

UCLA

UCLA Electronic Theses and Dissertations

Title

Incentivizing Truthful Language Models via Peer Elicitation Games

Permalink

<https://escholarship.org/uc/item/9qf018tz>

ISBN

9798265483089

Author

Chen, Baiting

Publication Date

2025-12-10

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Incentivizing Truthful Language Models via Peer Elicitation Games

A thesis submitted in partial satisfaction
of the requirements for the degree
Master of Science in Statistics

by

Baiting Chen

2025

© Copyright by

Baiting Chen

2025

ABSTRACT OF THE THESIS

Incentivizing Truthful Language Models via Peer Elicitation Games

by

Baiting Chen

Master of Science in Statistics

University of California, Los Angeles, 2025

Xiaowu Dai, Chair

Large Language Models (LLMs) have demonstrated strong generative capabilities but remain prone to inconsistencies and hallucinations. We introduce Peer Elicitation Games (PEG), a training-free, game-theoretic framework for aligning LLMs through a peer elicitation mechanism involving a generator and multiple discriminators instantiated from distinct base models. Discriminators interact in a peer evaluation setting, where utilities are computed using a determinant-based mutual information score that provably incentivizes truthful reporting without requiring ground-truth labels. We establish theoretical guarantees showing that each agent, via online learning, achieves sublinear regret in the sense their cumulative performance approaches that of the best fixed truthful strategy in hindsight. Moreover, we prove last-iterate convergence to a truthful Nash equilibrium, ensuring that the actual policies used by agents converge to stable and truthful behavior over time. Empirical evaluations across multiple benchmarks demonstrate significant improvements in factual accuracy. These results position PEG as a practical approach for eliciting truthful behavior from LLMs without supervision or fine-tuning.

The thesis of Baiting Chen is approved.

Oscar Leong

Yuhua Zhu

Xiaowu Dai, Committee Chair

University of California, Los Angeles

2025

To my parents, with gratitude for your endless love, trust and support.

Contents

Abstract	ii
List of Figures	1
1 Introduction	2
1.1 Background	2
1.2 Related Literature	5
2 Methodology	6
2.1 Truthfulness as Incentive Compatibility	6
2.2 Peer Elicitation Games (PEG)	8
3 Theoretical Guarantees	12
3.1 Dominant Truthfulness	12
3.2 Regret Analysis	17
3.3 Last-Iterate Convergence	19
4 Experiments	24
4.1 Experiment Setup	24
4.2 Experiment Results	25
5 Conclusions	29
Appendices	30
A Proofs of Theorem 1	31
B Experiments Details	37

List of Figures

1.1	Comparison of the consensus game and PEG: when multiple discriminator LLMs independently evaluate the generator’s output and are rewarded based on mutual agreement, their collective judgment aligns more closely with true answers.	4
2.1	Overview of our method: multiple discriminators independently evaluate the response provided by the generator, while each discriminator is rewarded based on mutual agreement with peers via PEG. This setup incentivizes truthful reporting for discriminators and aligns the generator without requiring ground-truth labels.	9
4.1	Accuracy comparison between original model outputs (D) and PEG majority vote answers.	27

Chapter 1

Introduction

1.1 Background

LLMs have achieved remarkable progress in natural language generation, reasoning, and few-shot learning [1, 32, 37, 62]. Despite these advances, they remain fundamentally limited by hallucination—the generation of outputs that are inconsistent with previous responses or factually incorrect [3, 25, 26, 40]. A common failure mode is inconsistency: LLMs may produce different answers to semantically equivalent prompts due to stochastic decoding or sensitivity to minor variations in phrasing [12, 19, 40, 61]. Another major challenge is lack of truthfulness, where outputs are not semantically accurate or aligned with verifiable facts, even when the relevant knowledge is implicitly encoded in the model [3, 17]. These limitations are particularly concerning in high-stakes domains such as scientific discovery, education, and decision-making, where outputs must be both consistent and truthful [25, 59]. This motivates a central question: *How can we reliably elicit consistent and truthful behavior from LLMs?*

To address this question, a growing line of work has explored post-training alignment techniques, such as supervised fine-tuning and reinforcement learning with human feedback [4, 52, 54]. While these methods can be effective, they are typically computationally intensive, require extensive human annotation, and lack theoretical guarantees of truthful behavior

[31, 47, 52]. Their dependence on model internals also limits their scalability and transferability across different LLMs [9, 61].

An alternative line of research draws on game-theoretic frameworks that aim to improve LLM reliability through structured multi-agent interactions [13, 27]. These approaches offer the advantage of training-free alignment using only black-box access to LLMs. For example, the *consensus game* aligns a generator and a discriminator by rewarding agreement between their outputs [27]. However, because the objective is based on mutual agreement, it can lead to uninformative or collusive equilibria, where agents reinforce each other’s responses even when those responses are factually incorrect.

We propose *Peer Elicitation Games* (PEG), a training-free, game-theoretic framework for aligning LLMs through structured peer evaluation. In PEG, a generator produces candidate responses to prompts, and multiple independently instantiated LLMs act as discriminators. Each discriminator assesses the generator’s output and is, in turn, evaluated by the other discriminators—who serve as peer referees. This mutual evaluation mechanism assigns utilities based on the level of agreement among discriminators, encouraging truthful reporting without relying on ground-truth labels. This incentive structure ensures that truthful behavior by each discriminator constitutes a Nash equilibrium, while discouraging collusion or uninformative consensus. To implement the framework, we apply the online mirror descent algorithm to iteratively update each discriminator’s policy, enabling the system to converge toward equilibrium through repeated, utility-driven interaction. The majority vote among discriminators is then returned to the generator, serving as a feedback signal for improving future generations. The design of PEG—introducing multiple LLM discriminators and rewarding them through peer evaluation—is inspired by renowned concepts in biology, including cognitive synergy [38, 55] and collective intelligence [7, 64], where diverse agents, each holding partial or noisy information, can collectively arrive at judgments that surpass what any individual could achieve alone. PEG enables language models to self-organize toward truthful and stable outputs using only local incentives. This collective dynamic offers

a scalable and supervision-free approach to building more trustworthy LLM systems.

To summarize, the main contribution of our work is as follows:

- We propose PEG, a training-free framework for eliciting truthful behavior from LLMs, without relying on ground-truth labels or model fine-tuning. The framework casts the interaction between a generator and multiple heterogeneous discriminators as a multi-agent peer evaluation game as depicted in Figure 1.1.
- We provide theoretical guarantees showing that PEG promotes truthful reporting as a Nash equilibrium. Furthermore, when each discriminator updates its policy using online mirror descent, the system achieves sublinear regret and converges to a truthful Nash equilibrium over repeated interactions.
- Through experiments on a range of benchmarks, including ARC, MMLU, and GPQA, we demonstrate that PEG improves factual accuracy by over 10% compared to existing methods. Additionally, our results show that smaller models (e.g., 7B parameters) using PEG can match or outperform much larger models (e.g., 65B).




Figure 1.1: Comparison of the consensus game and PEG: when multiple discriminator LLMs independently evaluate the generator’s output and are rewarded based on mutual agreement, their collective judgment aligns more closely with true answers.

1.2 Related Literature

This work relates to several lines of research. *First*, game-theoretic frameworks have been explored to improve reasoning and alignment in LLMs through structured interactions [13]. Most relevant is the Consensus Game framework [27], which promotes self-consistency by reconciling the outputs of a generator LLM and a discriminator LLM through game-theoretic interaction. Follow-up studies extend this perspective to decoding [11], federated learning with competing interests [65], and embodied tasks like vision-language navigation [67]. Our work builds on this foundation by introducing a multi-agent formulation with explicit peer interaction and a utility mechanism that promotes truthful reporting rather than mere agreement. *Second*, a growing body of research demonstrates that multi-agent systems of LLMs can collaborate through debate and cooperation to enhance factuality and task performance [15, 36, 44]. Our approach draws on this idea by coordinating multiple LLMs through structured peer evaluation. *Third*, our work connects to the literature on learning Nash equilibria in multi-agent systems, particularly no-regret learning in general-sum games [8, 24, 60] and preference-based learning frameworks such as Nash learning from human feedback [41] and direct preference optimization [58]. We contribute to this literature by showing how no-regret learning dynamics converge to truthful equilibria in a structured, incentive-compatible setting. *Finally*, PEG builds on peer prediction mechanisms [39, 43], which have been extended to handle complex settings [46, 49, 63, 68]. We adapt these ideas to LLM agents by extending multi-task peer prediction mechanisms [2, 33] to a setting where multiple LLMs jointly evaluate responses via mutual scoring and incentive-aligned interaction.

Chapter 2

Methodology

This section introduces the key components of PEG including a formal definition of truthfulness and an overview of the interactions between the generator and multiple discriminator agents.

