Novelty Assessment Report

Paper: Learning a Game by Paying the Agents
PDF URL: https://openreview.net/pdf?id=8yRtP2n8OK
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-08

Abstract

We study the problem of learning the utility functions of no-regret learning agents in a repeated normal-form game. Differing from most prior literature, we introduce a principal with the power to observe the agents playing the game, send agents signals, and give agents payments as a function of their actions. We show that the principal can, using a number of rounds polynomial in the size of the game, learn the utility functions of all agents to any desired precision $\varepsilon > 0$, for any no-regret learning algorithms of the agents. Our main technique is to formulate a zero-sum game between the principal and the agents, where the principal's strategy space is the set of all payment functions. Finally, we discuss implications for the problem of steering agents to a desired equilibrium: in particular, we introduce, using our utility-learning algorithm as a subroutine, the first algorithm for steering arbitrary no-regret learning agents without prior knowledge of their utilities.

Core Task Landscape

This paper addresses: **Learning Utility Functions of No-Regret Agents through Payments and Signals**

A total of **44 papers** were analyzed and organized into a taxonomy with **18 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Principal-Agent Incentive Design and Learning**
- **Agent Behavior Inference and Reward Learning**
- **Multi-Agent Coordination and Reward Design**
- **Applied Incentive Systems and Behavioral Studies**
- **Rational Choice Theory and Behavioral Foundations**

Complete Taxonomy Tree

- Learning Utility Functions of No-Regret Agents through Payments and Signals Survey Taxonomy
- Principal-Agent Incentive Design and Learning
  - Utility and Type Inference through Strategic Interaction ★ (3 papers)
  - [0] Learning a Game by Paying the Agents (Anon et al., 2026) View paper
  - [14] Estimating and incentivizing imperfect-knowledge agents with hidden rewards (Dogan, 2023) View paper
  - [34] Active Inference through Incentive Design in Markov Decision Processes (Wei Xin-yi, 2025) View paper
  - Incentive Contract Design under Information Asymmetry (5 papers)
  - [13] Optimal incentive contracts with imperfect information (M. Harris, 1979) View paper
  - [23] An equilibrium model of incentive contracts in the presence of information manipulation (Eitan Goldman, 2006) View paper
  - [29] The theory of incentives: the principal-agent model (Martimort, 2002) View paper
  - [31] Optimal incentive contracts with multiple agents (Joel Demski, 1984) View paper
  - [32] Dynamic Reward Design (Liu, 2025) View paper
  - Dynamic and Adaptive Incentive Mechanisms (3 papers)
  - [6] Adaptive Incentive Design for Markov Decision Processes with Unknown Rewards (H Ma, 2025) View paper
  - [24] Solving dynamic principal-agent problems with a rationally inattentive principal (Mu Tong, 2022) View paper
  - [38] An experimental study of a dynamic principalâ□agent relationship (Werner GÃ¼th, 1998) View paper
  - Information Signaling and Revelation Mechanisms (3 papers)
  - [3] Strategic incentives and the optimal sale of information (Rosina Rodriguez Olivera, 2024) View paper
  - [4] Information Signaling With Concurrent Monetary Incentives in Bayesian Congestion Games (Bryce L. Ferguson, 2024) View paper
  - [26] A Marketplace for Data: An Algorithmic Solution (Anish Agarwal, 2022) View paper
- Agent Behavior Inference and Reward Learning
  - Reward Function Learning from Observations (2 papers)
  - [1] CLIP-Motion: Learning Reward Functions for Robotic Actions Using Consecutive Observations (Xuzhe Dang, 2023) View paper
  - [15] On the challenges and practices of reinforcement learning from real human feedback (Ball Sarah, 2023) View paper
  - Causal Analysis of Agent Incentives (1 papers)
  - [18] Agent Incentives: A Causal Perspective (Carey Ryan, 2021) View paper
- Multi-Agent Coordination and Reward Design
  - Cooperative Multi-Agent Reward Shaping (3 papers)
  - [5] Lero: Llm-driven evolutionary framework with hybrid rewards and enhanced observation for multi-agent reinforcement learning (Wei Yuan, 2025) View paper

