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

Ian Calaway Delong Meng Farrell Wu

September 6, 2026

Abstract

We develop models of developmental tournaments in which competitions both select and produce talent. Participants differ in talent and exert costly effort to build skill; an organizer recognizes winners using a noisy assessment; and winning yields a direct prize plus a productive augmentation that scales with their skill. Effort is hump-shaped in talent: middle-ability participants work hardest, while the most talented coast and the least talented give up. We solve tractable versions of the organizer's problem over recognition breadth, assessment precision, and rewards. In deterministic benchmarks, optimal recognition is interior. At fixed breadth, larger budgets and winners' prior skill tilt rewards toward augmentation, while maximizing population skill calls for noisier assessment than maximizing winner quality. Using two decades of records from the Texas University Interscholastic League and a 2015 reassignment reform, we find that a one-standard-deviation move toward the advancement cutoff is followed by gains of roughly 0.05–0.09 standard deviations in team and average-entrant performance. Effects increase with the size of the move and grow over time, consistent with cumulative skill.

We thank Zhen Fang (SJTU), Jingsai Ni (SJTU), Sitta Tantikul (MIT) for excellent research assistance. The authors have no financial 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

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

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.

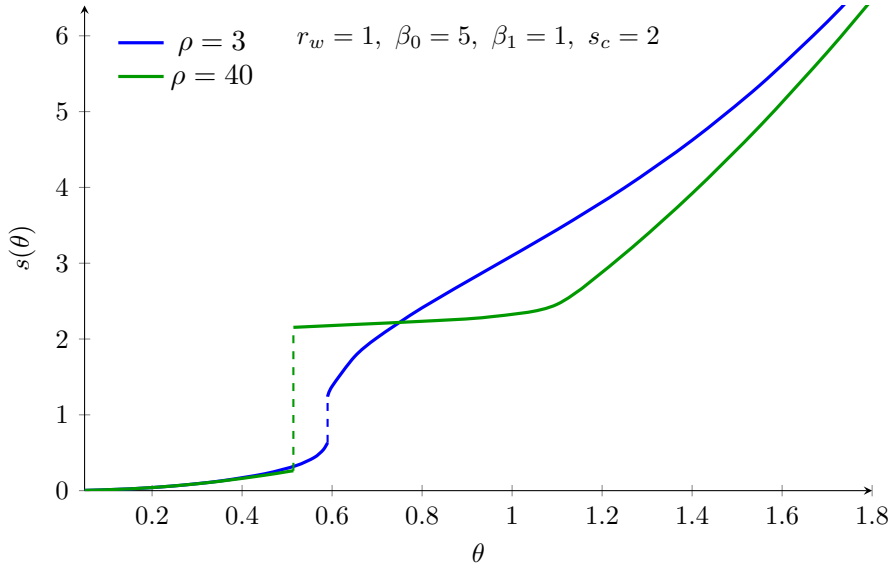


Figure 1: The Student's Choice of Skill

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 $\theta \to \infty$ 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.

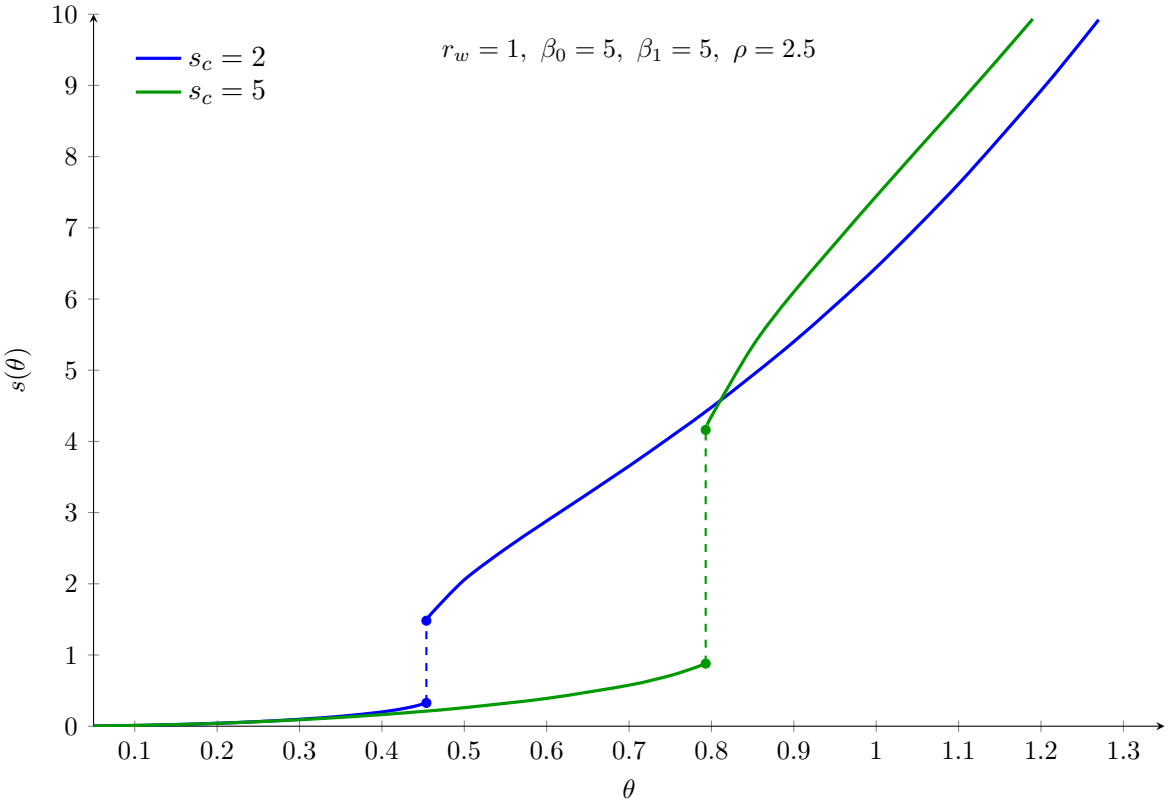


Figure 2: Changing the Cutoff Skill

4 Designer's Problem

Having characterized participant behavior, we now turn to the organizer's design problem. The organizer values post-competition skill but pays for the prizes and productive augmentation awarded to winners.

The designer chooses assessment sharpness ρ, the recognition threshold s_c, and the two winner benefits (β_0, β_1). Throughout this section, augmentation has unit-linear per-winner resource cost, $C_\beta(\beta_1) = \beta_1$. For any policy, let $s(\theta)$ denote type θ's equilibrium skill choice, and write

$$q(\theta) = \frac{s(\theta)^\rho}{s(\theta)^\rho + s_c^\rho}$$

for the resulting probability that type θ wins. The designer's general problem is

$$\begin{aligned} \max_{\rho, s_c, \beta_0, \beta_1} \quad & \int_0^\infty \Big[(1 - q(\theta))u(s(\theta)) + q(\theta)u\big((1+\beta_1)s(\theta)\big)\Big]\,\mathrm{d}F(\theta) \\ \text{s.t.} \quad & s(\theta) \in \arg\max_{s \geq 0} \left\{ r_w s + \frac{s^\rho}{s^\rho + s_c^\rho}\big[\beta_0 + \beta_1 r_w s\big] - \frac{s^2}{2\theta^2} \right\}, \\ & \beta_0 \int_0^\infty q(\theta)\,\mathrm{d}F(\theta) \leq B_0, \qquad \beta_1 \int_0^\infty q(\theta)\,\mathrm{d}F(\theta) \leq B_1. \end{aligned}$$

The incentive constraint requires $s(\theta)$ to be an equilibrium choice; the remaining constraints impose separate expected-expenditure budgets. Under a common award budget, they are replaced by

$$(\beta_0 + \beta_1) \int_0^\infty q(\theta)\, \mathrm{d}F(\theta) \le B.$$

Because s_c implements a winner share $\alpha = \mathbb{E}[q]$, we equivalently describe recognition breadth by α. The four design margins are therefore breadth α, sharpness ρ, the direct prize β_0, and productive augmentation β_1.

Rather than solve the full four-dimensional problem at once, we analyze three benchmark cases. First, deterministic recognition, $\rho = \infty$, together with $\beta_1 = 0$, isolates optimal recognition breadth. Second, restoring augmentation under deterministic recognition reveals how a common budget is divided between β_0 and β_1. Finally, fixing breadth and allowing finite ρ isolates optimal assessment sharpness under population-skill and winner-skill objectives.

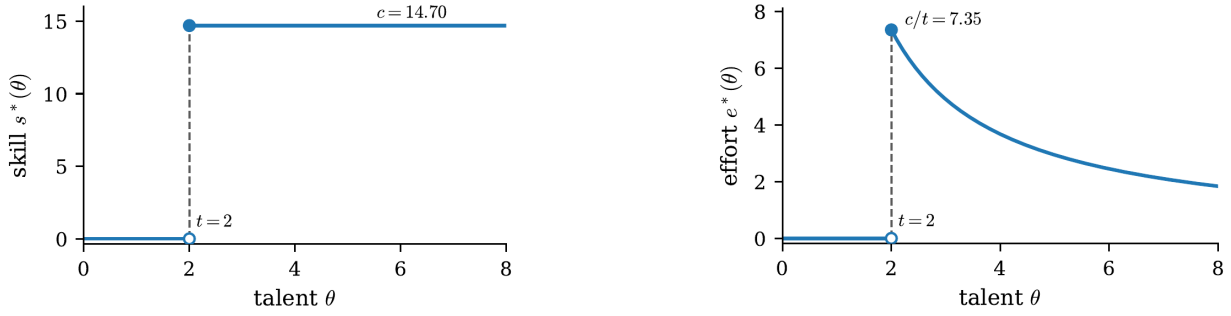
4.1 Deterministic direct-prize benchmarks

We first take the deterministic endpoint,

$$q_\infty(s; c) = \mathbf{1}\{s \ge c\},$$

and specialize the organizer's payoff to mean final skill, $u(s) = s$. Throughout this subsection, skill gains are measured relative to the no-tournament benchmark.

At the sharp corner $\beta_1 = r_w = 0$, losers choose zero skill and winners jump to a common cutoff. Figure 3(a)–(b) shows the resulting skill and effort schedules.



(a) Skill at $r_w = \beta_1 = 0$.