2.1 Truthfulness as Incentive Compatibility

We adopt the concept of *incentive compatibility* (IC) from mechanism design [42] to define the truthfulness. IC ensures that each agent maximizes its utility by reporting its true private information. Specifically, let c_i denote the true private information of agent i , and r_i a possible report. In our setting, the agents are discriminators, each independently evaluating a generated response. The true private information $c_i \in \{0, 1\}$ refers to the agent's truthful judgment on whether the response is factually correct ($c_i = 1$) or incorrect ($c_i = 0$). The report $r_i \in \{0, 1\}$ is the label that the agent chooses to submit to the system based on (or possibly deviating from) this truthful judgment. Let c_{-i} represent the truthful information of all other agents. The mechanism assigns utility $u_i(r_i, r_{-i})$ to agent i based on its own report and the reports of all other agents. A mechanism is said to be *incentive compatible (IC)* if, for all agents i , all c_i , c_{-i} , and all possible reports r_i ,

$$u_i(c_i, c_{-i}) \geq u_i(r_i, c_{-i}). \quad (2.1)$$

The consensus game framework [27] involves one LLM generator and one LLM discriminator, aiming to improve consistency by rewarding agreement between agents. However, this structure does not guarantee truthfulness, as agents are rewarded only for agreement rather than accuracy and may benefit from reporting false but mutually agreeable outputs. For example, the generator produces an incorrect answer to a question when tasked with generating a correct one. The discriminator, whose truthful evaluation would be to mark the answer as incorrect ($c_i = 0$), may instead receive a higher utility for agreeing with the generator by reporting it as correct ($r_i = 1$). In such a scenario, the utility of truthful reporting is lower than that of misreporting, i.e., $u_i(c_i, c_{-i}) < u_i(r_i, c_{-i})$, which creates a clear incentive for the discriminator to misreport. This violates the IC property defined in Eq. (2.1), highlighting a fundamental limitation of consensus-based approaches.

In contrast, recent studies have shown that multi-agent debates can better integrate the diverse perspectives of multiple models, leading to more accurate and reliable outputs [10, 15]. Motivated by these findings, we propose PEG in which multiple discriminator agents independently evaluate each generated response. In this setup, each discriminator is rewarded based on agreement with its peer discriminators. We show that when all other discriminators report truthfully, the best response for any individual discriminator is also to report truthfully—thus, truthful reporting becomes a Nash equilibrium. That is, PEG achieves IC property defined in (2.1), where no agent can improve its utility by deviating unilaterally. Figure 1.1 illustrates a key distinction between our PEG and the consensus game, where the generator and a single discriminator are incentivized to align with each other, which can lead to consistent but incorrect outputs. In contrast, PEG relies on independent mutual evaluations from multiple discriminators, promoting outputs that are not only consistent but also correct outputs.

2.2 Peer Elicitation Games (PEG)

We consider a generator G and a set of n discriminators $D_1, D_2, \dots, D_n$. Each generator and discriminator maintains a probabilistic policy. At each round $t \in \{1, 2, \dots, T\}$, the system assigns a set of tasks indexed by $k_t = 1, 2, \dots, K_t$, where K_t is the number of tasks in round t . For each task (t, k_t) , the generator receives an input question X_{t,k_t} along with a correctness label $V_{t,k_t} \in \{0, 1\}$, and produces a response Y_{t,k_t} according to its probabilistic policy. The generator’s policy is a conditional distribution over responses given the input and target label:

$$\pi_G(Y_{t,k_t} \mid X_{t,k_t}, V_{t,k_t}) = \mathbb{P}(Y_{t,k_t} \mid X_{t,k_t}, V_{t,k_t}; \theta_G),$$

where θ_G denotes the generator’s parameters.

Each discriminator D_i observes X_{t,k_t} and the corresponding generator response Y_{t,k_t} , and outputs a predicted label $V_{i,t,k_t} \in \{0, 1\}$ according to its policy:

$$\pi_{D_i}(V_{i,t,k_t} \mid X_{t,k_t}, Y_{t,k_t}) = \mathbb{P}(V_{i,t,k_t} \mid X_{t,k_t}, Y_{t,k_t}; \phi_i),$$

where ϕ_i denotes the parameters of discriminator D_i .

The goal is to generate responses that are both consistent and truthful. Our method achieves this by aligning the generator’s output with the majority judgment of discriminators, while incentivizing truthful evaluation from discriminators through peer evaluation. The overview of our method is illustrated in Figure 2.1. The left branch (in red) corresponds to the supervised alignment goal of imitating majority vote labels, while the right branch (in blue) depicts the PEG that ensures incentive compatibility among discriminators.

Step 1: Response Generation and Discriminator Evaluation. At each round, given an input question X_{t,k_t} and a correctness label V_{t,k_t} , the generator produces a response $Y_{t,k_t} \sim \pi_G(\cdot \mid X_{t,k_t}, V_{t,k_t})$. Each discriminator independently evaluates the response based on the question and outputs a correctness report $V_{i,t,k_t} \sim \pi_{D_i}(\cdot \mid X_{t,k_t}, Y_{t,k_t})$.




Figure 2.1: Overview of our method: multiple discriminators independently evaluate the response provided by the generator, while each discriminator is rewarded based on mutual agreement with peers via PEG. This setup incentivizes truthful reporting for discriminators and aligns the generator without requiring ground-truth labels.

Step 2: Peer Elicitation Games. To incentivize truthful reporting, the discriminator agents engage in a peer elicitation game, where the utility of each agent is based on a mutual evaluation. Importantly, this utility requires only reports from discriminators and does not rely on access to ground truth. Assume that in round t , all discriminators are assigned K_t tasks. For each task $k_t \in [K_t]$, discriminator i privately observes a signal $C_{i,t,k_t} \in \{0, 1\}$ and submits a report $V_{i,t,k_t} \in \{0, 1\}$. The task set $\{1, \dots, K_t\}$ is arbitrarily partitioned into two disjoint subsets $\mathcal{K}_{1,t}$ and $\mathcal{K}_{2,t}$. For every pair of distinct agents $i \neq j \in [n]$, we construct two 2×2 co-report matrices $\mathbf{M}_{1,t}^{ij}$ and $\mathbf{M}_{2,t}^{ij}$, one for each subset $\mathcal{K}_{\ell,t}$ where $\ell = 1, 2$. Each entry of $\mathbf{M}_{\ell,t}^{ij}(c, c')$ counts the number of tasks $k_t \in \mathcal{K}_{\ell,t}$ for which agents i and j reported the pair (c, c') , i.e.,

$$\mathbf{M}_{\ell,t}^{ij}(c, c') := \sum_{k_t \in \mathcal{K}_{\ell,t}} \mathbb{1}((V_{i,t,k_t}, V_{j,t,k_t}) = (c, c')).$$

The total payment to agent i in round t is defined as the sum over all other agents of the product of the determinants of these matrices:

$$p_{i,t} := \sum_{j \neq i} \det(\mathbf{M}_{1,t}^{ij}) \cdot \det(\mathbf{M}_{2,t}^{ij}).$$

In our setting, the discriminators are rewarded based on mutual evaluation. Meanwhile, the

generator is incentivized to produce responses that align with the majority consensus among high-confidence discriminators. The utility functions are computed over a batch of K_t tasks in round t . The overall utility functions are defined as:

$$\begin{aligned} u_{D_i}^{(t)} &= \mathbb{E}_{V_{i,t,k_t} \sim \pi_{D_i}(\cdot | X_{t,k_t}, Y_{t,k_t})} [p_{i,t}], \quad \forall i \in \{1, \dots, n\}, \\ u_G^{(t)} &= \mathbb{E}_{Y_{t,k_t} \sim \pi_G(\cdot | X_{t,k_t}, V_{t,k_t})} \left[\mathbb{1}(V_{t,k_t} = \hat{V}_{t,k_t}) \right], \end{aligned} \quad (2.2)$$

where $k_t \in \{1, \dots, K_t\}$ indexes the tasks in round t , and $\hat{V}_{t,k_t}$ denotes the majority vote label aggregated from the discriminator reports on task (t, k_t) .

Step 3: Policy Update. We assume that the policy space consists of conditional probability distributions over inputs, and focus on a localized subset around a reference policy π^* (e.g., a truthful reporting strategy) [50]. This constraint reflects the intuition that although policies may adapt to optimize for high utility, they should not obviate too far from truthful behavior to maintain semantic consistency and interpretability. In practical terms, this can be enforced by initializing the agent with a fine-tuned model that embodies π^* , and constraining updates to remain within a trust region [45, 69]. Specifically, we define the local policy neighborhood as: $\Pi_{\text{local}} = \{\pi(\cdot | \text{input}; \theta) \in \Pi \mid \|\pi(\cdot | \text{input}; \theta) - \pi^*(\cdot | \text{input})\| < \delta\}$. Our goal is to iteratively learn the policy that maximizes the utility function, aligning with the online learning framework that sequentially optimizes an objective function by searching for its critical point [22, 51]. Since each policy is a probability distribution, we adopt the Online Mirror Descent (OMD) algorithm with negative entropy as the Bregman divergence, which naturally preserves the probabilistic structure of the policy and regularizes each update by penalizing large deviations from the previous policy [16]. The update rules for the discriminators and generator are:

$$\begin{aligned} \pi_{D_i}^{(t'+1)}(V_{i,t,k_t} \mid X_{t,k_t}, Y_{t,k_t}) &\propto \pi_{D_i}^{(t')}(V_{i,t,k_t} \mid X_{t,k_t}, Y_{t,k_t}) \exp \left[\eta_t^{D_i} \nabla_{\pi_{D_i}^{(t')}(V_{i,t,k_t} \mid X_{t,k_t}, Y_{t,k_t})} u_{D_i}^{(t')} \right], \\ \pi_G^{(t'+1)}(Y_{t,k_t} \mid X_{t,k_t}, V_{t,k_t}) &\propto \pi_G^{(t')}(Y_{t,k_t} \mid X_{t,k_t}, V_{t,k_t}) \exp \left[\eta_t^G \nabla_{\pi_G^{(t')}(Y_{t,k_t} \mid X_{t,k_t}, V_{t,k_t})} u_G^{(t')} \right]. \end{aligned} \quad (2.3)$$

Intuitively, these updates guide each agent to adjust its policy to direction yields higher utility. After updating, the individual judgments $\{V_{1,t,k_t}, \dots, V_{n,t,k_t}\}$ from all discriminators on task (t, k_t) are aggregated via majority vote to form a consensus label $\hat{V}_{t,k_t}$, which serves as a proxy for correctness.

