On Benchmark Hacking in ML Contests: Modeling, Insights and Design*

Xiaoyun Qiu† Haifeng Xu‡ Yang Yu§

This draft: August 2025

Abstract

Benchmark hacking refers to tuning a machine learning model to score highly on certain evaluation criteria without improving true generalization or faithfully solving the intended problem. We study this phenomenon in a generic *machine learning contest*, where each contestant chooses two types of effort: *exploratory effort* that improves model capability as desired by the contest host, and *exploitative effort* that only improves the model's fitness to the particular task in contest without contributing to true generalization. We establish the existence of a symmetric monotone pure strategy equilibrium in this competition game. It also provides a natural definition of benchmark hacking in this strategic context by comparing a player's equilibrium effort allocation to that of a single-agent baseline scenario. Under our definition, contestants with types below certain threshold (*low* types) always engage in benchmark hacking, whereas those above the threshold do not. Furthermore, we show that more skewed reward structures (favoring top-ranked contestants) can elicit more desirable contest outcomes. We also provide empirical evidence to support our theoretical predictions.

Keywords— contests, benchmark hacking, gaming, machine learning

*Haifeng Xu is supported by AI2050 program at Schmidt Sciences (Grant G-24-66104), NSF through CCF-2303372 and ARO through W911NF-23-1-0030. We acknowledge Yingkai Li for helpful conversations at the early stage of this project. We also thank Ben Brooks, Kyohei Okumura, Abhishek Sarkar, Chenhao Zhang, Chunru Zheng and participants at 2024 DC IO Day for helpful comments and suggestions.

†Department of Economics, Northwestern University. Email: xiaoyun.qiu@u.northwestern.edu

‡Department of Computer Science, University of Chicago. Email: haifengxu@uchicago.edu

§Sloan School of Management & Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Email: yyu7@mit.edu

1 Introduction

With the pervasive application of machine learning techniques today, *machine learning (ML) contests* become a very popular way for firms or individuals to elicit high-quality ML methods from the machine learning community. Notable examples include platforms like Kaggle and AIcrowd. Moreover, ML contests are also popular among academics as a way to advance ML methodology development. The most notable example perhaps is the *ImageNet Challenge* (Deng et al., 2009), which attracted many computer vision researchers to compete on the common problem of image classification and sparked the revolutionary era of deep learning, started by the development of a novel ML model called *deep convolutional neural networks* that significantly beats all other ML models at the time (Krizhevsky et al., 2012). Finally, Ml content also happens in industry, though often implictly, such as the recent race in LLM model development (including the reported Llama 4's benchmark hacking (Ndzomga, 2025)).

Unlike other avenues of ML methodology development, machine learning contests raise the concern of *benchmark hacking* — the practice of optimizing only for contest rules through trial-and-error approaches but not truly improving a model's generalization ability or faithfully solving the underlying problem. From the perspective of the contest host (henceforth, the *principal*), she would want to elicit a model that delivers meaningful improvements, novel solutions, or even entirely new approaches to a problem that could potentially be applied to future problems.[1] However, in contests where prizes are awarded based on model performance on the current data set (due to the lack of perfect evaluation metrics), the contestants' direct incentives are often not the model's true capability, but rather how it fits the particular evaluation metric specified by the contest. This can be achieved not only through genuine model innovation but also through trial-and-error tactics, such as repeatedly tweaking hyperparameters and making numerous submissions to find the configuration that yields the highest score. Hence, compared to a single data analyst's ML model development, what is fundamentally different in ML contests is the mis-aligned incentives between the principal and contestants. As long as the principal cannot distinguish genuine innovation from benchmark hacking — a task typically infeasible given the high cost of manually reviewing hundreds or even thousands of scripts — this misalignment persists.

To study the above issue of benchmark hacking in machine learning contests, this paper puts forward a game-theoretic model for analyzing the potentially mis-aligned incentives between principal and contestants. Our model distinguishes two types of effort that a ML contestant could exert: (1) *exploratory effort* which delivers meaningful improvements and innovative solutions hence benefits the principal; for instance, exploratory effort could include the time a contestant spent on developing foundational innovation (e.g., proposing a new algorithmic framework) and creative

[1] There are typically two types of contests on Kaggle. In the first, the principal's primary objective is to solve a specific problem, focusing solely on eliciting a high-fitness model without concern for innovation or generalizability to other tasks. In the second, the true goal is to leverage the winning model for similar future tasks or to advance research within a broader community. This paper focuses on the latter type.

problem-specific adaptation (e.g., creating problem-specific features, designing a loss function that better reflects the nature of the problem, reformulating the problem to to capture its underlying structure); and (2) *exploitative effort* which improves the contest outcome through trial-and-error optimization using standard, plug-and-play machine learning techniques but does *not* benefit the principal; examples include the time spent on trying out default feature sets and model architectures without adapting them to the domain or data structure.[2][3] A benefit of modeling each contestant's effort allocation strategy is that it abstracts out concrete methodology development detail and allows us to focus our attention on analyzing the novel question of incentive conflict in the contest setup.

Let us use the famous ImageNet Challenge as an example contest to help readers better understand the two types of effort. From the very beginning, the goal of this challenge was clearly set as to advance computer vision methodology and achieve vision-level intelligence (Deng et al., 2009). The challenge tasked researchers with classifying objects (e.g., animal species) from a standardized dataset, the *ImageNet*, using accuracy as the common evaluation metric. In this challenge, exploitation efforts include hyperparameter tuning of known models (e.g., adjusting learning rates, batch sizes, regularization terms, or other hyperparameters for existing algorithms), combining off-the-shelf methods (e.g., using standard classification techniques such as random forests and logistic regression and combining their outputs via simple ensembling, without developing new architectures), and modifying image features superficially (e.g., resizing and rotating images). While these tactics earned top performers recognition, they contributed little to major innovations and fell short of the challenge's goals.

A turning point came in 2012, when Krizhevsky et al. (2012) introduced a highly original *deep convolutional neural network (CNN)* that boosted ImageNet accuracy to 85%, over 10% higher than the second-best of that year on the leaderboard (ImageNet, 2012). This type of effort that involves creativity and more generalizable is what we call exploratory effort. During the entire ImageNet Challenge spanning 8 years from 2010 to 2017, there were multiple milestones of novel methodological advances that overcame fundamental limitations of previous methods and led to more effective computer vision technologies, as the contest organizers had aimed for.

Developing a contest model to analyze benchmark hacking in ML competitions faces two main challenges. First, we need a realistic yet tractable game-theoretic framework to capture contestants' allocation of effort between exploration and exploitation. Second, we must formalize benchmark hacking and then study whether it arises at equilibrium. This is particularly subtle since benchmark hacking is a *relative* concept — some exploitation is natural and organic, but excessive focus on it at the expense of exploration harms innovation and generalizability — it is crucial to compare benchmark hacking outcomes against a well-defined baseline, which our model provides.

[2]Surveys show that many contestants do this by registering different accounts on the same platform, see https://image-net.org/challenges/LSVRC/announcement-June-2-2015.

[3]We also provide examples for exploratory and exploitative effort from scripts submitted to Kaggle in the empirical validation section below.

We develop a contest model for machine learning where contestants compete on the *performance* of their developed models for a prediction task.[4] Each player allocates effort between *exploration* and *exploitation*. While both efforts can improve contest performance on the specific prediction task and thus rewards, only exploration drives innovation and generalizability. For this reason, the principal only benefits from the exploration. Therefore, the principal faces a moral hazard problem (Grossman and Hart, 1992; Holmström, 1979) — that is, she cannot distinguish between the two types of contestant efforts and can only observe a signal about the overall *fitness* of the model to the particular contest task, which is an aggregated outcome of both efforts.

Each contestant has a private type, interpreted as the *capability* for model development. While exploitation improves any contestant's contest reward in the same way, exploration improves a contestant's reward differently for different types. In particularly, under the same level of exploration, more capable contestants develop a model with higher quality. This heterogeneous capability in model development leads to different effort allocation strategies for different types of contestants at the contest's Bayesian Nash equilibrium. Under reasonable assumptions, we prove the existence of a symmetric monotone pure strategy Nash equilibrium.

We then apply our model to study benchmark hacking in machine learning contests. Our fundamental hypothesis is that every contestant is mature enough such that if the learning task is appointed solely to this contestant alone, then this contestant's optimal choice of exploration and exploitation will be an organic baseline with no benchmark hacking. This hypothesis is based on the fact that most contestants on platforms like Kaggle and AIcrowd are mature data scientist and machine learning researcher, hence their optimal effort allocation should be appropriate if they were assigned to complete a task alone in absence of any competition. We say a contestant engages in benchmark hacking if his exploits more explores less in the contest setting compared to the baseline.

Under this new definition, we find that contestants with types below certain threshold (*low* types) always engage in benchmark hacking, whereas those above the threshold don't. As the contest reward skews towards top-ranked contestants, low types exploit more while high types explore more. We also inform ML contest design, suggesting that a more skewed (favoring top performers) reward structure can incentivize the development of winning models that are not only better-fitted to the contest task but also more innovative and generalizable to future applications.

We are motivated by the concern of benchmark hacking in ML contests. Hence our model and analysis tightly reflect this fundamental question. While some economic settings also involve two types of efforts — e.g., a manager choosing between efforts toward short-term and long-term goals (Narayanan, 1985) — these are typically single agent settings rather than contests that this paper focuses on. Therefore, we do not consider these other applications in this paper.

[4] In some cases, contestants may work on scalable algorithms for training existing models (e.g., convolutional neural networks (Krizhevsky et al., 2012)); we unify this under "producing models" to avoid confusion.

1.1 Literature Review

Manipulation in machine learning is a widely studied problem in AI. However, to our knowledge, most AI research thus far looks to design models robust to strategic behavior by downstream users or data providers (Ahmadi et al., 2021; Björkegren et al., 2020; Freeman et al., 2020; Hardt et al., 2016; Perdomo et al., 2020). While sharing this goal, we address a different problem — analyzing ML contests — and focus on contestants' incentives rather than those of end users.

Contests are widely used to deliver innovative products and solutions (e.g., Chesbrough et al., 2006), and prior research has examined their design features, such as prize structures (Clark and Riis, 1998), number of contestants (Bhattacharya, 2021), allocation rules (Fullerton and McAfee, 1999; Moldovanu and Sela, 2001), and information disclosure (Lemus and Marshall, 2021; Rieck, 2010), to optimize outcome quality. Our paper departs by analyzing players' effort allocation across multiple tasks and, uniquely, formalizes the notion of benchmark hacking arising from misaligned principal-agent incentives in a game-theoretic framework.

As for modeling, our model relates to the classic literature on multi-tasking (e.g., Bénabou and Tirole, 2016; Holmstrom and Milgrom, 1991), where agents allocate effort across multiple tasks. Multi-tasking has been studied in contexts such as firm management (Narayanan, 1985), academic publishing (Ellison, 2002), and pharmaceutical innovation (Bryan et al., 2022). Unlike this literature, which typically examines a single agent deciding between tasks that all benefit the principal, we study a multi-agent contest where only one type of effort benefits the principal. Our model also shares similarities with moral hazard problems in principal-agent theory (Grossman and Hart, 1992; Holmström, 1979), but differs in allowing agents to choose the levels of two distinct efforts rather than a single action.

Our paper also contributes to the literature on gaming behaviors. Brodeur et al. (2020) studies "p-hacking" in data manipulation,[5] while Ball (2022) and Frankel and Kartik (2019) examine gaming in signaling games, and Perez-Richet and Skreta (2022) and Li and Qiu (2023) analyze fraud-proof information design and gaming in contests. However, in our setting, manipulation arises in machine learning contests as contestants split effort between productive and gaming actions, capturing the phenomenon of benchmark hacking.

2 Model

Machine learning contest. Suppose there are I players (each of whom we will refer to as "he" hereafter) and one principal (hereafter, "she"). The I players can choose whether or not to participate in contest where they compete on developing data analysis approaches, referred conveniently as the *model* hereafter, for a learning task described in the contest. We call the game

[5] "p-hacking" occurs when researchers selectively collect or analyze data until nonsignificant results appear significant, representing exploitative effort in our model.

a *machine learning (ML) contest*. Each player i has a private type θ_i. It is known by player i but unknown to other players and the principal. We interpret θ_i as player i's *capability* in developing models. For any i, θ_i is drawn identically and independently from a commonly known distribution F with support $\Theta = [\underline{\theta}, \bar{\theta}]$ and a bounded and atomless density function f. Let the type profile for all players be $\boldsymbol{\theta} = (\theta_i)_{i=1}^I$ and $\boldsymbol{\theta}_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$ contains all players' types but i'th.

Actions, model fitness and performance. When building his *model*, each player i can engage in either *exploratory* or *exploitative effort*. Let the amount of exploratory (exploitative) effort be a_i (b_i). The effort pair chosen by player i, (a_i, b_i), determines the *fitness* μ_i of his model as follows

$$\mu_i = \nu(a_i, \theta_i) + \xi(b_i). \tag{1}$$

The effort choice of player i is privately known by him. Let the model fitness profile of all agents be $\boldsymbol{\mu} = (\mu_i)_{i=1}^I$.

We assume that $\nu(a_i, \theta_i)$ is strictly increasing in both a_i and θ_i, concave in a_i, twice differentiable, and $\nu(0, \theta_i) = 0$ for all $\theta_i \in [\underline{\theta}, \bar{\theta}]$. Moreover, $\xi(b_i)$ is strictly increasing, concave and twice differentiable in b_i, $\xi(0) = 0$, and $\lim_{b\to\infty} \xi'(b) = 0$. We also assume that the exploration is more productive when the player's type is higher.