(b) Effort at $r_w = \beta_1 = 0$.

Figure 3: Deterministic direct-prize benchmarks, $B = 1$ and $\gamma = 3$. Panels (a)–(b) evaluate the pure-prize optimum $\alpha^* = 1/27$, where $t = 2$, $c = 6\sqrt{6}$, and effort jumps to $3\sqrt{6}$. Open and filled circles distinguish the losing and winning branches at the indifferent marginal type.

Begin with the sharp corner $\beta_1 = r_w = 0$. If the target winner share is α, the marginal winner is $t = \alpha^{-1/\gamma} - 1$ and the binding prize budget gives $\beta_0 = B/\alpha$. All losers choose zero skill, while every

winner chooses the common cutoff $c = t\sqrt{2B/\alpha}$. Figure 3(a)–(b) makes the discontinuity explicit. The marginal type is indifferent between the two branches; the inclusive cutoff and greatest-best-response convention place that measure-zero type on the winning branch. Economically, the least-talented winner works hardest, while more talented winners need progressively less effort to reach the same threshold.

In this corner case, we conclude that there will always exist an interior skill maximizing $\alpha^* \in (0, 1)$, for which we can derive the closed form in terms of the Lomax parameter γ.

Proposition 3 (Pure-prize breadth)**.** *Suppose talent is Lomax, $F(\theta) = 1 - (1+\theta)^{-\gamma}$, with $\gamma > 2$. If $r_w = 0$, deterministic recognition produces*

$$S_\infty(\alpha) = \sqrt{2B}\left(\alpha^{1/2-1/\gamma} - \alpha^{1/2}\right), \tag{1}$$

and has the unique optimum

$$\alpha^*_\infty(\gamma) = \left(\frac{\gamma-2}{\gamma}\right)^\gamma. \tag{2}$$

Moreover, $\alpha^*_\infty(\gamma)$ is strictly increasing in γ.

The optimum is interior because both extremes destroy aggregate skill creation. As $\alpha \to 0$, the award per winner becomes large but the recognized mass vanishes; as $\alpha \to 1$, the marginal winner approaches zero talent and the cutoff collapses. A larger γ thins the right tail, so narrowing recognition buys less additional talent and the organizer recognizes a broader group. In the limits, $\alpha^*_\infty \to 0$ as $\gamma \to 2$ and $\alpha^*_\infty \to e^{-2}$ as $\gamma \to \infty$.

Now let $r_w > 0$. For a fixed target share α, define

$$t = \alpha^{-1/\gamma} - 1, \qquad c = r_w t^2 + t\sqrt{\frac{2B}{\alpha}}, \qquad h = \sqrt{\frac{c}{r_w}}.$$

Types below t remain at baseline and lose, types in $[t, h)$ build exactly c, and types above h already clear the cutoff. Tournament-created skill is therefore

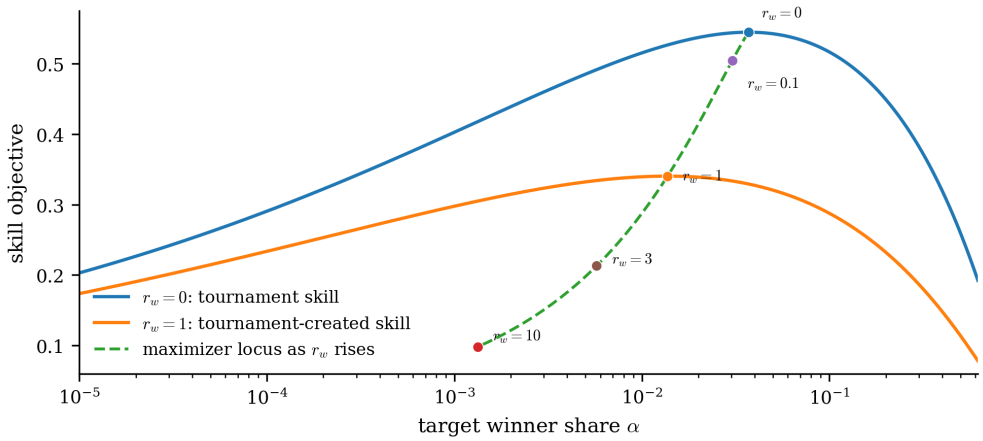
$$\Delta S_\infty(\alpha, r_w) = \int_t^h (c - r_w\theta^2)\,\mathrm{d}F(\theta), \qquad \frac{\partial \Delta S_\infty}{\partial r_w} = \int_t^h (t^2 - \theta^2)\,\mathrm{d}F(\theta) < 0.$$

Holding α fixed also fixes the marginal type t. Marginal indifference implies $\partial c/\partial r_w = t^2$, whereas the baseline skill of every threshold builder $\theta > t$ rises at the faster rate $\theta^2 > t^2$. Hence the incremental skill gain $c - r_w\theta^2$ shrinks with r_w for every threshold builder except the marginal type. Because this gain is already zero at the upper boundary h, movement in h contributes no boundary term. Thus

$$\frac{\partial \Delta S_\infty(\alpha, r_w)}{\partial r_w} < 0 :$$

at any fixed recognition breadth, a higher baseline return crowds out the skill induced by a direct prize.

Allowing the designer to adjust breadth produces a corresponding numerical pattern. In Figure 3(c), the objective remains single-peaked in α, and its numerical maximizer $\alpha^*_\infty(r_w)$ shifts downward as r_w rises. Thus, within the deterministic direct-prize benchmark, higher baseline returns are associated with narrower optimal recognition: as the prize has less incremental skill left to create, the designer concentrates the fixed budget on a smaller group of winners.



(c) Breadth objectives and the numerical maximizer locus.

Figure 3: Deterministic direct-prize benchmarks (continued). Panel (c) plots the $r_w = 0$ and $r_w = 1$ objectives and the dashed trajectory of their numerical maximizers as r_w rises from zero to ten.

For $B = 1$ and $\gamma = 3$, Figure 3(c) traces a steadily declining optimizer over the displayed range. The companion screen extends this exercise over $\gamma \in [2.05, 20]$ and $r_w/\sqrt{B} \in [10^{-4}, 10^4]$; every retained optimizer path declines and every retained objective has one numerical peak.

The decline is in tournament-created skill, not total skill. A larger r_w raises the skill participants would produce even without the tournament, but it leaves less for a direct prize to induce. At fixed α, $h^2 = t^2 + t\sqrt{2B/\alpha}/r_w$, so h falls with r_w and the share of winners who already clear the cutoff at baseline, $(1+h)^{-\gamma}/\alpha$, rises. Along the displayed maximizer locus, this baseline-qualified share increases from about 1.1 percent at $r_w = .1$ to 14.4 percent at $r_w = 1$, 33.5 percent at $r_w = 3$, and 59.0 percent at $r_w = 10$. The designer responds by recognizing a narrower group.

4.2 Productive augmentation and the deterministic reward mix

We now restore productive augmentation, retaining the unit-linear cost $C_\beta(\beta_1) = \beta_1$. At fixed winner share α, a binding common budget implies $\beta_0 + \beta_1 = R$, $R := \frac{B}{\alpha}$.

Write $x := \beta_1/R$ for augmentation's expenditure share. A direct prize primarily pays participants who must meet the cutoff. Augmentation applies to every winner: it raises the return to skill before the award and then multiplies final skill. It is therefore most attractive when winners already bring substantial prior skill.

Proposition 4 (Three-region allocation with productive augmentation)**.** *Suppose talent is Lomax with $\gamma > 2$. Fix $\alpha \in (0,1)$, $r_w > 0$, and awards $\beta_0, \beta_1 \geq 0$ that are not both zero. Let*

$$a := 1 + \beta_1, \qquad t := \alpha^{-1/\gamma} - 1, \qquad D := \sqrt{2\beta_0 + (a^2 - 1) r_w^2 t^2},$$

$$c := a r_w t^2 + tD, \qquad h := \sqrt{\frac{c}{a r_w}}.$$

Then $t < h$, the winning set is exactly $[t, \infty)$, and under the greatest-best-response convention the chosen and post-award skill schedules are

$$\big(s^*(\theta), s^+(\theta)\big) = \begin{cases} (r_w\theta^2, r_w\theta^2), & \theta < t, \\ (c, ac), & t \leq \theta < h, \\ (a r_w \theta^2, a^2 r_w \theta^2), & \theta \geq h. \end{cases}$$

Thus marginal winners bunch at the cutoff while sufficiently talented winners choose their unconstrained augmented skill.

The appendix proves the cutoff formula by a single-crossing comparison between losing at baseline and winning at the threshold.

Next, we calculate the integral of the total tournament worth, then deliver the first-dollar comparison.

Proposition 5 (First-dollar allocation at fixed breadth)**.** *Fix (α, γ) and $r_w > 0$. As $B \to 0$, the first award dollar goes to augmentation exactly when*

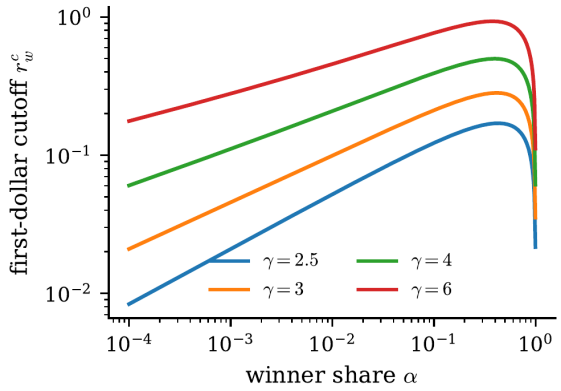
$$r_w > r_w^c(\alpha, \gamma) := \left[t^2 + \frac{4T_2(t)}{f(t)t} \right]^{-1/2},$$

where $t = \alpha^{-1/\gamma} - 1$, $f(t) = \gamma(1+t)^{-(\gamma+1)}$ is the Lomax density at the marginal winner, and $T_2(t) = \int_t^\infty \theta^2 \, \mathrm{d}F(\theta)$. At equality, the two instruments tie to first order.