Step 4: Generator Update. The utility of the generator is determined by whether its output aligns with the consensus label derived from the discriminators. Specifically, the generator receives utilities when its generated response matches the consensus label $\hat{V}_{t,k_t}$. This design enables learning without requiring supervised fine-tuning or access to explicit ground-truth correctness labels.

Chapter 3

Theoretical Guarantees

In this section, we present three main theoretical results: (i) the mechanism incentivizes dominantly truthful reporting in Section 3.1; (ii) both the generator and discriminators achieve sublinear regret under online learning dynamics in Section 3.2; and (iii) last-iterate converges to a truthful Nash equilibrium in Section 3.3. Here, last-iterate convergence means that the policy used in the final iteration converges to the equilibrium, rather than requiring averaging over past iterations [6].

3.1 Dominant Truthfulness

PEG satisfies the dominant truthfulness property, meaning that truthful reporting is a dominant strategy for each discriminator, as it yields the highest expected utility regardless of the strategies chosen by other agents. This ensures that truthful behavior is consistently incentivized. We now formally state this property of PEG.

Lemma 1. *Let n be the number of agents (e.g., discriminators) and K_t be the number of tasks assigned in round t as defined in Section 2.2. When $n \geq 2$ and $K_t \geq 4$, under mild assumptions, PEG is dominantly truthful and satisfies IC in Eq. (2.1). That is, for every agent i , the truthful reporting strategy maximizes their expected payment regardless of the*

strategies chosen by other agents.

The design of PEG is grounded in three principles to ensure this incentive guarantee. First, the use of determinant-based utility ensures information-monotonicity: the determinant achieves its maximum value when agents report truthfully, making truthful reporting the most rewarding strategy. Second, the utility function acts as an unbiased estimator of the joint distribution of agent reports, allowing the mechanism to approximate the true distribution without estimation error. Finally, to prevent negative payments, the set of tasks is divided into two disjoint subsets, and payments are computed using the product of determinants from these subsets. As both subsets serve as unbiased estimators of the distribution of agent reports, their determinants are expected to have the same sign. As a result, the payment is always non-negative.

Formally, we consider a binary-choice setting, where each task consists of two possible outcomes, denoted as $\{0, 1\}$. There are n agents. Each agent i is assigned K binary-choice tasks, where for each task k , agent i receives a private signal $c_i^k \in \{0, 1\}$. These private signals for all agents are drawn from a joint unknown prior distribution $U_{[n]}^k \in \Delta_{[n]}^2$, where $\Delta_{[n]}^2$ represents the set of all measurable distributions over $\{0, 1\}^n$. For the same task, the private signals of different agents are correlated. For different tasks, the private signals of the same agent independent. In our setting, discriminators do not incur effort to acquire their private signals.

For a multi-task peer evaluation mechanism, the agents do not know the specific realization of the prior distribution before receiving the private signals. After receiving the private signal c_i^k , each agent is required to report r_i^k , which may or may not reflect their true signal. The mechanism is designed to incentivize agents to truthfully report their signals. We first present three assumptions and corresponding justifications.

Assumption 1. [A Priori Similar Tasks] We assume that all tasks are drawn from a common unknown prior distribution $U_{[n]}$ such that $U_{[n]}^k = U_{[n]}$ for all tasks k .

This assumption states that all tasks are fundamentally similar, implying that the signals

across different tasks should follow the same joint distribution. Specifically, the signals $\mathbf{c}^k$ for each task k are assumed to be i.i.d.. While traditional single-task peer evaluation typically assumes a homogeneous prior, this multi-task setting allows for heterogeneity in agents' beliefs as long as the tasks themselves are similar. In our setting, this assumption implies that the problems are similar in the sense that the ground truths follow the same joint distribution within each batch at every round. We enforce this by ensuring that the same set of subjects and the same set of problems are presented in each round.

Definition 1 (Strategy). Each agent i 's strategy for reporting is a mapping from her private signal c_i^k to a distribution over possible reports r_i^k . Formally, a strategy can be represented as a function $S_i^k : \{0, 1\} \rightarrow \Delta(\{0, 1\})$, where $S_i^k(c_i^k)$ gives the probability distribution over the possible reports r_i^k conditioned on receiving the private signal c_i^k .

Every strategy S_i^k corresponds to a 2×2 transition matrix where $S_i^k(c_i^k, r_i^k)$ is the probability that agent i reports r_i^k given that she receives private signal c_i^k . A strategy is *truthful* if the agent always reports $r_i^k = c_i^k$. Agent i plays a *truthful strategy* if for every task k , S_i^k is an identity matrix.

Assumption 2 (Consistent Strategy). Each agent i plays the same strategy S_i for all tasks.

This assumption is reasonable because agents face structurally identical tasks drawn from the same distribution, with no task-specific information to adapt to.

Assumption 3 (Conditional Independence). We assume that agents' private signals $c_1, c_2, \dots, c_n$ are independent conditioning on ground truth. Since agents' strategies are independent, this also implies that agent's reports $\hat{c}_1, \hat{c}_2, \dots, \hat{c}_n$ are independent conditioning on ground truth.

In our case, the discriminators (agents) are instantiated independently and process inputs separately, so their reports are naturally conditionally independent given the underlying truth.

Definition 2 (Informative Peer). Agent i and agent j are considered each other's informative peers if the determinant of the joint distribution matrix over their private signals $\mathbf{c}_i$ and $\mathbf{c}_j$ is non-zero, i.e., $\det(U_{\mathbf{c}_i, \mathbf{c}_j}) \neq 0$, where $U_{\mathbf{c}_i, \mathbf{c}_j}$ represents the joint prior distribution of the signals $\mathbf{c}_i$ and $\mathbf{c}_j$, and $U_{\mathbf{c}_i, \mathbf{c}_j}$ is expressed in its matrix form.

Definition 2 captures whether two agents are "informative peers" by examining the structure of their shared information. If their private signals are sufficiently correlated (as captured by a non-zero determinant of the joint distribution matrix), they can serve as reliable references for each other in PEG.

Definition 3. Given two random variables X, Y which have the same support $\mathcal{C}$, we define the determinant mutual information (DMI) between X and Y as

$$\text{DMI}(X; Y) = |\det(\mathbf{U}_{X,Y})|.$$

Lemma 2 (Strict Information Monotonicity). For every two random variables X, Y with the same support $\mathcal{C}$, when X' is less informative than X , i.e., X' is independent of Y conditioning on X , it holds that

$$\text{DMI}(X'; Y) \leq \text{DMI}(X; Y).$$

The inequality is strict when $\det(\mathbf{U}_{X,Y}) \neq 0$ and $\mathbf{U}_{X'|X}$ is not a permutation matrix.

Lemma 3. Let $\mathbf{M}_\ell^{ij}$ be the co-occurrence matrix formed by agents i and j over the set of tasks $\mathcal{T}_\ell$, and define the score $\det(\mathbf{M}_\ell^{ij})$. Then the expectation of this score is an unbiased estimator of the determinant mutual information between $\hat{X}_i$ and $\hat{X}_j$:

$$\mathbb{E}_{\hat{X}_i, \hat{X}_j} [\det(\mathbf{M}_\ell^{ij})] = a_\ell \cdot \det(\mathbf{U}_{\hat{X}_i, \hat{X}_j}),$$

where $a_\ell = \binom{|\mathcal{T}_\ell|}{|\mathcal{C}|} \cdot |\mathcal{C}|!$ and $\mathbf{U}_{\hat{X}_i, \hat{X}_j}$ is the co-occurrence matrix of the random variables $\hat{X}_i$ and $\hat{X}_j$.

The proofs of Lemmas 2 and 3 follow the argument of Theorem 5.1 in [33]. Building on these results, we now present the proof of Lemma 1.

Proof. Let agent i truthfully report $\hat{X}_i = X_i$, and let other agents report signals $\hat{X}_j$. Consider the interaction between agent i and any peer agent j . The DMI between X_i and $\hat{X}_j$ satisfies:

$$\text{DMI}(\hat{X}_i; \hat{X}_j) \leq \text{DMI}(X_i; \hat{X}_j),$$

with equality if and only if $\hat{X}_i$ is a permutation of X_i . In particular, if $\hat{X}_j = X_j$ (agent j reports truthfully), then:

$$\text{DMI}(\hat{X}_i; X_j) \leq \text{DMI}(X_i; X_j),$$

and the inequality is strict if $\hat{X}_i$ is not a permutation of X_i , due to the strict information monotonicity of DMI. Therefore, the truthful strategy maximizes DMI against any truthfully reporting peer. From Lemma 3, we know that the utility between agents i and j , based on co-occurrence matrix $\mathbf{M}_\ell^{ij}$, satisfies:

$$\mathbb{E}[\det(\mathbf{M}_\ell^{ij})] = a_\ell \cdot \det(\mathbf{U}_{\hat{X}_i, \hat{X}_j}),$$

where a_ℓ is a constant depending on the number of tasks in round ℓ , and $\mathbf{U}_{\hat{X}_i, \hat{X}_j}$ is the empirical joint distribution matrix. Therefore, the expected squared utility is proportional to:

$$\{\mathbb{E}[\det(\mathbf{M}_\ell^{ij})]\}^2 \propto \text{DMI}^2(\hat{X}_i; \hat{X}_j),$$

and is maximized when $\hat{X}_i = X_i$, i.e., the agent reports truthfully. Since each agent's expected utility (aggregated over peers $j \neq i$) is a weighted sum of $\text{DMI}^2(\hat{X}_i; \hat{X}_j)$, and each term in the sum is maximized by truthful reporting, it follows that truth-telling is a dominant strategy. Furthermore, when all agents report truthfully, the resulting strategy profile forms an equilibrium, and any deviation leads to a strictly lower expected utility unless the deviation

