Fairness in Information Retrieval from an Economic Perspective

Chen Xu
Renmin University of China
Beijing, China
xc_chen@ruc.edu.cn

Clara Rus
University of Amsterdam
Amsterdam, The Netherlands
c.a.rus@uva.nl

Yuanna Liu
University of Amsterdam
Amsterdam, The Netherlands
y.liu8@uva.nl

Marleen de Jonge
University of Amsterdam
Amsterdam, The Netherlands
m.r.h.dejonge@uva.nl

Jun Xu
Renmin University of China
Beijing, China
junxu@ruc.edu.cn

Maarten de Rijke
University of Amsterdam
Amsterdam, The Netherlands
m.derijke@uva.nl

Abstract

Recently, fairness-aware information retrieval (IR) systems have been receiving much 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 re-organize fairness algorithms through 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 we demonstrate its benefits in industrial scenarios. 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.

CCS Concepts

• **Information systems → Information retrieval**.

Keywords

Recommender systems, Fairness, Economics

ACM Reference Format:
Chen Xu, Clara Rus, Yuanna Liu, Marleen de Jonge, Jun Xu, and Maarten de Rijke. 2025. Fairness in Information Retrieval from an Economic Perspective. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25), July 13–18, 2025, Padua, Italy*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3726302.3731694

SIGIR '25, Padua, Italy

ACM ISBN 979-8-4007-1592-1/2025/07
https://doi.org/10.1145/3726302.3731694

1 Motivation

1.1 Background of Fairness in IR

Information retrieval (IR) systems are designed to help users efficiently access information [35, 58]. However, IR systems such as search engines or recommender systems also influence and shape users' thoughts according to the given information [46]. This requires that IR systems should not only focus on accuracy, but also give attention to broader beyond-accuracy objectives such as fairness [30], bias mitigation [9], and novelty [23] to promote a healthier ecosystem [9, 30]. Among these factors, fairness is crucial for IR systems as it ensures that the system does not discriminate against certain user groups [44, 48, 53] and provides more support for the long tail of valuable creators or item categories [37, 39, 51].

Although numerous fairness-aware IR algorithms have been proposed, they are categorized into more than ten distinct levels [3, 14, 17, 28, 49, 61], including group vs. individual fairness [5], user vs. item fairness [27, 51], static vs. dynamic fairness [62], and short-term v.s. long-term fairness [57]. Moreover, the measures of fairness also vary, such as max-min fairness [51], gini index [12], and demographic parity [39]. This complexity in categorization stems from the diverse definitions of fairness itself [41] and the involvement of multiple stakeholders (e.g., users, items, platforms, creators) in IR [1], each with distinct goals. This complexity makes it challenging for the IR community to systematically summarize the existing work and identify clear directions for future research.

1.2 Economic Perspective on Fairness in IR

Inspired by literature published in 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 [56]. 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 [7]. 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 [24, 40]. 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, objects, and time of economic modeling to reorganize fairness-aware IR algorithms and evaluation methods:

• *Scale: macro vs. micro.* In economics, microeconomics focuses on the fair allocation among individuals, while macroeconomics is concerned with the aggregate outcome and fair distribution of societal resources across people. Micro-level fairness aligns with individual fairness, which emphasizes personalized user behaviors [5, 45]. In contrast, macro-level fairness corresponds to group or amortized fairness [37, 51], which focuses on aggregating individual preferences at the group level.

• *Objective: demand vs. supply.* Economists study the interplay between the demand and supply side of markets, where the supply side provides goods, and the demand side consumes goods. Demand-level fairness corresponds to user fairness, aligning with the principle of equity [41], which ensures that similar users receive comparable results. In contrast, supply-level fairness relates to provider fairness, reflecting the concept of equality [51], which emphasizes supporting weaker suppliers.

• *Time: long-term vs. short-term.* In economics, the value of goods is often measured in terms of their short-term and long-term value. Long-term economic fairness evaluates resources based on their future value, aligning with long-term or dynamic fairness [57] in IR. In contrast, short-term fairness considers only the immediate value of an item, corresponding to short-term or static IR fairness [39].

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 [38] and taxation theory [56], to rethink and tackle fairness challenges in IR systems.