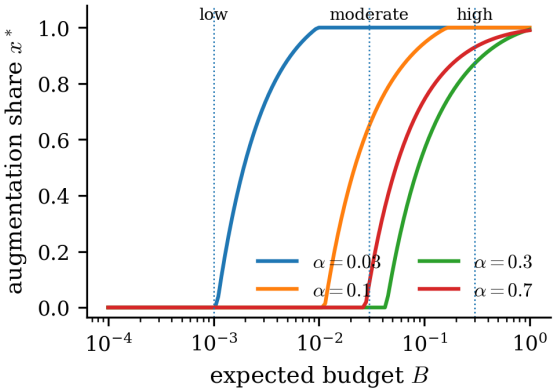
Figure 4(a) plots this first-dollar cutoff. It is low under very narrow recognition because selected winners are unusually skilled, and it falls again near universal recognition because augmentation then reaches almost the entire baseline skill stock. Direct prizes are relatively strongest at intermediate breadth. A thinner right tail raises the cutoff by reducing the upper-tail skill available to augment.

Panel (b) illustrates that non-monotonicity. At $r_w = .1$ and $\gamma = 3$, the broad policy $\alpha = .7$ starts using augmentation at $B \simeq .0275$, before $\alpha = .3$ at $B \simeq .0434$, because augmentation is again attractive when almost everyone is recognized. Yet $\alpha = .3$ reaches all augmentation first, at $B \simeq 1.18$ rather than 2.44. The reason is that the same expected budget delivers the larger per-winner amount $R = B/\alpha$ to the narrower group.

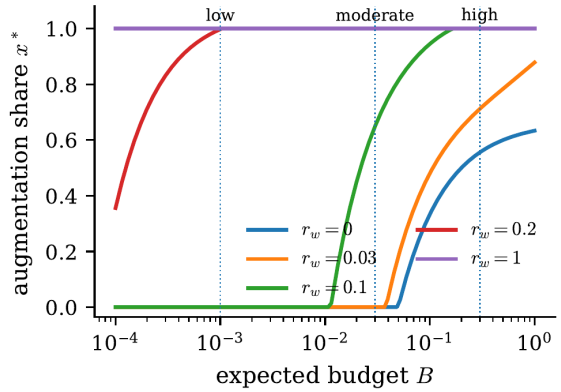
Panel (c) clarifies the role of the baseline return. Increasing r_w lowers the incremental value of a direct-prize tournament, as Figure 3(c) shows, but it simultaneously gives augmentation a larger



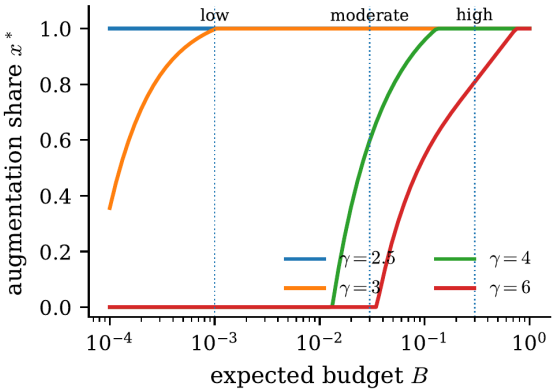
(a) First-dollar boundary.



(b) Varying α at $r_w = .1$, $\gamma = 3$.



(c) Varying r_w at $\alpha = .1$, $\gamma = 3$.



(d) Varying γ at $\alpha = .1$, $r_w = .2$.

Figure 4: Reward allocation with deterministic recognition and unit-linear augmentation cost. Panel (a) gives the first-dollar boundary; augmentation is preferred above each curve. Panels (b)–(d) report numerical global searches over $x \in [0, 1]$. Because $R = B/\alpha$, panel (b) varies both recognition breadth and reward intensity per winner, so its transition paths need not be ordered monotonically in α. The vertical guides mark representative low, moderate, and high expected budgets.

stock of pre-existing skill to scale. Hence the reward mix tilts toward β_1: at $r_w = 0$ a direct prize is needed to create the initial skill base, while at $r_w = 1$ the displayed solutions use only augmentation. Panel (d) gives the analogous tail comparison: a lower γ leaves more high-skill winners to augment and moves the switch toward β_1 to smaller budgets.

The calibration $(\alpha, r_w, \gamma) = (.1, .1, 3)$ displays all three regimes: all direct prize for $B \leq .01132$, a mixed award for $.01132 < B < .16474$, and all augmentation thereafter. At $B = .03$ the augmentation share is .649, and at $B = .10$ it is .945.

At fixed breadth, the figures suggest that the linear-cost high-budget limit is all augmentation when $r_w > 0$, whereas at $r_w = 0$ the augmentation share converges to 2/3.

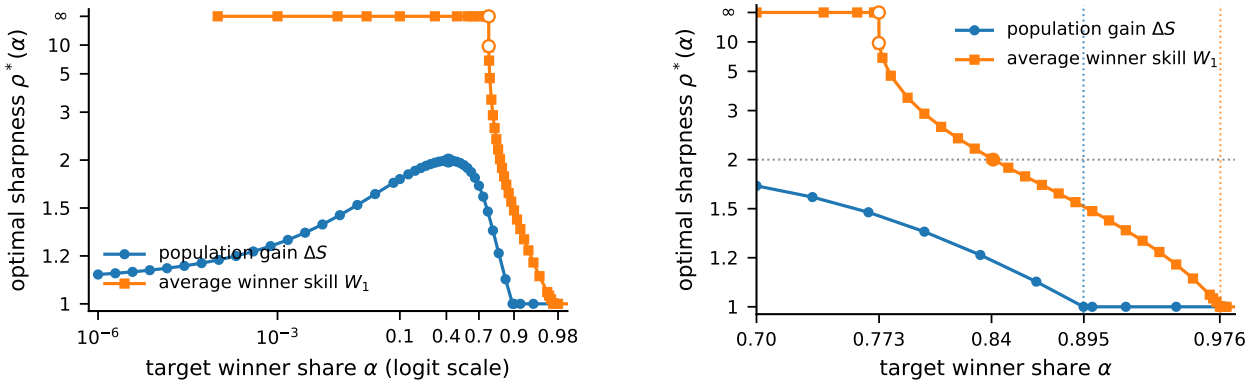
The fixed-breadth qualification is essential.[1] Thus Figure 4 is conditional on breadth. A well-posed joint design problem requires a lower bound on α, a cap or bounded augmentation technology, or a bounded skill payoff.

4.3 Assessment sharpness at fixed recognition breadth

We now return to the direct-prize case, $\beta_1 = 0$, and hold the target winner share α fixed. The binding budget sets $\beta_0 = B/\alpha$, while $s_c(\alpha, \rho)$ adjusts so that $\mathbb{E}[q] = \alpha$. We compare population skill creation with average final skill among realized winners:

$$\Delta S(\alpha, \rho) = \mathbb{E}\left[s^*(\theta) - r_w\theta^2\right], \qquad W_1(\alpha, \rho) = \frac{\mathbb{E}[q(\theta)s^*(\theta)]}{\alpha}.$$

The second objective rewards both induced development and the selection of high-skill participants into winner status.



(a) Full winner-share range.

(b) High-share comparison.

Figure 6A: Optimal assessment sharpness at fixed recognition breadth, $B = r_w = 1$ and $\gamma = 3$. Population-optimal sharpness peaks near $(\alpha, \rho) = (.419, 1.993)$ and reaches $\rho = 1$ near $\alpha = .895$. Average winner skill remains maximized by exact deterministic recognition until $\alpha = .773$, then switches to a finite rule with $\rho \simeq 9.63$, crosses $\rho = 2$ near $\alpha = .841$, and reaches $\rho = 1$ near $\alpha = .976$.

At $\alpha = .01$, ΔS peaks at $\rho = 1.4367$ while W_1 favors determinism; eventual nonwinners produce 78.7 percent of created skill at the population optimum. Thus ΔS values incentive reach, whereas W_1 also rewards selection. Figure 6A shows the divergence across α: population-optimal sharpness is hump-shaped, while the winner-skill organizer chooses a sharper rule for longer. The switch points are calibrated numerical conclusions. As $\alpha \to 1$, ΔS and $W_1 - r_w\mathbb{E}[\theta^2]/\alpha$ both equal $B(2-\rho)(1-\alpha)/(3r_w) + o(1-\alpha)$ for fixed $1 \le \rho < 2$, and are lower order for fixed $\rho \ge 2$. Under the corresponding uniform-argmax condition, both objectives therefore select $\rho^*(\alpha) \to 1$.

[1] *Proposition 6 (the infinitesimal-breadth edge case).* With $u(s) = s$, $r_w > 0$, and uncapped linear-cost augmentation, allowing $\alpha \to 0$ makes the joint problem unbounded; see the companion appendix for the proof.

5 Empirical Evidence

5.1 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.[2][3] 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 5.4.

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 10 displays the distribution of standardized scores for each event at the district level, pooling across all years and conferences.

[2]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 5.3.

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

5.2 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. The model's threshold s_c is the minimum score required to receive the prize. Empirically this threshold is revealed by the competitive margin: a team advances if and only if its score exceeds that of the strongest team it competes against. For each district-level competition cell c (a district-event-year), and for each school i in that cell, we define the skill threshold facing school i as the highest standardized score among all teams other than i:

$$s_{c,i} = \max_{\substack{j \in \mathcal{S}_c \\ j \neq i}} z_{jc}, \tag{3}$$

where $\mathcal{S}_c$ denotes the set of all teams in cell c and z_{jc} is team j's standardized score. Equivalently, $s_{c,i}$ is the score school i would have had to beat in order to advance, holding the rest of the district fixed. Excluding school i's own teams ensures that $s_{c,i}$ reflects only the competitive environment a school confronts and not its own performance; for a school that itself advanced, this leave-one-out construction returns the score of the runner-up it displaced. Figure 11 displays the distribution of $s_{c,i}$ across district-event-year cells. The substantial variation reflects differences in district competitiveness, which is key for our identification strategy in Section 5.3.

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:

$$\rho_a = \text{Corr}(z_{i,D,t},\, z_{i,R,t}), \tag{4}$$

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:

$$\rho_b = \text{Corr}(z_{i,e,t},\, z_{i,e,t-1}), \tag{5}$$

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 13 displays both measures.