Assumption 1 (Supermodularity)**.** *$\nu(a_i, \theta_i)$ is supermodular in a_i and θ_i: i.e., $\frac{\partial^2 \nu(a_i, \theta_i)}{\partial a_i \partial \theta_i} > 0$.*

The *performance* of player i's model is denoted by s_i. In practice, s_i is interpreted as the realized contest evaluation metric of player i's model. s_i is sampled from a continuous distribution $H_{\mu_i}(\cdot)$ with mean μ_i, support $\mathbb{R}_+$, and a bounded and atomless density function $h_{\mu_i}(\cdot)$. Furthermore,

Assumption 2 (FOSD)**.** *The distribution $H_\mu(s)$ first-order stochastically dominates $H_{\mu'}(s)$ for any $\mu > \mu'$. That is, $H_\mu(s) \leq (<) H_{\mu'}(s)$ for any(some) $s \geq 0$.*

This assumption formalizes the intuition that a better-fitted model is more likely to generates higher performance. Distributions including normal, exponential and Gumbel distribution satisfy this assumption. Denote the performance profile by $\boldsymbol{s} = (s_i)_{i=1}^I$ and the performance profile excluding i by $\boldsymbol{s}_{-i} = (s_j)_{j\neq i}$.

As previously discussed, each player could improve the performance of their model by spending either exploratory effort or exploitative effort. Moreover, exploratory effort requires experiences and analytical thinking, it naturally interact with each player's capability. However, exploitative effort is a mechanical searching problem hence it is reasonable to expect that its influences to the model performance is independent of the player's capability θ_i.

Contest awards, player costs and payoffs. The principal pre-specifies a prize vector $\boldsymbol{R} = (R_k)_{k=1}^I$ with $R_k \geq R_{k+1} \geq 0, 1 \leq k \leq I-1$, such that the player whose performance ranks kth would be rewarded the kth highest prize R_k. Denote the kth highest performance by $s_{(k)}$, so that $s_{(k)} \geq s_{(k+1)}$. Let i_k^* be the player whose performance ranks kth place, i.e., $s_{i_k^*} = s_{(k)}$.

The total effort spent by player i is $e_i \coloneqq a_i + b_i$. We let $c(e_i)$ denote the cost for exerting total effort e_i. Assume that c is twice differentiable, non-decreasing, convex in e_i, and $c(0) = 0$. Each player i's *payoff* for exerting effort (a_i, b_i) and achieving a performance level s_i is

$$\pi_i = \underbrace{\sum_{k=1}^{I} \mathbf{1}\left[i = i_k^*\right] \cdot R_k}_{\text{context gain}} + \underbrace{s_i}_{\text{private gain}} - c(a_i + b_i) \tag{2}$$

where $\mathbf{1}\left[i = i_k^*\right]$ is the (random) event that agent i wins the kth highest prize in the contest. Players benefit from building a better-fitted model through both *contest gain* and *private gain*. Contest gain captures the direct prize reward they win from the contest, while private gain is a reduced form approach to capture the personal gains in addition to contest prizes. These personal gains reflect the player's career concerns, where the market infers each player's type based on the observable signal (Cisternas, 2018). For example, winning or placing highly in a Kaggle competition can increase the participant's visibility in the community and showcase his skills in data science and machine learning to potential employers.

Timeline. This game is composed of three stages:

Stage 1: principal announces the prize vector.

Stage 2: each player observes his own type θ_i and chooses an effort allocation to achieve the fitnessμ_i.

Stage 3: Players performance are realized.[6] Rewards are allocated according to the contest rule and this ranking.

In reality, ML contests usually disclose players information on public leaderboard. However, in our paper, we consider a static game where there is no information disclosure about players status. This is because the main focus of our paper is to analyze players incentives to divide effort between exploratory and exploitative actions, whereas information disclosure and the contest architecture are of secondary order concerns.

Player optimization and equilibrium concept. Each player i's strategy β_i is a mapping from the type space Θ to an allocation of effort (a_i, b_i). The solution concept adopted is the standard *Bayesian Nash equilibrium*. A Bayesian Nash equilibrium is a strategy profile $\boldsymbol{\beta}^* = (\beta_i^*)_{i=1}^I$ such that for each player i, β_i^* maximizes player i's expected utility, given that other players strategies are $\boldsymbol{\beta}_{-i}^* = (\beta_j^*)_{j \neq i}$. In other words, β_i^* solves the following optimization problem.

$$\begin{aligned} \max_{\beta_i(\theta_i)} \quad & \mathbb{E}_{\boldsymbol{\theta}_{-i}, \boldsymbol{s}}\left[\pi_i(\beta_i, \theta_i, \boldsymbol{\beta}_{-i}^*) | \mu(\beta_i; \theta_i), \boldsymbol{\beta}_{-i}^*\right] \\ \text{s.t.} \quad & \mu(\beta_i; \theta_i) = \nu(a_i, \theta_i) + \xi(b_i) \end{aligned} \tag{$\mathcal{P}_i$}$$

[6] In our model, since the performance of each player is drawn from a continuous distribution, tie happens with probability 0.

Notably, the expectation above is over both other players' private types and the randomn performances given model fitness μ. Let $a_i^*(\theta_i)$, $b_i^*(\theta_i)$, and $e_i^*(\theta_i)$ be the equilibrium exploratory, exploitative and total effort, correspondingly, for the player i with type θ_i. Let the equilibrium fitness of player i be $\mu_i^*(\theta_i) := \mu(a_i^*(\theta_i), b_i^*(\theta_i); \theta_i)$. Finally, we say $\boldsymbol{\beta}^* = (\beta_i^*)_{i=1}^I$ is a *symmetric* Nash equilibrium if all players' strategies are the same, i.e., $\beta_i^* = \beta_j^*$ for any i, j.

3 Equilibrium Existence

Proposition 1. *Under Assumption 1 and 2, there exists a pure strategy Nash equilibrium that is*

- symmetric*: $\mu_i^*(\theta_i) = \mu^*(\theta_i)$ for all i;*
- *and* monotone*: $\mu^*(\theta_i)$ is non-decreasing in θ_i.*

Our proof is based on an alternative representation of the player's strategy space. Observe that in player i's problem $\mathcal{P}_i$ above, the effort on the two actions determines each agent's utility through two channels: (1) the model fitness μ_i and resultant probability distribution over the outcomes of game (i.e., ranking of players), (2) total effort. It turns out that we can transform each player's strategy space into an alternative and equivalent one in which each agent first chooses the model fitness μ_i to maximize his utility and then choosing the least-costly allocation of effort on the two actions to achieve the chosen fitness μ_i. This two-stage perspective allows us to disentangle their effects to the equilibrium and helps to simplify our analysis.

To describe this auxiliary game, we first introduce the cost of achieving fitness μ_i by type θ_i as the following function:

$$C(\mu_i, \theta_i) = \min_{a_i, b_i \geq 0} c(a_i + b_i) \qquad \text{s.t.} \quad \nu(a_i, \theta_i) + \xi(b_i) = \mu_i. \tag{3}$$

Now we consider the alternative game where each player i's utility is $\sum_{k=1}^I \mathbf{1}\left[i = i_k^*\right] \cdot R_k + \mu_i - C(\mu_i, \theta_i)$, when he has type θ_i and chooses fitness μ_i. Agent i's strategy in this alternative game is $(\mu_i, \boldsymbol{\phi})$, where μ_i is a mapping from his own type space Θ to the choice of fitness μ_i, and $\boldsymbol{\phi} = (\phi_1, \phi_2)$ maps from type and fitness space $\Theta \times \mathbb{R}_+$ to the allocation of two types of effort.

We then prove that the payoff function in the auxiliary game is bounded and continuous in μ_i for every $\theta_i \in \Theta$, and has increasing differences in (μ_i, θ_i), as formally proved in Lemma 2 and Lemma 4 in Appendix A. These then allow us to evoke Theorem 4.5 of Reny (2011) to show that there is a symmetric monotone equilibrium in the original game. The complete proof of Proposition 1 is provided in Appendix A. From now on, we will focus on symmetric monotone equilibrium, and drop index i when it is clear from the context.

4 The Baseline

We now introduce a baseline scenario, where the payoff of each agent comes only from the *private gain*. For notational convenience, we drop subscript i in this subsection. Each player's problem writes as

$$\max_{a,b\geq 0} \quad \pi^b(a,b;\theta) = \mu(a,b) - C(\mu(a,b),\theta). \tag{$\mathcal{P}_b$}$$

Let the optimal choice of effort be $(a^\dagger(\theta), b^\dagger(\theta)) = \arg\max_{a,b\geq 0} \pi^b(a,b;\theta)$, and the optimal total effort be $e^\dagger(\theta) = a^\dagger(\theta) + b^\dagger(\theta)$ in the baseline scenario.[7]

Assumption 3. *$\xi'(0) > c'(0)$, and $\lim_{\theta\to-\infty} \frac{\partial\nu(a,\theta)}{\partial a}\big|_{a=0} < c'(0)$, $\lim_{\theta\to+\infty} \frac{\partial\nu(a,\theta)}{\partial a}\big|_{a=0} > \xi'(0)$.*

The first half of this assumption indicates that marginal benefit of an infinitely small exploitative effort is greater than its marginal cost. It ensures that (1) all types exert positive efforts; and (2) not all types exert zero exploitative effort in the baseline scenario. The second half ensures that $\theta_1^\dagger$ and $\theta_2^\dagger$ are well-defined in Proposition 2.

Proposition 2. *Under Assumption 1, 2, and 3, $e^\dagger(\theta) > 0$ for any type θ. Moreover, there exists two thresholds $\theta_1^\dagger \leq \theta_2^\dagger$, such that*

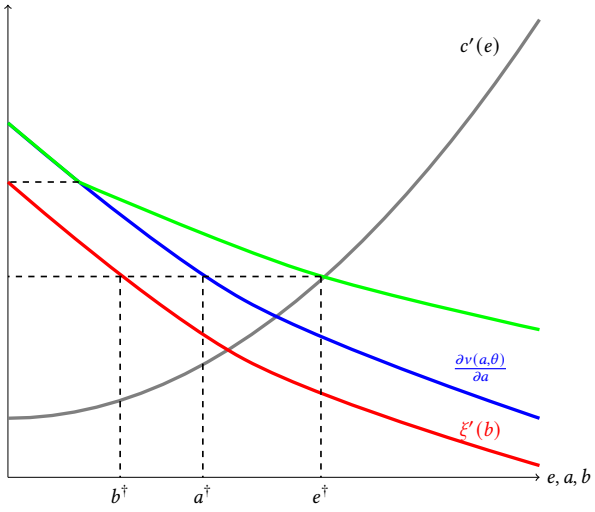
1. *Any type $\theta < \theta_1^\dagger$ only exploits.*
2. *Any type $\theta \in (\theta_1^\dagger, \theta_2^\dagger)$ both exploits and explores.*
3. *Any type above $\theta_2^\dagger$ only explores.*

Fig. 1 offers a graphical intuition for Proposition 2. The blue and red curves represent the marginal productivity of exploration and exploitation, respectively. The green curve is the horizontal sum of the blue and red curves, representing the marginal productivity of the total effort e. When θ is sufficiently large (small), the blue curve starts off above (below) the red curve. The kink on the green curve represents the switching point from only exploration (exploitation) to both efforts.

Intuitively, it's optimal to choose the level of exploration and exploitation where the marginal productivities of the two equalize. When θ is sufficiently high (low), the green curve crosses $c'(e)$ to the left of the kink, with the blue curve above (below) the red curve, then there is only exploration (exploitation). When θ falls within the intermediate range, both exploration and exploitation are positive.

[7] As a kind reminder, all the optimal choices for the baseline scenario are labeled with the "dagger" notation.

Figure 1: **Graphical illustration of Proposition 2**



5 Contest Rewards and Incentives

Definition 1 (A partial order on prize vectors)**.** *Consider two prize vectors* $\boldsymbol{R}, \boldsymbol{R}' \in \mathbb{R}_+^I$*. We say* $\boldsymbol{R} \geqq \boldsymbol{R}'$ *if* $R_k - R_{k+1} \geq R'_k - R'_{k+1}$ *for any* $k = 1, ..., I-1$*.*

This partial order compares the *skewness* of two prize vectors. The prize vector $\boldsymbol{R}$ is considered more "skewed" (towards top-ranked contestants) than $\boldsymbol{R}'$ if, for each step-up in ranks, the incremental increase in prize reward is greater under $\boldsymbol{R}$ than $\boldsymbol{R}'$. For example, let $I = 4$, then $(6, 4, 3, 0) \geqq (6, 5, 4, 3)$.

We focus on the symmetric monotone pure strategy equilibrium, $\mu^*(\theta, \boldsymbol{R})$.[8] The following main result of this section shows that an agent of any type chooses to produce a better-fitted model under a more skewed prize vector.

Proposition 3. *For any* $\theta \in \Theta$*, and any* $\boldsymbol{R} \geqq \boldsymbol{R}'$*,* $\mu^*(\theta, \boldsymbol{R}) \geq \mu^*(\theta, \boldsymbol{R}')$*.*

The proof follows a similar idea to the equilibrium existence argument, but instead of showing that the payoff function has increasing differences in (μ, θ), we now show that it has increasing differences in $(\mu, \boldsymbol{R})$. Such increasing difference comes from (1) the probability distribution of the final rank for a higher model fitness first-order stochastically dominates that of a lower model fitness, and (2) a more skewed prize vector allocates disproportionally greater rewards to top performers. Intuitively, the incremental gain from increasing the skewness of the prize vector is greater when a player is more likely to achieve a high rank (due to higher model fitness), as a more skewed prize vector disproportionately rewards top performers. Hence, a more skewed prize vector incentivizes any type to produce a better-fitted model.

[8] "monotone" only refers to that $\mu^*(\theta, \boldsymbol{R})$ is non-decreasing in θ but not necessarily $\boldsymbol{R}$.

Notice that the baseline scenario is equivalent to a contest where the prize vector is a zero vector $\mathbf{0} = (0, ..., 0)$. By Definition 1, any prize vector $\boldsymbol{R}$ satisfies $\boldsymbol{R} \geqq \mathbf{0}$. Thus, we derive the following Corollary that compares equilibrium fitness and the fitness in the baseline scenario.

