Strategic Generative AI for Machine Learning in Economic Environments

OMER MADMON, Technion - Israel Institute of Technology, Israel

MOSHE TENNENHOLTZ, Technion - Israel Institute of Technology, Israel

Generative AI (GenAI) is transforming the machine learning (ML) landscape by enabling the creation of high-quality synthetic data for training and evaluation. Yet when these synthetic datasets are used in *economic* or *multi-agent* environments—where learning systems interact with other decision-makers—the assumption that data can be generated independently of incentives often breaks down. In such settings, each agent's behavior depends on beliefs, payoffs, and strategic anticipation of others, making incentive structures an integral part of the data-generating process. This perspective advocates for an *incentive-aware* approach to GenAI, emphasizing the importance of embedding economic and strategic considerations into synthetic data generation. We review recent case studies in which incorporating incentive consistency has led to better-performing ML models in persuasion and competitive search environments. Finally, we outline a research agenda for developing *strategically aligned* generative models that integrate economic reasoning, mechanism design, and ML to ensure robustness and reliability in complex decision-making ecosystems.

CCS Concepts: • **General and reference → Surveys and overviews**; *Evaluation*; **Empirical studies**; • **Theory of computation →** Algorithmic game theory; • **Applied computing → Economics**.

Additional Key Words and Phrases: Generative AI, Machine Learning, Synthetic Data, Multi-Agent Systems, Game Theory, Economic Behavior, Decision-Making

1 Introduction

Generative AI (GenAI) is revolutionizing the way modern machine learning (ML) systems are designed, trained, and deployed. By leveraging large-scale foundation models capable of producing rich, diverse, and contextually grounded data, researchers and practitioners can now address challenges that were previously limited by data scarcity, bias, and privacy concerns. Among the many emerging applications, *synthetic data generation* stands out as one of the most transformative uses of GenAI—supporting the development of increasingly complex ML models and pipelines across domains such as computer vision, natural language processing, and multi-agent decision-making.

A growing body of work explores the use of generative models, particularly large language models (LLMs), to produce data that augment or replace real-world datasets for downstream learning [19, 20, 28, 29, 37]. This paradigm has proven especially valuable in domains where data collection is costly, sensitive, or limited by strategic or behavioral factors. However, when synthetic data are generated for *economic* or *strategic* environments—where ML models interact with other decision makers, including humans and AI agents—standard data generation methods can be misleading or even counterproductive. In these settings, each participant's incentives, goals, and beliefs jointly shape the outcomes of the system. A data generation process that ignores such incentive structures risks producing models that perform well in isolation but fail to generalize or behave appropriately in dynamic and strategic environments.

This perspective paper argues that *incentive-awareness* must be a first-class principle in the use of GenAI for modern ML, particularly when synthetic data are used to train agents that operate in multi-agent or socio-economic contexts [12, 16, 21, 22, 24, 27]. Generative models should not only simulate realistic data but also account for the strategic nature of the environments they aim to represent. This can be achieved through mechanisms informed by economic and game-theoretic reasoning [18, 23], or through adaptive procedures learned from experience [1, 38].

Authors' Contact Information: Omer Madmon, omermadmon@campus.technion.ac.il, Technion - Israel Institute of Technology, Haifa, Israel; Moshe Tennenholtz, moshet@technion.ac.il, Technion - Israel Institute of Technology, Haifa, Israel.

Although interaction-driven generation resembles ideas from self-play in reinforcement learning (RL) [9, 27, 45, 46], our perspective differs fundamentally in purpose and setting. Self-play is primarily a *learning mechanism* for improving policies within a fixed, fully specified environment, where the trajectories generated during training serve mainly as internal samples used to update the agent. In contrast, our perspective aims to *generate synthetic data* that captures the incentives, interactions, and feedback loops of economic environments, and that is intended to be used directly for training or evaluating other ML or LLM-based systems. Moreover, the environments we consider are typically *dynamic, partially observable, and not fully specified as games*: agents act through free-text communication or content generation, making the strategy space open-ended and richer than those typically studied in classical self-play [32, 44]. In this sense, our focus is on constructing strategically coherent *data distributions* themselves, rather than on improving a policy for a given environment.

We highlight the importance of *strategic and incentive-aware data generation* when ML models trained on synthetic data operate in economic environments with other decision-making entities. Two recent use cases illustrate this point. The first uses LLMs to generate behavioral data for predicting human choices in *persuasion games* [42], where incentive-consistent generation improved predictive validity. The second concerns *competitive search environments* [33], where LLM-guided fine-tuning of content-producing agents led to superior performance under realistic strategic dynamics. We conclude by outlining both challenges and future directions for *strategic GenAI*, where economic reasoning and data generation co-evolve toward more robust and strategically incentive-aligned ML systems.