5.3 Shifting Proximity to 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 competitive environment each school faced. Because the advancement threshold s_c defined above is a key lever governing equilibrium effort and skill production in the model, and because the model predicts that effort incentives are sharpest for competitors whose skill sits near s_c, the relevant object is not the level of the threshold alone but each school's proximity to it. The restructuring moved some schools closer to the threshold they had to clear to advance to the regional competitions and others farther away, providing identifying variation in exactly the quantity the model places at the center of the effort decision.

To measure this shift, we compare a counterfactual threshold to the school's own predetermined performance. For each school i and event e, let $\bar{z}_{ie}$ denote its average standardized team score over the pre-reform period (2004–2014), which we take as a measure of the school's position in the skill distribution. Using the leave-one-out threshold $s_{c,i}$ from the preceding subsection, we compute the average threshold $\bar{s}^{\text{pre}}_{c,ie}$ that school i actually faced among the schools it competed against before the reform, and the average threshold $\bar{s}^{\text{cf}}_{c,ie}$ it *would have faced* over the same pre-reform years had it competed against the schools it is matched with after the reform. Because both thresholds are built from pre-reform performance distributions, variation in their difference reflects only the change in competitive groupings and not any endogenous response by students or schools.

Our treatment is the reform-induced change in the school's distance to the threshold:

$$\Delta\hat{d}_{ie} = \left|\bar{z}_{ie} - \bar{s}^{\text{pre}}_{c,ie}\right| - \left|\bar{z}_{ie} - \bar{s}^{\text{cf}}_{c,ie}\right|, \tag{6}$$

where $\Delta\hat{d}_{ie} > 0$ indicates that the restructuring moved school i *closer* to the threshold it must clear to advance–the average highest score it must beat–and $\Delta\hat{d}_{ie} < 0$ indicates it moved farther away. Because $\bar{z}_{ie}$ is measured from predetermined performance and both thresholds use pre-reform distributions, $\Delta\hat{d}_{ie}$ isolates the exogenous change in proximity generated by the reassignment. The distribution of $\Delta\hat{d}_{ie}$ is roughly symmetric around zero (Figure 5). Identification therefore draws on schools moved away from the threshold and schools moved toward it in comparable measure, alongside a mass of cells near zero whose competitive environment was largely unchanged.

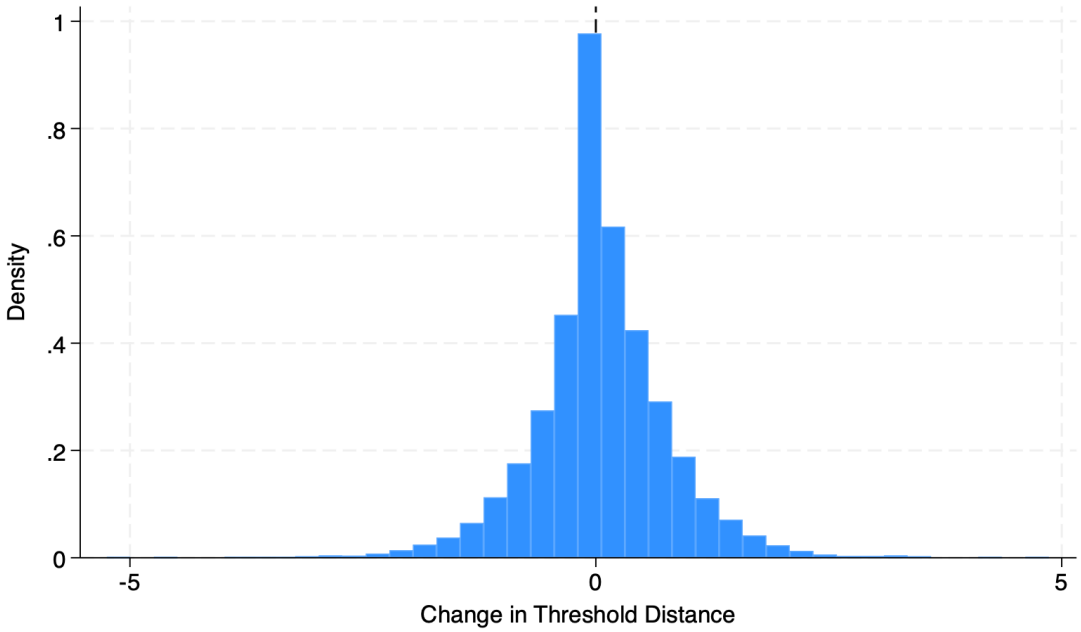


Figure 5: This figure plots the distribution of $\Delta\hat{d}_{ie}$ across school-event cells whose district composition changed in the 2015 restructuring. Positive values indicate that the reform moved a school closer to the threshold it must clear to advance. Both thresholds are leave-one-out and constructed from pre-reform scores, so variation reflects only the change in competitive groupings. The dashed line marks zero.

We estimate the effect of this proximity shift on school performance using the specification

$$y_{iet} = \beta \cdot \Delta\hat{d}_{ie} \cdot \mathbf{1}[t \geq 2015] + \gamma_{ie} + \delta_{et} + \varepsilon_{iet}, \tag{7}$$

where y_{iet} is one of four school-event-year outcomes (team advancement, the number of individual advancers, the team's standardized score, and the average standardized score of the school's individual entrants), γ_{ie} denotes school-by-event fixed effects, δ_{et} denotes event-by-year fixed effects, and standard errors are clustered by school. The sample runs from 2004 to 2022 and drops 2020 because no UIL meet was held. We restrict attention to schools whose district composition

actually changed in 2015. Both $\Delta\hat{d}_{ie}$ and the standardized outcomes are expressed in units of the within-event-year score distribution, so β is the effect of a one-unit move toward the threshold.

The two advancement outcomes verify that the treatment operates on the margin we claim: a school drawn closer to the threshold should advance more often and place more individual students at the regional competition. These responses are partly mechanical, and we read them as a check on the construction rather than as evidence of skill production. The performance outcomes carry the behavioral content. Scores are standardized within event-year across all districts rather than within district, so a change in district composition does not mechanically move a school's standardized score; movement in y_{iet} requires a change in what the school's students actually do. The model predicts $\beta > 0$ for these outcomes.

We also implement the event study analog:

$$y_{iet} = \sum_{k \neq 2014} \beta_k \cdot \Delta\hat{d}_{ie} \cdot \mathbf{1}[t = k] + \gamma_{ie} + \delta_{et} + \varepsilon_{iet}, \tag{8}$$

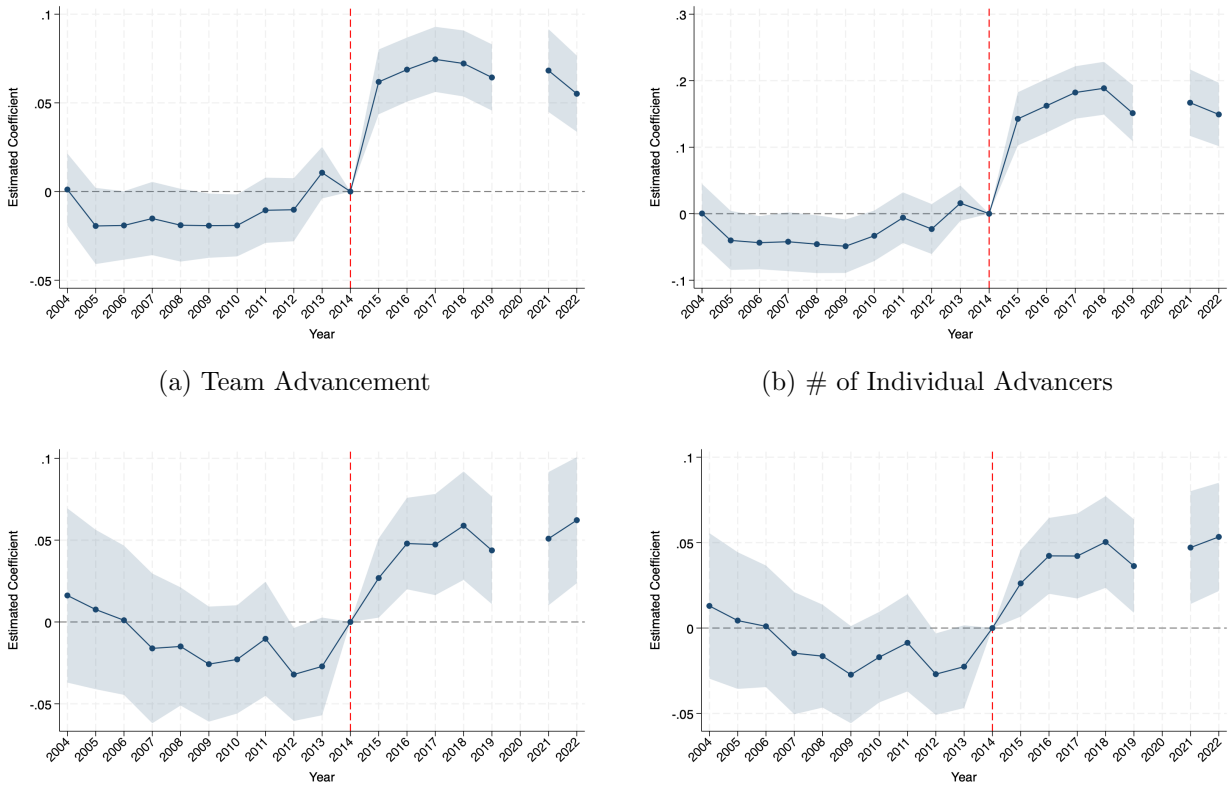
where k ranges over 2004 to 2022 excluding 2020, and 2014 is the omitted reference year: the last competition held under the old district assignments. The identifying assumption is that, absent the reform, schools facing larger moves toward their thresholds would have followed the same performance path as schools facing smaller moves, conditional on the fixed effects. A flat profile of β_k through 2014 supports it.

The estimated coefficients from Equation 7 are 0.053 (s.e. 0.011) for average individual performance and 0.059 (0.014) for team performance, so a one-unit move toward the threshold raises school performance by roughly 0.05 to 0.06 standard deviations of the within-event-year score distribution. The advancement outcomes move in the same direction, 0.077 (0.006) for team advancement and 0.186 (0.013) for the number of individual advancers, and all four estimates are significant at the one percent level. The event study shows a flat pre-trend and a level shift beginning in 2015 that persists through the end of the sample. One qualification is that the realignment was announced a few months before the 2014 competition, so the omitted year may already reflect an anticipatory response. If schools moved closer to their thresholds began exerting more effort in 2014, the pre-reform coefficients are displaced downward relative to a clean baseline and the post-reform estimates understate the effect of the reform.

