^{\rho-1}}{s_c^\rho}, \qquad \frac{s^\rho}{s^\rho + s_c^\rho} \approx \frac{s^\rho}{s_c^\rho}.$$

We obtain that

$$\begin{aligned} MB(s) &= r_w + \beta_0 \cdot \frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} + \beta_1 r_w \cdot \left(\frac{s^\rho}{s^\rho + s_c^\rho} + \frac{\rho s_c^\rho s^\rho}{(s^\rho + s_c^\rho)^2} \right). \\ &\approx r_w + \beta_0 \frac{\rho}{s_c^\rho} \cdot s^{\rho-1} + \beta_1 r_w \frac{1+\rho}{s_c^\rho} \cdot s^\rho. \end{aligned} \tag{7}$$

Notice if $\rho_1 < \rho_2$, then as s goes to 0 we have $\frac{s^{\rho_1 - 1}}{s^{\rho_2 - 1}} \to \infty$ and $\frac{s^{\rho_1}}{s^{\rho_2}} \to \infty$. As a result for sufficiently small s we have $MB_1(s) > MB_2(s)$. Since the MC is independent of ρ, we conclude that the choice of skill decreases (i.e. $s_1(\theta) > s_2(\theta)$).

Decrease in Skill for Sufficiently Large θ Suppose $\rho_1 < \rho_2$. We claim that for sufficiently large s we have $MB_1(s) > MB_2(s)$.

If s is sufficiently large, we have

$$\frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} \approx \frac{\rho s_c^\rho}{s^{\rho+1}}, \qquad \frac{s^\rho}{s^\rho + s_c^\rho} \approx 1 - \frac{s_c^\rho}{s^\rho}.$$

We deduce that

$$\begin{aligned} MB(s) &= r_w + \beta_0 \cdot \frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} + \beta_1 r_w \cdot \left(\frac{s^\rho}{s^\rho + s_c^\rho} + \frac{\rho s_c^\rho s^\rho}{(s^\rho + s_c^\rho)^2} \right). \\ &\approx r_w(1+\beta_1) + \beta_0 \frac{\rho s_c^\rho}{s^{\rho+1}} + \beta_1 r_w(\rho - 1) \cdot \frac{s_c^\rho}{s^\rho}. \end{aligned} \tag{8}$$

If $\rho_1 < \rho_2$, then as s goes to infinity we have $\frac{s^{\rho_2+1}}{s^{\rho_1+1}} \to 0$ and $\frac{s^{\rho_2}}{s^{\rho_1}} \to 0$. As a result for sufficiently large s we have $MB_1(s) > MB_2(s)$. Since the MC is independent of ρ, we conclude that the choice of skill decreases (i.e. $s_1(\theta) > s_2(\theta)$).

A.3 Proof of Proposition 2

Suppose $s_c < s_c^0$. We claim that if we increase s_c to s_c^0, then the chosen skill $s(\theta)$ decreases for sufficiently small θ and increases for sufficiently large θ.

Equivalently we need to show that as s_c increases to s_c^0 we have $MB(s)$ decreases for sufficiently small s, and $MB(s)$ increases for sufficiently large s. Indeed as Figure 3 shows for small and large θ we have MB and MC intersect only once, so increasing MB implies increasing $s(\theta)$.

If s is sufficiently small, Equation (7) shows that

$$MB(s) \approx r_w + \beta_0 \frac{\rho}{s_c^\rho} \cdot s^{\rho-1} + \beta_1 r_w \frac{1+\rho}{s_c^\rho} \cdot s^\rho.$$

As a result if s_c increases, then $MB(s)$ decreases.

Similarly if s is sufficiently large, Equation (8) shows that

$$MB(s) \approx r_w(1+\beta_1) + \beta_0 \frac{\rho s_c^\rho}{s^{\rho+1}} + \beta_1 r_w(\rho-1) \cdot \frac{s_c^\rho}{s^\rho}.$$

Consequently if s_c increases, then $MB(s)$ increases.