Fairness-aware applications using economic principles. We also present practical algorithms that implement these ideas in three real IR scenarios:

• *Recruitment search systems.* An application in the recruitment domain offers valuable insights into the complex interactions between various stakeholders. Unlike other settings, such as e-commerce, in a recruitment setting, there is a two-sided interaction [22, 60]: (i) candidates are searching for job offers and (ii) recruiters are searching for candidates given a job offer. This ecosystem allows us to apply economic principles, such as supply-demand theory, to better understand and analyze the complex interactions in the recruitment setting.

• *Next basket recommendation.* An e-commerce recommendation scenario where users exhibit both repetitive and exploratory purchase behaviors. Supply-side fairness takes into account the popularity and expected merits of the items. With the coexistence of repetitive and exploratory recommendation tasks, the evaluation and optimization of item fairness for next basket recommendation face unique challenges [26, 32].

• *Perzonalized financial product recommendations.* In banking and fintech, recommendation systems are increasingly used to match customers with financial products such as loans, credit cards, insurance plans, or investment portfolios, based on individual characteristics. Moreover, IR techniques, including ranking algorithms, are applied to tasks like credit scoring [6, 21]. These systems operate under real economic stakes—balancing financial risk, consumer protection, and regulatory fairness constraints.

1.3 Necessity and Timeliness of this Tutorial

Given the growing necessity and urgency of developing fair and trustworthy IR systems, we believe this is the right time to offer such a tutorial that helps researchers and industry practitioners summarize 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 gain a deeper understanding of these challenges while equipping them with economic insights to effectively address them in future research.

1.4 Qualification of Tutors

We have been working on fairness problems in information retrieval for a long time, underscored by a series of publications on the top-tier conferences and journals [26, 32, 33, 42–44, 51, 54–56]. Our team also has a strong educational background and working experience in an economic-related field. A highlight of our contributions includes two papers on IR fairness [51, 56], which were honored with the Spotlight-Best Paper Candidates at TheWebConf 2023 and the Best Paper Honorable Mention for SIGIR 2024. Moreover, our team has rich tutorial experience and has conducted more than 10 tutorials at various top-tier conferences, including SIGIR, TheWebConf, WSDM, KDD and RecSys [10, 25, 34, 59]. Our team also implemented a fairness-aware algorithm toolkit [52]. Thus we believe our tutorial will be comprehensive and insightful.

2 Objectives

Rooted in economic theory, this tutorial aims 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.

Additionally, this tutorial aims to equip attendees with economic insights to better understand and address 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 Relevance

This tutorial is acutely relevant to the core themes of SIGIR, with a specific focus on Fairness, Accountability, Transparency, Ethics, and Explainability (FATE) in IR, poised to inspire advancements in other trustworthy associated web applications.

In fairness-aware IR, several related tutorials have emerged, including **Recsys'19, SIGIR'19** [13], which consider fairness mainly from user study and evaluation perspective, **Recsys'20** [18], which proposes different tools and strategies to mitigate and evaluate unfairness in IR, **SIGIR'21** [31], which proposes a taxonomy on fairness-aware algorithms in recommender systems, and **CIKM'22** [16], which focuses on a fairness taxonomy for search systems from the machine learning perspective.

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 as powerful instruments for designing fair IR algorithms.

4 Format and Schedule

The outline for this tutorial is as follows:

00:00–00:20 Introduction (Maarten)
- 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 (Clara, Yuanna)
- Scale: macro vs. micro
  - Micro (individual) fairness [45].
  - Macro (group or amortized) fairness [5, 37, 51, 55].
  - Economic tools: game theory [2], risk theory [20].
- Objective: demand vs. supply
  - Demand (user) fairness [27, 29, 44, 50].
  - Supply (provider) fairness [4, 11, 51, 55, 56].
  - Economic tools: demand-supply theory [47].

Break, with Q&A

01:30–02:00 Economic-based Fairness Mitigation and Evaluation Strategies II (Clara, Yuanna)
- Time: long-term vs. short-term
  - Long-term or dynamic fairness [19, 36, 57].
  - Short-term or static fairness [39, 51, 56].
  - Economic tools: taxation [56], interests theory [8].

02:00–02:30 Application of economics-inspired IR (Marleen)
- Recruitment search systems [15, 42, 44].
- Next basket recommendation [26, 32, 33].
- Personalized financial product recommendations [6, 21].

