Economic Perspectives on Fairness in Information Retrieval

Chen Xu[2], Clara Rus[1](✉), Yuanna Liu[1], Marleen de Jonge[1], Jun Xu[2], and Maarten de Rijke[1]

[1] University of Amsterdam, Amsterdam, The Netherlands
{c.a.rus,y.liu8,m.r.h.dejonge,m.derijke}@uva.nl
[2] Renmin University of China, Beijing, China
{xc_chen,junxu}@ruc.edu.cn

Abstract. Fairness-aware information retrieval (IR) systems have been receiving more attention. Numerous fairness metrics and algorithms have been proposed. The complexity of fairness and IR systems makes it challenging to provide a systematic summary of the progress that has been made. This complexity calls for a more structured framework to navigate future fairness-aware IR research directions. The field of economics has long explored fairness, offering a strong theoretical and empirical foundation. Its system-oriented perspective enables the integration of IR fairness into a broader framework that considers societal and intertemporal trade-offs. In this tutorial, we first highlight that IR systems can be understood as a specialized economic market. Then, we reorganize fairness algorithms into an economic framework, which consists of three key economic dimensions: macro vs. micro, demand vs. supply, and short-term vs. long-term. We effectively view most fairness categories in IR from an economic perspective. Finally, we illustrate how this economic framework can be applied to various real-world IR applications and point out the future directions inspired by such a framework. Different from other fairness-aware tutorials, our tutorial not only provides a new and clear perspective to re-frame fairness-aware IR but also inspires the use of economic tools to solve fairness problems in IR. We hope this tutorial provides a fresh, broad perspective on fairness in IR, highlighting open problems and future research directions.

Keywords: Fairness · Economics · Information Retrieval

1 Motivation and Scope

1.1 Background on Fairness in IR

Information retrieval (IR) systems are designed to help users efficiently access information. However, IR systems such as search engines or recommender systems also influence and shape users' thoughts according to the given information [31]. This requires that IR systems should not only focus on accuracy, but also give attention to broader beyond-accuracy objectives such as fairness [22],

R. Campos et al. (Eds.): ECIR 2026, LNCS 16486, pp. 19–26, 2026.
https://doi.org/10.1007/978-3-032-21321-1_3

bias mitigation [5], and novelty [17] to promote a healthier ecosystem. Among these factors, fairness is crucial for IR systems as it ensures that the system does not discriminate against certain user groups [32] and provides more support for the long tail of valuable creators or item categories [33]. Although numerous fairness-aware IR algorithms have been proposed, they are categorized into more than ten distinct levels [11,20,39], including group vs. individual fairness [3], user vs. item fairness [32,33], static vs. dynamic fairness [40], and short-term v.s. long-term fairness [37]. Moreover, the measures of fairness also vary, such as max-min fairness [33], gini index [7], and demographic parity [29]. This complexity in categorization stems from the diverse definitions of fairness itself [28] and the involvement of multiple stakeholders, with distinct goals (e.g., users, items, platforms, creators) in IR [1]. This complexity makes it challenging for the IR community to systematically summarize the existing work and identify clear future research directions.

1.2 An Economic Perspective on Fairness in IR

Inspired by the field of economics, we can use established economic theory to systematically summarize and tackle complex fairness challenges in IR. In this tutorial, we first demonstrate that IR systems can be mapped to roles in a specialized economic market: users as consumers, items/documents as suppliers, and the platform as a central node, similar to the role of governments [36]. Specifically, in this market, users seek high-quality items, providers strive for maximum exposure of their products, and the platform aims to maximize profits and user satisfaction by delivering personalized services to users. Meanwhile, resources are limited (with a finite number of ranking slots), and the market price resembles the estimated ranking scores [4]. Given their shared structure, it is natural to bridge fairness issues in economics with those in IR systems.

Benefits of Using an Economic Perspective. The field of economics has long studied fairness, primarily focusing on how to allocate limited resources to best satisfy humans' unlimited desires [18,27]. Due to the established body of literature, an economic perspective on fairness offers a stronger theoretical and empirical foundation, allowing for a more structured analysis of complex fairness challenges in IR. Furthermore, the field of economics is concerned with interactions of different stakeholders in a system and analyzes the implications thereof over varying time spans. This systems-oriented thinking allows IR fairness to be embedded within a broader framework of societal and intertemporal trade-offs. By building on these well-established economic principles, fairness research in IR can benefit from greater coherence and avoid the proliferation of narrowly scoped or disconnected approaches.