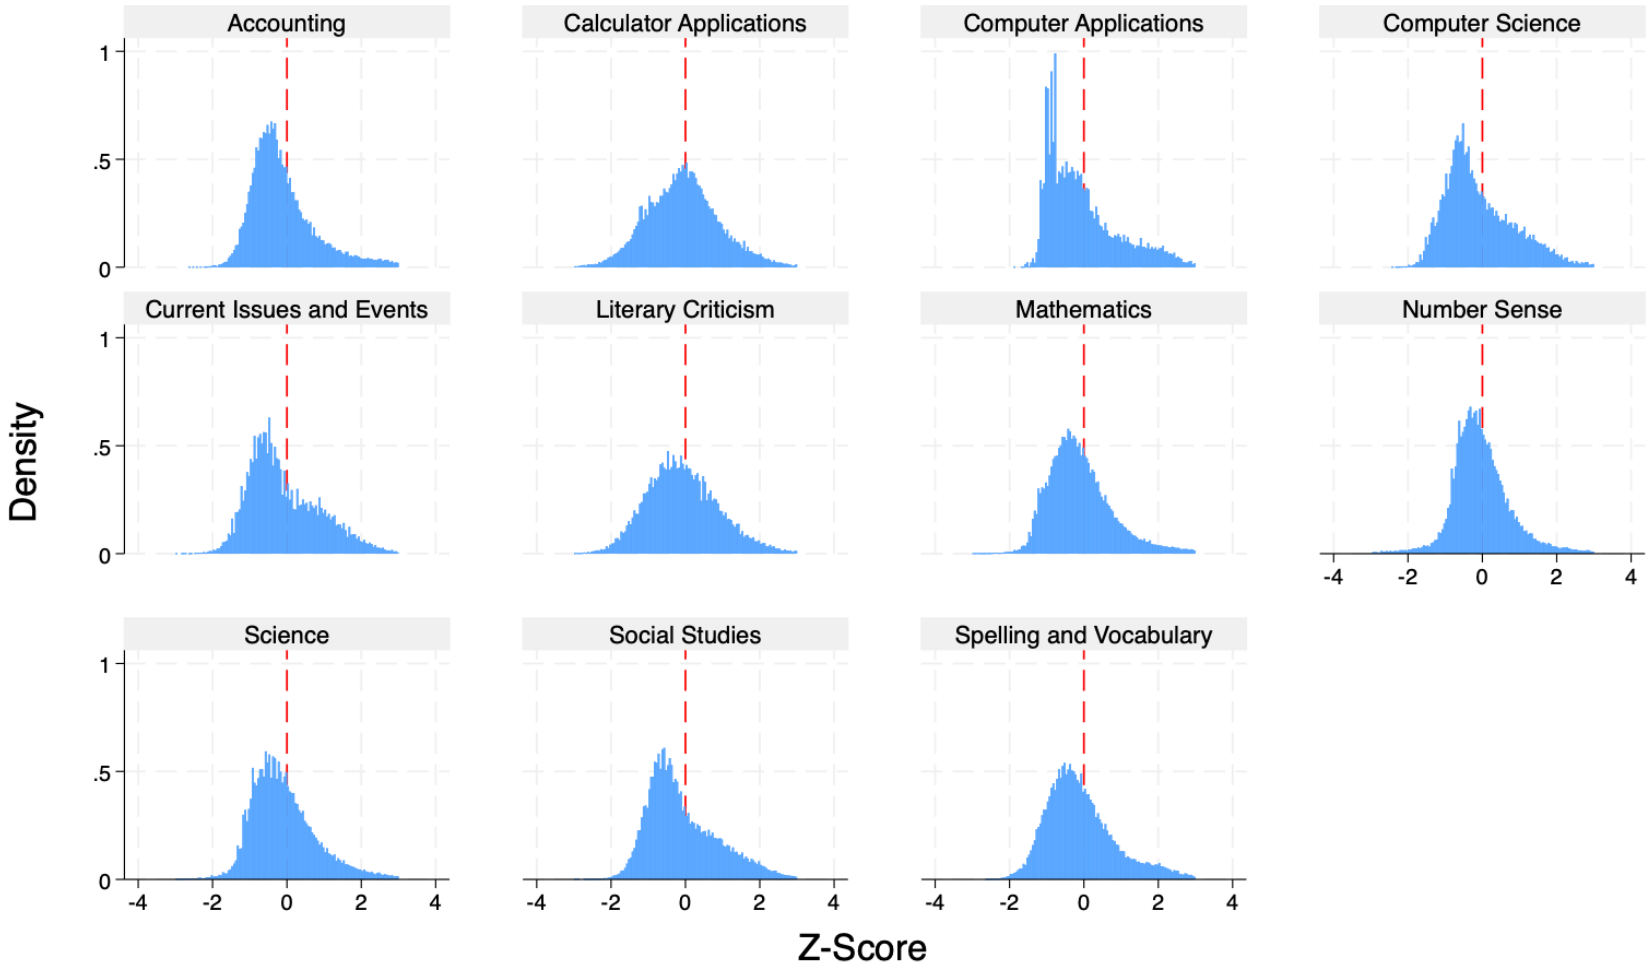
B Figures and Tables from the Empirical Section