- [11] Multi-agent cooperation policy gradient method based on enhanced exploration for cooperative tasks (Li-yang Zhao, 2024) View paper
- [17] Improving Multi-Agent Reinforcement Learning for Beer Game by Reward Design Based on Payment Mechanism (Masaaki HORI, 2023) View paper
- Market-Based Multi-Agent Coordination (1 papers)
- [9] Transactive local energy markets enable community-level resource coordination using individual rewards (Daniel May, 2024) View paper
- Crowdsourcing and Spatial Task Allocation Incentives (1 papers)
- [19] Incentive Mechanism for Spatial Crowdsourcing With Unknown Social-Aware Workers: A Three-Stage Stackelberg Game Approach (Yin Xu, 2022) View paper

- Applied Incentive Systems and Behavioral Studies
  - Incentives in Digital Platforms and Advertising (1 papers)
  - [35] IC Mechanisms for Risk-Averse Advertisers in the Online Advertising System (Wang, 2024) View paper
  - Fraud Detection and Payment Systems (2 papers)
  - [21] Application of deep reinforcement learning to payment fraud (Siddharth Vimal, 2021) View paper
  - [39] Minimum payments that reward honest reputation feedback (Radu Jurca, 2006) View paper
  - Organizational and Managerial Incentive Applications (2 papers)
  - [22] Interests, information, and incentives in higher education: Principal-agent theory and its potential applications to the study of higher education governance (Jason E. Lane, 2008) View paper
  - [28] Managerial influences on intraorganizational information technology use: A Principalâ□agent model (Anol Bhattacherjee, 1998) View paper
  - Supply Chain and Channel Coordination Incentives (1 papers)
  - [16] Balancing Supplier Channels: an Incentive Model for Online and Offline Sales Channels (N. Zivlak, 2023) View paper
  - Behavioral Experiments on Incentives and Cooperation (5 papers)
  - [8] Emergent cooperative decision-making in triadic Prisoner's Dilemmas: Effects of incentives and information. (Yinuo Du, 2025) View paper
  - [20] No-vaxxers are different in public good games (Marcello Basili, 2022) View paper
  - [33] Financial and symbolic incentives promote 'green' charging choices (Kacperski, 2025) View paper
  - [43] Combining payment for crop damages and reward for productivity to address wildlife conflict. (Poorva Joshi, 2022) View paper
  - [44] An informational versus monetary incentive in increasing physicians' response rates. (Alyssa N. Easton, 1998) View paper
- Rational Choice Theory and Behavioral Foundations
  - Foundations of Rational Choice and Incentive Theory (4 papers)
  - [2] Incentives: Motivation and the economics of information (Donald E. Campbell, 2018) View paper
  - [27] Rational herding in financial economics (Welch, 1996) View paper
  - [36] Rational choice (Hartmut Esser, 2010) View paper
  - [41] The rational choice theory of institutions: Implications for design (Randall L. Calvert, 1995) View paper
  - Behavioral Deviations from Rational Choice (3 papers)
  - [10] Rethinking rational choice (Frank, 2019) View paper
  - [30] A reâ□examination of probability matching and rational choice (Jd, 2002) View paper
  - [42] Economics of Mind, Body, Spirit â□ Experimental Games Reveal New Human Choices (Syed Shurid Khan, 2023) View paper
  - Rational Choice in Political and Social Contexts (4 papers)
  - [7] Rational choice and political imaging in mining areas: a case study of legislative elections in East Luwu (Rahmat Muhammad, 2025) View paper
  - [12] The Impact of Political Materialism on Voter Apathy in Nigeria: A Mixed-Methods Study (Chilaka F.C., 2024) View paper
  - [37] PERCEIVED FAIRNESS OF TAX INCENTIVES AND ITS EFFECT ON TAX ATTITUDES AMONG GIG WORKERS IN SOUTH-WEST NIGERIA (Abdulrasheed Taiwo Abdullahi, 2025) View paper
  - [40] A Case Study of Factors Influencing Default of Provincial Publicly-Funded Normal Studentsâ□â□Taking the Case of Default of the First Publicly-Funded Normal Students in Shandong Province as an Example (Xie Yuena, 2022) View paper
  - Institutional and Policy Design under Rational Agents (1 papers)
  - [25] An Institutionalist View on Experimentalist Governance: Local-level obstacles to policy-learning in European Union Cohesion Policy (Telle, 2021) View paper