Corollary 1. *Under Assumption 1, 2, and 3, $\mu^*(\theta) > \mu^\dagger(\theta)$ for any $\theta \in \Theta$.*

6 Benchmark Hacking

Consider a symmetric monotone equilibrium $(a^*(\theta), b^*(\theta))$ of the ML contest. Type θ is involved in *benchmark hacking* if $a^*(\theta) \leq a^\dagger(\theta)$ and $b^*(\theta) > b^\dagger(\theta)$, where $a^\dagger(\theta), b^\dagger(\theta)$ are the exploratory and exploitative effort in the baseline setting. In other words, a contestant is said to engage in benchmark hacking if he exhibits a bias toward exploitative effort during the contest. The following proposition compares the two types of effort in the baseline and in the contest.

Proposition 4. *Under Assumption 1, 2, and 3, there exists a threshold $\theta_1^* \in (-\infty, \theta_1^\dagger]$ such that*

1. *Any type $\theta < \theta_1^*$ only exploits, and $b^*(\theta) > b^\dagger(\theta) > 0$.*
2. *Any type $\theta \in (\theta_1^*, \theta_1^\dagger)$ exploits in both settings, and $b^*(\theta) > b^\dagger(\theta) > 0$; but explores only in the contest.*
3. *Any type $\theta \in (\theta_1^\dagger, \theta_2^\dagger)$ both explores and exploits, and each effort is higher in the contest.*

This proposition suggests that in the symmetric monotone equilibrium under contest, (1) *low* types achieve higher fitness level by choosing higher exploitative effort; (2) *high* types achieve higher fitness level by choosing higher effort on both actions.

Corollary 2. *Low types are involved in benchmark hacking and high types are not.*

This Corollary follow immediately from Proposition 4 that any type $\theta \leq \theta_1^*$ will exert zero exploratory effort in both contest and baseline case and higher exploitative effort in contest. In Appendix G.1, we provide a few concrete examples to illustrate the different possibilities of how types in each region allocate their effort.

The following Corollary is immediate from Proposition 3 and 4.

Corollary 3. *For any $\boldsymbol{R} \geqq \boldsymbol{R}'$,*

1. *it is never optimal to reward the lowest-ranked player.*
2. *Fewer types are involved in benchmark hacking under $\boldsymbol{R}$.*
3. *$\boldsymbol{R}$ encourages weakly higher model fitness for each type.*

This corollary informs the design of the ML contest. A more skewed prize vector elicits (1) better-fitted models (as is shown in Proposition 3), (2) more exploratory effort from the entire community, and (3) fewer types involved in benchmark hacking. However, it remains an open question on how the designer should tradeoff monetary rewards and machine learning models. We leave the more concrete characterization of optimal reward structure to future work.

7 Model Validation with Kaggle Data

In this section, we validate our model predictions using public data on featured competitions hosted by a popular ML contest platform *Kaggle* (Lemus and Marshall, 2021).[9] In a typical Kaggle competition, there is a prediction task and players compete on the performance of their designed ML models for it. Players are given a *training* dataset containing both features and labels, based on which they build an ML model. After that players submit the model to test its performance on the *preliminary* test dataset that is sampled from the same distribution as the training dataset but is not disclosed to the players. Rather, players are only told about the prediction accuracy of their submitted model on this test set. After seeing it, players may refine their model and submit again. At the end of the contest, each player is required to submit a final model, whose performance will be evaluated on the *main* test dataset that is from the same distribution as the training set and preliminary test set but contains different observations. The prediction accuracy on the main test set determines the rankings and the winner(s).

Data. To ensure that the prediction accuracy is comparable across contests, we select contests in which the prediction task is binary classification and the prizes are in U.S. dollars. Eventually, our dataset contains 26 contests. The unit of observation is a contest-player-submission pair, with contest indexed by j, player by i and submission by k. For each contest, the variables we observe include number of players (usually around 1000∼3000), number of prizes (usually around 1∼4), prize values (usually within 10K∼70K), contest year (within 2010∼2020) and duration (usually spanning 1∼3 months), etc. At player level, Kaggle has a feature called *Notebook* which counts how many upvotes a player receives. The number of upvotes is independent of the player's contest performance and is the best proxy for his expertise type (categorical from 0 (lowest) to 4 (highest)). In addition, we observe the date and prediction score (typically between 60% and 95%) for every submission made by each player in each contest. Detailed data statistics are reported in the appendices.

A key data challenge is to measure exploratory vs exploitative efforts as they are not observed directly from the data. It turns out that these two types of effort often lead to distinct model

[9]Featured competitions are the types of competitions in which players build a machine learning model to solve a (usually commercially-purposed) prediction task. Unlike other types of competitions on Kaggle, featured competitions offer medium- to high-stake cash prizes that can go as high as million dollars. https://www.kaggle.com/datasets/kaggle/meta-kaggle.

performance patterns over time. Exploratory effort involves step-by-step, creative, and problem-specific customizations. Participants usually cannot conceive all customizations upfront. They typically start with a baseline model, submit it and identify shortcomings from feedback, and iteratively develop tailored improvements. This feedback-driven process often results in a steady upward trend in accuracy over the course of the competition. In contrast, exploitative effort follows a trial-and-error approach, relying on plug-and-play machine learning techniques or well-known heuristics applied without deep understanding of the problem — often resulting in noisy and unpredictable performance fluctuations.

For example, in the *Driver Telematics Analysis* contest, participants were tasked with predicting the probability that a given trip was made by a specific driver, using a dataset of driving trips. Upon reviewing the coding summary shared by team *ice9* on Kaggle website (see Appendix for its screenshots), we observe that they employed several exploratory strategies (with some exploitative elements layered on top): decomposing trips into distances between turns and turn sizes, standardizing trip representations for meaningful comparison, and creating an adjacency tree to capture relationships between trips. Their accuracy scores generally improved over the course of the competition. Another example is the *RSNA-MICCAI Brain Tumor Radiogenomic Classification* contest, where participants were tasked with identifying tumor from MRI scans. The team *Aryaman Sharma* trained a 3D ResNet50 model—a standard choice for MRI data—using default settings and routine optimization techniques such as batch size tuning. This reflects a largely exploitative effort, and their performance showed fluctuating score patterns over time.

Hence, we infer the relative balance of exploratory versus exploitative effort from the monotonic trends in prediction scores at the contest-player level. Specifically, we use the Mann-Kendall test statistic on each participant's prediction scores over time as a proxy for this balance (mostly between -0.44 and 4.82).[10]

Empirical Evidence of Model Predictions. We generated a series of empirical facts from the data, from which we observed consistency with our model predictions as elaborated below. The following Table 1 contains some key estimated values for 3 out of the 5 player types that suffice for our illustration below; a full table with all 5 types, all estimates and their standard errors can be found in Table 5 in Appendix G.2 – in summary, all the conclusions below apply to the full table as well.

Test for Proposition 1. First, we test the monotonicity of the equilibrium, in particular, whether the equilibrium model fitness (measured by final prediction accuracy) is monotone in type. We estimate the impact of type on model fitness using regression $Fitness_{ij} = \eta_0 + TypeFE_i + ContestFE_j + \varepsilon_{ij}$, where $TypeFE$ is type fixed effect, $ContestFE$ is contest fixed effect and is included to account for variations in prediction accuracy specific to the distinct nature of different prediction tasks, ε_{ij}

[10]The Mann-Kendall test is a non-parametric method used to statistically assess whether there is a monotonic upward trend in the variable of interest over time (Mann, 1945)

captures noise in the prediction accuracy due to the limited sample size and the randomness in the splitting of training data for cross-validation. The result is shown in Table 1. For example, compared with type 2 players, type 3 players' model fitness is 1% higher, *ceteris paribus*. The coefficient of *TypeFE* is higher for higher types, suggesting that model fitness in equilibrium is monotone in player type, consistent with our finding in Proposition 1.[11]

Table 1: **Estimates of Coefficients (partial results for illustration purpose)**

TESTS	Prop. 1	Prop. 3	Prop. 4	Cor. 3
VARIABLES	Final score	Final score	MK stat	MK stat
T_2	6.389***	11.32***	1.969***	1.844***
T_3	7.380***	13.82***	2.453***	2.317***
$T_1 \times PV$		0.445***		0.0285***
$T_2 \times PV$		0.415***		0.0258***
$T_3 \times PV$		0.396***		0.0234***
$T_1 \times PS$		37.56***		2.318***
$T_2 \times PS$		37.09***		2.712***
$T_3 \times PS$		36.31***		2.927***
Constant	65.65***	22.83***	-0.238	-2.672***
Contest FE	Yes	Yes	Yes	Yes

Note: T_1 is a dummy that equals to 1 if the player's type is 1, same for T_2, T_3. PV and PS stand for total prize value (thousand USD) and prize skewness (favoring top-ranked players). ***p<0.01, **p<0.05, *p<0.1

Test for Proposition 3. Next, we test our comparative statics results established in Proposition 3, in particular, the impact of prize value and skewness on the choice of model fitness. Since we don't observe the split of the total prize in the data, we measure the prize skewness by the number of prizes. Specifically, prize skewness equals 1 if a contest has no more than 3 prizes, 0 otherwise. The rationale is that in a contest with fewer prizes, most of the rewards are captured by the top players. We use the following regression for estimation:

$$\begin{aligned} Fitness_{ij} = \rho_0 + \rho_1 \, TypeFE_i \times PV_j + \rho_2 \, TypeFE_i \times PS_j \\ + \, TypeFE_i + ContestFE_j + \epsilon_{ij}, \end{aligned} \tag{4}$$

where PV and PS stand for prize value and prize skewess, and the interaction terms $TypeFE_i \times PV_j$ and $TypeFE_i \times PS_j$ capture the impact of prize structure on the choice of model fitness *within* each type. We show the result in column (2) of Table 1. To interpret the coefficients of the interaction terms, take $T_1 \times PV$ for example, the coefficient of 0.445 means that type 1 players improve their model fitness by 0.445% per 1,000 USD increase in the total prize value, *ceteris*

[11] We also perform a Wald test to test the equality of fixed effects across types, and we reject the null hypothesis of equal effects at 1% significance level.

paribus. The coefficients of all interaction terms are significant and positive, consistent with our model predictions in Proposition 3 that a higher and more skewed prize vector leads to higher model fitness in equilibrium.

Test for Proposition 4. Proposition 4 suggests that low (high) type players tend to choose more exploitative (exploratory) effort. We test it using the regression $MK\ stat = \tau_0 + TypeFE_i + ContestFE_j + \omega_{ij}$, where *MK stat* is the Mann-Kendall test statistics. According to the results in column (3) of Table 1, higher types tend to exhibit higher MK statistics, indicating a greater likelihood of an upward trend in their prediction accuracy scores over time and, consequently, a stronger tendency toward exploratory effort.

Test for Corollary 3. Last, we test the impact of prize size and skewness on efforts using the following regression:

$$\begin{aligned} MK\ stat_{ij} = \kappa_0 + \kappa_1\, TypeFE_i \times PV_j + \kappa_2\, TypeFE_i \times PS_j \\ + TypeFE_i + ContestFE_j + \varsigma_{ij} \end{aligned} \tag{5}$$

The result is in column (4) of Table 1. The coefficients of (almost) all interaction terms are positive and significant, consistent with our findings in Corollary 3 that higher reward and skewness incentivize more exploratory effort.

8 Conclusion

This paper (1) develops a game-theoretic model to study misaligned incentives between contest hosts seeking high-quality ML solutions and contestants aiming to win prizes by strategically allocating exploratory and exploitative effort, (2) formally defines *benchmark hacking* in ML contests and examines its presence in equilibrium, and (3) tests the model's predictions using Kaggle competition data. Under this new definition, our results suggest that low-type contestants engage in benchmark hacking while high types do not. A more skewed prize structure encourages more exploration, less exploitation, and higher model fitness. These findings feature the heterogeneity among contestants regarding their tendency to benchmark hacking and the degree of mis-aligned incentives with contest hosts. We view this work as a step toward a more concrete characterization of optimal ML contest design that could better addressing misaligned incentives and producing more desirable outcomes for contest organizers.

References

Saba Ahmadi, Hedyeh Beyhaghi, Avrim Blum, and Keziah Naggita. The strategic perceptron. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 6–25, 2021.

Ian Ball. Scoring strategic agents. *arXiv preprint arXiv:1909.01888*, 2022.

Roland Bénabou and Jean Tirole. Bonus culture: Competitive pay, screening, and multitasking. *Journal of Political Economy*, 124(2):305–370, 2016.

Vivek Bhattacharya. An empirical model of r&d procurement contests: An analysis of the dod sbir program. *Econometrica*, 89(5):2189–2224, 2021.

Daniel Björkegren, Joshua E Blumenstock, and Samsun Knight. Manipulation-proof machine learning. *arXiv preprint arXiv:2004.03865*, 2020.

Abel Brodeur, Nikolai Cook, and Anthony Heyes. Methods matter: p-hacking and publication bias in causal analysis in economics. *American Economic Review*, 110(11):3634–60, November 2020. doi: 10.1257/aer.20190687. URL https://www.aeaweb.org/articles?id=10.1257/aer.20190687.

Kevin A Bryan, Jorge Lemus, and Guillermo Marshall. R&d competition and the direction of innovation. *International Journal of Industrial Organization*, 82:102841, 2022.

Henry Chesbrough, Wim Vanhaverbeke, and Joel West. *Open innovation: Researching a new paradigm*. Oxford university press, USA, 2006.

Gonzalo Cisternas. Career concerns and the nature of skills. *American Economic Journal: Microeconomics*, 10(2):152–189, 2018.

Derek J Clark and Christian Riis. Competition over more than one prize. *The American Economic Review*, 88(1):276–289, 1998.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

Glenn Ellison. Evolving standards for academic publishing: A q-r theory. *Journal of political economy*, 110(5):994–1034, 2002.

Alex Frankel and Navin Kartik. Muddled information. *Journal of Political Economy*, 127(4): 1739–1776, 2019.