Table 1: UIL Academic Event Descriptions

Abbr.	UIL academic event	Description
ACC	Accounting	Hour-long contest with 80–100 questions on elementary accounting principles and practices, including bookkeeping terminology, worksheets, financial statements, trial balances, journalizing, posting, bank reconciliation, payroll, and other basic accounting-cycle topics.
CAL	Calculator Applications	Thirty-minute contest with 70 calculator-based mathematical problems: numerical, stated, and geometry problems involving arithmetic, algebra, trigonometry, logarithms, calculus, statistics, matrix algebra, and related applied problem solving.
COM	Computer Applications	Contest focused on speed and accuracy in word processing, database, spreadsheet, and integrated computer applications; contestants produce scored printouts and complete a mandatory tiebreaker.
CSC	Computer Science	Contest covering computer science knowledge and programming; individual competition is based on a 45-minute written exam, while the team component combines written scores with a two-hour programming contest.
CIE	Current Issues and Events	Contest on current state, national, and world events and issues, consisting of 40 multiple-choice questions and an essay assessing students' ability to synthesize and evaluate information.
LIT	Literary Criticism	Ninety-minute contest assessing literary analysis, literary terms, literary history, annual reading-list works, and critical interpretation, with a required tiebreaking essay.
MTH	Mathematics	Forty-minute, 60-question multiple-choice exam testing Algebra I, Geometry, Algebra II, Pre-calculus, AP Calculus BC, AP Statistics, and selected director's-choice topics.
NUM	Number Sense	Ten-minute, 80-question mental mathematics test covering high school mathematics, emphasizing speed, shortcuts, and accuracy without calculator use.
SCI	Science	Two-hour objective test with 60 questions, divided equally across biology, chemistry, and physics, intended to test broad scientific knowledge, principles, experiments, and scientific understanding.
SOC	Social Studies	Contest with 45 objective questions and an essay; students study annual primary readings, selected documents, and general social studies concepts tied to a yearly topic.
SPV	Spelling & Vocabulary	Three-part contest based largely on UIL's annual Word Power list, including vocabulary/proofreading multiple-choice items and handwritten spelling of pronounced words plus a tiebreaker.