is a permutation of the truth. □

3.2 Regret Analysis

Next, we show that both the generator and the discriminators can progressively improve their behavior such that their cumulative performance asymptotically approaches that of the best truthful policy by performing online policy updates as defined in Eq. (2.3) in Section 2.2. This is formalized via the standard notion of *no-regret learning*, which measures the cumulative difference between the utility obtained by the learned policy over time and the utility that would have been achieved by the best fixed policy in hindsight [6, 18]. Since the generator and discriminators are updated independently, and their optimal strategy is to report truthfully regardless of others, we follow the setting in [27] and define the regret for each discriminator and the generator as:

$$\begin{aligned}\text{Regret}_{D_i}(T) &= \sum_{t'=1}^T u_{D_i}^{(t,t')}(\pi_{D_i}^{(t')}) - \max_{\pi_{D_i}^*} \sum_{t=1}^T u_{D_i}^{(t,t')}(\pi_{D_i}^*), \\ \text{Regret}_G(T) &= \sum_{t=1}^T \sum_{k_t=1}^{K_t} u_G^{(t,t',k_t)}(\pi_G^{(t')}) - \max_{\pi_G^*} \sum_{t'=1}^T \sum_{k_t=1}^{K_t} u_G^{(t,t',k_t)}(\pi_G^*),\end{aligned}$$

where $\pi_{D_i}^{(t')}$ and $\pi_G^{(t')}$ denote the policies of discriminator D_i and the generator at iteration t' , and $\pi_{D_i}^*$, π_G^* are their respective best fixed policies in hindsight. Each $u_G^{(t,t',k_t)}(\cdot)$ denotes the generator's utility on task k_t in round t .

We introduce the following Assumption 4 for theoretical guarantees.

Assumption 4. The utility function $u_{D_i}(\pi_{D_i})$ and $u_G(\pi_G)$ satisfies the following:

(Part 1: Local concavity). There exists a neighborhood $\mathcal{N}$ around a truthful reference policy such that $u_{D_i}(\pi_{D_i})$ and $u_G(\pi_G)$ is concave in π_{D_i} and π_G for all $\pi_{D_i}, \pi_G \in \mathcal{N}$.

(Part 2: Gradient boundedness). There exist constants $M_1, M_2 > 0$ such that for all π_{D_i} and π_G , the gradients are bounded in ℓ_2 -norm: $\|\nabla_{\pi_{D_i}} u_{D_i}(\pi_{D_i})\|_2 \leq M_1$, and $\|\nabla_{\pi_G} u_G(\pi_G)\|_2 \leq M_2$.

Part 1 of Assumption 4 assumes that the utility function is locally concave with respect to the policy within a restricted neighborhood around the truthful policy π^* , as defined in Section 2.2. This is justified because π^* is designed to be utility-maximizing under PEG, and policy updates are constrained to stay close to π^* in practice via deploying a fine-tuned model. Similar locality and curvature assumptions are standard in trust-region and online learning methods [20, 50]. Part 2 of Assumption 4 requires that the gradients of the utility function are bounded. These assumptions are standard in the convex optimization literature such as [6, 16] which also employ gradient-based optimization methods.

We now formalize the no-regret property of PEG. The following theorem shows that both the generator and each discriminator achieve sublinear regret when updated via mirror descent with appropriately chosen learning rates.

Theorem 1. *Under Assumption 4, and set the learning rate in Eq. (2.3) for the generator as $\eta^G := \sqrt{2D_{\text{KL}}(\pi_G^* \|\pi_G^{(1)})/(M_1^2 t_l)}$, and for each discriminator D_i as $\eta^{D_i} := \sqrt{2D_{\text{KL}}(\pi_{D_i}^* \|\pi_{D_i}^{(1)})/(M_2^2 t_l)}$ for $2^{t_l} \leq T < 2^{t_l+1}$. Then, the regrets of the generator and each discriminator are bounded by:*

$$\text{Regret}_G(T) \leq \frac{\sqrt{2}}{\sqrt{2}-1} M_1 \sqrt{2K_T D_{\text{KL}}(\pi_G^* \|\pi_G^{(1)}) \cdot T}, \text{ and } \text{Regret}_{D_i}(T) \leq \frac{\sqrt{2}}{\sqrt{2}-1} M_2 \sqrt{2K_T D_{\text{KL}}(\pi_{D_i}^* \|\pi_{D_i}^{(1)}) \cdot T},$$

respectively, where $D_{\text{KL}}(\cdot \|\cdot)$ denotes the KL divergence between the optimal policy π^ and the initial policy $\pi^{(1)}$, and K_T is the number of tasks at T iteration.*

Theorem 1 guarantees that both the generator and each discriminator achieve a regret bound of $O(\sqrt{T})$, which is consistent with standard results in online convex optimization and learning theory [e.g., 6, 30]. This sublinear regret implies that the policy update defined in Eq. (2.3) is *Hannan consistent*: their average regret vanishes as $T \rightarrow \infty$, meaning that each agent’s average performance converges to that of the best fixed policy [28]. More specifically, the regret bound depends on three key factors: (1) the bound on the gradient norm of the utility functions; (2) the number of tasks K ; and (3) the KL divergence between the initial policy and the optimal one, which measures how far the agent’s starting point is from the target behavior. A smaller KL divergence leads to a tighter regret bound, as the learning trajectory begins closer to the optimal policy. Importantly, regret analysis does not imply

convergence to the optimal policy, but ensures that the average utility gap to the best fixed policy vanishes as $T \rightarrow \infty$. While regret guarantees are well studied in online learning, applying them to a peer evaluation setting with multiple interdependent LLMs is novel. In this setting, agents influence each other only indirectly through their reported outputs, making this a theoretically grounded and promising direction for aligning LLMs.

3.3 Last-Iterate Convergence

Our third main theoretical result concerns the convergence behavior of agents to a Nash equilibrium. Specifically, we establish last-iterate convergence, a stronger guarantee that ensures that the actual sequence of policies converges to a fixed point [6].

In our PEG setup, each discriminator interacts with others repeatedly, aiming to report truthful evaluations based on shared signals. These interactions can be naturally modeled as a continuous multi-agent game $\mathcal{G} = (\mathcal{N}, \Pi, \{u_i\}_{i \in \mathcal{N}})$, where $\mathcal{N} = \{1, \dots, n\}$ denotes the set of agents, $\Pi = \prod_{i=1}^n \Pi_i$ is the joint space of stochastic policies, and $u_i : \Pi \rightarrow \mathbb{R}$ is the utility function for agent i , reflecting agreement with peers. We assume each u_i is continuous, differentiable in its own argument, and has Lipschitz continuous gradients.

A Nash equilibrium in this setting corresponds to a stable configuration of policies where no discriminator has an incentive to unilaterally deviate [23, 34]. Formally, a joint policy $\pi^* = (\pi_1^*, \dots, \pi_n^*) \in \Pi$ is a Nash equilibrium if for every $i \in \mathcal{N}$ and any alternative policy $\hat{\pi}_i \in \Pi_i$,

$$u_i(\pi_i^*, \pi_{-i}^*) \geq u_i(\hat{\pi}_i, \pi_{-i}^*). \quad (3.1)$$

Theorem 2 (Last-iterate Convergence to Nash). *Suppose each discriminator $i \in \mathcal{N}$ updates its policy using Eq. (2.3) with a decaying learning rate of order $O(1/t^p)$, for some $p \in (\frac{1}{2}, 1)$. Then, the sequence of joint policies π_t converges almost surely to the unique Nash equilibrium π^* of PEG.*

The proof of Theorem 2 builds on the Robbins–Siegmund Lemma. It is also inspired by the analysis in [16], though our setting is more specific—focusing on static games with perfect feedback. Theorem 2 establishes last-iterate convergence to the Nash equilibrium: the actual policies used by each discriminator at the end of training converge to the truthful equilibrium. This is in contrast to classical no-regret learning, which only ensures good average performance over time. Last-iterate convergence is particularly valuable in practice, as it guarantees that the learned policies are not just good on average but are inherently stable and reliable [6, 16]. In our case, this means PEG reliably leads to consistent and truthful evaluations, effectively enabling smaller language models to coordinate and reach human-aligned consensus without any supervised fine-tuning or distillation.

Lemma 4. *If the utility function is concave and differentiable, its gradient is monotone.*

Proof. By concavity, the first-order conditions for u at x and y give:

$$\begin{aligned} u(y) &\leq u(x) + \nabla u(x)^\top (y - x), \\ u(x) &\leq u(y) + \nabla u(y)^\top (x - y). \end{aligned}$$

Adding these inequalities:

$$u(y) + u(x) \leq u(x) + u(y) + \nabla u(x)^\top (y - x) + \nabla u(y)^\top (x - y).$$

Simplifying:

$$(\nabla u(x) - \nabla u(y))^\top (x - y) \leq 0.$$

Thus, ∇u is monotone:

$$(\nabla u(x) - \nabla u(y))^\top (x - y) \leq 0 \quad \forall x, y \in \mathcal{X}.$$

□

Lemma 5. Let $\pi_{i,t+1}$ be the update given by mirror descent:

$$\pi_{i,t+1} = \arg \min_{\pi_i \in \Pi_i} \{ \eta_t \langle -\nabla u_i(\pi_t), \pi_i \rangle + D_i(\pi_i, \pi_{i,t}) \},$$

where $D_i(\cdot, \cdot)$ is the Bregman divergence generated by a μ -strongly convex function h_i . Then for any $\pi_i^* \in \Pi_i$, we have

$$D_i(\pi_i^*, \pi_{i,t+1}) \leq D_i(\pi_i^*, \pi_{i,t}) - \eta_t \langle \nabla V_i(\pi_t), \pi_{i,t} - \pi_i^* \rangle + \frac{L_i \eta_t^2}{2},$$

where $L_i > 0$ is a constant that bounds $\|\nabla V_i(\pi_t)\|^2$ due to compactness of Π_i .