Rupert Freeman, David Pennock, Chara Podimata, and Jennifer Wortman Vaughan. No-regret and incentive-compatible online learning. In *International Conference on Machine Learning*, pages 3270–3279. PMLR, 2020.

Richard L Fullerton and R Preston McAfee. Auctionin entry into tournaments. *Journal of Political Economy*, 107(3):573–605, 1999.

Sanford J. Grossman and Oliver D. Hart. An analysis of the principal-agent problem. In *Foundations of insurance economics*, pages 302–340. Springer, 1992.

Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pages 111–122, 2016.

Bengt Holmström. Moral hazard and observability. *The Bell journal of economics*, pages 74–91, 1979.

Bengt Holmstrom and Paul Milgrom. Multitask principal–agent analyses: Incentive contracts, asset ownership, and job design. *The Journal of Law, Economics, and Organization*, 7(special_issue): 24–52, 1991.

ImageNet. Imagenet: Large scale visual recognition challenge 2012 (ilsvrc2012), 2012. https://www.image-net.org/challenges/LSVRC/2012/results.php.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

Jorge Lemus and Guillermo Marshall. Dynamic tournament design: Evidence from prediction contests. *Journal of Political Economy*, 129(2):383–420, 2021.

Yingkai Li and Xiaoyun Qiu. Contests as optimal mechanisms under signal manipulation. *arXiv preprint arXiv:2302.09168*, 2023.

Henry B Mann. Nonparametric tests against trend. *Econometrica: Journal of the econometric society*, pages 245–259, 1945.

Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001.

M P Narayanan. Managerial Incentives for Short-term Results. *Journal of Finance*, 40(5):1469–1484, December 1985. URL https://ideas.repec.org/a/bla/jfinan/v40y1985i5p1469-84.html.

FS Ndzomga. Llama 4 and benchmark lies. ai progress is overrated, 2025. https://medium.com/thoughts-on-machine-learning/llama-4-and-benchmark-lies-85f4445fac88.

Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pages 7599–7609. PMLR, 2020.

Eduardo Perez-Richet and Vasiliki Skreta. Test design under falsification. *Econometrica*, 90(3): 1109–1142, 2022.

Philip J Reny. On the existence of monotone pure-strategy equilibria in bayesian games. *Econometrica*, 79(2):499–553, 2011.

Thomas Rieck. Information disclosure in innovation contests. Technical report, Bonn Econ Discussion Papers, 2010.

A Equilibrium Existence

A.1 The auxiliary game.

An equilibrium of this auxiliary game is a strategy profile $((\mu_i^\ddagger)_i, \boldsymbol{\phi}^\ddagger)$ that satisfies the following two conditions. First, given other players strategy profile $\boldsymbol{\mu}_{-i}^\ddagger$, for each player i, $\mu_i^\ddagger$ solves the following problem

$$\begin{aligned} \max_{\mu_i} \quad & \mathbb{E}_{\boldsymbol{\theta}_{-i},\boldsymbol{s}}\left[\sum_{k=1}^{I} \mathbf{1}\left[i=i_k^*\right]\cdot R_k \middle| \mu_i, \boldsymbol{\mu}_{-i}^\ddagger\right] + \mu_i - C(\mu_i,\theta_i) \\ =\max_{\mu_i} \quad & g(\mu_i, \boldsymbol{\mu}_{-i}^\ddagger) + \mu_i - C(\mu_i,\theta_i), \end{aligned} \tag{$\mathcal{P}_i'$}$$

where

$$\begin{aligned} & g(\mu_i, \boldsymbol{\mu}_{-i}^\ddagger) \\ &= \mathbb{E}_{\boldsymbol{\theta}_{-i},\boldsymbol{s}}\left[\sum_{k=1}^{I} \mathbf{1}\left[i=i_k^*\right]\cdot R_k \middle| \mu_i, \boldsymbol{\mu}_{-i}^\ddagger\right] \\ &= \sum_{k=1}^{I} \mathbf{Pr}\left[i=i_k^* | \mu_i, \boldsymbol{\mu}_{-i}^\ddagger\right]\cdot R_k \end{aligned}$$

is the contest gain. Second, $\boldsymbol{\phi}^\ddagger(\mu,\theta)$ is the optimal effort allocation given fitness μ and type θ, i.e.,

$$\boldsymbol{\phi}^\ddagger(\mu,\theta) = \arg\min_{a,b\geq 0} c(a+b) \text{ s.t. } \nu(a,\theta)+\xi(b)=\mu. \tag{6}$$

The following lemma formally establishes the connection between our original game and the auxiliary game above. Armed with this equivalence result, we will then conduct our analysis on the auxiliary game in the following sections.

Lemma 1. *There is a one-to-one mapping from any equilibrium $(\beta_i^*)_i$ of the original game to an equilibrium $((\mu_i^\ddagger)_i, \boldsymbol{\phi}^\ddagger)$ in the auxiliary game, such that $\beta_i^*(\theta) = \boldsymbol{\phi}^\ddagger(\mu_i^\ddagger(\theta),\theta)$ for any $\theta\in\Theta$.*

The main reason is that in the original game, each player's payoff depends on other players' strategies only through the fitness. In other words, other players' fitness is a sufficient statistics of their strategies.

Proof of Lemma 1. Consider any player i and the original game. For any player $j\neq i$, player j's strategy can affect player i's payoff only by changing player i's ranking in the contest. Moreover, such influence is summarized by player j's fitness. Next, we argue that fixing other players fitness profile $\boldsymbol{\mu}_{-i}$, player i's choice of fitness is the same in the original game and the auxiliary game. This is because fixing other players fitness profile $\boldsymbol{\mu}_{-i}$, $\mathcal{P}_i$ and $\mathcal{P}_i'$ are essentially the same maximization problem and thus admit the same solution.

□

A.2 Existence of symmetric monotone equilibrium

In this subsection, we will show that there exists at least one symmetric monotone equilibrium in the auxiliary game, hence also in the original game due to Lemma 1. We start by characterizing properties of the contest gain function $g(\mu_i, \boldsymbol{\mu}_{-i})$.

Lemma 2. *Under Assumption 2, for any $\boldsymbol{\mu}_{-i}$, $g(\mu_i, \boldsymbol{\mu}_{-i})$ is non-negative, bounded, continuous and non-decreasing in μ_i.*

The fact that $g(\mu_i, \boldsymbol{\mu}_{-i})$ is non-decreasing in μ_i implies agent i will rank higher generally and wins higher prize in expectation by building a better fitted model, when fixing the strategy profile $\boldsymbol{\mu}_{-i}$ of other agents.

Proof of Lemma 2. Non-negativity and boundedness are straightforward. Continuity follows from continuity of the payoff function and continuity of the distribution function H_{μ_i}.

Non-decreasing: For notation convenience, we define $p_{\mu_i,\boldsymbol{\mu}_{-i}}(k) = \mathbf{Pr}\big[i = i_k | \mu_i, \boldsymbol{\mu}_{-i}\big]$, so that

$$g(\mu_i, \boldsymbol{\mu}_{-i}) = \sum_{k=1}^{I} p_{\mu_i,\boldsymbol{\mu}_{-i}}(k) \cdot R_k \tag{7}$$

Let $P_{\mu_i,\boldsymbol{\mu}_{-i}}(k) = \sum_{l=1}^{k} p_{\mu_i,\boldsymbol{\mu}_{-i}}(l)$ be the probability that player i ranks at least the kth place. Lemma 3 presents the properties of P.

Lemma 3. *Under Assumption 2, for any $\mu_i > \mu_i'$ and any $\boldsymbol{\mu}_{-i}$, $P_{\mu_i,\boldsymbol{\mu}_{-i}}(k) \geq P_{\mu_i',\boldsymbol{\mu}_{-i}}(k)$.*

Let $\mu_i > \mu_i'$. Lemma 3 has essentially established a first order stochastic dominance relation between the two distributions $P_{\mu_i',\boldsymbol{s}_{-i}}$ and $P_{\mu_i,\boldsymbol{s}_{-i}}$, both of which have support $\{1, \ldots, I\}$. Notice that the support of the two distributions consists of the ranking and $k = 1$ is the highest ranking, while $k = I$ is the lowest ranking. Hence we shall read Lemma 3 as $P_{\mu_i,\boldsymbol{s}_{-i}}$ first order stochastic dominates $P_{\mu_i',\boldsymbol{s}_{-i}}$. Recall that if one random variable stochastically dominates another in the first order, the expected value of a non-decreasing function of the dominating variable will be no less than the expected value of the same non-decreasing function of the variable it dominates. Using this property, we have

$$\mathbb{E}_{\boldsymbol{\theta}_{-i},\boldsymbol{s}_{-i}}\left[\sum_{k=1}^{I} p_{\mu_i',\boldsymbol{s}_{-i}}(k) \cdot R_k\right] \leq \mathbb{E}_{\boldsymbol{\theta}_{-i},\boldsymbol{s}_{-i}}\left[\sum_{k=1}^{I} p_{\mu_i,\boldsymbol{s}_{-i}}(k) \cdot R_k\right]$$

Hence $g(\mu_i, \boldsymbol{\mu}_{-i})$ is non-decreasing in μ_i.

□

Next, we characterize properties of the cost function $C(\mu_i, \theta_i)$.

Lemma 4. *Under Assumption 1, $C(\mu_i, \theta_i)$ has the following properties:*

1. $C(0, \theta_i) = 0 \ \forall \theta_i \in [\underline{\theta}, \bar{\theta}]$.

2. $C(\mu_i, \theta_i)$ *is continuous in* μ_i *and* θ_i.

3. $C(\mu_i, \theta_i)$ *is non-decreasing in* μ_i *and non-increasing in* θ_i.

4. $C(\mu_i, \theta_i)$ *is convex in* $\mu_i \ \forall \theta_i \in [\underline{\theta}, \bar{\theta}]$.

5. $C(\mu_i, \theta_i)$ *is continuously differentiable with respect to* $\mu_i \ \forall \theta_i \in [\underline{\theta}, \bar{\theta}]$.

6. $-C(\mu_i, \theta_i)$ *satisfies weak single crossing property. Moreover,* $-C(\mu_i, \theta_i)$ *has increasing differences, i.e., for any* $\mu_i' < \mu_i$ *and any* $\theta_i' < \theta_i$,

$$C(\mu_i, \theta_i) - C(\mu_i', \theta_i) \leq C(\mu_i, \theta_i') - C(\mu_i', \theta_i')$$

The most important property is the weak single crossing property, which is primarily attributed to the supermodarity assumption in ν. This property suggests that the incremental increase in total effort when improving the model fitness from μ_i to μ_i' is lesser for a higher type agent.

Proof of Lemma 4. Since cost minimization is a single-agent problem, we drop subscript i for notation convenience. The first property is straightforward. Continuity follows from the continuity of ν, ξ and effort cost function c. The monotonicity in μ stems from the monotonicity of model fitness in both efforts. Likewise, the monotonicity in θ is from the monotonicity of model fitness in θ.

Convexity: since c only depends on the total effort and is monotone in it, $C(\mu, \theta) = c(e(\mu, \theta))$ where

$$e(\mu, \theta) = \min_{a \geq 0, b \geq 0} a + b \quad \text{s.t. } \nu(a, \theta) + \xi(b) = \mu \tag{8}$$

Let $(a(\mu, \theta), b(\mu, \theta))$ be the optimal effort choices in Eq. (8). Note that it coincides with $\boldsymbol{\phi}^{\ddagger}(\mu, \theta)$.

We first show that $e(\mu, \theta)$ is convex in μ. For any $\mu > \mu'$ and any $\theta \in \Theta$, let $(a, b) = (a(\mu, \theta), b(\mu, \theta))$ and $(a', b') = (a(\mu', \theta), b(\mu', \theta))$. For any $t \in (0, 1)$, we construct an effort allocation $\hat{a} = ta + (1-t)a'$, $\hat{b} = tb + (1-t)b'$ that produces a model with fitness $\mu(\hat{a}, \hat{b}; \theta)$.

$$\begin{aligned}
\mu(\hat{a}, \hat{b}; \theta) &= \nu(ta + (1-t)a', \theta) + \xi(tb + (1-t)b') \\
&\geq t\nu(a, \theta) + (1-t)\nu(a', \theta) + t\xi(b) + (1-t)\xi(b') \\
&= t\mu + (1-t)\mu'
\end{aligned}$$

The inequality is by the concavity of ν and ξ. Since we have found an effort allocation that produces model fitness at least $t\mu + (1-t)\mu'$, the minimum cost of producing a model of fitness

$t\mu + (1-t)\mu'$ cannot be more than that. Hence,

$$\begin{aligned} e(t\mu + (1-t)\mu', \theta) &\leq \hat{a} + \hat{b} \\ &= t(a+b) + (1-t)(a'+b') \\ &= te(\mu, \theta) + (1-t)e(\mu', \theta) \end{aligned}$$

Therefore, $e(\mu, \theta)$ is convex in μ. Since convexity is preserved under convex transformation, $C(\mu, \theta)$ is also convex in μ.

Continuously differentiable: Since $C(\mu, \theta)$ is convex in μ, then it is continuously differentiable with respect to μ almost everywhere. The possible non-differentiable points are such that the effort choice switches from only exerting one type of effort to exerting both types of efforts, or vice versa. However, when switching, the marginal productivity of both efforts should be the same (otherwise, the player shouldn't switch efforts). Hence, as μ is approached from either the left or the right, the left-hand derivative and the right-hand derivative both exist and are equal in the limit. Hence, $C(\mu, \theta)$ is continuously differentiable with respect to μ everywhere.

Weak Single Crossing: For any $\mu \geq 0$ and any $\theta > \theta'$, due to the monotonicity property, $-C(\mu, \theta) - (-C(\mu, \theta')) \geq 0$. Hence, the weak single crossing property is trivial.