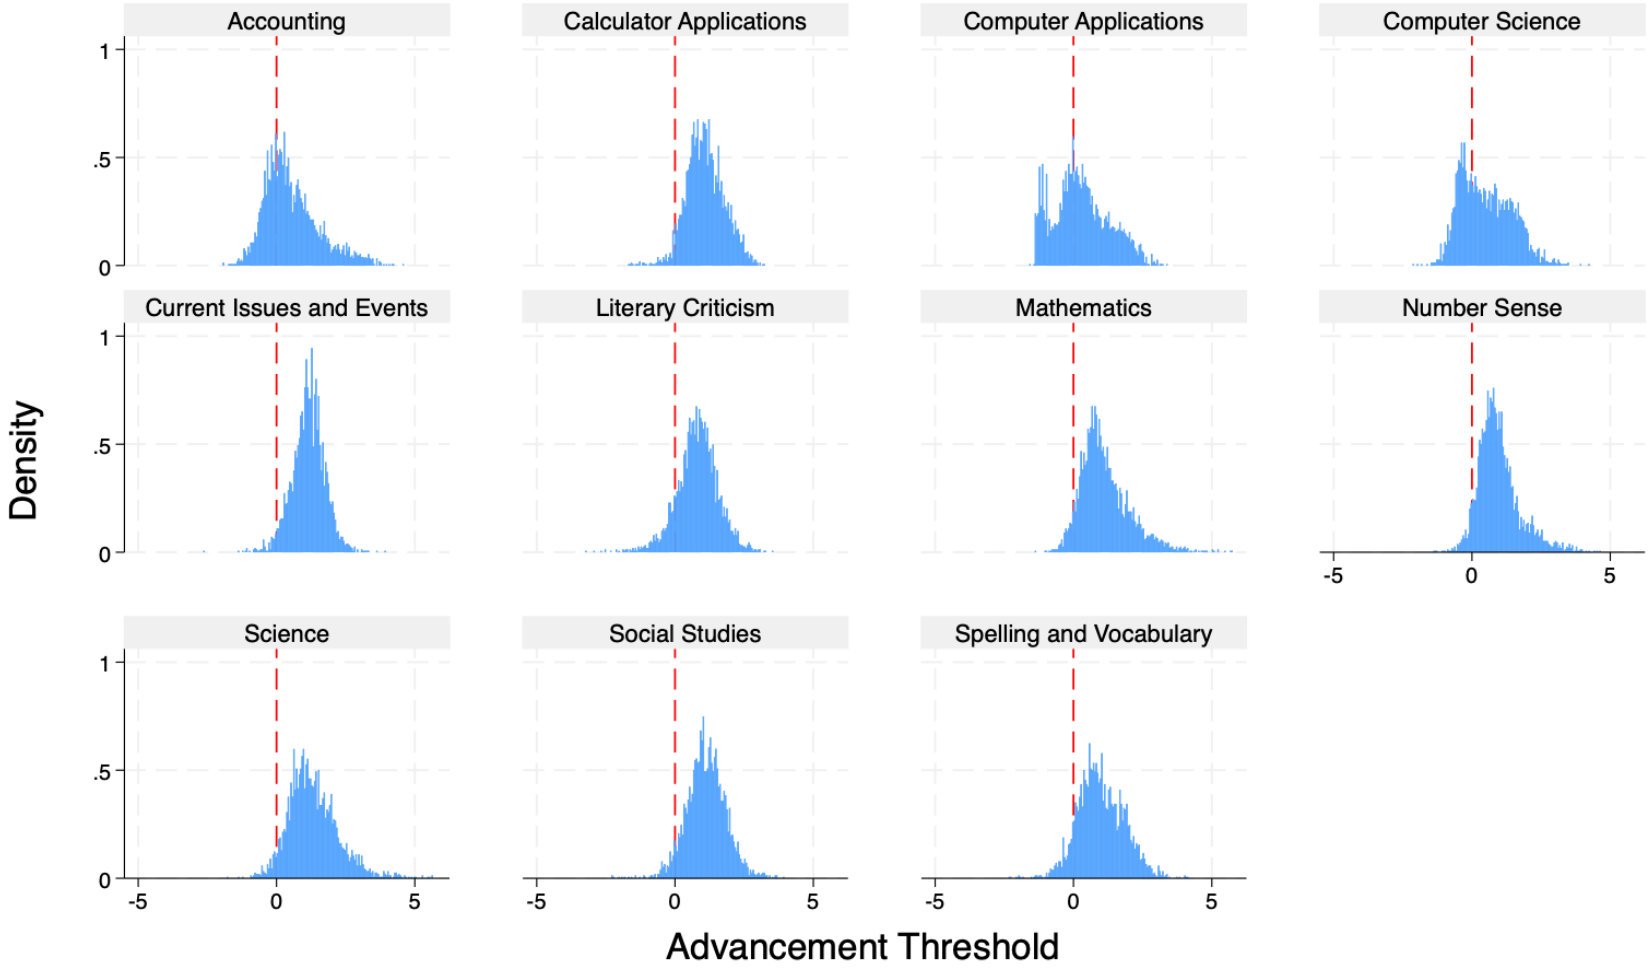
This table describes each of the high school UIL academic events analyzed in this paper, including the format, duration, item count, and content coverage of each contest.

