



Recommender System Fairness for Small Sellers on Multi-Sided Online Marketplaces

Sona Balasubramanian^{1*}, Kavita Park², Shanice Girard³

¹Department of Computer Science, Ural Graduate Institute of Computer Science, Hungary

²Department of Computer Science, Limpopo Center for Computer Science, Tanzania

³Department of Computer Science, Mekong Institute for Advanced Computer Science, United Kingdom

Article Info

Volume 1, Issue 1, 2026

doi: [10.62715/ijecis.2026.v1.i1.pp8-14](https://doi.org/10.62715/ijecis.2026.v1.i1.pp8-14)

Received: 05 February 2026

Accepted: 03 April 2026

Published: 20 May 2026

Abstract

Online marketplaces depend on recommender and ranking systems to decide which listings a shopper sees, and for a small seller that decision largely determines whether a sale is possible at all. Research on fairness in recommendation, which began with concern for the people receiving recommendations, has widened to include those being recommended. This narrative review examines that provider-side literature from the standpoint of small sellers on multi-sided marketplaces. It describes the mechanisms that concentrate exposure on established sellers, including popularity bias, position bias, cold start and feedback loops in training data. It then compares the main definitions of provider fairness, the algorithmic methods proposed to achieve them, and the evidence on trade-offs with consumer relevance and platform revenue. Problems of evaluation are discussed, as is the regulatory context in the European Union, where ranking transparency and self-preferencing are now subject to law. The review concludes that technical methods are mature relative to the understanding of what small sellers are owed, of how sellers respond strategically, and of long-run effects on the marketplace, and it identifies these as priorities for research.

Keywords: *Recommender systems, Provider fairness, Fairness of exposure, Online marketplaces, Small sellers, Multi-sided platforms*

© 2026 Sona Balasubramanian et al. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

1. Introduction

A marketplace with millions of listings is usable only because software chooses a few dozen of them for each shopper. Search ranking, recommendation carousels, panels of related products and the default seller attached to a product page all perform this selection. For consumers the result is convenience. For sellers it is something closer to a licence to trade: a listing that is never shown cannot be bought, however good the product or the price. Because attention falls steeply with position, small differences in predicted relevance become large differences in sales.

Early work on algorithmic fairness in recommendation asked whether different groups of users received recommendations of equal quality. Burke (2017) pointed out that recommender systems serve

* Corresponding author: Sona Balasubramanian, Department of Computer Science, Ural Graduate Institute of Computer Science, Hungary.
E-mail: sona.balasubramanian@ural.edu

© 2026 Sona Balasubramanian et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

several parties at once and distinguished fairness toward consumers, fairness toward providers, and the combination of the two. The idea was developed into a broader programme of multistakeholder recommendation ([Abdollahpouri *et al.*, 2020](#)), and provider-side fairness is now a recognised topic in surveys of the field ([Ekstrand *et al.*, 2022](#)).

This review reads that literature with one group of providers in mind: small sellers, meaning independent merchants with few listings, short sales histories and no bargaining power with the platform. Their position differs from that of the protected groups usually considered in fairness research, because smallness is an economic condition and not a legally protected attribute, and because the platform may itself compete with them. The review asks what is known about why small sellers receive little exposure, what fairness toward them could mean, which methods exist, and what remains unresolved.

2. Scope and Approach

The article is a narrative review of published work in recommender systems, information retrieval, and the economics and law of platforms. It concentrates on contributions that introduced a definition, a method or an empirical regularity on which later work has built, and it does not attempt a complete census of a literature that is growing quickly. Material on ranking is included alongside recommendation, since on a marketplace the two are implemented by similar learning-to-rank models and raise the same questions of exposure. Fairness toward consumers, and fairness across demographic groups of providers, are discussed only where they bear on sellers of different size.

3. Why Small Sellers Receive Little Exposure

Four mechanisms, each well documented, work against the small seller. The first is popularity bias. Collaborative filtering learns from past interactions, and items with many interactions are both better modelled and more often recommended, so that recommendation lists over-represent what is already popular ([Abdollahpouri *et al.*, 2017](#)). The second is the cold-start problem: a new listing or a new seller has no interaction history from which to infer relevance ([Schein *et al.*, 2002](#)). Small sellers are disproportionately new, and their listings spend more of their lives in this state.

The third mechanism is position bias. Users examine results roughly from the top down and often stop early, so a click is evidence of position as well as of relevance ([Craswell *et al.*, 2008](#)). A model trained on raw clicks therefore learns that items shown high were good, whether or not they were. The fourth mechanism joins the others into a loop, shown in Figure 1. The system's own output determines what users can interact with, those interactions become the next round of training data, and the model's earlier choices are confirmed. Simulation work has shown that such algorithmic confounding homogenises user behaviour without improving utility ([Chaney *et al.*, 2018](#)).