Increasing difference: We first analyze whether $-e(\mu, \theta)$ exhibits increasing difference. Consider the Lagrangian of the minimization problem in Eq. (8),

$$\mathcal{L}(a, b, \lambda, \eta_1, \eta_2) = a + b + \lambda(\mu - \nu(a, \theta) - \xi(b)) - \eta_1 a - \eta_2 b \tag{9}$$

where λ and $\eta_1, \eta_2 \geq 0$ are Lagrange multipliers. Let $(a^*, b^*, \lambda^*, \eta_1^*, \eta_2^*)$ be the solution to the minimization problem. Then it satisfies the following necessary conditions, including the first order conditions,

$$\frac{\partial \mathcal{L}(a, b, \lambda, \eta_1, \eta_2)}{\partial a} = 1 - \lambda \frac{\partial \nu(a, \theta)}{\partial a} - \eta_1 = 0, \tag{FOC-a}$$

and

$$\frac{\partial \mathcal{L}(a, b, \lambda, \eta_1, \eta_2)}{\partial b} = 1 - \lambda \xi'(b) - \eta_2 = 0, \tag{FOC-b}$$

the equality constraint (budget balanced condition)

$$\nu(a, \theta) + \xi(b) = \mu, \tag{BC}$$

and the comparative statics conditions

$$\eta_1 a = 0, \tag{CS-a}$$

$$\eta_2 b = 0. \tag{CS-b}$$

There are three cases.

Case 1: $a = 0$ **and** $b \neq 0$**.** By CS-b, $\eta_2 = 0$. By FOC-b and FOC-a

$$\xi'(b^*) = \frac{1}{\lambda^*} \geq \frac{\partial \nu(a,\theta)}{\partial a}|_{a=0} \tag{$C1$}$$

Intuitively, it is optimal to devote all effort to exploit if the marginal product from exploiting is still higher than the marginal product from exploration even after achieving μ with only exploiting effort.

Case 2: $a \neq 0$ **and** $b \neq 0$**.** By CS-a, $\eta_1 = 0$ and by CS-b, $\eta_2 = 0$. By FOC-b and FOC-a,

$$\xi'(b)\big|_{b=b^*} = \frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^*} = \frac{1}{\lambda^*}. \tag{$C2$}$$

In **case 1** and **case 2**, we have

$$\frac{\partial \frac{1}{\lambda^*}}{\partial \theta} = \frac{\partial b^*}{\partial \theta}\xi''(b)\Big|_{b=b^*} \geq 0$$

$\frac{\partial b^*}{\partial \theta} \leq 0$ because when θ increases, the marginal product of exploratory effort increases. Since the two efforts are substitutes, to achieve the same μ, the agent will substitute away from exploiting. $\xi''(b) \leq 0$ because ξ is assumed to be concave. Hence,

$$\frac{\partial \lambda^*}{\partial \theta} \leq 0$$

Therefore, under either **Case 1** or **Case 2**, by Envelop Theorem, $\frac{\partial^2 e^*(\mu,\theta)}{\partial \mu \partial \theta} = \frac{\partial \lambda^*}{\partial \theta} \leq 0$.

Case 3: $a \neq 0$ **and** $b = 0$**.** By CS-a, $\eta_1 = 0$. By FOC-b and FOC-a

$$\frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^*} = \frac{1}{\lambda^*} \geq \xi'(0) \tag{$C3$}$$

Under **Case 3**, we have the following lemma:

Lemma 5. *Let $\bar{a}(\mu,\theta)$ be the solution to $\nu(a,\theta) = \mu$. Under Assumption 1, $\frac{\partial^2 \bar{a}(\mu,\theta)}{\partial \mu \partial \theta} \leq 0$.*

See Proof in Appendix E. In **Case 3**, since an agent devotes all efforts to exploration, we have $e(\mu,\theta) = a(\mu,\theta) = \bar{a}(\mu,\theta)$. By Lemma 5, $\frac{\partial^2 e(\mu,\theta)}{\partial \mu \partial \theta} = \frac{\partial^2 \bar{a}(\mu,\theta)}{\partial \mu \partial \theta} \leq 0$.

So far, we have shown $-e(\mu,\theta)$ has increasing difference under each case. However, to show that $-e(\mu,\theta)$ has increasing difference globally, we need to consider scenarios where the effort choice

switches from one case to another when θ or μ change. While we have shown that $C(\mu,\theta)$ is continuously differentiable with respect to μ, so does $e(\mu,\theta)$ because $C(\mu,\theta)$ is a convex transformation of it, in the following, we provide a more general proof of the increasing difference property, even when $e(\mu,\theta)$ is not necessarily differentiable at the points where the effort choice switches.

An algorithm to decide the optimal choice of effort. Now let $\psi_{\theta,\mu}(y) := \xi'(x(y))$, where $x(y) = \xi^{-1}(\mu - \nu(y,\theta))$ solves BC. Let $x_{\min} = 0$, $x_{\max} = \xi^{-1}(\mu)$, $y_{\min} = 0$, $y_{\max}(\theta,\mu)$ solves $\nu(y,\theta) = \mu$. Then $\psi_{\theta,\mu}(y)$ is defined on $[y_{\min}, y_{\max}]$ and it is decreasing in μ but increasing in y,θ. Moreover, $y_{\max}(\theta,\mu)$ is increasing in μ but decreasing in θ.

Lemma 6 (An algorithm to decide optimal effort choice)**.** *The optimal effort choice can be determined in the following way:*

1. *If* $\psi_{\theta,\mu}(0) \geq \frac{\partial\nu(y,\theta)}{\partial y}\Big|_{y=0}$*, then the optimal effort choice belongs to* ***case 1****, i.e., it is optimal to set* $a = 0$*.*
2. *If* $\psi_{\theta,\mu}(y_{max}(\theta,\mu)) \leq \frac{\partial\nu(y,\theta)}{\partial y}\Big|_{y=y_{max}(\theta,\mu)}$*, then the optimal effort choice belongs to* ***case 3****, i.e., it is optimal to set* $b = 0$*.*
3. *Otherwise if* $\psi_{\theta,\mu}(y)$ *and* $\frac{\partial\nu(y,\theta)}{\partial y}$ *intersect once within the interval* $[y_{min}, y_{max}]$*, the optimal effort choice falls into* ***case 2****.*

Conversely, the following is also true.

1. *If the optimal effort choice belongs to* ***case 1****, i.e., it is optimal to set* $a = 0$*, then* $\psi_{\theta,\mu}(0) \geq \frac{\partial\nu(y,\theta)}{\partial y}\Big|_{y=0}$*.*
2. *If the optimal effort choice belongs to* ***case 3****, i.e., it is optimal to set* $b = 0$*, then* $\psi_{\theta,\mu}(y_{max}(\theta,\mu)) \leq \frac{\partial\nu(y,\theta)}{\partial y}\Big|_{y=y_{max}(\theta,\mu)}$*.*
3. *If the optimal effort choice falls into* ***case 2****, then* $\psi_{\theta,\mu}(y)$ *and* $\frac{\partial\nu(y,\theta)}{\partial y}$ *intersect once within the interval* $[y_{min}, y_{max}]$*.*

Consider any $\mu > \mu'$, and $\theta > \theta'$. Our goal is to show that

$$e(\mu,\theta) - e(\mu',\theta) \leq e(\mu,\theta') - e(\mu',\theta') \tag{10}$$

We first state the following two lemmas.

Lemma 7. *For any* θ*, and any* $\mu > \mu'$*, the following combination of optimal effort choice is not possible:*

(i) the optimal effort choice under (μ, θ) *falls into* ***Case 3*** *and the optimal effort choice under* (μ', θ) *falls into* ***Case 2****.*

(ii) the optimal effort choice under (μ, θ) *falls into* ***Case 1*** *and the optimal effort choice under* (μ', θ) *falls into* ***Case 2****.*

(iii) the optimal effort choice under (μ, θ) *falls into* ***Case 1*** *and the optimal effort choice under* (μ', θ) *falls into* ***Case 3****.*

(iv) the optimal effort choice under (μ, θ) *falls into* ***Case 3*** *and the optimal effort choice under* (μ', θ) *falls into* ***Case 1****.*

See proof in Appendix F.

Lemma 8. *For any* μ*, and any* $\theta > \theta'$*, the following combination of optimal effort choice is not possible:*

(i) the optimal effort choice under (μ, θ) *falls into* ***Case 1*** *and the optimal effort choice under* (μ, θ') *falls into* ***Case 2****.*

(ii) the optimal effort choice under (μ, θ) *falls into* ***Case 1*** *and the optimal effort choice under* (μ, θ') *falls into* ***Case 3****.*

(iii) the optimal effort choice under (μ, θ) *falls into* ***Case 2*** *and the optimal effort choice under* (μ, θ') *falls into* ***Case 3****.*

See proof in Appendix G.

For any $\theta > \theta'$, $\mu > \mu'$, Table 2 presents all the possible case combinations. By Lemma 7 and 8, all the combinations marked with a capital letter are possible combinations, and impossible otherwise. Hence we only need to show for each possible combination, Eq. (10) holds.

Table 2: **Possible combinations**

	(1,1)	(2,2)	(3,3)	(1,2)	(3,2)	(1,3)	(2,1)	(2,3)	(3,1)
(1,1)	A	×	×	×	×	×	×	×	×
(2,2)	B	C	×	D	×	×	×	×	×
(3,3)	E	F	G	H	I	×	×	×	×
(1,2)	J	×	×	K	×	×	×	×	×
(3,2)	L	M	×	N	O	×	×	×	×
(1,3)	×	×	×	×	×	×	×	×	×
(2,1)	×	×	×	×	×	×	×	×	×
(2,3)	×	×	×	×	×	×	×	×	×
(3,1)	×	×	×	×	×	×	×	×	×

Note: The tuples in the first column present the cases of the optimal effort choice under (μ', θ) and (μ, θ). For example, $(1, 2)$ means that the optimal effort choices fall in to **Case 1** under (μ', θ) and **Case 2** under (μ, θ). Similarly, the tuples in the first row present the cases of the optimal effort choice under (μ', θ') and (μ, θ').

We prove in the following order: (1) A, C, G; (2) B, E, J, L; (3) F, K ,O; (4) M, I, D, H, N.

Step 1. A, C, G Note that the effort choices in A, C, G do not switch across cases. Hence, we have already shown above that for these combinations Eq. (10) is satisfied.

Step 2. B, E, J, L Since the effort choices fall into **Case 1** under (μ', θ') and (μ, θ'), $e(\mu', \theta') = b(\mu', \theta')$ and $e(\mu, \theta') = b(\mu, \theta')$. For type θ to attain μ, consider a (not necessarily optimal) effort choices of $a(\mu', \theta)$ and $b(\mu', \theta) + b(\mu, \theta') - b(\mu', \theta')$. The model fitness of this effort choice by type θ is

$$\begin{aligned}
&\nu(a(\mu', \theta), \theta) + \xi(b(\mu', \theta) + b(\mu, \theta') - b(\mu', \theta')) \\
=&\nu(a(\mu', \theta), \theta) + \xi(b(\mu', \theta)) + \int_{b(\mu', \theta)}^{b(\mu', \theta) + b(\mu, \theta') - b(\mu', \theta')} \xi'(b) db \\
\geq&\mu' + \int_{b(\mu', \theta')}^{b(\mu', \theta') + b(\mu, \theta') - b(\mu', \theta')} \xi'(b) db \\
=&\mu' + \xi(b(\mu, \theta')) - \xi(b(\mu', \theta')) \\
=&\mu' + \mu - \mu' \\
=&\mu
\end{aligned}$$

The first inequality is because 1) $b(\mu', \theta) \leq b(\mu', \theta')$ (otherwise type θ will produce a model fitness higher than μ'), and 2) ξ is concave so that $\xi'(b)$ is decreasing in b. Hence, the effort choice we

consider can attain fitness at least μ for type θ.

Since the optimal effort choice under (μ, θ) must be no more costly than any other effort choice that achieves μ,

$$\begin{aligned} e(\mu, \theta) &\leq a(\mu', \theta) + b(\mu', \theta) + b(\mu, \theta') - b(\mu', \theta') \\ &= e(\mu', \theta) + e(\mu, \theta') - e(\mu', \theta') \end{aligned}$$

Hence, Eq. (10) is satisfied for these cases.

Step 3. F, K, O We start with F. By the definition of F, $a(\mu, \theta), a(\mu', \theta) > 0$, $b(\mu, \theta) = b(\mu', \theta) = 0$, $a(\mu, \theta'), a(\mu', \theta')$, $b(\mu, \theta'), b(\mu', \theta') \geq 0$. Suppose Eq. (10) is false, that is,

$$a(\mu, \theta) - a(\mu', \theta) > a(\mu, \theta') - a(\mu', \theta') + b(\mu, \theta') - b(\mu', \theta') \tag{11}$$

Then

$$\begin{aligned} \mu - \mu' &= \int_{a(\mu', \theta)}^{a(\mu, \theta)} \frac{\partial \nu(a, \theta)}{\partial a} da \\ &> \int_{a(\mu', \theta)}^{a(\mu, \theta)} \xi'(0) da \\ &> \int_{a(\mu', \theta')}^{a(\mu, \theta')} \xi'(0) da + \int_{b(\mu', \theta')}^{b(\mu, \theta')} \xi'(0) db \\ &> \int_{a(\mu', \theta')}^{a(\mu, \theta')} \frac{\partial \nu(a, \theta')}{\partial a} da + \int_{b(\mu', \theta')}^{b(\mu, \theta')} \xi'(b) db \\ &= \mu - \mu' \end{aligned}$$

The first inequality is by Eq. ($C3$). The second inequality is by Eq. (11). The third inequality is by Eq. ($C2$) plus ξ is concave and thus ξ' is decreasing. Hence, we derive a contradiction.

Combinations K and O can be proved using similar method as Step 2.

Step 4. M, I, D, H, N These combinations can be proved by finding a marginal type. Take M for example.

Consider the marginal type θ_0 such that $\frac{\nu(a, \theta_0)}{a}\Big|_{a=a(\mu', \theta_0)} = \xi'(0)$. So that the optimal choices fall into either **Case 2** or **Case 3** under (μ', θ_0), and **Case 2** under (μ, θ_0). Moreoever, $\theta' \leq \theta_0 \leq \theta$.