Proof. Let $x = \pi_{i,t}$, $x^+ = \pi_{i,t+1}$, $z = \pi_i^*$, and $g_i = \nabla u_i(\pi_t)$. The mirror descent update is:

$$x^+ = \arg \min_{\pi_i \in \Pi_i} \{ -\eta_t \langle g_i, \pi_i \rangle + D_i(\pi_i, x) \}.$$

By the first-order optimality condition for convex minimization over Π_i , we have:

$$\langle -\eta_t g_i + \nabla h_i(x^+) - \nabla h_i(x), z - x^+ \rangle \geq 0,$$

which gives:

$$\eta_t \langle g_i, z - x^+ \rangle \leq \langle \nabla h_i(x^+) - \nabla h_i(x), z - x^+ \rangle.$$

Using the three-point identity for Bregman divergence:

$$\langle \nabla h_i(x^+) - \nabla h_i(x), z - x^+ \rangle = D_i(z, x) - D_i(z, x^+) - D_i(x^+, x),$$

we substitute and rearrange:

$$D_i(z, x^+) \leq D_i(z, x) - \eta_t \langle g_i, z - x \rangle + \eta_t \langle g_i, x^+ - x \rangle - D_i(x^+, x).$$

Now apply Young's inequality:

$$\eta_t \langle g_i, x^+ - x \rangle \leq \frac{1}{2\mu} \|x^+ - x\|^2 + \frac{\mu\eta_t^2}{2} \|g_i\|^2.$$

By strong convexity, $D_i(x^+, x) \geq \frac{\mu}{2} \|x^+ - x\|^2$, so:

$$\eta_t \langle g_i, x^+ - x \rangle \leq D_i(x^+, x) + \frac{\eta_t^2 \|g_i\|^2}{2\mu}.$$

Substituting back, the $D_i(x^+, x)$ term cancels, and we obtain:

$$D_i(z, x^+) \leq D_i(z, x) - \eta_t \langle g_i, z - x \rangle + \frac{\eta_t^2 \|g_i\|^2}{2\mu}.$$

Let $L_i := \frac{1}{\mu} \sup_{\pi \in \mathcal{K}} \|\nabla u_i(\pi)\|^2$, which is finite due to compactness of Π_i . Therefore:

$$D_i(\pi_i^*, \pi_{i,t+1}) \leq D_i(\pi_i^*, \pi_{i,t}) - \eta_t \langle \nabla u_i(\pi_t), \pi_i^* - \pi_{i,t} \rangle + \frac{L_i \eta_t^2}{2}.$$

□

Next we can move to the proof of Theorem 2.

Proof. We define an energy function that tracks the distance to the Nash equilibrium using Bregman divergences. For each player $i \in \mathcal{N}$, let D_i denote the associated Bregman divergence:

$$D_i(\pi_i^*, \pi_{i,t}) := h_i(\pi_i^*) - h_i(\pi_{i,t}) - \langle \nabla h_i(\pi_{i,t}), \pi_i^* - \pi_{i,t} \rangle.$$

We define the total energy function as

$$E_t := \sum_{i \in \mathcal{N}} D_i(\pi_i^*, \pi_{i,t}).$$

Each player updates their strategy via mirror descent:

$$\pi_{i,t+1} = \arg \min_{\pi_i \in \Pi_i} \{ \eta_t \langle -\nabla u_i(\pi_t), \pi_i \rangle + D_i(\pi_i, \pi_{i,t}) \}.$$

With Lemma 5, we have:

$$D_i(\pi_i^*, \pi_{i,t+1}) \leq D_i(\pi_i^*, \pi_{i,t}) - \eta_t \langle \nabla u_i(\pi_t), \pi_{i,t} - \pi_i^* \rangle + \frac{L_i \eta_t^2}{2}.$$

Summing over all players, we obtain:

$$E_{t+1} \leq E_t - \eta_t \langle \nabla u(\pi_t), \pi_t - \pi^* \rangle + C \eta_t^2,$$

for some constant $C > 0$, where $\nabla V(\pi_t) = (\nabla u_1(\pi_t), \dots, \nabla u_N(\pi_t))$.

Because the utility map ∇u is strictly monotone, there exists $\mu > 0$ such that

$$\langle \nabla u(\pi_t), \pi_t - \pi^* \rangle \geq \mu \|\pi_t - \pi^*\|^2.$$

Substituting into the previous inequality yields:

$$E_{t+1} \leq E_t - \mu \eta_t \|\pi_t - \pi^*\|^2 + C \eta_t^2.$$

This recursion is in the form required by the Robbins–Siegmund Lemma:

$$E_{t+1} \leq E_t - a_t + b_t,$$

where $a_t = \mu \eta_t \|\pi_t - \pi^*\|^2 \geq 0$ and $b_t = C \eta_t^2$, with $\sum_{t=1}^{\infty} b_t < \infty$ since $\eta_t = O(1/t^p)$ with $2p > 1$. Therefore, the Robbins–Siegmund Lemma implies that E_t converges almost surely to a finite random variable E_{∞} , and $\sum_t \eta_t \|\pi_t - \pi^*\|^2 < \infty$. Since $\sum_t \eta_t^2 < \infty$ and $\sum_t \eta_t = \infty$, $\|\pi_t - \pi^*\| \rightarrow 0$ almost surely. $\square$

Chapter 4

Experiments

We evaluate PEG on question-answering (QA) tasks. In this setup, a generator LLM agent produces answers, while a group of discriminator LLM agents independently evaluate each response, serving as peers to promote both accuracy and consistency.

4.1 Experiment Setup

Models. The main experiment employs the following models: DeepSeek-R1-Distill-Qwen-7B (deepseek-Qwen-7b), deepseek-ai/deepseek-llm-7b-chat (deepseek-Llama-7b), and Qwen/Qwen2.5-7B-Instruct (OQwen-7b) as discriminators, with Qwen/Qwen2.5-7B-Instruct also serving as the generator. We consider Gemma-7B [53], Mistral-7B [29], Ai-Yi-9B [66] and OpenChat-7B [57] as candidate discriminators. Unless otherwise specified, we set the learning rate $\eta = 0.1$ for all experiments. The PEG mechanism between discriminators is run for 10 iterations for 8 tasks.

Prompts. We evaluate our models with zero-shotting following the format described in [21]. By default, conditioning the P_{LM} on $(x, \text{correct})$ corresponds to the standard zero-shot prompt. Conditioning on $(x, \text{incorrect})$ uses the same structure, except that the phrase “Answer:” is replaced with “Incorrect Answer:” in the prompt B.0.1.

Baselines. We incorporate several baseline methods, each representing a different

approach to generating and selecting responses. To evaluate against our method, we apply each baseline to obtain responses from the discriminators, and then compare their outputs using majority voting [61].

- **Generative Ranking (G):** A standard baseline that ranks candidate answers by the probability $P_{LM}(y \mid x, \text{correct})$ and selects the top candidate [56].
- **Discriminative Ranking (D):** This method employs a discriminator π_D to estimate $P(\text{correct} \mid x, y)$ and ranks responses accordingly [27].
- **Mutual Information Ranking (MI):** This method reweights each candidate by the product of the forward and reverse likelihoods, $P_{LM}(y \mid x, \text{correct}) \cdot P_{LM}(\text{correct} \mid x, y)$ [35].
- **Equilibrium Ranking Discriminator (ER-D):** Based on the Consensus Game framework of [27], this method formulates the interaction between a generator and a discriminator as a signaling game. The discriminator iteratively updates its estimates to maximize agreement with the generator. Each query–candidate pair (x, y) is reweighted by $\pi_D^*(\text{correct} \mid x, y)$, encouraging consistency in the final ranking.

Datasets. We conduct our evaluations using four diverse datasets: ARC-Easy, ARC-Challenge, Massive Multitask Language Understanding (MMLU) and Graduate-Level Google-Proof Q&A Benchmark (GPQA) [14, 21, 48]. Details of the datasets are in Appendix B.0.2. Each dataset presents unique challenges, allowing us to test PEG under various knowledge domains.

4.2 Experiment Results

PEG Improves Accuracy. Results in Table 4.1 show that PEG consistently outperforms all baselines across the evaluated datasets. Notably, it achieves more than a 10% improvement in accuracy on the most challenging benchmarks, ARC-Challenge and MMLU, compared

Table 4.1: Accuracy (%) of majority vote answers across benchmark datasets for each method. **Bold** indicates the best performance in each row.

Dataset	G	D	MI	ER-D	PEG
ARC-Easy	88.19	84.18	88.61	88.57	91.78
ARC-Challenge	77.01	70.68	78.03	77.52	87.01
MMLU	59.98	50.75	59.85	59.66	70.73
GPQA-Main	18.08	9.15	16.29	16.52	22.54

to the strongest baseline. This performance gain is due to two key factors. First, PEG leverages the complementary strengths and diverse reasoning capabilities of multiple LLM discriminators. Second, unlike agreement-based methods such as MI and ER-D, PEG promotes truthful reporting through mutual information signals, which more effectively elicit the latent capabilities of the models. Table 4.2 further supports this, showing that each individual discriminator consistently improves in accuracy following policy updates under PEG.