The economic literature gives a mixed picture of the aggregate result. Online retail made a far wider range of products available than physical shops could stock, and consumers gained substantially from that variety ([Brynjolfsson *et al.*, 2003](#)), an observation often discussed under the name of the long tail. Yet availability is not exposure. [Fleder and Hosanagar \(2009\)](#) showed analytically and by simulation that common recommender designs can reduce the diversity of sales in aggregate even while individual users are led to items new to them, because the system steers many users toward the same items. For a small seller the practical meaning is that being listed in the tail does not ensure being found.

4. What Fairness Toward Sellers Could Mean

Definitions of provider fairness differ in what is to be equalised, among whom, and over what period. Table 1 summarises the principal notions. Most treat exposure, meaning the attention a ranking position is expected to receive, as the resource being distributed.

[Singh and Joachims \(2018\)](#) gave the formulation that most later work adopts. They represented a ranking policy as a probability distribution over rankings, defined exposure through a position-discount model, and expressed fairness as linear constraints, which lets the platform maximise expected utility to

users subject to a chosen exposure rule. Their constraints include equal exposure across groups and exposure in proportion to the groups' relevance. Biega et al. (2018) argued that fairness within a single ranking is unattainable when several items are nearly equally relevant, since only one can be first, and proposed that attention be made proportional to relevance cumulatively over a sequence of rankings.



Other definitions set floors instead of proportions. The FA*IR algorithm guarantees that the share of a protected group in every prefix of a top-k list does not fall significantly below a target (Zehlike et al., 2017). FairRec, designed for two-sided platforms, guarantees each producer a minimum level of exposure while keeping the allocation nearly envy-free among customers (Patro et al., 2020). Underlying all of these is the older distinction between group fairness and individual fairness, the latter requiring that similar individuals be treated similarly (Dwork et al., 2012). A survey by Zehlike et al. (2022) classifies ranking methods by the normative position each embodies.

Table 1: Principal notions of provider-side fairness in ranking and recommendation, and the difficulty each meets when applied to small sellers

Notion	Requirement	Source	Difficulty for small sellers
Equal exposure across groups	Groups of providers receive the same average exposure	Singh and Joachims, 2018	Disregards differences in relevance and quality
Exposure proportional to merit	Exposure of each group tracks its relevance to users	Singh and Joachims, 2018	Merit estimates are biased by past exposure
Equity of attention	Attention proportional to relevance, accumulated over many rankings	Biega et al., 2018	New sellers appear in too few rankings to accumulate
Minimum share in the top-k	Every prefix of the list holds a minimum share of a group	Zehlike et al., 2017	Needs a group boundary and a chosen quota
Guaranteed minimum exposure	Each provider receives at least a floor of exposure	Patro et al., 2020	A floor may be too low to yield sales
Individual fairness	Similar providers are treated similarly	Dwork et al., 2012	Rests on a contested measure of similarity

Applied to small sellers, each definition meets a difficulty of its own. Proportional rules depend on a measure of merit, and predicted relevance is contaminated by the exposure history it is meant to correct. Group rules need a boundary, and any threshold of size invites sellers to arrange themselves on its favourable side. Floors can be met with exposure in placements that seldom convert. And none of the definitions says why size should matter. Possible grounds include equal opportunity to compete, the long-run value of a diverse seller base to consumers, and the dependence of sellers on a platform they cannot readily leave. These are different reasons and they support different rules.

5. Methods for Allocating Exposure

Methods are usually grouped by where they act. Pre-processing approaches correct the training data. Inverse propensity weighting treats the logging policy as an intervention and reweights observed interactions by the probability that they could have been observed, which removes selection bias from ratings ([Schnabel et al., 2016](#)) and position bias from clicks ([Joachims et al., 2017](#)). These techniques were developed for accuracy, not fairness, but they address the contamination of merit estimates directly and are arguably the first step that any seller-fair system should take.

In-processing approaches change the learning objective. A regularisation term can penalise lists drawn only from popular items ([Abdollahpouri et al., 2017](#)), and provider constraints can be built into the optimisation that generates recommendations ([Sürer et al., 2018](#)). Post-processing approaches leave the model alone and re-rank its output, which is how FA*IR and FairRec operate and why they are attractive to platforms that are reluctant to retrain production models. Dynamic methods go further and treat fairness as a control problem: the FairCo controller adjusts rankings over time so that amortised exposure converges to the target while relevance estimates are debiased as data arrive ([Morik et al., 2020](#)).