2 Predicting Human Choice in Persuasion Games

Background. Persuasion games are a cornerstone of economic theory, capturing how information asymmetry shapes strategic interaction between an informed *sender* and an uninformed *receiver* [14, 17, 25]. The sender observes a private state of the world and strategically decides how much information to *reveal* to influence the receiver's beliefs and actions. This framework highlights the power of information control in shaping decisions and outcomes, with a wide range of applications, including reputation systems [4, 10], market segmentation [3, 8, 15], AI alignment [5], and beyond. While theoretical models typically assume fully rational behavior, real-world persuasion often involves *human decision makers* whose behavior exhibits a spectrum of patterns. The empirical study of human decision making in persuasion games was initiated by Apel et al. [2], who demonstrated the effectiveness of ML models in predicting choices within a repeated, language-based persuasion game, where a hotel expert bot interacted with human participants by revealing selected reviews to promote hotel adoption. Subsequent studies extended this line of work by optimizing the sender's strategy against human behavior [39] and by developing rule-based simulation frameworks to enhance the effectiveness of learning models [41]. A recent study by Shapira et al. [42] explored the potential of LLMs to generate synthetic training data for human choice prediction, showing that LLM-generated data can successfully replace or augment human data for training predictive models in repeated, language-based persuasion games.

Game setup and learning task. In the experimental setting of Shapira et al. [42], the *expert* (sender) is a rule-based bot that knows the true quality of a hotel and strategically selects which textual review to reveal to human *decision makers* (DMs) in order to persuade them to "adopt" (book) the hotel. The interaction unfolds over multiple rounds: after observing the expert's chosen message and the cumulative history of past interactions, each human DM decides whether to adopt or reject the hotel. The learning task is to predict these human decisions—specifically, to estimate the probability that a DM will adopt in the current round given the sequence of revealed reviews and past actions.

The setup therefore combines linguistic, sequential, and strategic components, making it an ideal testbed for studying human decision making and choice prediction in complex, language-based environments.

Synthetic data generation. To investigate whether LLMs can substitute for human participants in such settings, Shapira et al. [42] designed a pipeline in which LLMs act as decision makers under varying personas [11], and games are simulated under multiple expert strategies to capture diverse interaction patterns. Each synthetic agent observes the same review and interaction history as the human counterpart and outputs a binary adoption decision, thereby generating synthetic behavioral trajectories. This *behavioral* data generation approach is contrasted with a *sentiment-based baseline*, where the decision is derived only from the review's sentiment score predicted by an LLM, ignoring historical and strategic context.

Results and implications. The authors train several predictive models on three data sources: (i) human-choice data, (ii) LLM-generated behavioral data, and (iii) the sentiment-based baseline. When evaluated on held-out human data, models trained solely on LLM-generated behavioral data achieve predictive accuracy comparable to, and sometimes exceeding, that of models trained on real human data. In contrast, the sentiment-based baseline performs substantially worse, indicating that the advantage of LLM-generated data stems from its ability to capture the sequential and incentive-aware structure of the interaction, rather than mere textual sentiment. Overall, Shapira et al. [42] demonstrate that LLMs can serve as *strategic simulators* capable of producing *incentive-consistent synthetic data* that embeds the behavioral dependencies induced by the underlying game. This insight exemplifies a broader principle of *incentive-aware synthetic data generation*: synthetic agents should not merely mimic linguistic signals, but also internalize the strategic and motivational structure of the environment that shapes human behavior.

3 Fine-Tuning a Competitive Search Agent

Background. Competitive search environments capture the strategic interplay between content creators (or *publishers*) who compete for exposure in ranking systems [26, 36]. In such settings, each publisher strategically adjusts or generates content to maximize visibility, anticipating both the behavior of the ranking algorithm and the responses of competing publishers. This paradigm has been studied extensively in both theoretical and applied literatures, encompassing models of game-theoretic analysis of search and recommendation ecosystems [7, 40, 47–50], adaptive behavior and learning dynamics [30, 31, 35], and recent advances in AI-driven content generation [6, 13, 33, 43]. These studies reveal that modern search and recommendation ecosystems are inherently strategic, and that designing robust, socially aligned mechanisms requires understanding learning dynamics when feedback and incentives are mediated by algorithms.