PEG Enables Coordination Among Heterogeneous Agents. Despite substantial differences in architecture and baseline accuracy, Table 4.2 shows that all discriminators benefit from PEG. Remarkably, even the weakest model, OQwen-7B, which initially has the lowest accuracy among all agents, becomes the most accurate after participating in PEG. This suggests that PEG encourages cross-agent learning, where each model learns to align with truthful, high-confidence signals provided by its peers. Figure 4.1 reports the accuracy of each individual discriminator, as well as the majority vote, after applying PEG. The majority vote achieves a notable accuracy gain over any single model, validating the benefits of collaborative learning among heterogeneous agents.

PEG with Varying Number of Discriminators. We further evaluate PEG on the ARC-Challenge dataset with extended discriminator settings: 5-discriminators (adding Gemma-7B [53] and Mistral-7B [29]) and 7-discriminators (further adding Ai-Yi-9B [66] and OpenChat-7B [57]). The results consistently highlight the strong impact of PEG on both individual and collective performance. Notably, the weakest models with the lowest

Table 4.2: Accuracy (%) of each model before (D) and after applying the PEG mechanism (PEG) across four benchmark datasets. **Bold** highlights the best PEG result per dataset.

Model	ARC-Easy		ARC-Challenge		MMLU		GPQA	
	D	PEG	D	PEG	D	PEG	D	PEG
OQwen-7B	90.43	90.89	82.91	85.04	62.72	68.19	16.07	22.77
deepseek-Qwen-7B	70.48	91.73	57.18	86.15	42.38	69.04	15.18	20.09
deepseek-LLaMA-7B	76.30	91.44	59.83	85.98	46.23	68.71	17.41	23.66




Figure 4.1: Accuracy comparison between original model outputs (D) and PEG majority vote answers.

Table 4.3: Accuracy (%) of individual discriminators before and after applying PEG in the 5-discriminator setting.

5 Discriminators	OQwen-7B	deepseek-Qwen-7B	deepseek-Llama2-7B	Gemma-7B	Mistral-7B
Original	82.91	57.18	59.83	69.32	70.26
After PEG	86.84	84.27	83.76	84.27	82.74

initial accuracy benefit the most from PEG. For example, in Tables 4.3 and 4.4, Qwen-7B and LLaMA2-7B gain over 20% improvement after applying PEG, whereas initially stronger models such as OQwen-7B and Ai-Yi-9B exhibit only marginal gains. Nevertheless, all models, regardless of their initial performance, converge toward coordinated outcomes after PEG. These results demonstrate that agents indeed learn from one another and achieve coordination through PEG. Moreover, when varying the number of discriminators, PEG

consistently outperforms both initial discriminator majority vote D and ER-D, especially with 3 and 5 discriminators, where it achieves over 10% improvement, as shown in Table 4.5.

Table 4.4: Accuracy (%) of individual discriminators before and after applying PEG in the 7-discriminator setting.

7 Discriminators	OQwen-7B	Qwen-7B	Llama2-7B	Gemma-7B	Mistral-7B	Ai-Yi-9B	OpenChat-7B
Original	82.91	57.18	59.83	69.32	70.26	81.28	79.06
After PEG	83.85	75.30	73.08	77.61	79.06	81.62	82.99

Table 4.5: Overall accuracy (%) comparison of initial discriminator majority vote D, ER-D and PEG under 3-, 5-, and 7-discriminator settings. **Bold** indicates the best performance in each row.

Setting	D	ER-D	PEG
3 Discriminators	70.68	77.52	87.01
5 Discriminators	71.71	76.32	86.75
7 Discriminators	76.50	81.54	81.97

Finally, we observe a slight decrease in accuracy when expanding to seven discriminators. This can be attributed to a mild violation of the conditional independence assumption (Assumption 3), which ensures that the mutual information-based utility remains informative and non-degenerate. Intuitively, when discriminators are highly similar, their outputs become redundant, i.e., observing one provides little new information beyond the others. As a result, the mutual-information-based utility degenerates to zero, so truthfulness no longer maximizes the utility. Consequently, the learning dynamics may converge to a suboptimal equilibrium where discriminators are not incentive-compatible in (2.1), resulting in a reduction in accuracy. Overall, our results indicate that PEG remains robust with up to five discriminators; see, Table 4.5. As a practical guideline, we recommend using three to five heterogeneous LLMs, which provide a strong balance between performance, stability, and computational efficiency. Expanding beyond this should be done cautiously and only when sufficient diversity and independence among LLMs can be ensured.

Chapter 5

Conclusions

In this paper, we propose a training-free, game-theoretic framework for aligning LLMs through a multi-agent peer elicitation game. Through mutual evaluations among agents, our peer elicitation game facilitates interactions between a generator and multiple discriminators in a way that provably incentivizes truthful behavior. We theoretically show that truthful reporting is a dominant strategy for each discriminator. Furthermore, using online mirror descent, each agent achieves sublinear regret, ensuring that its average performance approaches that of the best fixed truthful strategy. The agents’ strategies also converge in the last iterate to a truthful Nash equilibrium. Empirically, our framework significantly improves factual accuracy across a range of benchmarks and performs competitively with much larger models, highlighting a practical direction for deploying lightweight models in resource-constrained environments.

There are several promising directions for future research. One is to extend PEG to high-stakes settings such as medical decision support, scientific fact verification, and policy-relevant summarization, where truthful and consistent outputs are essential for safety and reliability. Another is to incorporate concepts from game theory and economics, such as reputation systems, repeated interactions, and budget-aware mechanisms, to further enhance alignment and robustness among LLM agents, particularly in open-ended or adversarial environments.

Appendices

Appendix A

Proofs of Theorem 1

Definition 4 (Bregman divergence). Let $\varphi : \Omega \rightarrow \mathbb{R}$ be a differentiable and μ -strongly convex ($\mu > 0$) function with respect to a norm $\|\cdot\|$, that is, satisfy

$$\varphi(x) \geq \varphi(y) + \langle \nabla \varphi(y), x - y \rangle + \frac{\mu}{2} \|x - y\|^2, \quad \forall x, y \in \Omega.$$

The Bregman divergence centered in $y \in \Omega$ is the function $D_\varphi(x \parallel y)$ defined as

$$D_\varphi(x \parallel y) := \varphi(x) - \varphi(y) - \langle \nabla \varphi(y), x - y \rangle.$$

Definition 5. Let $f : \Omega \rightarrow \mathbb{R}$ be a convex function, and let D_φ be a Bregman divergence. The *proximal mapping* (or *proximal step*) for a point x_t with step size $\eta > 0$ is defined as:

$$\text{Prox}_\varphi(\eta \nabla f(x_t), x_t) = \arg \min_{x \in \Omega} \{ \eta \langle \nabla f(x_t), x \rangle + D_\varphi(x \parallel x_t) \}.$$

The *mirror descent algorithm* is then given by the iterative update:

$$x_{t+1} := \text{Prox}_\varphi(\eta \nabla f(x_t), x_t).$$

Lemma 6. Let $\pi_\theta(y|x)$ be a policy and $R(y)$ be a utility function independent of θ . Define

the objective function:

$$J(\theta) = \mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [R(y)].$$

Then, the gradient of $J(\theta)$ with respect to θ is:

$$\nabla_\theta J(\theta) = \mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [R(y) \nabla_\theta \log \pi_\theta(y|x)].$$

Proof.

$$\begin{aligned} \nabla_\theta J(\theta) &= \nabla_\theta \int \pi_\theta(y|x) R(y) dy = \int \nabla_\theta \pi_\theta(y|x) R(y) dy \\ &= \int \pi_\theta(y|x) \nabla_\theta \log \pi_\theta(y|x) \cdot R(y) dy \\ &= \mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [R(y) \nabla_\theta \log \pi_\theta(y|x)]. \end{aligned}$$

□

Corollary 1 (Gradient of utility Functions). *Using Lemma 6, the gradients of the utility functions with respect to the discriminator and generator parameters are given by: For each discriminator D_i :*

$$\nabla_{\theta_{D_i}} u_{D_i} = \mathbb{E}_{r_{ik} \sim \pi_{D_i}(\cdot|x_k, y_k)} \left[\left(\sum_{j \neq i} \det(M_{ij}^1) \cdot \det(M_{ij}^2) \right) \nabla_{\theta_{D_i}} \log \pi_{D_i}(r_{ik}|x_k, y_k) \right].$$

For the generator G :

$$\nabla_{\theta_G} u_G = \mathbb{E}_{y_k \sim \pi_G(\cdot|x_k, v_k)} [\mathbf{1}(v_k = \hat{v}_k) \nabla_{\theta_G} \log \pi_G(y_k|x_k, v_k)].$$

Lemma 7. *When $\Omega = \Delta^n$ is the set of full-support distributions over n objects and the φ is set to the negative entropy function, which is 1-strongly convex with respect to the ℓ_1 norm $\|\cdot\|_1$, the corresponding Bregman divergence is the Kullback-Leibler (KL) divergence.*

Lemma 8. *Let $x' = \text{Prox}_\varphi(g, x)$ be the proximal update. Then, for all $y \in \Omega$, we have:*

$$\langle -g, y - x' \rangle \leq -D_\varphi(y \| x') + D_\varphi(y \| x) - D_\varphi(x' \| x).$$

By setting $y = x$, this three-point inequality simplifies to

$$\langle -g, x - x' \rangle \leq -D_\varphi(x \| x') - D_\varphi(x' \| x).$$

Proof. The objective function of the proximal step problem is given by

$$h(z) := \langle g, z \rangle + D_\varphi(z \| x), \quad z \in \Omega.$$

The first-order optimality conditions applied to the solution $z = x'$ are therefore

$$\begin{aligned} -\nabla h(x') \in \mathcal{N}_\Omega(x') &\iff -g - \nabla \varphi(x') + \nabla \varphi(x) \in \mathcal{N}_\Omega(x') \\ &\iff \langle -g - \nabla \varphi(x') + \nabla \varphi(x), y - x' \rangle \leq 0 \quad \forall y \in \Omega \\ &\iff \langle -g, y - x' \rangle \leq \langle \nabla \varphi(x') - \nabla \varphi(x), y - x' \rangle \quad \forall y \in \Omega. \end{aligned}$$