Figure 5: UIL Performance Distributions



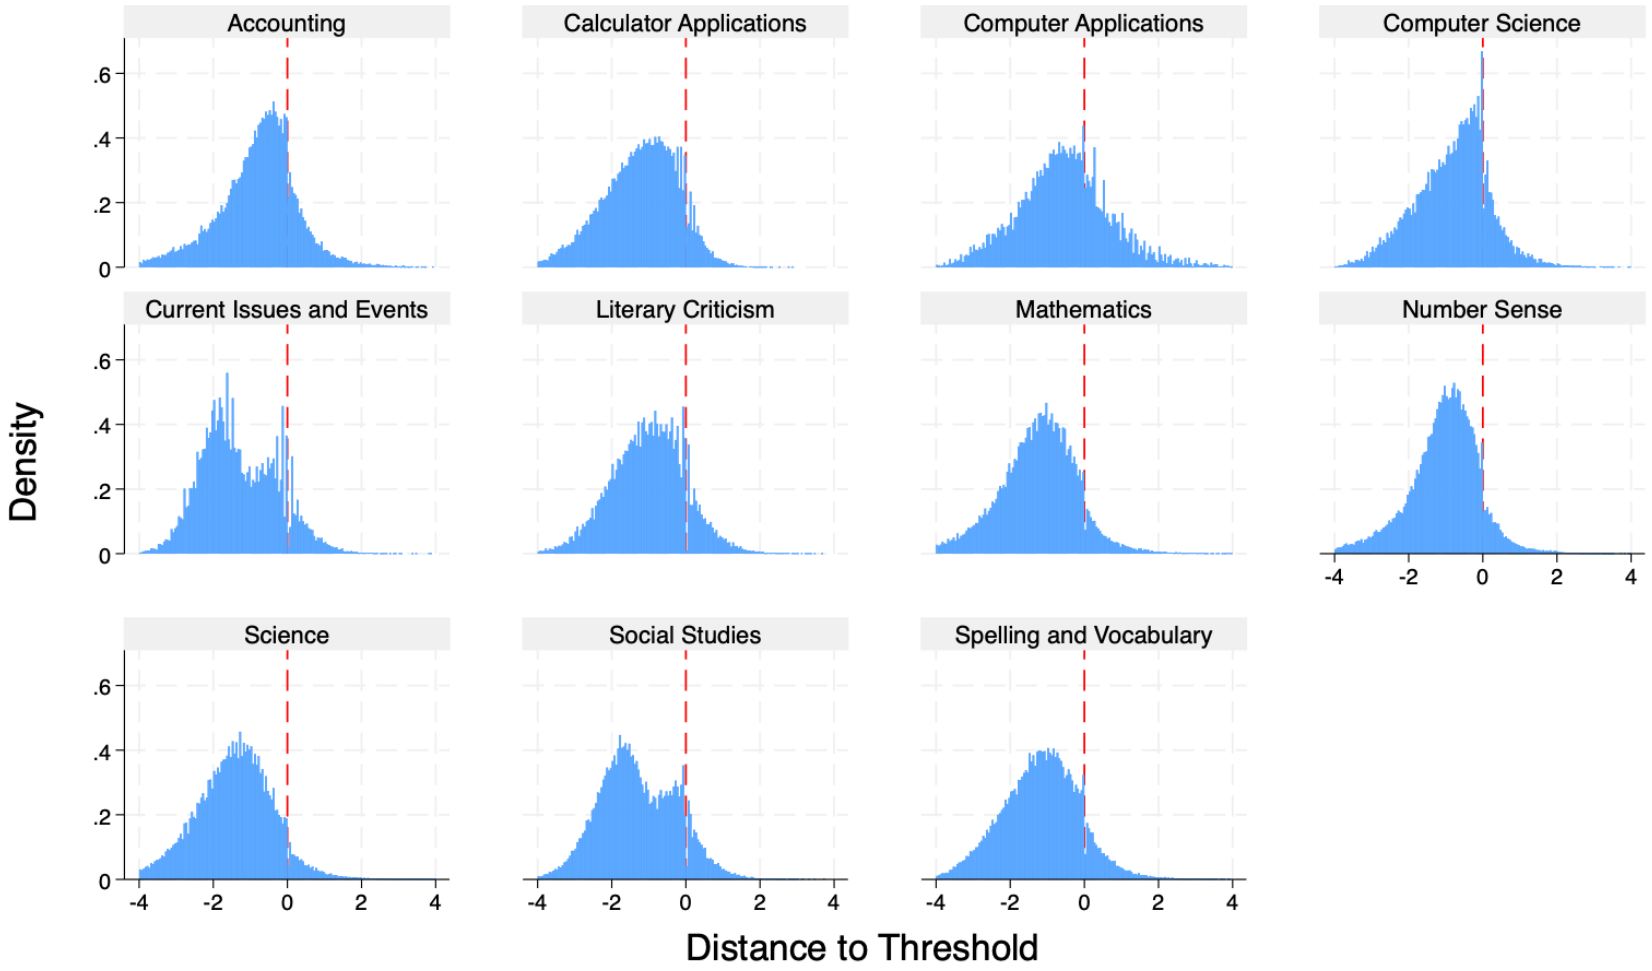
This figure plots the distribution of standardized district-level scores for each event, pooling across all years and conferences. Scores are standardized within each event-year cell. The dashed line marks the mean.