Narrative

Core task: learning utility functions of no-regret agents through payments and signals. The field structure reflects a broad landscape spanning theoretical foundations and practical applications of incentive design. Principal-Agent Incentive Design and Learning encompasses classical contract theory and modern adaptive mechanisms that infer agent preferences through strategic interaction, often drawing on foundational work such as Optimal Incentive Contracts[13] and Theory of Incentives[29]. Agent Behavior Inference and Reward Learning focuses on extracting hidden preferences from observed actions, addressing challenges like imperfect knowledge and human feedback noise, as seen in Imperfect Knowledge Hidden Rewards[14] and Human Feedback Challenges[15]. Multi-Agent Coordination and Reward Design examines how incentives shape collective behavior in settings ranging from congestion games to cooperative dilemmas, while Applied Incentive Systems and Behavioral Studies explore real-world domains such as transactive energy markets, spatial crowdsourcing, and gig economy taxation. Rational Choice Theory and Behavioral Foundations provides the underlying decision-theoretic models, including critiques and extensions of classical rationality assumptions.

A particularly active line of work investigates how principals can learn agent types or utilities by strategically offering payments or information signals, balancing exploration of unknown preferences with exploitation of learned models. Learning Game Paying Agents[0] sits squarely within this vein, focusing on utility inference for no-regret learners through carefully designed payment schemes. It shares thematic ground with Active Inference Incentive[34], which also considers how agents update beliefs and respond to incentives, and with Strategic Incentives Information Sale[3], which examines information provision as a lever for influencing strategic behavior. Compared to Imperfect Knowledge Hidden Rewards[14], which addresses hidden reward structures in single-agent settings, Learning Game Paying Agents[0] emphasizes the interactive, game-theoretic dimension where the principal must adapt to agent learning dynamics. This positioning highlights ongoing questions about how to efficiently elicit preferences when agents themselves are adapting, and how payment mechanisms can serve dual roles as both incentives and informative signals.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. Estimating and incentivizing imperfect-knowledge agents with hidden rewards

Authors: Dogan, Ilgin, IlgÄ±n DoÄan, Shen, Zuo-Jun Max, et al. (11 authors total) | **Year/Venue**: 2023 | **URL**: View paper

Abstract

In practice, incentive providers (i.e., principals) often cannot observe the reward realizations of incentivized agents, which is in contrast to many principal-agent models that have been previously studied. This information asymmetry challenges the principal to consistently estimate the agent's unknown rewards by solely watching the agent's decisions, which becomes even more challenging when the agent has to learn its own rewards. This complex setting is observed in various real-life scenarios ...

Relationship Analysis

Both papers belong to the same taxonomy category focusing on utility and type inference through strategic interaction, where principals learn about agents through payments and signals. They overlap in studying how principals can learn agent utility functions in repeated games with strategic agents, using payment mechanisms to influence behavior and extract information. However, the original paper focuses on learning utility functions of no-regret learning agents in normal-form games through a zero-sum game formulation with polynomial-time guarantees, while the candidate paper addresses learning from imperfect-knowledge agents solving multi-armed bandit problems with hidden reward realizations, using a slack-minimization estimator and regret analysis that accounts for the agent's own learning process.

---

2. Active Inference through Incentive Design in Markov Decision Processes

Authors: Wei Xin-yi, Xinyi Wei, Shi, Chongyang, Chongyang Shi, et al. (16 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

We present a method for active inference with partial observations in stochastic systems through incentive design, also known as the leader-follower game. Consider a leader agent who aims to infer a follower agent's type given a finite set of possible types. Different types of followers differ in either the dynamical model, the reward function, or both. We assume the leader can partially observe a follower's behavior in the stochastic system modeled as a Markov decision process, in which the fol...