The specification above imposes that the response is linear in $\Delta\hat{d}_{ie}$, which the model does not require. The marginal incentive to build skill depends on the probability of winning through the term $q(1-q)$, so it is small for schools far from the threshold in either direction and largest for schools near it, and a linear treatment averages over that curvature. To relax the functional form, we partition school-event cells into five bins of $\Delta\hat{d}_{ie}$ and estimate

$$y_{iet} = \sum_{b \neq 3} \sum_{k \neq 2014} \beta_{bk} \cdot \mathbf{1}[(i,e) \in \mathcal{B}_b] \cdot \mathbf{1}[t = k] + \gamma_{ie} + \delta_{et} + \varepsilon_{iet}, \tag{9}$$

where the bins are $\mathcal{B}_1 = (-\infty, -0.75]$, $\mathcal{B}_2 = (-0.75, -0.25]$, $\mathcal{B}_3 = (-0.25, 0.25]$, $\mathcal{B}_4 = (0.25, 0.75]$,



(a) Team Advancement

(b) # of Individual Advancers

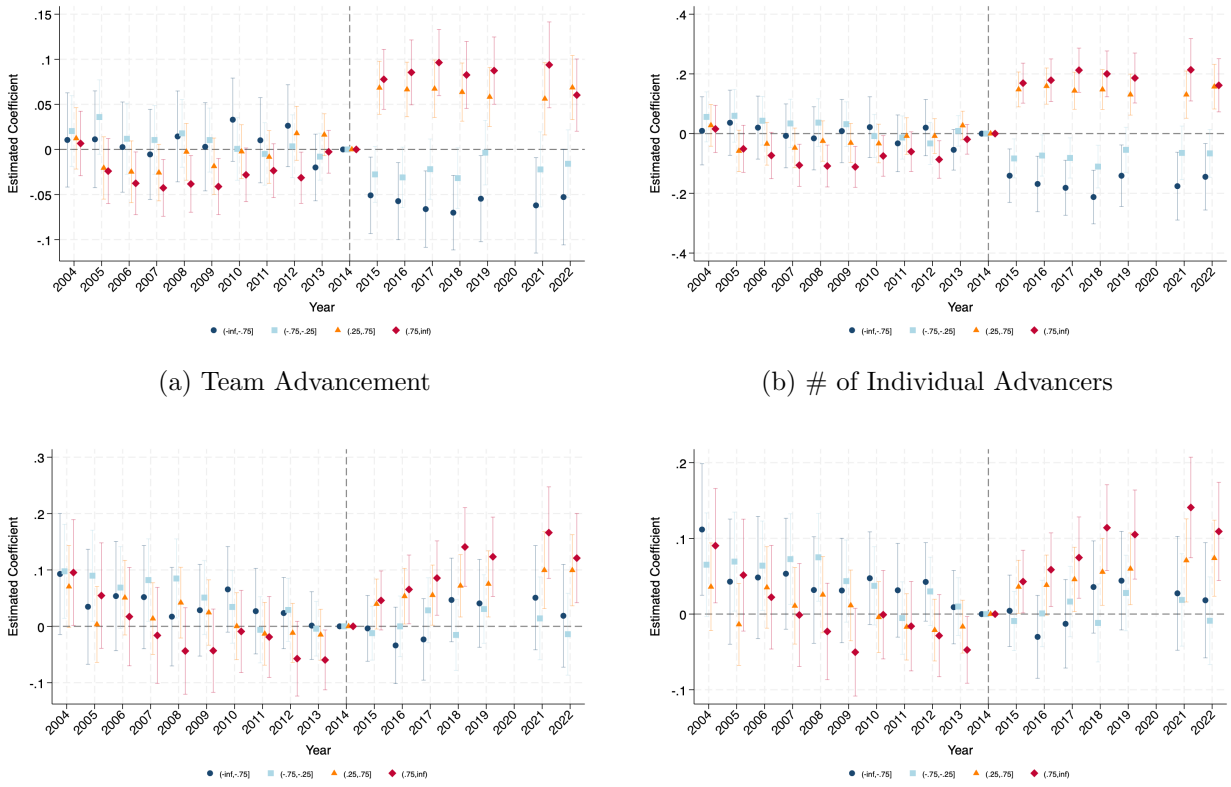
(c) Team Performance

(d) Average Individual Performance

Figure 6: This figure plots β_k from the event-study specification, interacting $\Delta\hat{d}_{ie}$ with each competition year. The sample is school-event-year observations from 2004 to 2022 for schools whose district composition changed in 2015; 2020 is omitted because no UIL meet was held. Specifications include school-by-event and event-by-year fixed effects, with standard errors clustered by school. The omitted reference year is 2014, the last competition held under the old district assignments. Coefficients are per one-unit move toward the threshold in units of the within-event-year score distribution. Bands are 95 percent confidence intervals.

and $\mathcal{B}_5 = (0.75, \infty)$, and $\mathcal{B}_3$ is omitted. Cells in $\mathcal{B}_3$ experienced essentially no change in their competitive environment, so each β_{bk} measures the performance of schools in bin b relative to schools whose reassignment left their proximity to the threshold roughly unchanged.

Figure 7 reports the binned estimates. The post-reform coefficients are monotone in the bins for all four outcomes, so the linear specification is not concealing a non-monotonicity, and the dose-response pattern appears without being imposed. The responses also grow over the post-reform period rather than jumping to a constant level, consistent with skill accumulating over repeated years of competition. Schools moved farther away are indistinguishable from the reference bin in the performance outcomes while declining in advancement, which is what the model implies when nearly all schools sit below the threshold: $q(1-q)$ is small and flat far below the cutoff, so additional distance changes effort little, and the advancement decline is mechanical. The pre-reform coefficients are not exactly flat; for the performance outcomes the largest positive bin declines toward zero as the reference year approaches, a movement opposite in sign to the post-reform



(a) Team Advancement

(b) # of Individual Advancers

(c) Team Performance

(d) Average Individual Performance

Figure 7: This figure plots β_{bk} from the binned event-study specification, which replaces the continuous treatment with indicators for four bins of $\Delta\hat{d}_{ie}$ interacted with competition year. The omitted bin is $(-0.25, 0.25]$, comprising schools whose proximity to the threshold was largely unchanged by the reform. Coefficients are therefore differences relative to those schools rather than per-unit effects. Sample, fixed effects, clustering, and reference year are as in Figure 6. Markers are offset horizontally within each year so that confidence intervals do not overlap. Bars are 95 percent confidence intervals.

response.

5.4 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

$$z_{iet} = \sum_k \beta_k \cdot \mathbf{1}[\text{bin}_k(\hat{s}_{c,i} - z_{i,\text{pre}})] \cdot \mathbf{1}[t \geq 2023] + \gamma_i + \delta_{et} + \varepsilon_{iet}, \tag{10}$$

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.

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 $s(\theta) \approx r_w \theta^2$.

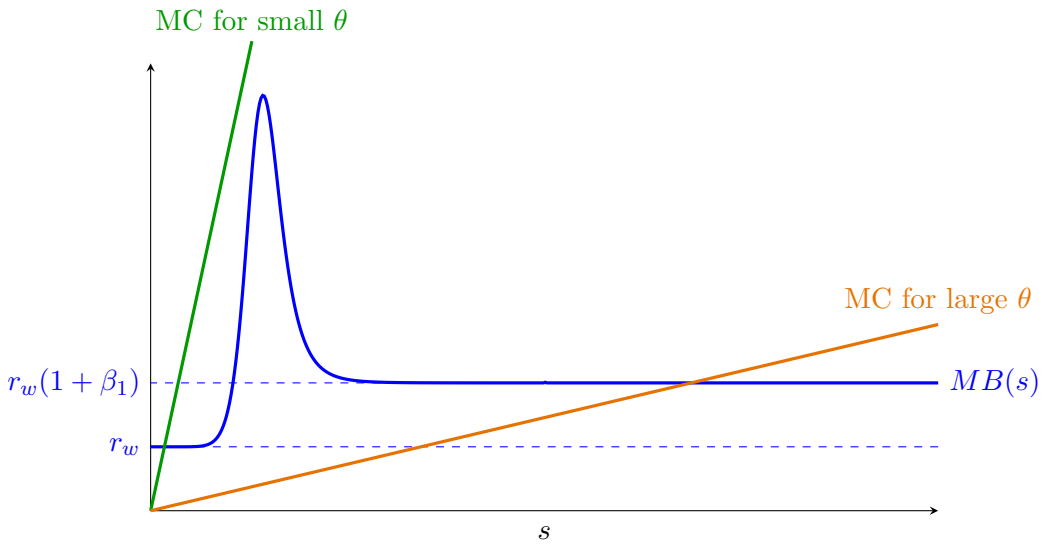


Figure 8: Changing θ: MB is fixed; the slope of MC changes.

- Suppose $\theta \to \infty$. 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 $\theta \to \infty$ the intersect of the marginal cost and the marginal benefit goes to infinity, which means $s(\theta) \to \infty$.

  As $s \to \infty$ 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 $\theta \to \infty$ 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 8 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 9 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).