Game setup and learning task. Mordo et al. [33] study a multi-agent version of this problem, in which a set of *LLM-based publisher agents* compete for exposure in a search engine. Each round begins with a query (drawn from a diverse set that induces distinct competitive games) and a set of initial documents. Agents can modify or regenerate their documents before submitting them to the ranker, which evaluates all documents and returns a ranking. The resulting positions determine feedback signals, which serve as rewards for the agents. The overarching objective is to train publisher agents who learn to strategically improve their content and adapt to competitors' behavior in order to achieve better rankings. In this framework, training the agents amounts to fine-tuning language models to generate documents that perform well under a competitive ranking process.

Synthetic data generation. A central methodological contribution of Mordo et al. [33] is the introduction of two distinct synthetic data generation regimes: *Static Generation (SG)* and *Dynamic Generation (DG)*. In the static regime, data is collected from isolated document-generation episodes, where each agent produces candidate texts independently, without considering the responses of other agents or the evolving state of competition. In contrast, the dynamic regime constructs datasets from simulated multi-agent interactions [34, 51], where multiple publisher agents repeatedly compete and adapt to one another. These simulations capture the iterative feedback loops of real competitive environments, where agents refine their strategies based on past ranking outcomes and observed rival behaviors. The resulting dynamic datasets therefore embed the incentive structure of the environment.

Results and implications. Mordo et al. [33] show that models fine-tuned on dynamically generated data outperform those trained on static data in both in-distribution and out-of-distribution evaluations. Agents trained with *Reinforcement Learning from Ranker Feedback (RLRF)* using DG data consistently achieve higher win-rates compared to thosed trained using SG data, demonstrating that synthetic data generated through interactive, incentive-consistent pipelines leads to improved strategic behavior in multi-agent environments. This finding reinforces a broader lesson for incentive-aware synthetic data generation: when the downstream task involves strategic competition or feedback-driven adaptation, effective training requires *data generation pipelines that capture the dynamic incentives shaping behavior.*

4 Strategic Generative AI: Challenges, Limitations, and Vision

Challenges. Embedding strategic reasoning into GenAI introduces fundamental conceptual tensions. Standard ML frameworks assume a passive environment where data distributions are exogenous and fixed, while economic systems are inherently *endogenous*: each agent's data-generating process depends on others' incentives and beliefs. As a result, defining meaningful training objectives and evaluation metrics requires reconciling the probabilistic modeling view of GenAI with equilibrium-based reasoning from game theory. In such hybrid systems, equilibrium concepts interact with statistical ones, raising open questions about convergence, identifiability, and robustness under adaptive feedback.

Limitations. Even when synthetic data are generated dynamically within incentive-aware paradigms, significant robustness challenges remain. In realistic environments, the underlying incentives themselves evolve: agents' preferences may shift, competitors may change, action spaces may expand or shrink as new tools emerge, and market conditions may fluctuate. These dynamics imply that models trained on strategically generated data can still fail when deployed under altered incentive landscapes. Addressing these challenges requires developing methods that support *adaptive behavior* not only during training but also at evaluation and inference time. This calls for combining tools from RL and online learning with the strategic synthetic data paradigm, enabling systems to continually adjust to evolving equilibria and to determine when and how retraining should occur as interactions progress.

Vision. A key step forward is to develop a *unified framework for synthetic data generation in multi-agent environments*. Such a GenAI-driven framework should formalize how incentives, interactions, and information structures across diverse simulated games jointly determine the synthetic data distribution. It would allow researchers to model diverse domains (e.g., persuasion, search, voting and auctions) within a single generative schema that supports both *simulation* (dynamic data generation) and *learning* (synthetic data utilization). Concretely, this entails defining standard interfaces for (i) specifying agent objectives and observables, (ii) generating data through equilibrium or learning dynamics, and (iii) evaluating resulting datasets in terms of strategic fidelity, stability, and welfare impact. A unified approach would not only consolidate existing applications but also enable systematic development of learning systems trained (and potentially even evaluated) on strategically coherent synthetic data—capturing how incentives, competition, and feedback shape the data distribution itself in economic environments. Ultimately, the goal of Strategic GenAI is to establish a principled foundation for *synthetic data generation tailored to ML systems that operate in dynamic, incentive-driven economic environments*, ensuring that models learn, adapt, and generalize in alignment with the strategic feedback loops that govern real-world behavior.

References

[1] Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. 2025. Playing repeated games with large language models. *Nature Human Behaviour* (2025), 1–11.