Relationship Analysis

Both papers belong to the category of utility and type inference through strategic interaction, where a principal learns about agents through strategic mechanisms. They overlap in addressing how a principal can infer agent characteristics (utility functions vs. types) through strategic payments in settings with learning agents. The key difference is that the original paper focuses on learning complete utility functions of no-regret agents in repeated normal-form games using payments and signals, while the candidate paper addresses active inference of agent types in Markov decision processes through incentive design, using information-theoretic objectives (conditional entropy) and partial observations in a leader-follower framework.

Contributions Analysis

Overall novelty summary. The paper contributes a polynomial-time algorithm for learning utility functions of no-regret agents through strategic payments and signals, alongside a zero-sum game formulation for the principal-agent interaction. It resides in the 'Utility and Type Inference through Strategic Interaction' leaf, which contains only three papers total, indicating a relatively sparse research direction within the broader taxonomy. This leaf sits under 'Principal-Agent Incentive Design and Learning', distinguishing work where principals actively design incentives to infer agent characteristics from purely observational approaches.

The taxonomy reveals neighboring research directions that provide context. 'Incentive Contract Design under Information Asymmetry' (five papers) focuses on optimal contracts without learning objectives, while 'Dynamic and Adaptive Incentive Mechanisms' (three papers) examines time-varying incentive systems. The 'Agent Behavior Inference and Reward Learning' branch addresses preference learning without principal intervention, including reward function learning from observations. The paper's approach bridges these areas by combining strategic payment design with utility inference, positioning it at the intersection of mechanism design and learning theory in repeated games.

Among 26 candidates examined across three contributions, the analysis reveals varied novelty profiles. The polynomial-time utility learning algorithm (6 candidates examined, 0 refutable) and zero-sum game formulation (10 candidates examined, 0 refutable) show no clear prior work overlap within the limited search scope. The steering algorithm contribution (10 candidates examined, 1 refutable) appears to have more substantial prior work, with at least one candidate providing overlapping methods. This suggests the core utility learning mechanism may represent the more distinctive technical contribution, though the limited search scale means potentially relevant work outside the top-26 semantic matches remains unexamined.

Based on the top-26 semantic matches and taxonomy structure, the work appears to occupy a relatively underexplored niche combining no-regret learning dynamics with principal-designed payment mechanisms. The sparse population of its taxonomy leaf and limited refutable prior work suggest novelty, though the analysis cannot rule out relevant contributions beyond the examined candidate set. The steering algorithm's partial overlap with prior work indicates this application may be more incremental than the core learning framework.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Polynomial-time algorithm for learning utility functions of no-regret agents via payments

Description: The authors introduce an algorithm that enables a principal to learn the utility functions of agents playing a repeated normal-form game by providing payments and signals. The algorithm works for arbitrary no-regret learning agents and achieves learning in polynomially many rounds with respect to game size.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Game Manipulators--the Strategic Implications of Binding Contracts

URL: View paper

Brief Assessment

Game Manipulators Binding Contracts[56] focuses on commitment devices and binding contracts in game theory, analyzing how external manipulators can extract fees by distorting payoff matrices. It does not address the problem of learning utility functions of no-regret agents through payments and signals in repeated normal-form games.

---

2. Economics and the Theory of Games

URL: View paper

Brief Assessment

Economics Theory Games[55] appears to be a foundational economics text on game theory. The provided context contains only fragmentary text about rational players and preferences, with no discussion of learning algorithms, no-regret learning, or computational complexity for utility function learning.

---

3. Algorithmic Monetary Policies for Blockchain Participation Games

URL: View paper

Brief Assessment

Blockchain Participation Policies[57] addresses monetary policy design for blockchain participation games with type and stake parameters, not utility function learning of no-regret agents in repeated normal-form games via payments and signals.