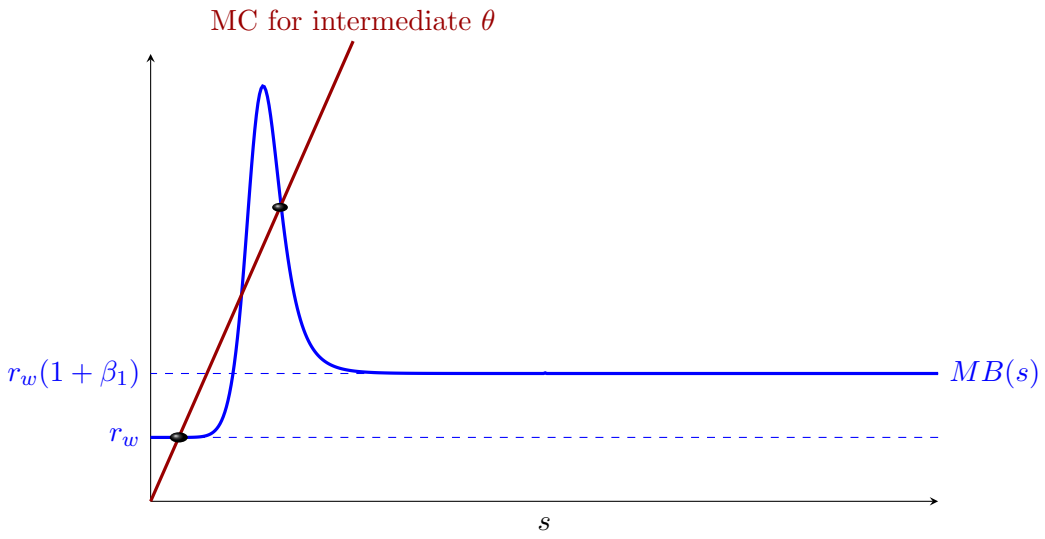


Figure 9: Two Possible Optimal Skill Choices

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

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

Consequently we have

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

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

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

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

$$\begin{aligned} MB(s) &= r_w + \beta_0 \cdot \frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} + \beta_1 r_w \cdot \left(\frac{s^\rho}{s^\rho + s_c^\rho} + \frac{\rho s_c^\rho s^\rho}{(s^\rho + s_c^\rho)^2} \right). \\ &\approx r_w + \beta_0 \frac{\rho}{s_c^\rho} \cdot s^{\rho-1} + \beta_1 r_w \frac{1+\rho}{s_c^\rho} \cdot s^\rho. \end{aligned} \tag{11}$$

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

$$\begin{aligned} MB(s) &= r_w + \beta_0 \cdot \frac{\rho s_c^\rho s^{\rho-1}}{(s^\rho + s_c^\rho)^2} + \beta_1 r_w \cdot \left(\frac{s^\rho}{s^\rho + s_c^\rho} + \frac{\rho s_c^\rho s^\rho}{(s^\rho + s_c^\rho)^2} \right). \\ &\approx r_w(1+\beta_1) + \beta_0 \frac{\rho s_c^\rho}{s^{\rho+1}} + \beta_1 r_w(\rho - 1) \cdot \frac{s_c^\rho}{s^\rho}. \end{aligned} \tag{12}$$

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 $s_c < s_c'$. We claim that if we increase s_c to s_c', 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' we have $MB(s)$ decreases for sufficiently small s, and $MB(s)$ increases for sufficiently large s. Indeed as Figure 8 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 (11) 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 (12) 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.

A.4 Proof of Proposition 3

Set $r_w = \beta_1 = 0$ and let $\beta_0 = B/\alpha$. A loser maximizes $-s^2/(2\theta^2)$ and therefore chooses zero. Conditional on winning, the cheapest feasible action is $s = c$, so a type wins precisely when

$$\beta_0 - \frac{c^2}{2\theta^2} \geq 0.$$

This gain is increasing in talent. Under the greatest-best-response convention the winning set is the top α tail, whose marginal type is

$$t = t_\alpha = \alpha^{-1/\gamma} - 1.$$

Marginal indifference gives

$$c = t\sqrt{\frac{2B}{\alpha}}.$$

Hence the selected schedules are

$$s^*(\theta) = \begin{cases} 0, & \theta < t, \\ c, & \theta \geq t, \end{cases} \qquad e^*(\theta) = \begin{cases} 0, & \theta < t, \\ c/\theta, & \theta \geq t. \end{cases}$$

Both schedules have a genuine jump. At $\theta = t$ the lower and upper actions tie; changing the tie-breaking convention changes that one action but no aggregate object.

Every winner produces c, so

$$S_\infty(\alpha) = \alpha c = \sqrt{2B}\left(\alpha^{1/2-1/\gamma} - \alpha^{1/2}\right).$$

Differentiation yields

$$S'_\infty(\alpha) = \sqrt{2B}\,\alpha^{-1/2}\left[\frac{\gamma-2}{2\gamma}\alpha^{-1/\gamma} - \frac{1}{2}\right].$$

The bracket is strictly decreasing, positive near zero, and negative near one. Its unique zero is

$$\alpha^*_\infty(\gamma) = \left(\frac{\gamma-2}{\gamma}\right)^\gamma.$$

To establish the tail comparative static, differentiate its logarithm:

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\log\alpha^*_\infty(\gamma) = \log\left(1 - \frac{2}{\gamma}\right) + \frac{2}{\gamma-2}.$$

Putting $z = 2/(\gamma-2) > 0$ turns the right side into $z - \log(1+z) > 0$. Thus $\mathrm{d}\alpha^*_\infty/\mathrm{d}\gamma > 0$. The limits follow directly from the closed form:

$$\alpha^*_\infty(\gamma) \to 0 \quad (\gamma \to 2), \qquad \alpha^*_\infty(\gamma) \to e^{-2} \quad (\gamma \to \infty).$$

A.5 Proof of Proposition 4

A type's deterministic payoff is

$$U(s;\theta) = r_w s - \frac{s^2}{2\theta^2} + \mathbf{1}\{s \geq c\}\,[\beta_0 + \beta_1 r_w s].$$

Assume $r_w > 0$ and $(\beta_0, \beta_1) \neq (0, 0)$. Conditional on losing, the unconstrained optimum is the baseline skill $r_w\theta^2$. Conditional on winning, the objective is strictly concave and its constrained optimum is

$$s_W(\theta) = \max\{c, (1 + \beta_1)_w\theta^2\}.$$

Set $h = \sqrt{c/(ar_w)}$ and $a = 1 + \beta_1 \geq 1$. For $\theta < h$, baseline skill lies below the cutoff, and the utility gain from winning at exactly c rather than remaining at baseline is

$$\Phi(\theta) = \beta_0 + ar_w c - \frac{c^2}{2\theta^2} - \frac{r_w^2\theta^2}{2}.$$

On $(0, h)$,

$$\Phi'(\theta) = \frac{c^2}{\theta^3} - r_w^2\theta > 0,$$

because $\theta < h$ implies $r_w^2\theta^4 < c^2/a^2 \leq c^2$. Moreover $\Phi(\theta) \to -\infty$ as $\theta \to 0$, whereas

$$\Phi(h) = \beta_0 + \frac{a^2 - 1}{2a} r_w c > 0.$$

Hence the comparison below h has a unique crossing. For $\theta \geq h$, the winning optimum $ar_w\theta^2$ is feasible and exceeds the unconstrained no-award value by

$$\beta_0 + \frac{(a^2 - 1)r_w^2\theta^2}{2} > 0;$$

any losing action is no better than that no-award value. The winning set is therefore an upper tail. Target-share clearing sets its marginal type at

$$t = \alpha^{-1/\gamma} - 1 < h.$$

Marginal indifference, $\Phi(t) = 0$, gives

$$c^2 - 2ar_w t^2 c + r_w^2 t^4 - 2\beta_0 t^2 = 0.$$

The relevant root is the one satisfying $c > ar_w t^2$:

$$D := \sqrt{2\beta_0 + (a^2 - 1)r_w^2 t^2}, \qquad c = ar_w t^2 + tD, \qquad h = \sqrt{\frac{c}{ar_w}}. \tag{13}$$

Because $D > 0$, this root indeed gives $t < h$. Types below t choose baseline skill and lose; types in $[t, h)$ win by choosing exactly c; and types above h choose the unconstrained winning optimum $ar_w\theta^2$. Multiplying winners' chosen skill by a gives the post-award schedule stated in the proposition.

This completes the proof.

A.6 Proof of Proposition 5

Fix (α, r_w, γ) with $\alpha \in (0,1)$ and $r_w > 0$, and write $b = \beta_0$, and $a = 1 + \beta_1$. For Lomax talent, $f(z) = \gamma(1+z)^{-(\gamma+1)}$. From (13),

$$D^2 = 2\beta_0 + (a^2 - 1)r_w^2 t^2, \qquad h^2 = t^2 + \frac{tD}{ar_w}.$$

Let $\bar{F}(x) = (1+x)^{-\gamma}$, and write $T_2(x) = \int_x^\infty \theta^2 \, \mathrm{d}F(\theta)$. Integrating the three branches and subtracting no-tournament skill yields

$$V(\beta_0, \beta_1; \alpha) = ac\big[\alpha - \bar{F}(h)\big] - r_w T_2(t) + a^2 r_w T_2(h). \tag{14}$$

For Lomax talent,

$$T_2(x) = \frac{\gamma}{\gamma - 2}(1+x)^{2-\gamma} - \frac{2\gamma}{\gamma - 1}(1+x)^{1-\gamma} + (1+x)^{-\gamma}.$$

Using $ac = a^2 r_w h^2$, the exact objective $V(\beta_0, \beta_1; \alpha) = ac\big[\alpha - (1+h)^{-\gamma}\big] - r_w T_2(t) + a^2 r_w T_2(h)$ can be rewritten as

$$V(\beta_0, \beta_1; \alpha) = (a^2 - 1) r_w T_2(t) + a^2 r_w \int_t^h (h^2 - \theta^2) \, \mathrm{d}F(\theta). \tag{15}$$

As $(\beta_0, \beta_1) \to (0,0)$ in the nonnegative quadrant, $D \to 0$ and $h \to t$. Continuity of the Lomax density at t gives

$$\int_t^h (h^2 - \theta^2) \, \mathrm{d}F(\theta) = f(t) t (h-t)^2 + o\big((h-t)^2\big).$$

Also

$$h - t = \frac{D}{2r_w} + o(D), \qquad a^2 - 1 = 2\beta_1 + o(\beta_1).$$

Substituting these expansions and then using the expression for D^2 yields

$$V(\beta_0, \beta_1; \alpha) = \underbrace{\frac{f(t)t}{2r_w}}_{\mu_0} \beta_0 + \underbrace{r_w \left[\frac{f(t)t^3}{2} + 2T_2(t) \right]}_{\mu_1} \beta_1 + o(\beta_0 + \beta_1). \tag{16}$$

Under the unit-linear common budget, let $R = B/\alpha$ and set

$$\beta_0 = (1-x)R, \qquad \beta_1 = xR, \qquad x \in [0,1].$$

Because the remainder in (16) is little-o in $\beta_0 + \beta_1 = R$, uniformly over this compact simplex,

$$V(R, x) = R\big[(1-x)\mu_0 + x\mu_1\big] + o(R) \quad \text{uniformly in } x \in [0,1].$$

Thus a strict inequality $\mu_1 > \mu_0$ sends every small-budget maximizer toward $x = 1$, while $\mu_1 < \mu_0$ sends it toward $x = 0$; at equality the two instruments tie to first order. Finally,

$$\mu_1 > \mu_0 \quad \Longleftrightarrow \quad r_w^2 \left[t^2 + \frac{4T_2(t)}{f(t)t} \right] > 1,$$

which, since $r_w > 0$, is equivalent to

$$r_w > \left[t^2 + \frac{4T_2(t)}{f(t)t} \right]^{-1/2}.$$

This proves the stated threshold.