We first consider θ_0 and θ. Type θ falls into **Case 3** under μ' and **Case 2** under μ. Since type θ_0 is marginal type between **Case 2** and **3** under μ', we think as if she falls into **Case 3** under μ' and **Case 2** under μ without loss of generality. Then the combination falls into Combination O

proved in Step 3. Hence,

$$e(\mu,\theta)-e(\mu',\theta)\leq e(\mu,\theta_0)-e(\mu',\theta_0)$$

Then we consider θ_0 and θ'. Instead we think as if type θ_0 falls into **Case 2** under both μ' and μ. By the same logic, the combination falls into Combination C proved in Step 1.

$$e(\mu,\theta_0)-e(\mu',\theta_0)\leq e(\mu,\theta')-e(\mu',\theta')$$

Hence Eq. (10) is straightforward. Other combinations can be proved in similar ways.

Since $c(\cdot)$ is convex, it's easy to show that

$$c(e(\mu,\theta))-c(e(\mu',\theta))\leq c(e(\mu,\theta'))-c(e(\mu',\theta'))$$

Hence, $-C(\mu,\theta)$ has increasing difference.

□

Proof of Proposition 1. To evoke the sufficient condition for the existence of symmetric monotone pure strategy equilibrium in Reny (2011), we first show that the payoff function 1) is bounded and continuous in μ_i for every $\theta_i\in\Theta$, and 2) has increasing differences in (μ_i,θ_i).

Continuity follows from continuity of the contest gain $g(\mu_i,\boldsymbol{\mu}_{-i})$ and the cost function $C(\mu_i,\theta_i)$. Boundedness follows from continuity of the payoff function and boundedness of μ_i. μ_i is bounded because the cost function is convex, hence for any type θ_i, the choice of μ_i should be finite. Increasing difference is immediate from Lemma 4.

By Theorem 4.5, the sufficient condition for the existence of symmetric monotone pure strategy equilibrium, in Reny (2011), there exists a pure strategy Nash equilibrium that is symmetric and monotone.

□

B Baseline

Proof of Proposition 2. **Optimization problem.** Each agent solves the following maximization problem

$$\begin{aligned}\max_{a,b}\quad & \mu-c(e)\\ \text{s.t.}\quad & a+b=e\\ & a,b\geq 0.\end{aligned}\tag{12}$$

Set up the Lagrangian

$$\tilde{\mathcal{L}}(a, b, \eta_1, \eta_2) = \nu(a, \theta) + \xi(b) - c(a + b) - \eta_1 a - \eta_2 b, \tag{13}$$

where $\eta_1, \eta_2 \geq 0$.

The solution satisfies the following conditions

$$\frac{\partial \tilde{\mathcal{L}}}{\partial a} = c'(e) + \eta_1, \tag{FOC-$a^\dagger$}$$

and

$$\frac{\partial \tilde{\mathcal{L}}}{\partial b} = c'(e) + \eta_2, \tag{FOC-$b^\dagger$}$$

the equality constraint (budget balanced condition)

$$a + b = e, \tag{BC†}$$

and the comparative statics conditions

$$\eta_1 a = 0, \tag{CS-$a^\dagger$}$$

$$\eta_2 b = 0. \tag{CS-$b^\dagger$}$$

There are three cases.

Case 1: $a = 0$ **and** $b \neq 0$**.** By CS-$b^\dagger$, $\eta_2 = 0$. By FOC-$b^\dagger$ and FOC-$a^\dagger$

$$\frac{\partial \nu(a, \theta)}{\partial a}|_{a=0} < \xi'(b^\dagger) = c'(b^\dagger) \tag{$B1$}$$

Intuitively, it is optimal to devote all effort to exploit exiting methods if the marginal product from exploitation is still higher than the marginal product from exploration. And the marginal product from exploitation should equal to the marginal cost.

Case 2: $a \neq 0$ **and** $b \neq 0$**.** By CS-$a^\dagger$, $\eta_1 = 0$ and by CS-$b^\dagger$, $\eta_2 = 0$. By FOC-$b^\dagger$ and FOC-$a^\dagger$,

$$\xi'(b^\dagger) = \frac{\partial \nu(a, \theta)}{\partial a}|_{a=a^\dagger} = c'(b^\dagger + a^\dagger). \tag{$B2$}$$

Case 3: $a \neq 0$ **and** $b = 0$**.** By CS-$a^\dagger$, $\eta_1 = 0$. By FOC-$b^\dagger$ and FOC-$a^\dagger$

$$\xi'(0) < \frac{\partial \nu(a, \theta)}{\partial a}|_{a=a^\dagger} = c'(a^\dagger) \tag{$B3$}$$

Interval structure. Suppose for some type θ', it is optimal to devote all effort to exploit exiting methods, i.e., **case 1** is true for type θ'. Then for any $\theta'' < \theta'$, by supermodularity of ν, $\frac{\partial \nu(a,\theta'')}{\partial a}|_{a=0} < \frac{\partial \nu(a,\theta')}{\partial a}|_{a=0}$. Hence, **case 1** is also true for type θ'', i.e., it is optimal to devote all effort to exploit exiting methods for type θ''. Define $\theta_1^\dagger = \sup\{\theta : \frac{\partial \nu(a,\theta)}{\partial a}|_{a=0} < \xi'(b^\dagger)\}$, where $b^\dagger > 0$ solves $\xi'(b^\dagger) = c'(b^\dagger)$.

Similarly, suppose for some type θ', it is optimal to devote all effort to exploration, i.e., **case 3** is true for type θ'. Then for any $\theta'' > \theta'$, by supermodularity of ν, $\frac{\partial \nu(a,\theta'')}{\partial a}|_{a=a^\dagger(\theta')} > \frac{\partial \nu(a,\theta')}{\partial a}|_{a=a^\dagger(\theta')}$. Hence, **case 3** is also true for type θ'', i.e., it is optimal to devote all effort to exploration for type θ''. Define $\theta_2^\dagger = \inf\{\theta : \xi'(0) < \frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^\dagger}\}$, where $a^\dagger > 0$ solves $\frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^\dagger} = c'(a^\dagger)$.

Next, we show that $\theta_2^\dagger > \theta_1^\dagger$. Since $\nu(a,\theta)$ is supermodular, $\frac{\partial \nu(a,\theta'')}{\partial a}$ is increasing in θ. Moreover $\nu(a,\theta)$ is twice differentiable and Assumption 3 together imply that

$$\begin{aligned} \theta_1^\dagger &= \max\{\theta : \frac{\partial \nu(a,\theta)}{\partial a}|_{a=0} < \xi'(b^\dagger)\} \\ \theta_2^\dagger &= \min\{\theta : \xi'(0) < \frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^\dagger}\}. \end{aligned} \tag{14}$$

Since $\xi()$ is concave, we have $\xi'(b^\dagger) < \xi'(0)$. Supermodularity of $\nu(a,\theta)$ therefore implies that $\theta_2^\dagger > \theta_1^\dagger$.

Moreover, for any $\theta_1^\dagger < \theta < \theta_2^\dagger$, only **case 2** can be true. This is directly implied by definitions of $\theta_1^\dagger$ and $\theta_2^\dagger$.

□

C Contest Rewards and Incentives

Proof of Proposition 3. For any given $\boldsymbol{\mu}_{-i}$ and θ_i, we write the payoff function as

$$\pi(\mu_i, \boldsymbol{R}) = \sum_{k=1}^{I} p_{\mu_i, \boldsymbol{\mu}_{-i}}(k) \cdot R_k + \mu_i - C(\mu_i, \theta_i),$$

recall that $p_{\mu_i, \boldsymbol{\mu}_{-i}}(k)$ is defined as the probability that player i ranks the kth place given his model fitness μ_i and other players' model fitnesses $\boldsymbol{\mu}_{-i}$. For any $\mu_i \geq \mu_i'$ and $\boldsymbol{R} \geqq \boldsymbol{R}'$,

$$
\begin{aligned}
&\big(\pi(\mu_i, \boldsymbol{R}) - \pi(\mu_i', \boldsymbol{R})\big) - \big(\pi(\mu_i, \boldsymbol{R}') - \pi(\mu_i', \boldsymbol{R}')\big) \\
=&\left(\sum_{k=1}^{I} p_{\mu_i, \boldsymbol{\mu}_{-i}}(k) \cdot R_k + \mu_i - C(\mu_i, \theta_i)\right) \\
-&\left(\sum_{k=1}^{I} p_{\mu_i', \boldsymbol{\mu}_{-i}}(k) \cdot R_k + \mu_i' - C(\mu_i', \theta_i)\right) \\
-&\left(\sum_{k=1}^{I} p_{\mu_i, \boldsymbol{\mu}_{-i}}(k) \cdot R_k' + \mu_i - C(\mu_i, \theta_i)\right) \\
+&\left(\sum_{k=1}^{I} p_{\mu_i', \boldsymbol{\mu}_{-i}}(k) \cdot R_k' + \mu_i' - C(\mu_i', \theta_i)\right) \\
=&\sum_{k=1}^{I} p_{\mu_i, \boldsymbol{\mu}_{-i}}(k) \cdot (R_k - R_k') - \sum_{k=1}^{I} p_{\mu_i', \boldsymbol{\mu}_{-i}}(k) \cdot (R_k - R_k') \\
\geq& 0
\end{aligned}
$$

The inequality follows from 1) the first order stochastic dominance relation between the two distributions $P_{\mu_i, \boldsymbol{\mu}_{-i}}(k)$ and $P_{\mu_i', \boldsymbol{\mu}_{-i}}(k)$ as is established in Lemma 3, and 2) the non-decreasing property of $R_k - R_k'$ that is immediate from the definition of partial order on reward vectors in Definition 1.

Hence, the payoff function $\pi(\mu_i, \boldsymbol{R})$ has increasing differences in $(\mu_i, \boldsymbol{R})$. Then we can evoke Theorem 4.5 in Reny (2011) again and conclude that there exists a symmetric pure strategy Nash equilibrium that is monotone in $\boldsymbol{R}$.

□

D Benchmark Hacking

Proof of Proposition 4. First we show that for any $\theta' < \theta$, if $a^*(\theta) = 0$, then $a^*(\theta') = 0$.

By Lemma 6, $a^*(\theta) = 0$ implies that $\psi_{\theta, \mu^*(\theta)}(0) \geq \left.\frac{\partial \nu(a,\theta)}{\partial a}\right|_{a=0}$.

In any symmetric monotone equilibrium, we have $\mu^*(\theta') \leq \mu^*(\theta)$. This implies that $\psi_{\theta', \mu^*(\theta')}(0) \geq \psi_{\theta, \mu^*(\theta)}(0)$, where $\psi_{\theta,\mu}(0) = \xi'(\xi^{-1}(\mu))$ is constant in θ and decreasing in μ. Moreover, by the supermodularity of ν, we have $\left.\frac{\partial \nu(a,\theta)}{\partial a}\right|_{a=0} > \left.\frac{\partial \nu(a,\theta')}{\partial a}\right|_{a=0}$. Hence, $\psi_{\theta', \mu^*(\theta')}(0) \geq \left.\frac{\partial \nu(a,\theta')}{\partial a}\right|_{a=0}$. By Lemma 6, we have $a^*(\theta') = 0$.

Define $\theta_1^* = \sup\{\theta : \psi_{\theta, \mu^*(\theta)}(0) \geq \left.\frac{\partial \nu(a,\theta)}{\partial a}\right|_{a=0}\} = \sup\{\theta : \xi'(\xi^{-1}(\mu^*(\theta))) \geq \left.\frac{\partial \nu(a,\theta)}{\partial a}\right|_{a=0}\}$. Notice that θ_1^* is well-defined because as θ increases, the LHS is decreasing while the RHS is increasing. Alternatively, since $b^*(\theta) = \xi^{-1}(\mu^*(\theta))$ when such a type only exerts exploitative effort, we can also

write $\theta_1^* = \max\{\theta : \xi'(b^*(\theta)) \geq \frac{\partial \nu(a,\theta)}{\partial a}\Big|_{a=0}\}$. Hence, the first bullet point is implied by Corollary 1 and that exploratory effort is zero.

Similarly, we want to show for any $\theta' < \theta$, if $b^*(\theta') = 0$, then $b^*(\theta) = 0$. To do this, we first define $a_{\max}(\theta)$ such that $\frac{\partial \nu(a,\theta)}{\partial a}\Big|_{a=a_{\max}(\theta)} = \xi'(0)$. That is, $a_{\max}(\theta)$ is the maximum possible exploratory effort a rational player of θ can exert without incurring any exploitative effort. Hence, the maximum possible fitness the player can attain without making exploitative effort is $\nu(a_{\max}(\theta), \theta)$. Moreover, since ν is concave in a and supermodular in (a, θ), $a_{\max}(\theta)$ is increasing in θ.