---

4. Compensatory transfers in two-player decision problems

URL: View paper

Brief Assessment

Compensatory Transfers Decision[58] focuses on two-player decision problems with monetary transfers in a static setting, not on learning utility functions of no-regret agents in repeated games through strategic payments and signals.

---

5. No-regret Learning and a Mechanism for Distributed Convex Optimisation and Coordination

URL: View paper

Brief Assessment

Distributed Convex Optimization[60] focuses on distributed convex optimization through no-regret learning in resource allocation problems, not on learning utility functions of agents in repeated normal-form games via payments and signals.

---

6. Unidirectional substitutes and complements

URL: View paper

Brief Assessment

Unidirectional Substitutes Complements[59] addresses matching markets with complementarities between worker groups, not learning utility functions of agents in repeated games through payments and signals.

Contribution 2: Zero-sum game formulation between principal and agents for utility learning

Description: The core technical contribution is a novel formulation where the principal and agents play a zero-sum game. The principal chooses payment functions while agents choose actions to maximize rewards. This formulation enables the principal to learn utility functions through convergence to equilibrium.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Robust mechanisms: the curvature case

URL: View paper

Brief Assessment

Robust Mechanisms Curvature[67] focuses on mechanism design with curvature constraints in robust settings, not on utility learning through zero-sum game formulation with payments and no-regret learning agents.

---

2. Fairness and incentives in a multiâtask principalâagent model

URL: View paper

Brief Assessment

Fairness Multi Task Agent[62] focuses on fairness and incentives in multi-task principal-agent models with explicit monetary transfers, not on zero-sum game formulations for utility learning through payment functions and equilibrium convergence.

---

3. Review of Incentive Mechanisms of Differential Privacy Based Federated Learning Protocols: From the Economics and Game Theoretical Perspectives

URL: View paper

Brief Assessment

Differential Privacy Federated Learning[68] focuses on incentive mechanisms for federated learning protocols with differential privacy, examining reward and compensation mechanisms. It does not address zero-sum game formulations for principal-agent utility learning through payment functions.

---

4. More than privacy: Adopting differential privacy in game-theoretic mechanism design

URL: View paper

Brief Assessment

Differential Privacy Mechanism Design[63] focuses on privacy-preserving mechanism design in game-theoretic settings, not on utility learning through zero-sum game formulation between principal and agents with payment functions.

---

5. Paying to do better: Games with payments between learning agents

URL: View paper

Brief Assessment

Paying Better Games[61] focuses on payments between learning agents in repeated games to influence equilibrium outcomes, not on a principal learning agent utility functions through zero-sum game formulation with payment functions.

---

6. Game theory and business applications

URL: View paper

Brief Assessment

Game Theory Business Applications[65] discusses principal-agent models and incentive compensation in general business contexts, not a zero-sum game formulation for utility learning through payment functions and convergence to equilibrium.

---

7. The Huntâvitell general theoryof marketing ethics: can it enhance our understanding of principal-agent relationships in channels of distribution?

URL: View paper

Brief Assessment

Hunt Vitell Marketing Ethics[69] focuses on marketing ethics theory and relational exchanges in distribution channels, not on game-theoretic utility learning algorithms or principal-agent payment mechanisms for learning utility functions.

---

8. Principal-Agent Reward Shaping in MDPs

URL: View paper

Brief Assessment

Principal Agent Reward Shaping[70] addresses reward shaping in MDPs where the principal modifies agent rewards through budget-constrained payments, not utility learning through zero-sum games. The candidate focuses on Stackelberg games for incentive alignment, not on learning unknown utility functions through adversarial equilibrium convergence.

---

9. Optimal profit-loss sharing contracts with symmetric and asymmetric information (principal-agent model approach)

URL: View paper

Brief Assessment

Profit Loss Sharing Contracts[64] focuses on optimal profit-loss sharing contracts in Islamic banking using principal-agent models, not on utility learning through zero-sum game formulation. The candidate discusses cooperative game structures for profit-sharing rather than adversarial zero-sum games for learning agent utilities.