Figure 6: UIL District Advancement Thresholds



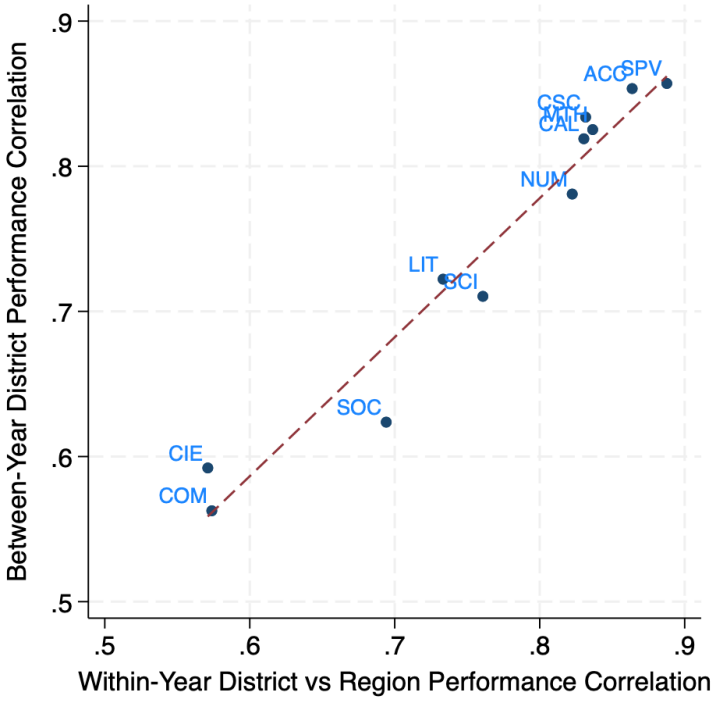
This figure plots the distribution of the implied advancement threshold $\hat{s}_c$ across district-year-event cells, defined as the standardized score of the lowest-performing advancer in each cell. Variation reflects differences in district competitiveness and participation rates.

Figure 7: Participant Distance to Advancement Threshold



This figure plots the distribution of each participant's standardized score relative to the advancement threshold, $z_{ict} - \bar{s}_c$, by event. The dashed line marks the threshold.

Figure 8: Sharpness Across Events and Methods



This figure compares the two sharpness measures across events. The horizontal axis plots the same-year district-to-regional performance correlation (ρ_a); the vertical axis plots the between-year within-student correlation (ρ_b).