[2] Reut Apel, Ido Erev, Roi Reichart, and Moshe Tennenholtz. 2022. Predicting decisions in language based persuasion games. *Journal of Artificial Intelligence Research* 73 (2022), 1025–1091.

[3] Itai Arieli, Yakov Babichenko, Omer Madmon, and Moshe Tennenholtz. 2025. Robust price discrimination. *Games and Economic Behavior* 154 (2025), 377–395.

[4] Itai Arieli, Omer Madmon, and Moshe Tennenholtz. 2024. Reputation-based persuasion platforms. *Games and Economic Behavior* 147 (2024), 128–147.

[5] Fengshuo Bai, Mingzhi Wang, Zhaowei Zhang, Boyuan Chen, Yinda Xu, Ying Wen, and Yaodong Yang. 2024. Efficient model-agnostic alignment via bayesian persuasion. *arXiv preprint arXiv:2405.18718* (2024).

[6] Niv Bardas, Tommy Mordo, Oren Kurland, and Moshe Tennenholtz. 2025. Automatic Document Editing for Improved Ranking. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*. Association for Computing Machinery, New York, NY, USA, 2779–2783. doi:10.1145/3726302.3730168

[7] Omer Ben-Porat and Moshe Tennenholtz. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems* 31 (2018).

[8] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2015. The limits of price discrimination. *American Economic Review* 105, 3 (2015), 921–957.

[9] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680* (2019).

[10] James Best and Daniel Quigley. 2024. Persuasion for the long run. *Journal of Political Economy* 132, 5 (2024), 1740–1791.

[11] Hyeong Kyu Choi and Yixuan Li. 2024. Picle: Eliciting diverse behaviors from large language models with persona in-context learning. *arXiv preprint arXiv:2405.02501* (2024).

[12] Lee Cohen, Saeed Sharifi-Malvajerdi, Kevin Stang, Ali Vakilian, and Juba Ziani. 2024. Bayesian strategic classification. *Advances in Neural Information Processing Systems* 37 (2024), 111649–111678.

[13] Erica Coppolillo, Federico Cinus, Marco Minici, Francesco Bonchi, and Giuseppe Manco. 2024. Engagement-Driven Content Generation with Large Language Models. *arXiv preprint arXiv:2411.13187* (2024).

[14] Vincent P Crawford and Joel Sobel. 1982. Strategic information transmission. *Econometrica: Journal of the Econometric Society* (1982), 1431–1451.

[15] Rachel Cummings, Nikhil R Devanur, Zhiyi Huang, and Xiangning Wang. 2020. Algorithmic price discrimination. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2432–2451.

[16] Caoyun Fan, Jindou Chen, Yaohui Jin, and Hao He. 2024. Can large language models serve as rational players in game theory? a systematic analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 17960–17967.

[17] Joseph Farrell and Matthew Rabin. 1996. Cheap talk. *Journal of Economic perspectives* 10, 3 (1996), 103–118.

[18] Denizalp Goktas, Amy Greenwald, Takayuki Osogami, Roma Patel, Kevin Leyton-Brown, Grant Schoenebeck, Daphne Cornelisse, Constantinos Daskalakis, Ian Gemp, John Horton, et al. 2025. Strategic Foundation Models. (2025).

[19] Mandeep Goyal and Qusay H Mahmoud. 2025. An llm-based framework for synthetic data generation. In *2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, 00340–00346.

[20] Xu Guo and Yiqiang Chen. 2024. Generative ai for synthetic data generation: Methods, challenges and the future. *arXiv preprint arXiv:2403.04190* (2024).

[21] Gillian K Hadfield and Andrew Koh. 2025. An economy of AI agents. *arXiv preprint arXiv:2509.01063* (2025).

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

[23] Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, et al. 2024. Game-theoretic llm: Agent workflow for negotiation games. *arXiv preprint arXiv:2411.05990* (2024).

[24] Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzman, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. *Advances in neural information processing systems* 31 (2018).

[25] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.

[26] Oren Kurland and Moshe Tennenholtz. 2022. Competitive search. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2838–2849.

[27] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems* 30 (2017).

[28] Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. Synthetic data generation with large language models for text classification: Potential and limitations. *arXiv preprint arXiv:2310.07849* (2023).

[29] Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. 2024. On llms-driven synthetic data generation, curation, and evaluation: A survey. *arXiv preprint arXiv:2406.15126* (2024).