Notice that in the symmetric equilibrium of the contest, μ^* solves each agent's maximization problem Eq. ($\mathcal{P}_i'$), i.e., it solves the following equation for every θ and every i:

$$\frac{\partial g(\mu_i, \boldsymbol{\mu}_{-i})}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)} + 1 = \frac{\partial C(\mu_i, \theta)}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)}. \tag{15}$$

By Lemma 4, we know that $\frac{\partial C(\mu_i,\theta)}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)} = c'(e^*)\lambda^*$, where λ^* is the langrangian multiplier defined in Eq. (9). Moreover, we know that when agents only explore, Eq. ($C3$) is true, i.e., $\frac{1}{\lambda^*} = \frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^*}$. Hence we have

$$\frac{\partial C(\mu_i, \theta)}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)} = \frac{c'(e^*)}{\frac{\partial \nu(a,\theta)}{\partial a}|_{a=a^*}}, \tag{16}$$

For any $\theta' < \theta$, $\mu^*(\theta') < \mu^*(\theta)$. Since $\frac{\partial g(\mu_i,\boldsymbol{\mu}_{-i})}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)}$ is decreasing in $\mu^*(\theta)$, combining with Eq. (15), we have

$$\begin{aligned} \frac{\partial C(\mu_i, \theta')}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta')} &= \frac{\partial g(\mu_i, \boldsymbol{\mu}_{-i})}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta')} + 1 \\ > \frac{\partial g(\mu_i, \boldsymbol{\mu}_{-i})}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)} + 1 &= \frac{\partial C(\mu_i, \theta)}{\partial \mu_i}\Big|_{\mu_i=\mu^*(\theta)}. \end{aligned} \tag{17}$$

To show that $b^*(\theta) = 0$, we only need to show $\mu^*(\theta) < \nu(a_{\max}(\theta), \theta)$. We prove by contradiction. Suppose $\mu^*(\theta) > \nu(a_{\max}(\theta), \theta)$, then

$$
\begin{aligned}
\left.\frac{\partial C(\mu_i,\theta)}{\partial \mu_i}\right|_{\mu_i=\mu^*(\theta)} &> \left.\frac{\partial C(\mu_i,\theta)}{\partial \mu_i}\right|_{\mu_i=\nu(a_{\max}(\theta),\theta)} \\
&= \frac{c'(a_{\max}(\theta))}{\frac{\partial \nu(a,\theta)}{\partial a}|_{a=a_{\max}(\theta)}} \\
&= \frac{c'(a_{\max}(\theta))}{\xi'(0)} \\
&> \frac{c'(a_{\max}(\theta'))}{\xi'(0)} \\
&= \frac{c'(a_{\max}(\theta'))}{\frac{\partial \nu(a,\theta')}{\partial a}|_{a=a_{\max}(\theta')}} \\
&= \left.\frac{\partial C(\mu_i,\theta')}{\partial \mu_i}\right|_{\mu_i=\nu(a_{\max}(\theta'),\theta')} \\
&> \left.\frac{\partial C(\mu_i,\theta')}{\partial \mu_i}\right|_{\mu_i=\mu^*(\theta')}.
\end{aligned} \tag{18}
$$

The first inequality is because $C(\mu_i,\theta)$ is convex in μ_i (derived in Lemma 4), the first equality is by Eq. (16), the second equality is by the definition of $a_{\max}$, the second inequality is because $a_{\max}(\theta)$ is increasing in θ and c is convex, the last inequality is because $\xi(\theta')=0$ (then it has to be $\nu(a_{\max}(\theta'),\theta') > \mu^*(\theta')$) and $C(\mu_i,\theta)$ is convex in μ_i.

Since Eq. (18) contradicts Eq. (17), it has to be $\nu(a_{\max}(\theta),\theta) > \mu^*(\theta)$. Then we also have $b^*(\theta)=0$.

Next we argue that $\theta_1^* \le \theta_1^\dagger$. Notice that in the baseline, we have $\xi(b^\dagger)=\mu^\dagger$ for any $\theta \le \theta_1^\dagger$, while in the contest, we have $\xi(b^*(\theta))=\mu^*(\theta)$ for any $\theta \le \theta_1^*$. By Corollary 1, we know that $\mu^*(\theta) \ge \mu^\dagger$. Hence, we have $b^*(\theta) \ge b^\dagger$. This proves the second bullet point.

Notice that

$$
\left.\frac{\partial \nu(a,\theta_1^\dagger)}{\partial a}\right|_{a=0} = \xi'(b^\dagger) \ge \xi'(b^*(\theta)) = \left.\frac{\partial \nu(a,\theta_1^*)}{\partial a}\right|_{a=0}
$$

Comparing the definitions of $\theta_1^\dagger$ and θ_1^*, we can conclude that $\theta_1^\dagger \ge \theta_1^*$.

The last bullet point follows from Corollary 1 and that exploitative effort is zero in both settings.

□

E Proof of Lemma 5

Proof. Let $\theta_1 > \theta_2$. For any $\mu \ge 0$, choose a_1, a_2 such that $\nu(a_1,\theta_1)=\nu(a_2,\theta_2)=\mu$. By the definition of $\bar{a}(\mu,\theta)$, we have $a_1=\bar{a}(\mu,\theta_1), a_2=\bar{a}(\mu,\theta_2)$. Because $\nu(a,\theta)$ is increasing in both a and θ, and $\theta_1 \ge \theta_2$, then it's straightforward that $a_1 \le a_2$. Hence,

$$
\left.\frac{\partial \nu(a,\theta_1)}{\partial a}\right|_{a=a_1} \ge \left.\frac{\partial \nu(a,\theta_1)}{\partial a}\right|_{a=a_2} \ge \left.\frac{\partial \nu(a,\theta_2)}{\partial a}\right|_{a=a_2}
$$

The first inequality is because ν is concave in a, the second inequality is because ν is supermodular in a and θ. Then take the inverse of the inequalities above,

$$\begin{aligned}
&\frac{\partial a}{\partial \nu(a,\theta_1)}\Big|_{a=a_1} \leq \frac{\partial a}{\partial \nu(a,\theta_2)}\Big|_{a=a_2} \\
\Rightarrow\ &\frac{\partial \bar{a}(\mu,\theta_1)}{\partial \mu} \leq \frac{\partial \bar{a}(\mu,\theta_2)}{\partial \mu} \\
\Rightarrow\ &\frac{\partial^2 \bar{a}(\mu,\theta)}{\partial \mu \partial \theta} \leq 0
\end{aligned}$$

□

F Proof of Lemma 7

Proof. We prove by contradiction.

Case i: Since the optimal effort allocation falls into **Case 2** under (μ',θ), by Eq. (C2),

$$\xi'(b)\big|_{b=b(\mu',\theta)} = \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=a(\mu',\theta)} > \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=\bar{a}(\mu',\theta)},$$

where $a(\mu',\theta), b(\mu',\theta) > 0$ and $\bar{a}(\mu',\theta) > a(\mu',\theta)$. The inequality is because ν is concave in a and $\bar{a}(\mu',\theta) > a(\mu',\theta)$.

Since the optimal effort allocation falls into **Case 3** under (μ,θ), by Eq. (C3),

$$\frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=\bar{a}(\mu,\theta)} \geq \xi'(b)\big|_{b=0}$$

Since $\xi(b)$ is concave, we have $\xi'(b)\big|_{b=0} > \xi'(b)\big|_{b=b(\mu',\theta)}$. Hence,

$$\begin{aligned}
&\frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=\bar{a}(\mu,\theta)} \\
&\geq \xi'(b)\big|_{b=0} > \xi'(b)\big|_{b=b(\mu',\theta)} \\
&= \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=a(\mu',\theta)} \\
&> \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=\bar{a}(\mu',\theta)}
\end{aligned}$$

Since ν is concave in a, and $\bar{a}(\mu,\theta) > \bar{a}(\mu',\theta)$, we have

$$\frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=\bar{a}(\mu,\theta)} < \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=\bar{a}(\mu',\theta)}$$

By contradiction, there won't be such a senario.

Case ii: By the same method of proof as **Case i**, **Case ii** is impossible.

Case iii: Since the optimal effort choice under (μ, θ) falls into **Case 1**, by Eq. (*C*1),

$$\xi'(b)\big|_{b=b(\mu,\theta)} \geq \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=0}$$

Since the optimal effort choice under (μ', θ) falls into **Case 3**, by Eq. (*C*3),

$$\frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=a(\mu',\theta)} \geq \xi'(b)\big|_{b=0}$$

Since ξ is concave and $b(\mu, \theta) > 0$, we have

$$\frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=a(\mu',\theta)} \geq \xi'(b)\big|_{b=0} > \xi'(b)\big|_{b=b(\mu,\theta)} \geq \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=0}$$

Since ν is concave in a and $a(\mu', \theta) > 0$, we have

$$\frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=a(\mu',\theta)} < \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=0}$$

By contradiction, there won't be such a scenario.

Case iv: By the same method of proof as **Case iii**, **Case iv** is impossible. □

G Proof of Lemma 8

Proof. Let $\bar{b}(\mu)$ be such that $\xi(\bar{b}(\mu)) = \mu$. So that $\bar{b}(\mu)$ is the effort to achieve fitness μ if only exploitative effort is exerted.

Case i: If the optimal effort choice under (μ, θ) falls into **Case 1** and the optimal effort choice under (μ, θ') falls into **Case 2**,

$$\begin{aligned}
&\xi'(b)\big|_{b=\bar{b}(\mu)} \\
&\geq \frac{\partial \nu(a,\theta)}{\partial a}\big|_{a=0} > \frac{\partial \nu(a,\theta')}{\partial a}\big|_{a=0} > \frac{\partial \nu(a,\theta')}{\partial a}\big|_{a=a(\mu,\theta')} \\
&= \xi'(b)\big|_{b=b(\mu,\theta')} \\
&> \xi'(b)\big|_{b=\bar{b}(\mu)}
\end{aligned}$$

The first inequality is because the optimal effort choice under (μ, θ) falls into **Case 1**, the second inequality is because $\nu(a, \theta)$ is supermodular, the third inequality is because $\nu(a, \theta)$ is concave in a

and $a(\mu, \theta') > 0$, the first equality is because the optimal effort choice under (μ, θ') falls into **Case 2**, the last inequality is because $\xi(b)$ is concave in b and $b(\mu, \theta') < \bar{b}(\mu)$.

From above, we derive $\xi'(b)\big|_{b=\bar{b}(\mu)} > \xi'(b)\big|_{b=\bar{b}(\mu)}$, which is incorrect. Hence, **Case i** is impossible.

Case ii, iii: These two cases can be proved impossible using similar method. □

G.1 Examples

We present two examples and identify the set of agents who overfit in each of the example. In the first example, there is no *middle* types, and only *high* and *low* types. *High* types don't overfit, while *low* types do. In the second example, there are only *middle* types, where all types always spend positive effort on both actions, and all overfit. In the third example, there are no *high* types, and *middle* types engage in both actions and don't overfit. In the forth example, three types co-exist and *middle* types overfit as well.

Example 1 (No *middle* types)**.** *Consider $I \geq 2$ symmetric players. Let $\Theta = [0, 3]$, assume that $\mu = \theta a + b$ and $c(e) = \frac{1}{2}e^2$. Then $\theta_1^\dagger = \theta_2^\dagger = \theta_1^* = \theta_2^* = 1$.*

In this example, an agent engages solely in exploration if $\theta > 1$ and exploitation if $\theta < 1$. The cost of attaining fitness μ by type $\theta \geq 1$ is $C(\mu, \theta) = \frac{1}{2}\left(\frac{\mu}{\theta}\right)^2$ and by $\theta < 1$ is $C(\mu, \theta) = \frac{1}{2}\mu^2$.

In the baseline scenario, an agent's optimized choices of (μ, a, b) are $(\theta^2, \theta, 0)$ if $\theta > 1$ and $(1, 0, 1)$ if $\theta < 1$. In the contest, agents with $\theta > 1$ only dedicates to exploration and $\theta < 1$ only to exploitation. The intuition is, as the two efforts are perfect substitutes, agents always opt for the action where they exhibit higher productivity, given the total efforts exerted. Applying Proposition 4, any types $\theta < 1$ overfit and any types $\theta > 1$ don't.

Example 2 (Only *middle* types)**.** *Consider $I \geq 2$ symmetric players. Let $\Theta = [0, \infty)$. Assume that $\mu = \theta(a)^{\frac{1}{2}} + (b)^{\frac{1}{2}}$ and $c(e) = e$. Then $\theta_1^\dagger = \theta_1^* = 0$, $\theta_2^\dagger = \theta_2^* = \infty$.*

In the baseline, an agent's optimized choices of (μ, a, b) are $(\frac{1}{2}\theta^2 + \frac{1}{2}, \frac{1}{4}\theta^2, \frac{1}{4})$. Here is a symmetric monotone equilibrium under contest: any agent i with any type $\theta \in (0, \infty)$ engages in both of the two actions. Moreover, the ratio of the two efforts, $\frac{a^*(\theta)}{b^*(\theta)}$, is θ^2, coinciding with the ratio in the baseline. This fixed ratio property is attributed to the Cobb-Douglas functional form of μ. Formally, for any $\theta \in [0, \infty)$, $a^*(\theta) \geq a^\dagger(\theta)$ and $b^*(\theta) \geq b^\dagger(\theta)$. This is because the agent has to improve both efforts in order to attain higher model fitness in the contest while keep the effort ratio fixed. Hence in this example, all types overfit.

In the next example, middle types do not overfit.

Example 3 (No *high* types)**.** *Consider $I \geq 2$ symmetric players. Let $\Theta = [0, 3]$, assume that $\mu = \theta a + \sqrt{b}$ and $c(e) = \frac{1}{2}e^2$. To achieve fitness μ, type θ player chooses $a = 0$ if $\theta \leq \frac{1}{2\mu}$ and*

$a = \frac{\mu}{\theta} - \frac{1}{2\theta^2}$ *if* $\theta > \frac{1}{2\mu}$; $b = \mu^2$ *if* $\theta \leq \frac{1}{2\mu}$ *and* $b = \frac{1}{4\theta^2}$ *if* $\theta > \frac{1}{2\mu}$. *Hence*

$$C(\mu, \theta) = \begin{cases} \frac{1}{2}\mu^4 \text{ if } \theta \leq \frac{1}{2\mu} \\ \frac{1}{2}(\frac{\mu}{\theta} - \frac{1}{4\theta^2})^2 \text{ if } \theta > \frac{1}{2\mu} \end{cases}$$

In the baseline scenario, $\theta_1^\dagger = (\frac{1}{2})^{\frac{2}{3}}$ *and* $\theta_2^\dagger = \infty$. $\mu^\dagger(\theta) = (\frac{1}{2})^{\frac{1}{3}}$ *if* $\theta \leq \theta_1^\dagger$.

In the symmetric equilibrium under contest, middle types do not overfit, because the effort on exploitation for middle types only depends on type but not the fitness.

In the next example, middle types overfit as well.