The statement now follows from using the identity

$$\langle \nabla \varphi(x') - \nabla \varphi(x), y - x' \rangle = -D_\varphi(y \| x') + D_\varphi(y \| x) - D_\varphi(x' \| x),$$

which can be checked directly from the definition of Bregman divergence. □

Lemma 9. *Let $f : \Omega \rightarrow \mathbb{R}$ be convex. Each step of the mirror descent algorithm satisfies*

$$f(x_t) \leq f(y) + \langle \nabla f(x_t), x_t - x_{t+1} \rangle - \frac{1}{\eta_t} D_\varphi(y \| x_{t+1}) + \frac{1}{\eta_t} D_\varphi(y \| x_t) - \frac{1}{\eta_t} D_\varphi(x_{t+1} \| x_t).$$

Proof. Using the linear lower bound property of convex functions, we can write

$$f(x_t) \leq f(y) - \langle \nabla f(x_t), y - x_t \rangle = f(y) + \langle \nabla f(x_t), x_t - x_{t+1} \rangle - \langle \nabla f(x_t), y - x_{t+1} \rangle.$$

On the other hand, from Lemma 8 applied to the mirror descent step (that is, for the choices $g = \eta_t \nabla f(x_t)$, $x' = x_{t+1}$, $x = x_t$), we have

$$-\eta_t \langle \nabla f(x_t), y - x_{t+1} \rangle \leq -D_\varphi(y \| x_{t+1}) + D_\varphi(y \| x_t) - D_\varphi(x_{t+1} \| x_t).$$

Hence, dividing by η_t and plugging into the previous inequality, we obtain the statement. $\square$

Lemma 10. *Let $\|\cdot\|$ be the norm with respect to which the DGF φ is 1-strongly convex, and $\|\cdot\|_*$ be the dual norm. If all functions $f_t : \Omega \rightarrow \mathbb{R}$ are convex, the regret of online mirror descent is bounded by*

$$R_T := \sum_{t=1}^T f_t(x_t) - \min_{x \in \Omega} \sum_{t=1}^T f_t(x) \leq \frac{1}{\eta} D_\varphi(x \| x_1) + \frac{\eta}{2} \sum_{t=1}^T \|\nabla f_t(x_t)\|_*^2.$$

In particular, assuming that all dual gradient norms are upper bounded by G , and setting

$$\eta := \sqrt{\frac{2D_\varphi(x \| x_1)}{M^2 \cdot 2^i}} \quad \text{for } 2^i \leq T < 2^{i+1}.$$

we find

$$R_T \leq M \sqrt{2D_\varphi(x \| x_1) \cdot T} \frac{\sqrt{2}}{\sqrt{2} - 1}.$$

Proof. Since functions f_t are convex, we can use Lemma 9:

$$f_t(x_t) \leq f_t(x) + \langle \nabla f_t(x_t), x_t - x_{t+1} \rangle - \frac{1}{\eta_t} D_\varphi(x \| x_{t+1}) + \frac{1}{\eta_t} D_\varphi(x \| x_t) - \frac{1}{\eta_t} D_\varphi(x_{t+1} \| x_t) \quad \forall x \in \Omega.$$

Using the Cauchy-Schwarz inequality, we can bound the right-hand side by

$$f_t(x_t) \leq f_t(x) + \|\nabla f_t(x_t)\|_* \cdot \|x_t - x_{t+1}\| - \frac{1}{\eta_t} D_\varphi(x\|x_{t+1}) + \frac{1}{\eta_t} D_\varphi(x\|x_t) - \frac{1}{\eta_t} D_\varphi(x_{t+1}\|x_t).$$

Using Young's inequality, as well as the 1-strong convexity of the KL divergence, which implies

$$\begin{aligned} \frac{1}{\eta_t} D_\varphi(x_{t+1}\|x_t) &\geq \frac{1}{2\eta_t} \|x_{t+1} - x_t\|^2, \\ f_t(x_t) &\leq f_t(x) + \frac{\eta_t}{2} \|\nabla f_t(x_t)\|^2 + \frac{1}{2\eta_t} \|x_t - x_{t+1}\|^2 - \frac{1}{\eta_t} D_\varphi(x\|x_{t+1}) + \frac{1}{\eta_t} D_\varphi(x\|x_t) - \frac{1}{2\eta_t} \|x_{t+1} - x_t\|^2. \end{aligned}$$

Rearranging terms, we obtain:

$$f_t(x_t) \leq f_t(x) + \frac{\eta_t}{2} \|\nabla f_t(x_t)\|^2 - \frac{1}{\eta_t} D_\varphi(x\|x_{t+1}) + \frac{1}{\eta_t} D_\varphi(x\|x_t).$$

Summing over $t = 1, \dots, T$:

$$\sum_{t=1}^T (f_t(x_t) - f_t(x)) \leq \frac{1}{2} \sum_{t=1}^T \eta_t \|\nabla f_t(x_t)\|_2^2 + \sum_{t=1}^T \frac{1}{\eta_t} (D_\varphi(x\|x_t) - D_\varphi(x\|x_{t+1})).$$

The regret incurred by the algorithm is upper bound by the regret incurred in each of the intervals $2^i \leq T < 2^{i+1}$. Suppose the algorithm has been run until time $2^i \leq T < 2^{i+1}$. Hence, the regret is upper bounded by

$$\begin{aligned} \text{Reg}_T &\leq \left(M \sqrt{\frac{D_\varphi(x\|x_1)}{2}} \sum_{i=0}^I (\sqrt{2})^i \right) + \left(\sum_{i=0}^{I-1} \frac{\sqrt{D_\varphi(x\|x_1)}}{2M^2 2^i} M^2 2^i \right) + \left(\sqrt{\frac{2D_\varphi(x\|x_1)}{2M^2 2^I}} \right) M^2 (T - 2^I) \\ &= M \sqrt{\frac{D_\varphi(x\|x_1)}{2}} \left(\sum_{i=0}^I (\sqrt{2})^i + \sum_{i=0}^{I-1} (\sqrt{2})^i \right) + \sqrt{\frac{D_\varphi(x\|x_1)}{2M^2 2^I}} M^2 (T - 2^I) \end{aligned}$$

In particular, since $T < 2^{I+1}$,

$$\begin{aligned}
\text{Reg}_T &\leq M \sqrt{\frac{D_\varphi(x\|x_1)}{2}} \left(\sum_{i=0}^I (\sqrt{2})^i + \sum_{i=0}^{I-1} (\sqrt{2})^i \right) + \sqrt{\frac{D_\varphi(x\|x_1)}{2M^2 2^I}} M^2 2^I \\
&= M \sqrt{\frac{D_\varphi(x\|x_1)}{2}} \left(\sum_{i=0}^I (\sqrt{2})^i + \sum_{i=0}^I (\sqrt{2})^i \right) \\
&= M \sqrt{2D_\varphi(x\|x_1)} \sum_{i=0}^I (\sqrt{2})^i \\
&= M \sqrt{2D_\varphi(x\|x_1)} \cdot \frac{(\sqrt{2})^{I+1} - 1}{\sqrt{2} - 1} \\
&\leq M \sqrt{2D_\varphi(x\|x_1)T} \cdot \frac{\sqrt{2}}{\sqrt{2} - 1}.
\end{aligned}$$

□

Lemma 11. *Suppose the utility is defined over a set of K policies, and each policy is updated independently via a no-regret algorithm with individual regret bounded by $O(\sqrt{T})$. Then, the total regret of the system with respect to the best fixed policy in hindsight is bounded by $O(\sqrt{KT})$. Therefore, when both the generator and each discriminator maintain and update K independent policies, their regret bounds will incur an additional $\sqrt{K}$ factor, yielding an overall regret of $O(\sqrt{KT})$ per agent.*

This result that the regret grows as $O(\sqrt{KT})$ when updating and aggregating over K policies is a classical result in online learning and multi-armed bandits, as established in [5, 51]. Therefore, by combining Lemmas 10 and 11, we complete the proof of Theorem 1

Appendix B

Experiments Details

B.0.1 Prompt

Generator Prompt

The following are multiple choice questions (with answers) about Geology. Question:

Which property of a mineral can be determined just by looking at it?

- A) luster
- B) mass
- C) weight
- D) hardness

Answer:

Discriminator Prompt

You are an expert evaluator of questions about Geology. Determine if the proposed answer is correct. Output ONLY 'A' or 'B'.

Question: Luster is the property of a mineral can be determined just by looking at it.

Proposed Answer: luster

Is this answer correct? Respond ONLY with: A. Correct B. Incorrect

Answer:

B.0.2 Datasets

Datasets from different domains are used in the experiments. Details are summarized below.

- **ARC** dataset [14] evaluates scientific reasoning, includes both the 'easy' and 'challenge' sets. It consists of 7787 science questions, all nondiagrams, multiple choice (typically 4-way multiple choice). The experiment applies a zero-shot setting for ARC, assessing how well the model navigates science-related questions requiring logical reasoning rather than memorized knowledge.
- **MMLU** [21] is a benchmark with questions across humanities, STEM, and social sciences. It requires the model to demonstrate broad general knowledge. For MMLU, we evaluate our models following the format described in setting [21].
- **GPQA** [48] is a challenging dataset designed to evaluate LLM capabilities and scalable oversight mechanisms. It consists of 448 multiple choice questions that cover biology, physics, and chemistry. These questions are intentionally designed to be high-quality and extremely difficult. The experiment will apply the zero-shot setting for GPQA.