---

10. Supermodularity and Monotonicity in Economics

URL: View paper

Brief Assessment

Supermodularity Monotonicity Economics[66] is an editorial overview of papers using supermodular game theory and monotone methods. It does not address principal-agent utility learning through zero-sum game formulations with payment functions.

Contribution 3: First steering algorithm for no-regret agents without prior utility knowledge

Description: The authors present the first algorithm that can steer no-regret learning agents toward desired equilibria without requiring prior knowledge of the agents' utility functions. This is achieved by combining their utility-learning algorithm with a steering procedure, and they characterize the optimal achievable value through correlated equilibrium with payments.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Steering no-regret learners to optimal equilibria

URL: View paper

Prior Art Analysis

Steering No Regret Learners[45] demonstrates that steering no-regret learning agents to optimal equilibria is possible without prior knowledge of utility functions. The candidate paper presents algorithms for steering agents to desired equilibria using payments and signals, establishing both theoretical results and experimental validation. While the original paper claims to be 'the first algorithm for steering arbitrary no-regret learning agents without prior knowledge of their utilities,' the candidate paper (published earlier, as indicated by submission to ICLR 2024 vs. ICLR 2026) already addresses the steering problem for no-regret learners, including scenarios where equilibria must be computed online without full prior knowledge of the game structure.

Evidence

Evidence 1 - **Rationale**: The candidate paper establishes fundamental results on steering no-regret learners to equilibria, demonstrating that this problem can be solved with vanishing payments, which directly addresses steering without requiring complete prior utility knowledge. - **Original**: we introduce, using our utility-learning algorithm as a subroutine, the first algorithm for steering arbitrary no-regret learning agents without prior knowledge of their utilities. - **Candidate**: we consider the problem of steering no-regret-learning agents to play desirable equilibria via nonnegative payments. we first show that steering is impossible if the total budget (across all iterations) is finite, both in normaland extensive-form games. however, we establish thatvanishing average paym…

Evidence 2 - **Rationale**: The candidate's online steering algorithm learns and steers simultaneously without requiring precomputed equilibria or full utility knowledge, directly challenging the novelty claim of being 'first' to steer without prior utility knowledge. - **Original**: our result 2 is the first in steering any no-regret agents without prior knowledge of their utilities, and the first exact characterization of the optimal value achievable by the principal in the steering problem. - **Candidate**: algorithm 6.4(online steer). the mediator runs a regret minimization algorithmr0 over its own strategy spacex0, which we assume has regret at mostr0(t) after t rounds. on each round, the mediator does the following: • get a strategyμ(t) from r0. playμ(t), and setp(t) i as defined in (2) in μ(t) . • ...

---

2. Is rlhf more difficult than standard rl? a theoretical perspective

URL: View paper

Brief Assessment

RLHF Difficulty Theoretical[50] focuses on converting preference feedback to reward signals for standard RL algorithms, not on steering no-regret learning agents toward equilibria without utility knowledge. The paper addresses a different problem: whether preference-based RL is more difficult than reward-based RL through reduction approaches.

---

3. No-Regret Learning and Equilibrium Computation in Quantum Games

URL: View paper

Brief Assessment

Quantum Games No Regret[46] focuses on quantum game equilibria and no-regret learning in quantum systems, not on steering classical no-regret agents or learning utility functions in normal-form games without prior knowledge.

---

4. No-Regret Learning for Stackelberg Equilibrium Computation in Newsvendor Pricing Games

URL: View paper

Brief Assessment

Newsvendor Pricing Stackelberg[51] focuses on a specific two-player supply chain game with stochastic demand learning, not on steering arbitrary no-regret agents in general normal-form games without utility knowledge. The candidate addresses Stackelberg equilibrium computation in a newsvendor pricing context, which is a specialized economic application rather than a general framework for steering no-regret learners across arbitrary games.