A.7 Proof of Proposition 6

Fix $B > 0$, $r_w > 0$, and $\gamma > 2$, and consider the feasible pure-augmentation sequence

$$\beta_0 = 0, \qquad \beta_1 = \frac{B}{\alpha}, \qquad \alpha \to 0.$$

Let $a = 1 + B/\alpha \to \infty$ and $t = \alpha^{-1/\gamma} - 1$. With $b = 0$,

$$D = r_w t\sqrt{a^2 - 1}, \qquad \frac{c}{a r_w t^2} = 1 + \sqrt{1 - a^{-2}} \longrightarrow 2, \qquad \frac{h}{t}, \frac{1+h}{1+t} \longrightarrow \sqrt{2}.$$

Regular variation of the Lomax tail gives

$$\frac{\alpha - \bar{F}(h)}{\alpha} = 1 - \frac{\bar{F}(h)}{\bar{F}(t)} \longrightarrow 1 - 2^{-\gamma/2},$$
$$\frac{T_2(h)}{\alpha t^2} \longrightarrow \frac{\gamma}{\gamma - 2} 2^{1-\gamma/2}.$$

Dividing (14) by $a^2 r_w \alpha t^2$ therefore yields

$$\frac{V(0, B/\alpha; \alpha)}{a^2 r_w \alpha t^2} \longrightarrow K_\gamma := 2\big(1 - 2^{-\gamma/2}\big) + \frac{\gamma}{\gamma - 2} 2^{1-\gamma/2} > 0.$$

Finally,

$$a^2 \alpha t^2 \sim B^2 \alpha^{-1-2/\gamma}.$$

Hence

$$V(0, B/\alpha; \alpha) \sim K_\gamma r_w B^2 \alpha^{-1-2/\gamma} \to \infty,$$

which proves the unboundedness result.

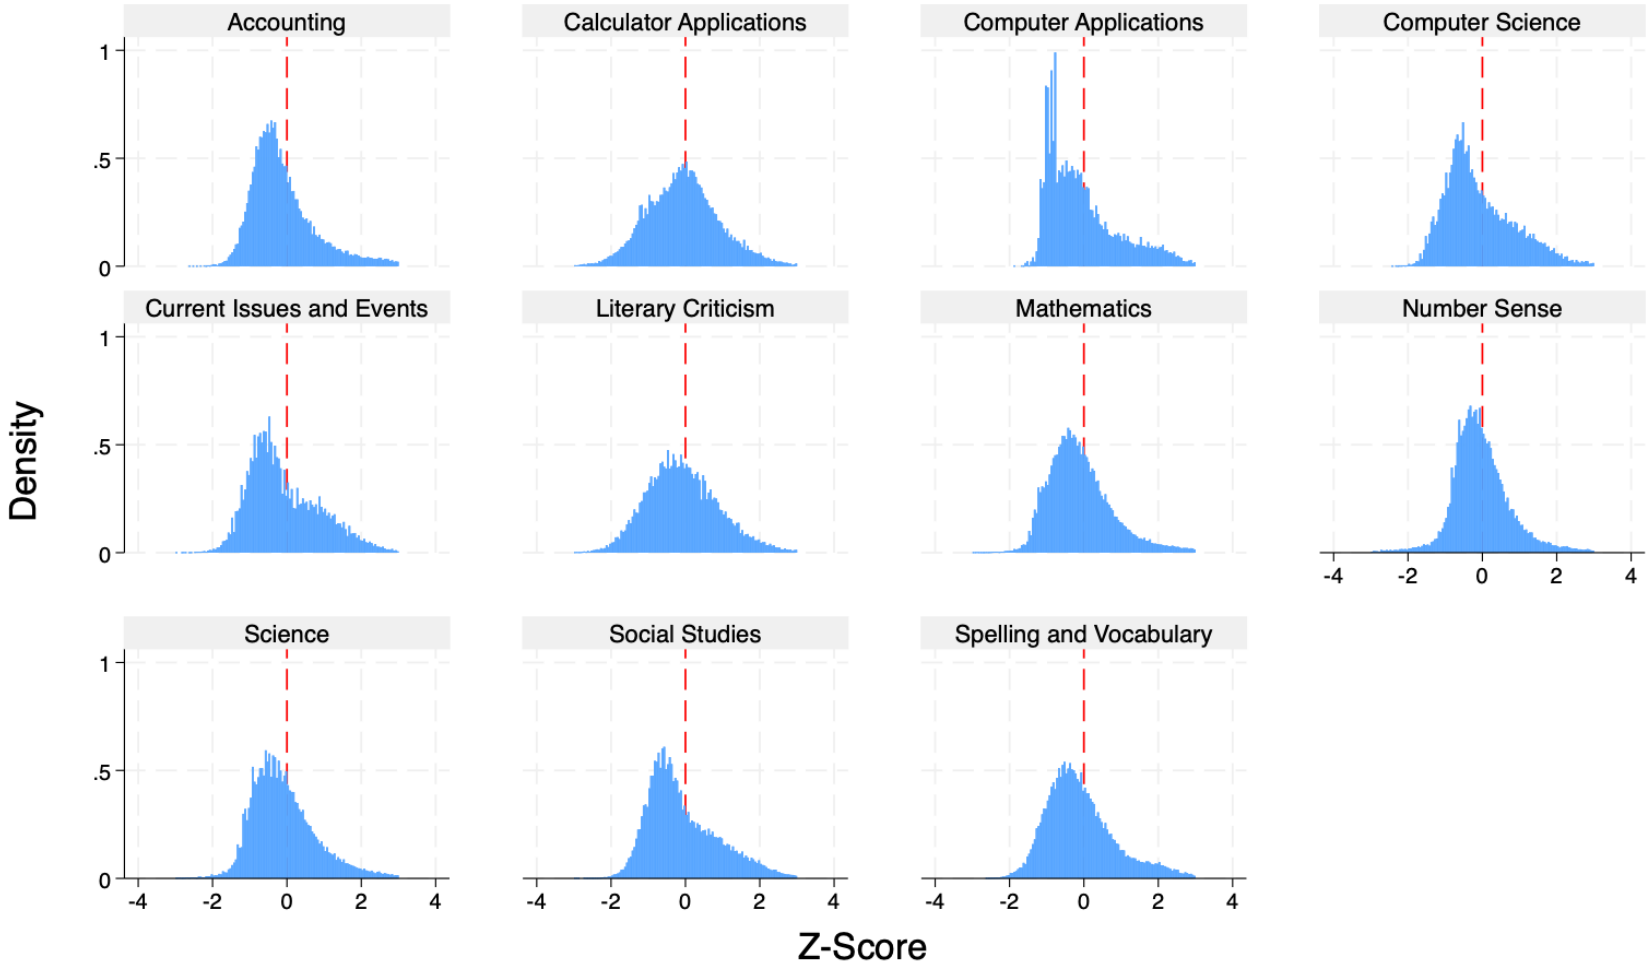
B Figures and Tables from the Empirical Section

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.