02:30–02:50 Open Problems, Future Directions, and Conclusions (Chen, Maarten)
- Ignoring fairness problems in the IR markets.
- Insights from economic perspective.
- Benchmarks and evaluation.
- Conclusions.

02:50–03:00 Q&A

5 Materials

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

Bibliography. 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 [52], 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 agreement No 101070212. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

[1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder Recommendation: Survey and Research Directions. *User Modeling and User-Adapted Interaction* 30, 1 (2020), 127–158.

[2] 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).

[3] Nolwenn Bernard and Krisztian Balog. 2025. A Systematic Review of Fairness, Accountability, Transparency, and Ethics in Information Retrieval. *Comput. Surveys* 57, 6 (2025), 1–29.

[4] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H Chi, et al. 2019. Fairness in recommendation Ranking through Pairwise Comparisons. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2212–2220.

[5] Asia J. Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing Individual Fairness in Rankings. In *The 41st international ACM Sigir Conference on Research & Development in Information Retrieval*. 405–414.

[6] Bonnie G Buchanan. 2019. Artificial Intelligence in Finance. (2019).

[7] Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. 2024. Economics of Recommender Systems. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) *(RecSys '24)*. 1279–1280.

[8] Joseph W Conard. 1963. *An introduction to the theory of interest*. Univ of California Press.

[9] Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2024. 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* (Barcelona, Spain) *(KDD '24)*. 6437–6447.

[10] Sunhao Dai, Chen Xu, Shicheng Xu, Liang Pang, Zhenhua Dong, and Jun Xu. 2025. Unifying Bias and Unfairness in Information Retrieval: New Challenges in the LLM Era. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining (WSDM '25)*. 998–1001.

[11] Virginie Do, Sam Corbett-Davies, Jamal Atif, and Nicolas Usunier. 2021. Two-sided Fairness in Rankings via Lorenz Dominance. *Advances in Neural Information Processing Systems* 34 (2021), 8596–8608.

[12] Virginie Do and Nicolas Usunier. 2022. Optimizing Generalized Gini Indices for Fairness in Rankings. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 737–747.

[13] Michael D. Ekstrand, Robin Burke, and Fernando Diaz. 2019. Fairness and Discrimination in Retrieval and Recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) *(SIGIR'19)*. 1403–1404.

[14] Michael D. Ekstrand, Anubrata Das, Robin Burke, Fernando Diaz, et al. 2022. Fairness in information access systems. *Foundations and Trends in Information Retrieval* 16, 1-2 (2022), 1–177.

[15] Alessandro Fabris, Nina Baranowska, Matthew J Dennis, David Graus, Philipp Hacker, Jorge Saldivar, Frederik Zuiderveen Borgesius, and Asia J Biega. 2025. Fairness and bias in algorithmic hiring: A multidisciplinary survey. *ACM Transactions on Intelligent Systems and Technology* 16, 1 (2025), 1–54.

[16] Yi Fang, Hongfu Liu, Zhiqiang Tao, and Mikhail Yurochkin. 2022. Fairness of Machine Learning in Search Engines. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 5132–5135.

[17] Yi Fang, Ashudeep Singh, Zhiqiang Tao, et al. 2024. Fairness in Search Systems. *Foundations and Trends in Information Retrieval* 18, 3 (2024), 262–416.

[18] Ruoyuan Gao and Chirag Shah. 2020. Counteracting Bias and Increasing Fairness in Search and Recommender Systems. In *Proceedings of the 14th ACM Conference on Recommender Systems* (Virtual Event, Brazil) *(RecSys '20)*. 745–747.

[19] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. 2021. Towards Long-term Fairness in Recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*. 445–453.

[20] Jan Grandell. 2012. *Aspects of Risk Theory*. Springer Science & Business Media.

[21] Janin Karoli Hentzen, Arvid Hoffmann, Rebecca Dolan, and Erol Pala. 2022. Artificial Intelligence in Customer-facing Financial Services: A Systematic Literature Review and Agenda for future research. *International Journal of Bank Marketing* 40, 6 (2022), 1299–1336.

[22] Xiao Hu, Yuan Cheng, Zhi Zheng, Yue Wang, Xinxin Chi, and Hengshu Zhu. 2023. Boss: A Bilateral Occupational-Suitability-Aware Recommender System for Online Recruitment. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4146–4155.