---

5. No-regret sample-efficient Bayesian optimization for finding Nash equilibria with unknown utilities

URL: View paper

Brief Assessment

Bayesian Nash Optimization[53] focuses on finding Nash equilibria through Bayesian optimization with Gaussian processes in games with unknown utility functions, but does not address the steering problem of guiding no-regret learning agents toward desired equilibria through payments and signals as described in the original paper.

---

6. Near-Optimal No-Regret Learning in General Games

URL: View paper

Brief Assessment

Near Optimal General Games[54] focuses on achieving polylogarithmic regret bounds for optimistic hedge in multi-player games and convergence to coarse correlated equilibria, not on steering agents toward desired equilibria without prior utility knowledge through payments and signals.

---

7. Regret minimization and convergence to equilibria in general-sum markov games

URL: View paper

Brief Assessment

Regret Minimization Markov Games[49] focuses on decentralized regret minimization in general-sum Markov games where all agents adopt the same learning algorithm, achieving convergence to correlated equilibria. The original paper addresses a different problem: steering no-regret agents toward desired equilibria without prior utility knowledge through payments and signals, which is not the focus of the candidate paper.

---

8. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games

URL: View paper

Brief Assessment

Correlated Equilibria No Regret[48] focuses on convergence rates to correlated equilibria through no-regret learning dynamics (OMWU, internal/swap regret algorithms), not on a principal actively steering agents or learning utility functions without prior knowledge.

---

9. No-regret learning and mixed Nash equilibria: They do not mix

URL: View paper

Brief Assessment

No Regret Mixed Nash[52] focuses on the stability and convergence properties of no-regret learning dynamics to Nash equilibria, not on steering agents or learning utility functions without prior knowledge.

---

10. No-Regret Learning in Stackelberg Games with an Application to Electric Ride-Hailing

URL: View paper

Brief Assessment

Electric Ride Hailing Stackelberg[47] focuses on single-leader multi-follower Stackelberg games with learning at the lower level, not general repeated normal-form games with payments and signals. The candidate addresses a different problem structure (hierarchical Stackelberg vs. peer game with principal intervention).

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Learning a Game by Paying the Agents View paper
- [1] CLIP-Motion: Learning Reward Functions for Robotic Actions Using Consecutive Observations View paper
- [2] Incentives: Motivation and the economics of information View paper
- [3] Strategic incentives and the optimal sale of information View paper
- [4] Information Signaling With Concurrent Monetary Incentives in Bayesian Congestion Games View paper
- [5] Lero: Llm-driven evolutionary framework with hybrid rewards and enhanced observation for multi-agent reinforcement learning View paper
- [6] Adaptive Incentive Design for Markov Decision Processes with Unknown Rewards View paper
- [7] Rational choice and political imaging in mining areas: a case study of legislative elections in East Luwu View paper
- [8] Emergent cooperative decision-making in triadic Prisoner's Dilemmas: Effects of incentives and information. View paper
- [9] Transactive local energy markets enable community-level resource coordination using individual rewards View paper
- [10] Rethinking rational choice View paper
- [11] Multi-agent cooperation policy gradient method based on enhanced exploration for cooperative tasks View paper
- [12] The Impact of Political Materialism on Voter Apathy in Nigeria: A Mixed-Methods Study View paper
- [13] Optimal incentive contracts with imperfect information View paper
- [14] Estimating and incentivizing imperfect-knowledge agents with hidden rewards View paper
- [15] On the challenges and practices of reinforcement learning from real human feedback View paper
- [16] Balancing Supplier Channels: an Incentive Model for Online and Offline Sales Channels View paper
- [17] Improving Multi-Agent Reinforcement Learning for Beer Game by Reward Design Based on Payment Mechanism View paper
- [18] Agent Incentives: A Causal Perspective View paper
- [19] Incentive Mechanism for Spatial Crowdsourcing With Unknown Social-Aware Workers: A Three-Stage Stackelberg Game Approach View paper
- [20] No-vaxxers are different in public good games View paper
- [21] Application of deep reinforcement learning to payment fraud View paper
- [22] Interests, information, and incentives in higher education: Principal-agent theory and its potential applications to the study of higher education governance View paper
- [23] An equilibrium model of incentive contracts in the presence of information manipulation View paper
- [24] Solving dynamic principal-agent problems with a rationally inattentive principal View paper
- [25] An Institutionalist View on Experimentalist Governance: Local-level obstacles to policy-learning in European Union Cohesion Policy View paper
- [26] A Marketplace for Data: An Algorithmic Solution View paper
- [27] Rational herding in financial economics View paper
- [28] Managerial influences on intraorganizational information technology use: A Principalâ□agent model View paper