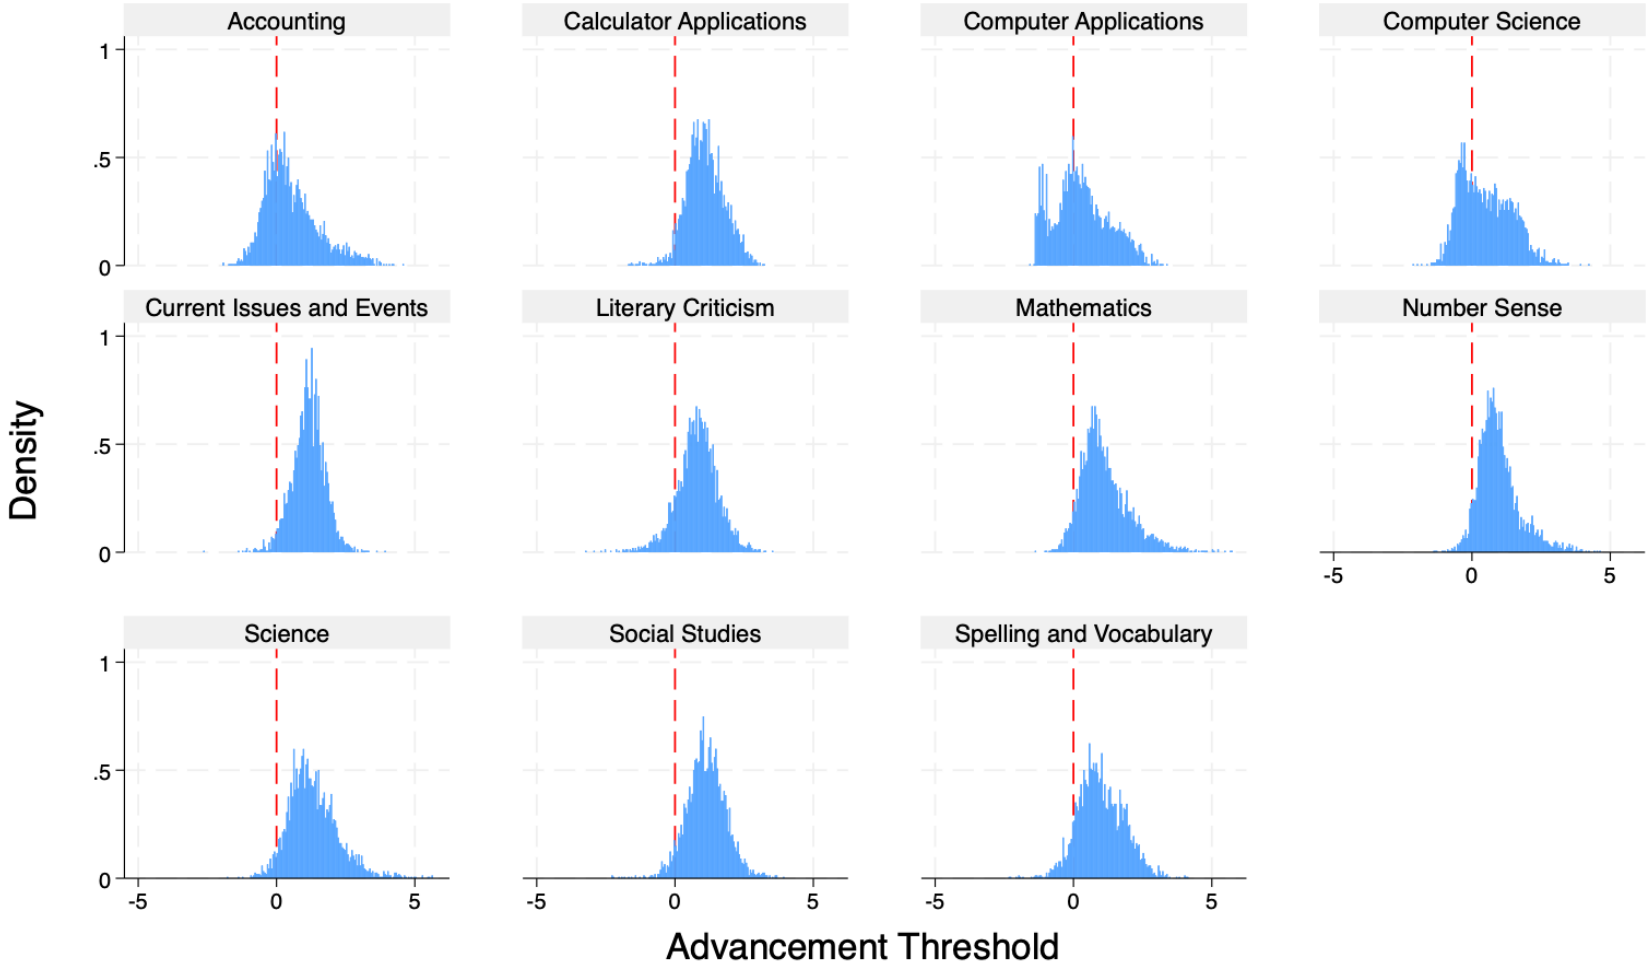
This table describes each of the high school UIL academic events analyzed in this paper, including the format, duration, item count, and content coverage of each contest.

Figure 10: UIL Performance Distributions



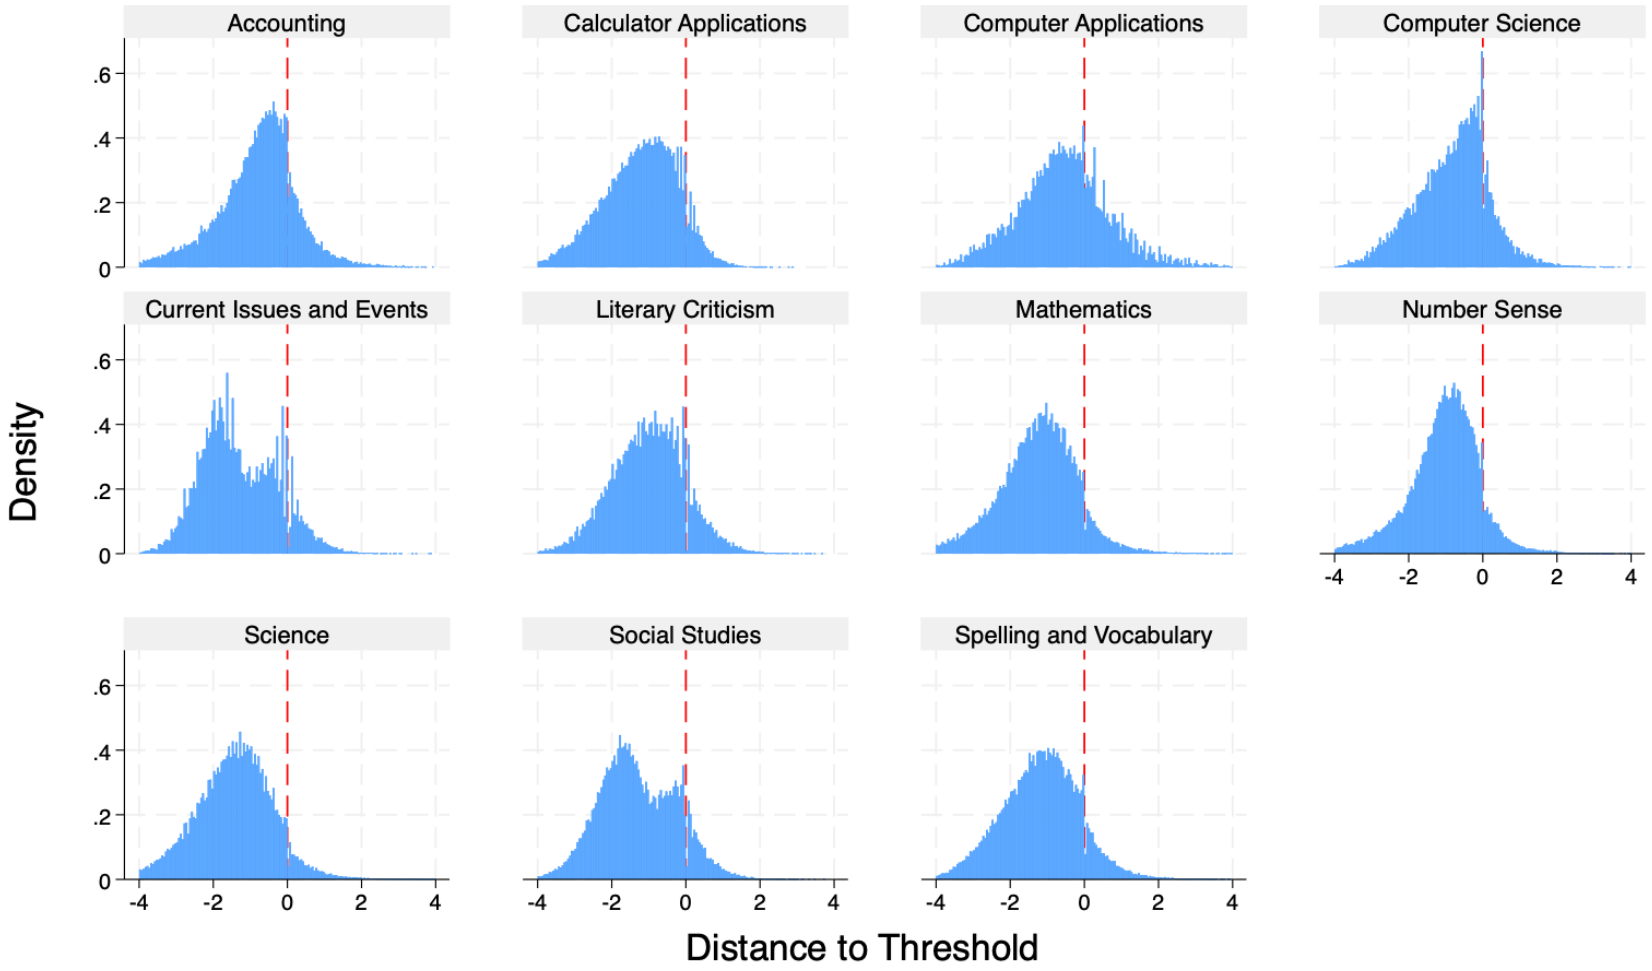
This figure plots the distribution of standardized district-level scores for each event, pooling across all years and conferences. Scores are standardized within each event-year cell. The dashed line marks the mean.

Figure 11: UIL District Advancement Thresholds



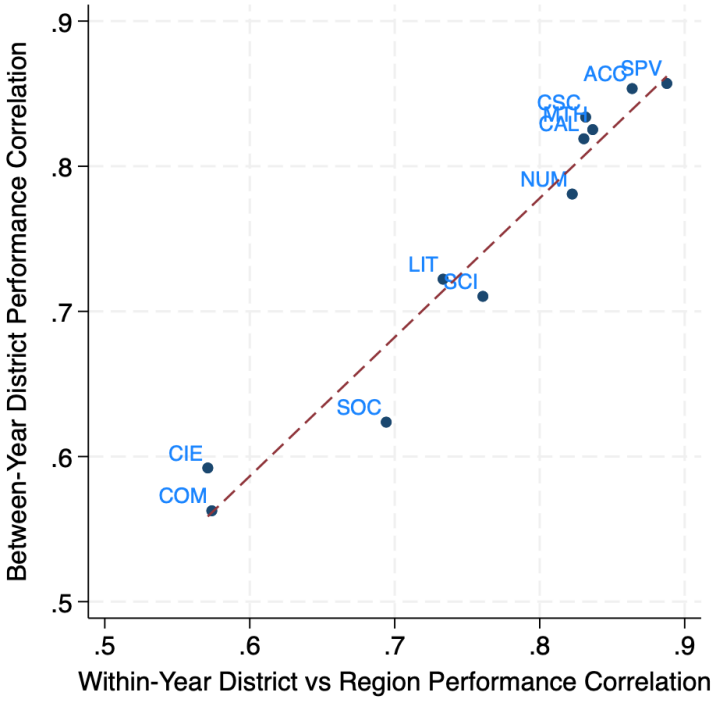
This figure plots the distribution of the implied advancement threshold $\hat{s}_c$ across district-year-event cells, defined as the standardized score of the lowest-performing advancer in each cell. Variation reflects differences in district competitiveness and participation rates.

Figure 12: Participant Distance to Advancement Threshold



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

Figure 13: 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).