IR Fairness Framework from an Economic Perspective. In this tutorial, we will elaborate on the following three economic dimensions: scale (macro vs. micro), objects (demand vs. supply), and time (long-term vs. short-term) of economic modeling to re-organize fairness-aware IR algorithms and evaluation methods. This perspective not only provides a novel and structured approach to reframing fairness-aware IR but also underscores the potential of applying economic principles and methodologies, such as game theory [25] and taxation theory [36], to rethink and tackle fairness challenges in IR systems.

2 Learning Objectives

Grounded in economic theory, this tutorial seeks to introduce and summarize fairness-aware IR from an economic perspective. By leveraging the well-established economic literature on fairness, we will systematically categorize and analyze fairness-related data, algorithms, and evaluation methodologies in IR, providing a unique perspective that not only deepens understanding but also identifies key open research directions for future exploration. Furthermore, this tutorial aims to equip attendees with economic insights to better comprehend and tackle broader trustworthy IR challenges beyond fairness, such as novelty, diversity, and interpretability. By drawing parallels between IR systems and economic markets, we hope to provide a fresh perspective that enables participants to design more equitable and transparent IR systems.

3 Scope

The tutorial will cover basic notions of fairness proposed in the IR community, economic perspectives on (un)fairness mitigation and evaluation, and build bridges between the two, with a special focus on teaching participants about conceptual tools and resources from economics and their potential in IR. The tutorial will *not* cover the long histories of fairness in either field nor offer a comprehensive side-by-side comparison of the many notions of fairness proposed in either field.

4 Relevance

Given the increasing importance and urgency of developing fair and trustworthy IR systems, we believe this is the right time to present a tutorial that helps researchers and industry practitioners consolidate current advancements and explore future directions in fairness-aware IR systems, especially in the era of LLMs. Moreover, our tutorial offers a fresh, well-structured perspective on emerging fairness issues in IR, enabling participants to develop a deeper understanding of these challenges while gaining economic insights to address them effectively in future research. This tutorial is highly relevant to the core themes of ECIR, with a specific focus on Fairness, Accountability, Transparency, Ethics, and Explainability (FATE) in IR, poised to inspire advancements in other trustworthy associated web applications.

5 Format

The tutorial will be a mixture of lectures, interactive demonstrations based on the FairDiverse environment [34], and Q&A sessions.

6 Length

The tutorial is 3 h (half day).

7 Detailed Outline

The outline for this tutorial is as follows:

00:00–00:20 Introduction
- Introduction of information retrieval systems.
- Introduction of fairness definition.
- Taxonomy for fairness in IR.
- Organization of the tutorial.

00:20–00:50 An Economic View on Fairness in IR (Chen)
- Introduction of economics.
- Introduction of fairness in economics.
- Relating IR systems to the economic markets.
- Re-framing fairness in IR through economics.

• **00:50–01:00** Q&A

• **01:00–01:30** Economic-based Fairness Mitigation and Evaluation Strategies I
- Objective: demand vs. supply
  - Demand (user) fairness [19,21,29].
  - Supply (provider) fairness [6,33,35,36].
  - Economic tools: taxation [36].

Break, with Q&A

01:30–02:00 Economic-based Fairness Mitigation and Evaluation Strategies II
- Scale: macro vs. micro
  - Micro (individual) fairness [30].
  - Macro (group or amortized) fairness [3,24,33,35].
  - Economic tools: game theory [2].

02:00–02:30 Economic-based Fairness Mitigation and Evaluation Strategies III
- Time: long-term vs. short-term
  - Long-term or dynamic fairness [13,37].
  - Short-term or static fairness [26,33].
  - Economic tools: risk theory [14].

02:30–02:50 Open Problems, Future Directions, and Conclusions
- Ignoring fairness problems in the IR markets.
- Benchmarks and evaluation.
- Conclusions.

02:50–03:00 Q&A

8 Target Audience

This tutorial explores fairness in IR from an economic standpoint, making it valuable for researchers, practitioners, and students interested in fairness, bias, and trustworthiness in IR systems such as search and recommendation, as well as for those investigating the economic dimensions of IR. It assumes only a basic understanding of IR and fairness, with no prior knowledge of economics required. The tutorial introduces essential economic concepts, such as taxation and interest rates, through clear, intuitive examples to ensure accessibility and understanding. Our goal is for the audience to learn how to apply economic frameworks and tools to understand fairness issues, and to extend this perspective to broader challenges such as bias and misinformation.