Bibliography

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Floren-
cia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat,
et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] Arpit Agarwal, Debmalya Mandal, David C Parkes, and Nisarg Shah. Peer prediction
with heterogeneous users. *ACM Transactions on Economics and Computation*, 8(1):1–34,
2020.
- [3] Amos Azaria and Tom Mitchell. The internal state of an LLM knows when it’s lying.
arXiv preprint arXiv:2304.13734, 2023.
- [4] Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and
Mike Zheng Shou. Hallucination of multimodal large language models: A survey. *arXiv
preprint arXiv:2404.18930*, 2024.
- [5] Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and non-
stochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*,
5(1):1–122, 2012.
- [6] Yang Cai and Weiqiang Zheng. Doubly optimal no-regret learning in monotone games.
In *International Conference on Machine Learning*, pages 3507–3524. PMLR, 2023.
- [7] José A Carrillo, Shi Jin, Lei Li, and Yuhua Zhu. A consensus-based global optimization

- method for high dimensional machine learning problems. *ESAIM: Control, Optimisation and Calculus of Variations*, 27:S5, 2021.
- [8] Philippe Casgrain, Brian Ning, and Sebastian Jaimungal. Deep Q-learning for Nash equilibria: Nash-DQN. *Applied Mathematical Finance*, 29(1):62–78, 2022.
- [9] Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomek Korbak, David Lindner, Pedro Freire, et al. Open problems and fundamental limitations of reinforcement learning from human feedback. *Transactions on Machine Learning Research*, 2023.
- [10] Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. Chateval: Towards better llm-based evaluators through multi-agent debate. In *The Twelfth International Conference on Learning Representations*, 2024.
- [11] Sijin Chen, Omar Hagrass, and Jason M Klusowski. Decoding game: On minimax optimality of heuristic text generation strategies. *arXiv preprint arXiv:2410.03968*, 2024.
- [12] Yangyi Chen, Lifan Yuan, Ganqu Cui, Zhiyuan Liu, and Heng Ji. A close look into the calibration of pre-trained language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1343–1367, 2023.
- [13] Pengyu Cheng, Tianhao Hu, Han Xu, Zhisong Zhang, Yong Dai, Lei Han, Xiaolong Li, et al. Self-playing adversarial language game enhances llm reasoning. *Advances in Neural Information Processing Systems*, 37:126515–126543, 2024.
- [14] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

- [15] Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*, 2023.
- [16] Benoit Duvocelle, Panayotis Mertikopoulos, Mathias Staudigl, and Dries Vermeulen. Multiagent online learning in time-varying games. *Mathematics of Operations Research*, 48(2):914–941, 2023.
- [17] Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. Measuring and improving consistency in pretrained language models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031, 2021.
- [18] Xiand Gao, Xiaobo Li, and Shuzhong Zhang. Online learning with non-convex losses and non-stationary regret. In *International Conference on Artificial Intelligence and Statistics*, pages 235–243. PMLR, 2018.
- [19] Zhixiong Han, Yaru Hao, Li Dong, Yutao Sun, and Furu Wei. Prototypical calibration for few-shot learning of language models. In *The Eleventh International Conference on Learning Representations*, 2023.
- [20] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [21] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2020.
- [22] Steven CH Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *Neurocomputing*, 459:249–289, 2021.

- [23] Charles A Holt and Alvin E Roth. The nash equilibrium: A perspective. *Proceedings of the National Academy of Sciences*, 101(12):3999–4002, 2004.
- [24] Junling Hu and Michael P Wellman. Nash q-learning for general-sum stochastic games. *Journal of Machine Learning Research*, 4(Nov):1039–1069, 2003.
- [25] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025.
- [26] Xiaowei Huang, Wenjie Ruan, Wei Huang, Gaojie Jin, Yi Dong, Changshun Wu, Saddek Bensalem, Ronghui Mu, Yi Qi, Xingyu Zhao, et al. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *Artificial Intelligence Review*, 57(7):175, 2024.
- [27] Athul Paul Jacob, Yikang Shen, Gabriele Farina, and Jacob Andreas. The consensus game: Language model generation via equilibrium search. In *The Twelfth International Conference on Learning Representations*, 2023.
- [28] Amir Jafari, Amy Greenwald, David Gondek, and Gunes Ercal. On no-regret learning, fictitious play, and nash equilibrium. In *ICML*, volume 1, pages 226–233, 2001.
- [29] Albert Qiaochu Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mistral 7b. *ArXiv*, abs/2310.06825, 2023.
- [30] Matthew Jones, Huy Nguyen, and Thy Nguyen. An efficient algorithm for fair multi-agent multi-armed bandit with low regret. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.

- [31] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [32] Ryuto Koike, Masahiro Kaneko, and Naoaki Okazaki. Outfox: Llm-generated essay detection through in-context learning with adversarially generated examples. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.
- [33] Yuqing Kong. Dominantly truthful multi-task peer prediction with a constant number of tasks. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2398–2411. SIAM, 2020.
- [34] David M Kreps. Nash equilibrium. In *Game Theory*, pages 167–177. Springer, 1989.
- [35] Jiwei Li and Dan Jurafsky. Mutual information and diverse decoding improve neural machine translation. *arXiv preprint arXiv:1601.00372*, 2016.
- [36] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. Encouraging divergent thinking in large language models through multi-agent debate. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904. Association for Computational Linguistics, November 2024.
- [37] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2024.
- [38] Andrea I Luppi, Pedro AM Mediano, Fernando E Rosas, Negin Holland, Tim D Fryer, John T O’Brien, James B Rowe, David K Menon, Daniel Bor, and Emmanuel A Stamatakis. A synergistic core for human brain evolution and cognition. *Nature Neuroscience*, 25(6):771–782, 2022.

- [39] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.
- [40] Eric Mitchell, Joseph Noh, Siyan Li, Will Armstrong, Ananth Agarwal, Patrick Liu, Chelsea Finn, and Christopher D Manning. Enhancing self-consistency and performance of pre-trained language models through natural language inference. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1754–1768, 2022.
- [41] Remi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiegl, et al. Nash learning from human feedback. In *International Conference on Machine Learning*, pages 36743–36768. PMLR, 2024.
- [42] Roger B Myerson. Incentive compatibility and the bargaining problem. *Econometrica: journal of the Econometric Society*, pages 61–73, 1979.
- [43] Drazen Prelec. A Bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.
- [44] Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. ChatDev: Communicative agents for software development. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186, Bangkok, Thailand, August 2024. Association for Computational Linguistics.
- [45] Guannan Qu, Adam Wierman, and Na Li. Scalable reinforcement learning of localized policies for multi-agent networked systems. In *Learning for Dynamics and Control*, pages 256–266. PMLR, 2020.

- [46] Goran Radanovic and Boi Faltings. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2014.
- [47] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36:53728–53741, 2023.
- [48] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.
- [49] Grant Schoenebeck and Fang-Yi Yu. Two strongly truthful mechanisms for three heterogeneous agents answering one question. *ACM Transactions on Economics and Computation*, 10(4):1–26, 2023.
- [50] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897. PMLR, 2015.
- [51] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- [52] Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liangyan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 13088–13110, 2024.
- [53] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love,

- et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- [54] SM Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313*, 6, 2024.
- [55] Giulio Tononi, Olaf Sporns, and Gerald M Edelman. A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences*, 91(11):5033–5037, 1994.
- [56] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [57] Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. Openchat: Advancing open-source language models with mixed-quality data. *arXiv preprint arXiv:2309.11235*, 2023.
- [58] Mingzhi Wang, Chengdong Ma, Qizhi Chen, Linjian Meng, Yang Han, Jiancong Xiao, Zhaowei Zhang, Jing Huo, Weijie J Su, and Yaodong Yang. Magnetic preference optimization: Achieving last-iterate convergence for language model alignment. *arXiv preprint arXiv:2410.16714*, 2024.
- [59] Tianlong Wang, Xianfeng Jiao, Yinghao Zhu, Zhongzhi Chen, Yifan He, Xu Chu, Junyi Gao, Yasha Wang, and Liantao Ma. Adaptive activation steering: A tuning-free llm truthfulness improvement method for diverse hallucinations categories. In *Proceedings of the ACM on Web Conference 2025*, pages 2562–2578, 2025.
- [60] Xiaofeng Wang and Tuomas Sandholm. Reinforcement learning to play an optimal nash

- equilibrium in team markov games. *Advances in Neural Information Processing Systems*, 15, 2002.
- [61] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2022.
- [62] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- [63] Jens Witkowski and David Parkes. A robust bayesian truth serum for small populations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2012.
- [64] Anita Williams Woolley, Christopher F Chabris, Alex Pentland, Nada Hashmi, and Thomas W Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688, 2010.
- [65] TaeHo Yoon, Sayantan Choudhury, and Nicolas Loizou. Multiplayer federated learning: Reaching equilibrium with less communication. *arXiv preprint arXiv:2501.08263*, 2025.
- [66] Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin Wang, Heng Li, Jiangcheng Zhu, Jianqun Chen, et al. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*, 2024.
- [67] Bangguo Yu, Yuzhen Liu, Lei Han, Hamidreza Kasaei, Tingguang Li, and Ming Cao. Vln-game: Vision-language equilibrium search for zero-shot semantic navigation. *arXiv preprint arXiv:2411.11609*, 2024.
- [68] Peter Zhang and Yiling Chen. Elicitability and knowledge-free elicitation with peer

prediction. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, pages 245–252, 2014.

- [69] Yizhou Zhang, Guannan Qu, Pan Xu, Yiheng Lin, Zaiwei Chen, and Adam Wierman. Global convergence of localized policy iteration in networked multi-agent reinforcement learning. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(1):1–51, 2023.