Example 4. *Consider* $I \geq 2$ *symmetric players. Let* $\Theta = [0, 9]$. *Assume that* $\mu = \theta(-\exp(-a) + 1) + b$, $c(e) = \frac{1}{4}e^2$. *Then effort on exploration is* $a = -\ln(1 - \frac{\mu}{\theta})$ *if* $\theta > \mu + 1$, $a = \ln\theta$ *if* $1 < \theta \leq \mu + 1$, *and* $a = 0$, *if* $\theta \leq 1$. *The effort on exploitation is* $b = \mu$ *if* $\theta \leq 1$, $b = \mu - \theta + 1$ *if* $1 < \theta \leq \mu + 1$ *and* $b = 0$ *if* $\theta > \mu + 1$.

$$C(\mu, \theta) = \begin{cases} \frac{1}{4}\mu^2 \text{ if } \theta < 1 \\ \frac{1}{4}(\mu - \theta + 1 + \ln\theta)^2 \text{ if } 1 \leq \theta < 1 + \mu \\ \frac{1}{4}(\ln\frac{\theta}{\theta - \mu})^2 \text{ if } \theta > \mu + 1 \end{cases}$$

In the baseline scenario, $\mu^\dagger(\theta) = 2$ *if* $\theta < 1$, $\mu^\dagger(\theta) = 1 + \theta - \ln\theta$ *if* $\theta < e^2$, $2(\theta - \mu^\dagger(\theta)) = \ln\frac{\theta}{\theta - \mu^\dagger(\theta)}$ *if* $\theta \geq e^2$. *Hence* $\theta_1^\dagger = 1$, *i.e., player with a type below* 1 *only exert effort on exploitation, and* $\theta_2^\dagger = e^2$, *i.e., player with a type above* e^2 *exerts all effort on exploration, while player with a type in between exert positive effort on both actions.*

In the symmetric equilibrium under contest, middle types overfit because the optimal choice of effort for middle types on exploitation increases as fitness increases and the fitness under contest is higher than the finitess in the baseline scenario.

G.2 More Details on Empirical Evaluations

Data Set Description. We select contests in which the prediction task is binary classification and the prizes are in U.S. dollars.[12] Eventually, there are 26 contests in our sample. The unit of observation is a contest-player pair, and player is indexed by i and contest by j. Table 3 describes the variables we observe.

A player's type comes from an independent data source, thus is *not* determined by competition outcomes. Specifically, each player on Kaggle is evaluated on two categories of data science exper-

[12] We only include binary classification contests to ensure that the prediction accuracy is comparable across contests.

Table 3: **Variable Description**

Variables	Description
A. Contest level	
NumPlayer	Number of players that submit at least once
NumPrize	Number of prizes
PrizeValue	Monetary value of all prizes
Duration	Number of days a contest lasts
Year	Year in which a contest occurs
Ratio	The size of the preliminary test set to all test sets
MaxSubmission	Maximum number of submissions allowed per player per day
B. Player level	
Type	Categorical type ranging from 0 (lowest) to 4 (highest)
C. Contest-player level	
Fitness	A player's prediction accuracy on the main test set of a given contest
D. Contest-player-submission level	
Date	The date a submission occurs
Accuracy score	A player's prediction accuracy ranging from 0 to 100

tise: *Competition* and *Notebook*. The expertise on *Competition* is (solely) determined by a player's performance in Kaggle-hosted competitions, but the expertise on *Notebook* is determined *not* by how well a player did in competitions, but by the popularity (measured by the number of upvotes) of all public, mostly non-competition, scripts he has created and shared with the Kaggle community.[13] Since most, if not all, of Kaggle voters are experienced and knowledgeable data scientists capable of assessing script quality, we consider *Notebook* expertise to be a reasonable measure of a player's exploratory capability and in the meantime is not a result of competition performance. Hence, we use a player's expertise in *Notebook* as a proxy for his type.[14]

Table 4 presents the summary statistics.

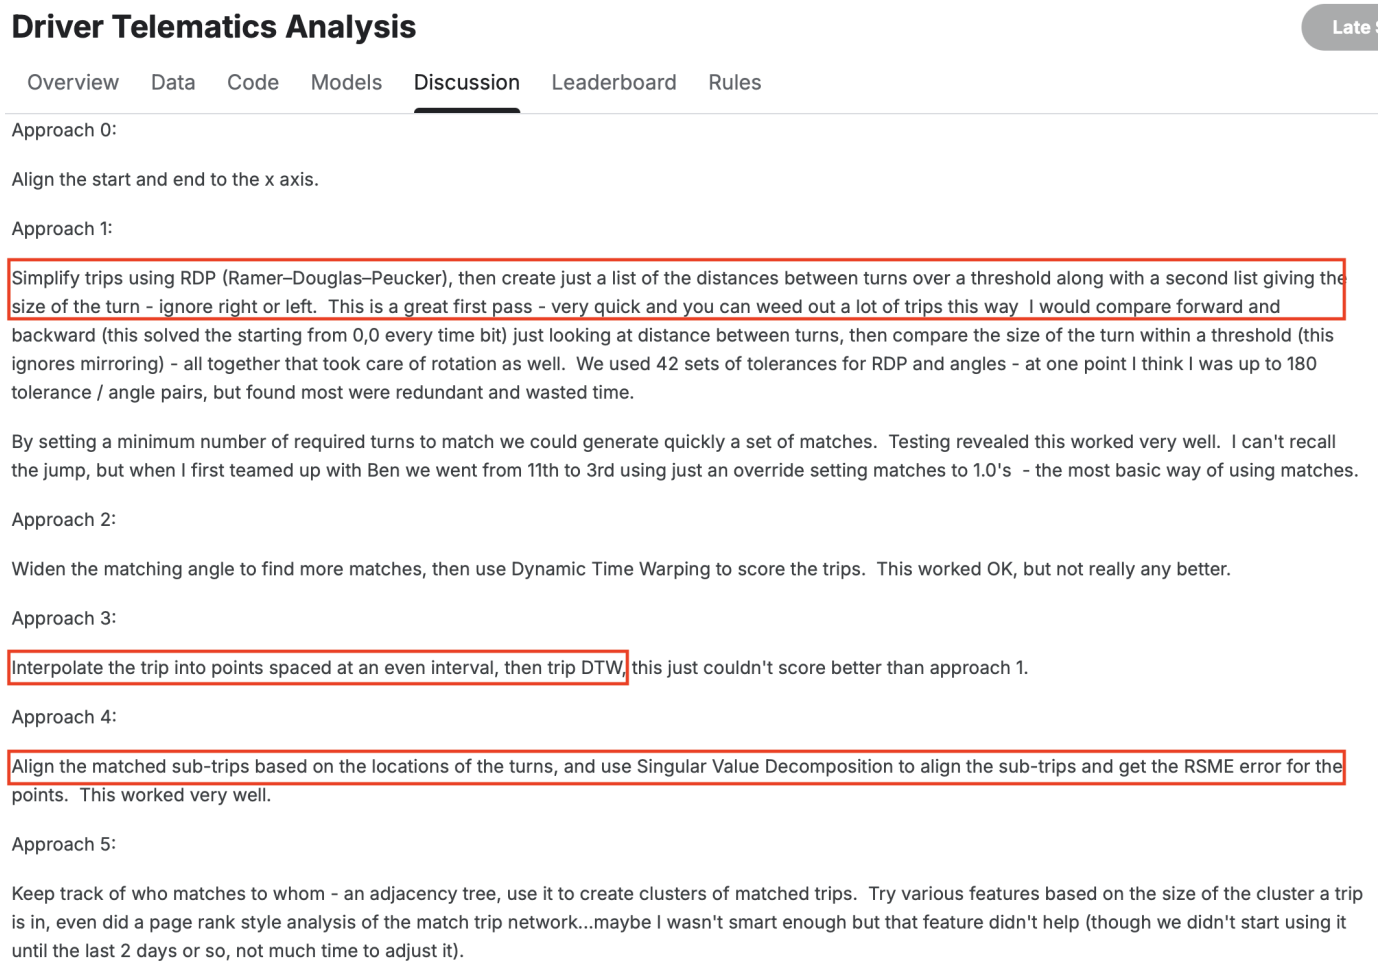
Coding Summary Examples. The following Figure 2 and 3 provide an example of exploration and exploitation from Kaggle.

Data Estimate Results. The following Table 5 describes full estimate data, part of which are presented in Table 1 of the main body.

[13]See details at https://www.kaggle.com/progression

[14]In Kaggle competitions, players usually form teams to participate (though are not required). We treat each team as a single player, and in case of multiple members in a team, we use the teamleader's type for the team's type.

Figure 2: **Exploratory Effort**



Driver Telematics Analysis

Late

Overview Data Code Models Discussion Leaderboard Rules

Approach 0:

Align the start and end to the x axis.

Approach 1:

Simplify trips using RDP (Ramer–Douglas–Peucker), then create just a list of the distances between turns over a threshold along with a second list giving the size of the turn - ignore right or left. This is a great first pass - very quick and you can weed out a lot of trips this way I would compare forward and backward (this solved the starting from 0,0 every time bit) just looking at distance between turns, then compare the size of the turn within a threshold (this ignores mirroring) - all together that took care of rotation as well. We used 42 sets of tolerances for RDP and angles - at one point I think I was up to 180 tolerance / angle pairs, but found most were redundant and wasted time.

By setting a minimum number of required turns to match we could generate quickly a set of matches. Testing revealed this worked very well. I can't recall the jump, but when I first teamed up with Ben we went from 11th to 3rd using just an override setting matches to 1.0's - the most basic way of using matches.

Approach 2:

Widen the matching angle to find more matches, then use Dynamic Time Warping to score the trips. This worked OK, but not really any better.

Approach 3:

Interpolate the trip into points spaced at an even interval, then trip DTW, this just couldn't score better than approach 1.

Approach 4:

Align the matched sub-trips based on the locations of the turns, and use Singular Value Decomposition to align the sub-trips and get the RSME error for the points. This worked very well.

Approach 5:

Keep track of who matches to whom - an adjacency tree, use it to create clusters of matched trips. Try various features based on the size of the cluster a trip is in, even did a page rank style analysis of the match trip network...maybe I wasn't smart enough but that feature didn't help (though we didn't start using it until the last 2 days or so, not much time to adjust it).

Note: This screenshot was taken from https://www.kaggle.com/competitions/axa-driver-telematics-analysis/discussion/12850. The highlighted texts exemplify exploratory effort.

Figure 3: **Exploitative Effort**

RSNA-MICCAI Brain Tumor Radiogenomic Classification

Overview Data Code Models Discussion Leaderboard Rules

255th place solution - Using 3D Resnet50 model

Although my model didn't generalise i'd still like to summarise my approach, this being my first kaggle competition.

Data preprocessing steps

- Load T2w Images
- Removed Blank slices
- Crop slices
- SIZ [1] depth 64
- Normalise
- Zero centre

Training

- Trained a Resnet50 3D with a LR scheduler (factor 10, tolerance 5) Initial LR=1e-2, batch size = 16
- Single fold cross validation
  Generally the results appeared like:
  Screenshot-from-2021-10-20-12-20-00.png
  Chose best 5 models from each training cycle and created a final ensemble for submission

[1] https://link.springer.com/chapter/10.1007%2F978-3-030-59354-4_15

Binary Classification Computer Vision Intermediate

Note: The highlighted texts exemplify exploitative effort.

Table 4: **Summary Statistics**

Variables	**Count**	**Mean**	**Std**	**Min**	**Max**
A.					
NumPlayer	26	1811.54	2233.00	26	8751
NumPrize	26	2.58	1.30	1	5
PrizeValue	26	27.32	30.64	0.15	100
Duration	26	68.23	25.67	1	121
Year	26	2014	3.18	2010	2020
Ratio, %	26	30.08	13.57	2	62
MaxSubmission	26	5.42	7.28	2	40
B.					
Type	35026	0.67	0.55	0	4
C.					
Fitness	46848	79.69	14.22	1.89	99.77
D.					
Accuracy	767844	81.39	1340.78	1.50	99.74

Note: Fitness ranges from 0 to 100 and represents the percentage of correct predictions. The maximum number of submissions each player can make is capped by the product of *MaxDailySubmission* and *Duration*. In the data, only 0.2% of observations hit the cap, so we assume that the submission cap doesn't bind.

Table 5: **Estimates of Coefficients**

TESTS VARIABLES	Prop. 1 Fitness	Prop. 3 Fitness	Prop. 4 MK stat	Cor. 3 MK stat
T_1	3.629***	6.652***	0.819***	0.787***
	(0.134)	(0.529)	(0.0315)	(0.124)
T_2	6.389***	11.32***	1.969***	1.844***
	(0.323)	(1.302)	(0.0724)	(0.292)
T_3	7.380***	13.82***	2.453***	2.317***
	(0.670)	(2.702)	(0.150)	(0.600)
T_4	8.761***	13.52***	3.100***	4.361***
	(0.839)	(3.451)	(0.184)	(0.752)
$T_0 \times PV$		0.475***		0.0316***
		(0.0207)		(0.00476)
$T_1 \times PV$		0.445***		0.0285***
		(0.0198)		(0.00453)
$T_2 \times PV$		0.415***		0.0258***
		(0.0240)		(0.00545)
$T_3 \times PV$		0.396***		0.0234***
		(0.0356)		(0.00796)
$T_4 \times PV$		0.408***		0.0118
		(0.0421)		(0.00930)
$T_0 \times PS$		39.65***		1.969***
		(1.173)		(0.272)
$T_1 \times PS$		37.56***		2.318***
		(1.153)		(0.267)
$T_2 \times PS$		37.09***		2.712***
		(1.409)		(0.321)
$T_3 \times PS$		36.31***		2.927***
		(2.174)		(0.489)
$T_4 \times PS$		37.64***		1.645***
		(2.748)		(0.601)
Constant	65.65***	22.83***	-0.238	-2.672***
	(0.847)	(2.058)	(0.193)	(0.472)
Observations	45,103	45,103	38,102	38,102
R-squared	0.328	0.329	0.069	0.072

Note: ***p<0.01, **p<0.05, *p<0.1