- [29] The theory of incentives: the principal-agent model View paper
- [30] A reâ□□examination of probability matching and rational choice View paper
- [31] Optimal incentive contracts with multiple agents View paper
- [32] Dynamic Reward Design View paper
- [33] Financial and symbolic incentives promote 'green' charging choices View paper
- [34] Active Inference through Incentive Design in Markov Decision Processes View paper
- [35] IC Mechanisms for Risk-Averse Advertisers in the Online Advertising System View paper
- [36] Rational choice View paper
- [37] PERCEIVED FAIRNESS OF TAX INCENTIVES AND ITS EFFECT ON TAX ATTITUDES AMONG GIG WORKERS IN SOUTH-WEST NIGERIA View paper
- [38] An experimental study of a dynamic principalâ□□agent relationship View paper
- [39] Minimum payments that reward honest reputation feedback View paper
- [40] A Case Study of Factors Influencing Default of Provincial Publicly-Funded Normal Studentsâ□□â□□Taking the Case of Default of the First Publicly-Funded Normal Students in Shandong Province as an Example View paper
- [41] The rational choice theory of institutions: Implications for design View paper
- [42] Economics of Mind, Body, Spirit â□□ Experimental Games Reveal New Human Choices View paper
- [43] Combining payment for crop damages and reward for productivity to address wildlife conflict. View paper
- [44] An informational versus monetary incentive in increasing physicians' response rates. View paper
- [45] Steering no-regret learners to optimal equilibria View paper
- [46] No-Regret Learning and Equilibrium Computation in Quantum Games View paper
- [47] No-Regret Learning in Stackelberg Games with an Application to Electric Ride-Hailing View paper
- [48] Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games View paper
- [49] Regret minimization and convergence to equilibria in general-sum markov games View paper
- [50] Is rlhf more difficult than standard rl? a theoretical perspective View paper
- [51] No-Regret Learning for Stackelberg Equilibrium Computation in Newsvendor Pricing Games View paper
- [52] No-regret learning and mixed Nash equilibria: They do not mix View paper
- [53] No-regret sample-efficient Bayesian optimization for finding Nash equilibria with unknown utilities View paper
- [54] Near-Optimal No-Regret Learning in General Games View paper
- [55] Economics and the Theory of Games View paper
- [56] Game Manipulators--the Strategic Implications of Binding Contracts View paper
- [57] Algorithmic Monetary Policies for Blockchain Participation Games View paper
- [58] Compensatory transfers in two-player decision problems View paper
- [59] Unidirectional substitutes and complements View paper
- [60] No-regret Learning and a Mechanism for Distributed Convex Optimisation and Coordination View paper
- [61] Paying to do better: Games with payments between learning agents View paper
- [62] Fairness and incentives in a multiâ□□task principalâ□□agent model View paper
- [63] More than privacy: Adopting differential privacy in game-theoretic mechanism design View paper
- [64] Optimal profit-loss sharing contracts with symmetric and asymmetric information (principal-agent model approach) View paper
- [65] Game theory and business applications View paper
- [66] Supermodularity and Monotonicity in Economics View paper
- [67] Robust mechanisms: the curvature case View paper
- [68] Review of Incentive Mechanisms of Differential Privacy Based Federated Learning Protocols: From the Economics and Game Theoretical Perspectives View paper
- [69] The Huntâ□□vitell general theoryof marketing ethics: can it enhance our understanding of principal-agent relationships in channels of distribution? View paper
- [70] Principal-Agent Reward Shaping in MDPs View paper