[23] Neil Hurley and Mi Zhang. 2011. Novelty and diversity in top-n recommendation–analysis and evaluation. *ACM Transactions on Internet Technology (TOIT)* 10, 4 (2011), 1–30.

[24] Leonid Hurwicz. 1973. The Design of Mechanisms for Resource Allocation. *The American Economic Review* 63, 2 (1973), 1–30.

[25] Wenqiang Lei, Chongming Gao, and Maarten de Rijke. 2021. RecSys 2021 Tutorial on Conversational Recommendation: Formulation, Methods, and Evaluation. In *RecSys 2021: 15th ACM Conference on Recommender Systems*. ACM, 824–844.

[26] Ming Li, Yuanna Liu, Sami Jullien, Mozhdeh Ariannezhad, Andrew Yates, Mohammad Aliannejadi, and Maarten de Rijke. 2024. Are We Really Achieving Better Beyond-Accuracy Performance in Next Basket Recommendation?. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. 924–934.

[27] Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2021. User-oriented Fairness in Recommendation. In *Proceedings of the Web Conference*. 624–632.

[28] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, Juntao Tan, Shuchang Liu, and Yongfeng Zhang. 2023. Fairness in recommendation: Foundations, methods, and applications. *ACM Transactions on Intelligent Systems and Technology* 14, 5 (2023), 1–48.

[29] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2021. Towards Personalized Fairness based on Causal Notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1054–1063.

[30] Yingji Li, Mengnan Du, Rui Song, Xin Wang, and Ying Wang. 2023. A Survey on Fairness in Large Language Models. *arXiv peprint arXiv:2308.10149* (2023).

[31] Yunqi Li, Yingqiang Ge, and Yongfeng Zhang. 2021. 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* (Virtual Event, Canada) *(SIGIR '21)*. 2654–2657.

[32] Yuanna Liu, Ming Li, Mohammad Aliannejadi, and Maarten de Rijke. 2025. Repeat-bias-aware Optimization of Beyond-accuracy Metrics for Next Basket Recommendation. *arXiv e-prints* (2025), arXiv–2501.

[33] Yuanna Liu, Ming Li, Mozhdeh Ariannezhad, Masoud Mansoury, Mohammad Aliannejadi, and Maarten de Rijke. 2024. Measuring Item Fairness in Next Basket Recommendation: A Reproducibility Study. In *European Conference on Information Retrieval*. Springer, 210–225.

[34] Yu-An Liu, Ruqing Zhang, Jiafeng Guo, and Maarten de Rijke. 2024. Robust Information Retrieval. In *SIGIR 2024: 47th international ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 3009–3012.

[35] Christopher D. Manning, Hinrich Schütze, and Prabhakar Raghavan. 2009. *An Introduction to Information Retrieval*. Cambridge university press.

[36] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the Trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 2243–2251.

[37] Mohammadmehdi Naghiaei, Hossein A Rahmani, and Yashar Deldjoo. 2022. CPFair: Personalized Consumer and Producer Fairness Re-ranking for Recommender Systems. *arXiv preprint arXiv:2204.08085* (2022).

[38] Guillermo Owen. 2013. *Game theory*. Emerald Group Publishing.

[39] Gourab K. Patro, Arpita Biswas, Niloy Ganguly, Krishna P. Gummadi, and Abhijnan Chakraborty. 2020. FairRec: Two-sided Fairness for Personalized Recommendations in Two-sided Platforms. In *Proceedings of the Web Conference*. 1194–1204.

[40] Frank P. Ramsey. 1927. A Contribution to the Theory of Taxation. *The Economic Journal* 37, 145 (1927), 47–61.

[41] John Rawls. 1958. Justice as Fairness. *The Philosophical Review* 67, 2 (1958), 164–194.

[42] Clara Rus, Maarten de Rijke, and Andrew Yates. 2023. Counterfactual Representations for Intersectional Fair Ranking in Recruitment.. In *HR@ RecSys*.

[43] Clara Rus, Gabrielle Poerwawinata, Andrew Yates, and Maarten de Rijke. 2024. AnnoRank: A Comprehensive Web-Based Framework for Collecting Annotations and Assessing Rankings. In *CIKM 2024: 33rd ACM International Conference on Information and Knowledge Management*. ACM, 5400–5404.