9 Tutorial History

We presented an earlier version of the tutorial at SIGIR 2025, with more than 30 people attending. The tutorial slides and materials are provided on the website https://economic-fairness-ir.github.io/. Compared to this earlier version, the current tutorial has been updated to include the latest research on unfairness [15,16,38] presented at recent conferences such as SIGIR and KDD. We have also added more economic examples to clarify key concepts and algorithms.

10 Related Tutorials

For fairness-aware IR, several related tutorials have emerged, including **Recsys'19, SIGIR'19** [9], considering fairness mainly from user study and evaluation perspective, **Recsys'20** [12], proposing different strategies to mitigate and evaluate unfairness in IR, **SIGIR'21** [23], proposing a taxonomy on fairness-aware algorithms in recommender systems, **CIKM'22** [10], focusing on a fairness taxonomy for search systems from the machine learning perspective and **ECIR'25** [8] providing a concise roadmap to fairness in information retrieval. The key distinction of our tutorial is that we organize the body of knowledge on fairness in IR through the structured lens of economics. This unique perspective not only provides a more systematic way to understand fairness in IR but also introduces economic tools for designing fair IR systems.

11 Materials

Slides and Bibliography. The slides and a bibliography file will be released on the tutorial website https://economic-fairness-ir.github.io/.

Related Benchmark. We provide a benchmark that introduced fairness-aware IR algorithms, evaluation metrics, and datasets [34], available on https://github.com/XuChen0427/FairDiverse.

Acknowledgments. This work was funded by the National Natural Science Foundation of China (No. 62472426). This research was (partially) supported by the Dutch Research Council (NWO), under project numbers 024.004.022, NWA.1389.20.183, and KICH3.LTP.20.006, and the European Union's Horizon Europe program under grant agreements No. 101070212 (FINDHR) and No. 101201510 (UNITE). All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdollahpouri, H., et al.: Multistakeholder recommendation: survey and research directions. User Model. User-Adap. Inter. **30**(1), 127–158 (2020). https://doi.org/10.1007/s11257-019-09256-1
2. Ben-Porat, O., Tennenholtz, M.: A game-theoretic approach to recommendation systems with strategic content providers. Adv. Neural Inf. Process. Syst. **31** (2018)
3. Biega, A.J., Gummadi, K.P., Weikum, G.: Equity of attention: amortizing individual fairness in rankings. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, pp. 405–414 (2018)
4. Calvano, E., Calzolari, G., Denicolo, V., et al.: Economics of recommender systems. In: Proceedings of the 18th ACM Conference on Recommender Systems, RecSys '24, pp. 1279–1280 (2024)
5. Dai, S., Xu, C., Xu, S., et al.: Bias and unfairness in information retrieval systems: new challenges in the LLM era. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pp. 6437–6447 (2024)
6. Do, V., Corbett-Davies, S., Atif, J., et al.: Two-sided fairness in rankings via lorenz dominance. Adv. Neural. Inf. Process. Syst. **34**, 8596–8608 (2021)
7. Do, V., Usunier, N.: Optimizing generalized gini indices for fairness in rankings. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 737–747 (2022)
8. Ekstrand, M.D.: Fairness in information access conceptual foundations and new directions. In: Advances in Information Retrieval: 47th European Conference on Information Retrieval, Part V, pp. 262–265. Springer, Heidelberg (2025). https://doi.org/10.1007/978-3-031-88720-8_41
9. Ekstrand, M.D., Burke, R., Diaz, F.: Fairness and discrimination in retrieval and recommendation. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1403–1404 (2019)
10. Fang, Y., Liu, H., Tao, Z., et al.: Fairness of machine learning in search engines. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, pp. 5132–5135 (2022)
11. Fang, Y., Singh, A., Tao, Z., et al.: Fairness in search systems. Found. Trends Inf. Retr. **18**(3), 262–416 (2024)

