

Economics of Content Creation With Generative AI

Boaz Taitler

Technion — Israel Institute of Technology, Israel
Boaztaitler@campus.technion.ac.il

Abstract

Generative AI (GenAI) has become an integral part of our lives, serving both as a tool for creating content and as a platform that interacts with users. In doing so, it not only shapes user experiences but also competes with other actors within broader economic ecosystems. This dissertation examines how GenAI interacts with users and competitors within these ecosystems. Its goal is to deepen our understanding of the implications of integrating GenAI and to design approaches that increase social welfare. The contributions include developing theoretical models of these interactions and proposing solutions through algorithmic methods, mechanism design, and regulatory frameworks that ensure welfare-improving outcomes.

Introduction

GenAI systems have revolutionized content creation by producing high-quality outputs in seconds, a task that previously required significant time and effort. Trained on vast datasets, these systems generate content across diverse domains in both textual and visual forms, transforming practices for firms and individuals alike.

This revolution has introduced new competitive dynamics, both between GenAI and traditional content creators and between GenAI platforms and rival firms. For example, while GenAI can produce content rapidly and at low cost, it depends on continuous access to fresh data. As GenAI reduces competitors' engagement and revenue, it may also weaken their ability to generate the very data on which GenAI relies, creating a feedback loop with potentially adverse effects.

This proposal, situated at the intersection of machine learning and game theory, seeks to advance our understanding of GenAI's role in competitive ecosystems. First, it characterizes GenAI's impact on the various participants in these ecosystems and analyzes their strategic responses. Second, it proposes mechanisms and regulatory strategies to integrate GenAI in ways that mitigate negative effects and enhance overall welfare.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Related Work

This work contributes to the growing literature on the strategic and societal aspects of foundation models and machine learning (Boehmer, Fish, and Procaccia 2025; Duetting et al. 2024; Raghavan 2024). In this line of research, the work most closely related to ours is that of Yao et al. (2024), who model the competition between classical content creators and those using GenAI to compete for user engagement. Conceptually, Yao et al. (2024) views GenAI as a tool and models the strategic behavior and utilities of entities that adopt GenAI versus those who do not. In contrast, my work can be divided into two conceptual directions that differ from Yao et al. (2024). The first treats GenAI as a strategic entity, focusing on the effects of its strategic behavior on the surrounding environment. The second examines the use of GenAI in interactions that involve both competitive and cooperative aspects among users.

Research Plan

The overarching aim of this proposal is to understand how to better integrate GenAI into competitive ecosystems to maximize social welfare. It makes two main contributions. First, it models and analyzes the dynamics between GenAI and the broader ecosystems in which it operates. Second, it proposes solutions for the effective management of GenAI through algorithmic approaches, mechanism design, and regulatory frameworks.

Our first work (Ben-Porat et al. 2024), accepted to AAAI24, develops a principal-agent framework for steering autonomous agents through incentives. We study how a principal can teach a sequential policy by shaping an agent's rewards subject to a limited budget. The agent operates in a Markov decision process with its own reward function, while the principal has a distinct reward function and a finite budget that can be used to augment the agent's rewards. The goal is to design reward adjustments so the agent's optimal response policy maximizes the principal's payoff.

Our second work (Taitler and Ben-Porat 2025a), accepted to AAAI25 examines the tension between GenAI and online forums. Prior work (del Rio-Chanona, Laurentsyeva, and Wachs 2023) documented a decline in new posts on Stack Overflow following the release of ChatGPT. Motivated by this, we model competition between ChatGPT and Stack Overflow for user engagement in the Q&A domain. We

capture features such as Stack Overflow’s network effects and ChatGPT’s reliance on fresh training data. The analysis shows that ChatGPT’s strategic decisions about when to re-train can reduce users’ long-term social welfare, potentially making them worse off than in a world without ChatGPT.

Building on this line of research, our third work (Taitler and Ben-Porat 2025c), accepted to ICML 2025, investigates whether GenAI should provide complete answers to users. We model the process of data generation on the forum and capture GenAI’s accuracy as a function of its available data. Our analysis reveals an inherent exploration–exploitation tradeoff: answering yields immediate revenue, while not answering generates additional data that may improve future revenue. In particular, we show that offering partial answers or refraining from answering altogether can enhance both users’ social welfare and GenAI’s revenue. Finally, we propose an approximation algorithm that determines when to provide answers in order to optimize GenAI’s revenue.