[44] Clara Rus, Andrew Yates, and Maarten de Rijke. 2024. A Study of Pre-processing Fairness Intervention Methods for Ranking People. In *European Conference on Information Retrieval*. Springer, 336–350.

[45] Yuta Saito and Thorsten Joachims. 2022. 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*. 1514–1524.

[46] Jonathan Stray, Alon Halevy, Parisa Assar, Dylan Hadfield-Menell, Craig Boutilier, Amar Ashar, Chloe Bakalar, Lex Beattie, Michael Ekstrand, Claire Leibowicz, et al. 2024. Building human values into recommender systems: An interdisciplinary synthesis. *ACM Transactions on Recommender Systems* 2, 3 (2024), 1–57.

[47] John Von Neumann and Oskar Morgenstern. 1947. *Theory of Games and Economic Behavior* (2nd rev ed.). Princeton University Press.

[48] Lequn Wang and Thorsten Joachims. 2021. 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*. 23–41.

[49] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Transactions on Information Systems* 41, 3 (2023), 1–43.

[50] Le Wu, Lei Chen, Pengyang Shao, Richang Hong, Xiting Wang, and Meng Wang. 2021. Learning Fair Representations for Recommendation: A Graph-based Perspective. In *Proceedings of the Web Conference 2021*. 2198–2208.

[51] Chen Xu, Sirui Chen, Jun Xu, Weiran Shen, Xiao Zhang, Gang Wang, and Zhenhua Dong. 2023. P-MMF: Provider Max-min Fairness Re-ranking in Recommender System. In *Proceedings of the ACM Web Conference 2023*. 3701–3711.

[52] Chen Xu, Zhirui Deng, Clara Rus, Xiaopeng Ye, Yuanna Liu, Jun Xu, Zhicheng Dou, Ji-Rong Wen, and Maarten de Rijke. 2025. FairDiverse: A Comprehensive Toolkit for Fair and Diverse Information Retrieval Algorithms. *arXiv preprint arXiv:2502.11883* (February 2025).

[53] Chen Xu, Wenjie Wang, Yuxin Li, Liang Pang, Jun Xu, and Tat-Seng Chua. 2023. Do LLMs Implicitly Exhibit User Discrimination in Recommendation? An Empirical Study. *arXiv* (2023).

[54] Chen Xu, Wenjie Wang, Yuxin Li, Liang Pang, Jun Xu, and Tat-Seng Chua. 2024. A Study of Implicit Ranking Unfairness in Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 7957–7970.

[55] Chen Xu, Jun Xu, Yiming Ding, Xiao Zhang, and Qi Qi. 2024. FairSync: Ensuring Amortized Group Exposure in Distributed Recommendation Retrieval. In *Proceedings of the ACM Web Conference 2024* (Singapore, Singapore) *(WWW '24)*. 1092–1102.

[56] Chen Xu, Xiaopeng Ye, Wenjie Wang, Liang Pang, Jun Xu, and Tat-Seng Chua. 2024. A Taxation Perspective for Fair Re-ranking. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. 1494–1503.

[57] Chen Xu, Xiaopeng Ye, Jun Xu, Xiao Zhang, Weiran Shen, and Ji-Rong Wen. 2023. LTP-MMF: Towards Long-term Provider Max-min Fairness Under Recommendation Feedback Loops. *arXiv preprint arXiv:2308.05902* (2023).

[58] Jun Xu, Xiangnan He, and Hang Li. 2018. Deep Learning for Matching in Search and Recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. 1365–1368.

[59] Jun Xu, Xiangnan He, and Hang Li. 2018. Deep Learning for Matching in Search and Recommendation. In *Proceedings of the 2018 World Wide Web Conference*.

[60] Chen Yang, Yupeng Hou, Yang Song, Tao Zhang, Ji-Rong Wen, and Wayne Xin Zhao. 2022. Modeling Two-way Selection Preference for Person-Job Fit. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 102–112.

[61] Meike Zehlike, Ke Yang, and Julia Stoyanovich. 2022. Fairness in Ranking, Part II: Learning-to-rank and Recommender Systems. *Comput. Surveys* 55, 6 (2022), 1–41.

[62] Ziwei Zhu, Yun He, Xing Zhao, and James Caverlee. 2021. Popularity Bias in Dynamic Recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2439–2449.