12. Gao, R., Shah, C.: Counteracting bias and increasing fairness in search and recommender systems. In: Proceedings of the 14th ACM Conference on Recommender Systems, pp. 745–747 (2020)
13. Ge, Y., Liu, S., Gao, R., Xian, Y., Li, Y., Zhao, X., et al.: Towards long-term fairness in recommendation. In: Proceedings of the 14th ACM International Conference on Web Search and Data Mining, pp. 445–453 (2021)
14. Grandell, J.: Aspects of Risk Theory. Springer, Heidelberg (2012)
15. Guo, H., Sun, Z., Wang, D., et al.: Enhancing new-item fairness in dynamic recommender systems. In: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1707–1716 (2025)
16. Gupta, A., Henzinger, T.A., Kueffner, K., et al.: Monitoring robustness and individual fairness. In: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining, vol. 2, pp. 790–801 (2025)
17. Hurley, N., Zhang, M.: Novelty and diversity in top-n recommendation-analysis and evaluation. ACM Trans. Internet Technol. **10**(4), 1–30 (2011)
18. Hurwicz, L.: The design of mechanisms for resource allocation. Am. Econ. Rev. **63**(2), 1–30 (1973)
19. Li, Y., Chen, H., Fu, Z., et al.: User-oriented fairness in recommendation. In: Proceedings of the Web Conference, pp. 624–632 (2021)
20. Li, Y., et al.: Fairness in recommendation: foundations, methods, and applications. ACM Trans. Intell. Syst. Technol. **14**(5), 1–48 (2023)
21. Li, Y., Chen, H., Xu, S., Ge, Y., Zhang, Y.: Towards personalized fairness based on causal notion. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1054–1063 (2021)
22. Li, Y., Du, M., Song, R., Wang, X., Wang, Y.: A survey on fairness in large language models. arXiv peprint arXiv:2308.10149 (2023)
23. Li, Y., Ge, Y., Zhang, Y.: Tutorial on fairness of machine learning in recommender systems. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2654–2657 (2021)
24. Naghiaei, M., Rahmani, H.A., Deldjoo, Y.: Cpfair: personalized consumer and producer fairness re-ranking for recommender systems. arXiv preprint arXiv:2204.08085 (2022)
25. Owen, G.: Game Theory. Emerald Group Publishing (2013)
26. Patro, G.K., Biswas, A., Ganguly, N., et al.: FairRec: two-sided fairness for personalized recommendations in two-sided platforms. In: Proceedings of the Web Conference, pp. 1194–1204 (2020)
27. Ramsey, F.P.: A contribution to the theory of taxation. Econ. J. **37**(145), 47–61 (1927)
28. Rawls, J.: Justice as fairness. Phil. Rev. **67**(2), 164–194 (1958)
29. Rus, C., Yates, A., de Rijke, M.: A study of pre-processing fairness intervention methods for ranking people. In: European Conference on Information Retrieval, pp. 336–350. Springer, Heidelberg (2024). https://doi.org/10.1007/978-3-031-56066-8_26
30. Saito, Y., Joachims, T.: Fair ranking as fair division: impact-based individual fairness in ranking. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pp. 1514–1524 (2022)
31. Stray, J., Halevy, A., Assar, P., et al.: Building human values into recommender systems: an interdisciplinary synthesis. ACM Trans. Recommend. Syst. **2**(3), 1–57 (2024)

32. Wang, L., Joachims, T.: User fairness, item fairness, and diversity for rankings in two-sided markets. In: Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval, pp. 23–41 (2021)
33. Xu, C., Chen, S., Xu, J., et al.: P-MMF: provider max-min fairness re-ranking in recommender system. In: Proceedings of the ACM Web Conference 2023, pp. 3701–3711 (2023)
34. Xu, C., Deng, Z., Rus, C., et al.: FairDiverse: a comprehensive toolkit for fair and diverse information retrieval algorithms (2025)
35. Xu, C., Xu, J., Ding, Y., et al.: FairSync: ensuring amortized group exposure in distributed recommendation retrieval. In: Proceedings of the ACM Web Conference 2024, WWW '24, pp. 1092–1102 (2024)
36. Xu, C., Ye, X., Wang, W., et al.: A taxation perspective for fair re-ranking. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1494–1503 (2024)
37. Xu, C., Ye, X., Xu, J., et al.: LTP-MMF: towards long-term provider max-min fairness under recommendation feedback loops. arXiv preprint arXiv:2308.05902 (2023)
38. Xu, C., Zhao, J., Wang, W., et al.: Understanding accuracy-fairness trade-offs in re-ranking through elasticity in economics. In: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 539–548 (2025)
39. Zehlike, M., Yang, K., Stoyanovich, J.: Fairness in ranking, part ii: learning-to-rank and recommender systems. ACM Comput. Surv. **55**(6), 1–41 (2022)
40. Zhu, Z., He, Y., Zhao, X., et al.: Popularity bias in dynamic recommendation. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, pp. 2439–2449 (2021)