Our fourth work (Taitler et al. 2025), currently under review, examines data sharing between a GenAI platform and a rival firm offering similar services. In our model, the platform may acquire data either from the firm or from external experts to improve accuracy and attract users, while the firm decides how much of its data to share. Using a mechanism design framework, we characterize the equilibrium for different prices of the firm’s data. We show that, in equilibrium, the firm may even be willing to pay the GenAI platform to purchase its data - a case we refer to as the ‘costly equilibrium.’ Finally, we identify the equilibria that yield Pareto improvements in both parties’ revenues relative to the case of no sharing, and we characterize the conditions under which the costly equilibrium is Pareto improving.

Our fifth work (Taitler and Ben-Porat 2025b), also under review, studies the use of GenAI in team projects. We model a manager with multiple tasks who can either hire employees for a share of the revenue or use GenAI to perform them. Employees then decide how much effort to invest and whether to rely on GenAI. We show that in a fixed team, some employees who would have exerted effort without access to GenAI may instead adopt GenAI without contributing effort, even when GenAI’s contribution is minimal. We further show that any employee whose marginal contribution exceeds that of GenAI can be critical to the team’s productivity. Moreover, as the manager relies on GenAI for more tasks, the marginal value of the remaining employees decreases, leading to either very large or very small teams.

Future works

There are two future research directions. The first is motivated by the growing availability of GenAI as a tool for copying and generating new content. With GenAI, content creators can adopt and modify the work of others, enabling them to produce new content more quickly. This dynamic introduces a challenge: creating content to attract users can, in fact, strengthen competitors and thereby reduce one’s own exposure through recommendation systems. Our preliminary results show that traditional Tullock-contest mechanisms fail to guarantee the existence of equilibrium. Therefore, our goal is to study how to incentivize content cre-

ators to invest effort under recommendation-based exposure, where each creator’s quality depends on the efforts of others. We propose a new mechanism that guarantees equilibrium existence and design an algorithm that achieves, in expectation, a fraction of the optimal total quality. Our next step is to identify axioms that uniquely characterize this mechanism. We aim to complete this project by February 2026.

The second direction examines the role of search engines within AI chats, focusing on the trade-offs between search engine optimization (SEO) and generative engine optimization (GEO). For websites, choosing which engine to optimize for is non-trivial: optimizing for one provides greater exposure there, but the type of exposure differs and may influence downstream outcomes. For example, SEO often leads to clicks and direct visits, while GEO produces mentions of the site as an information source. This distinction is crucial, as websites that rely on advertising cannot generate revenue without actual user visits. The goal is to study how the tension between SEO and GEO shapes websites’ incentives and the strategic decisions of their creators.

Acknowledgments

This research was supported by the Israel Science Foundation (ISF; Grant No. 3079/24).

References

- Ben-Porat, O.; Mansour, Y.; Moshkovitz, M.; and Taitler, B. 2024. Principal-agent reward shaping in mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9502–9510.
- Boehmer, N.; Fish, S.; and Procaccia, A. D. 2025. Generative Social Choice: The Next Generation. In *Forty-second International Conference on Machine Learning*.
- del Rio-Chanona, M.; Laurensyeva, N.; and Wachs, J. 2023. Are large language models a threat to digital public goods? evidence from activity on stack overflow. *arXiv preprint arXiv:2307.07367*.
- Duetting, P.; Mirrokni, V.; Paes Leme, R.; Xu, H.; and Zuo, S. 2024. Mechanism design for large language models. In *Proceedings of the ACM Web Conference 2024*, 144–155.
- Raghavan, M. 2024. Competition and diversity in generative ai. *arXiv preprint arXiv:2412.08610*.
- Taitler, B.; and Ben-Porat, O. 2025a. Braess’s paradox of generative ai. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 14139–14147.
- Taitler, B.; and Ben-Porat, O. 2025b. Collaborating with GenAI: Incentives and Replacements. *arXiv preprint arXiv:2508.20213*.
- Taitler, B.; and Ben-Porat, O. 2025c. Selective Response Strategies for GenAI. In *Forty-second International Conference on Machine Learning*.
- Taitler, B.; Madmon, O.; Tennenholtz, M.; and Ben-Porat, O. 2025. Data Sharing with a Generative AI Competitor. *arXiv preprint arXiv:2505.12386*.
- Yao, F.; Li, C.; Nekipelov, D.; Wang, H.; and Xu, H. 2024. Human vs. generative AI in content creation competition: symbiosis or conflict? *arXiv preprint arXiv:2402.15467*.