[30] Omer Madmon, Idan Pipano, Itamar Reinman, and Moshe Tennenholtz. 2025. On the Convergence of No-Regret Dynamics in Information Retrieval Games with Proportional Ranking Functions. In *The Thirteenth International Conference on Learning Representations*. https://openreview.net/forum?id=jJXZvPe5z0

[31] Omer Madmon, Idan Pipano, Itamar Reinman, and Moshe Tennenholtz. 2025. The search for stability: Learning dynamics of strategic publishers with initial documents. *Journal of Artificial Intelligence Research* 83 (2025).

[32] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.

[33] Tommy Mordo, Sagie Dekel, Omer Madmon, Moshe Tennenholtz, and Oren Kurland. 2025. RLRF: Competitive Search Agent Design via Reinforcement Learning from Ranker Feedback. *arXiv preprint arXiv:2510.04096* (2025).

[34] Tommy Mordo, Tomer Kordonsky, Haya Nachimovsky, Moshe Tennenholtz, and Oren Kurland. 2025. LEMSS: LLM-Based Platform for Multi-Agent Competitive Search Simulation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3595–3605.

[35] Haya Nachimovsky and Moshe Tennenholtz. 2025. On the Power of Strategic Corpus Enrichment in Content Creation Games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 14019–14026.

[36] Haya Nachimovsky, Moshe Tennenholtz, and Oren Kurland. 2025. A Multi-Agent Perspective on Modern Information Retrieval. *arXiv preprint arXiv:2502.14796* (2025).

[37] Mihai Nadăș, Laura Dioșan, and Andreea Tomescu. 2025. Synthetic data generation using large language models: Advances in text and code. *IEEE Access* (2025).

[38] Chanwoo Park, Xiangyu Liu, Asuman Ozdaglar, and Kaiqing Zhang. 2024. Do llm agents have regret? a case study in online learning and games. *arXiv preprint arXiv:2403.16843* (2024).

[39] Maya Raifer, Guy Rotman, Reut Apel, Moshe Tennenholtz, and Roi Reichart. 2022. Designing an automatic agent for repeated language–based persuasion games. *Transactions of the Association for Computational Linguistics* 10 (2022), 307–324.

[40] Nimrod Raifer, Fiana Raiber, Moshe Tennenholtz, and Oren Kurland. 2017. Information retrieval meets game theory: The ranking competition between documents' authors. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 465–474.

[41] Eilam Shapira, Omer Madmon, Reut Apel, Moshe Tennenholtz, and Roi Reichart. 2025. Human choice prediction in language-based persuasion games: Simulation-based off-policy evaluation. *Transactions of the Association for Computational Linguistics* 13 (2025), 980–1006.

[42] Eilam Shapira, Omer Madmon, Roi Reichart, and Moshe Tennenholtz. 2024. Can LLMs Replace Economic Choice Prediction Labs? The Case of Language-based Persuasion Games. *arXiv preprint arXiv:2401.17435* (2024).

[43] Amit Sharma, Neha Patel, and Rajesh Gupta. 2022. Leveraging Reinforcement Learning and Natural Language Processing for Enhanced Social Media Content Optimization. *European Advanced AI Journal* 11, 8 (2022).

[44] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.

[45] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. 2017. Mastering the game of go without human knowledge. *nature* 550, 7676 (2017), 354–359.

[46] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *nature* 575, 7782 (2019), 350–354.

[47] Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. 2023. How Bad is Top-K Recommendation under Competing Content Creators?. In *International Conference on Machine Learning*. PMLR, 39674–39701.

[48] Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. 2024. Human vs. Generative AI in Content Creation Competition: Symbiosis or Conflict? *arXiv preprint arXiv:2402.15467* (2024).

[49] Fan Yao, Chuanhao Li, Karthik Abinav Sankararaman, Yiming Liao, Yan Zhu, Qifan Wang, Hongning Wang, and Haifeng Xu. 2024. Rethinking Incentives in Recommender Systems: Are Monotone Rewards Always Beneficial? *Advances in Neural Information Processing Systems* 36 (2024).

[50] Fan Yao, Yiming Liao, Mingzhe Wu, Chuanhao Li, Yan Zhu, James Yang, Jingzhou Liu, Qifan Wang, Haifeng Xu, and Hongning Wang. 2024. User welfare optimization in recommender systems with competing content creators. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3874–3885.

[51] Xiaopeng Ye, Chen Xu, Zhongxiang Sun, Jun Xu, Gang Wang, Zhenhua Dong, and Ji-Rong Wen. 2025. LLM-Empowered Creator Simulation for Long-Term Evaluation of Recommender Systems Under Information Asymmetry. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 201–211.