Exploration deserves separate mention because it speaks to the cold-start condition in which small sellers so often find themselves. A system that occasionally shows uncertain items learns their true appeal and gives new listings a chance to earn a history. From the seller's side, a guaranteed period of exploratory exposure may matter more than any steady-state fairness constraint, since it determines whether a listing ever leaves the cold-start state. The fairness literature and the literature on exploration in recommendation have developed largely apart, and the connection between them is under-studied.

6. Trade-offs With Consumers and the Platform

Any reallocation of exposure away from the items a model ranks highest has a cost in predicted relevance, and the size of that cost is the central empirical question. [Mehrotra et al. \(2018\)](#), working with data from a music streaming service, used counterfactual estimation to study the trade-off between relevance, supplier fairness and user satisfaction, and considered policies that adapt the weight placed on fairness to each user's tolerance for less familiar suppliers. The general lesson is that the trade-off is not a fixed exchange rate. It depends on how many alternatives of nearly equal relevance exist, and in broad product categories there are often many.

The platform's own interest is harder to read from outside. In the theory of two-sided markets a platform gains from participation on both sides and sets its terms to balance them ([Rochet and Tirole, 2003](#)), and a marketplace in the sense of [Hagiu and Wright \(2015\)](#) leaves control of the sale to independent sellers. That gives the platform a reason to keep small sellers viable. It also earns commission and advertising revenue that rise with the sales of large sellers, and it may sell its own products alongside theirs. Whether seller fairness aligns with platform profit is therefore an empirical matter that varies by platform, and published evidence from marketplaces themselves is scarce.

7. Problems of Evaluation

Evaluating provider fairness is more difficult than defining it. Offline metrics computed on logged data inherit the biases of the logging policy, so an exposure measure may look fair on historical data and behave differently once deployed. Exposure itself is modelled, typically with a logarithmic or geometric discount by position, and real attention on a mobile grid or a carousel may not follow the assumed curve.

Online experiments face a further obstacle: sellers compete for the same finite attention, so treating some sellers changes outcomes for the rest, and the usual assumption that experimental units do not interfere with one another fails.

There is also the question of what to measure. Exposure is an input. Sellers care about sales, revenue and survival, and the mapping from exposure to those outcomes differs with price, reviews and fulfilment. A small seller given more impressions but lacking the reviews that reassure buyers may convert few of them, and a system judged fair by exposure may leave outcomes unchanged. Long-run measures, such as the rate at which new sellers reach a sustainable volume or the concentration of sales over time, would be more informative and are seldom reported.

8. Regulation and Self-Preferencing

Law has begun to address ranking on platforms, at least in the European Union. Regulation (EU) 2019/1150 obliges online intermediation services to set out in their terms the main parameters determining ranking and the reasons for their relative importance, and to describe any differentiated treatment they give to their own goods (European Union, 2019). The Digital Markets Act goes beyond transparency: designated gatekeepers may not rank their own products and services more favourably than those of third parties and must apply transparent, fair and non-discriminatory conditions to ranking (European Union, 2022). The concern behind these provisions, that a platform acting as both marketplace and merchant has the means and motive to disadvantage dependent sellers, was set out influentially by Khan (2017).

These rules concern the platform's treatment of rivals and the disclosure of ranking criteria. They do not require any particular distribution of exposure among third-party sellers, and a ranking that is transparent and free of self-preferencing can still concentrate attention on a few large merchants through the mechanisms described earlier. The technical literature and the legal one thus address adjacent problems. A useful bridge would be auditing methods that let a regulator or a seller association test a disclosed ranking rule against observed exposure.

9. Gaps and Future Work

Three gaps stand out. One is normative. The field has many formal definitions and little argument about which of them a small seller has a claim to, or on what ground. Work that connects exposure rules to explicit principles, such as opportunity for new entrants or protection of dependent trading partners, would make the choice among definitions less arbitrary. A second gap is strategic behaviour. Sellers are not passive items. They respond to ranking rules by splitting or merging accounts, duplicating listings, buying advertising and adjusting prices, and a fairness rule keyed to size or novelty creates incentives of its own. Mechanism-design analysis of fairness interventions is at an early stage.

The third gap is evidence from the field. Many published evaluations use public datasets from domains such as film, music or hiring, where providers are not merchants and the unit economics of a sale are absent. Studies on marketplace data, with outcomes measured in seller revenue and survival, are needed, as are studies of sponsored listings, which occupy many of the most visible positions and sit outside most fairness formulations. Small sellers who cannot afford to advertise may be affected more by the growth of paid placement than by any property of the organic ranking model.

10. Conclusion

Research on recommender fairness now recognises sellers as parties with a stake in the ranking. The mechanisms that concentrate exposure, namely popularity bias, cold start, position bias and feedback through training data, are well understood, and a range of methods can redistribute exposure at a cost in relevance that depends on how many close substitutes exist. What the literature does not yet supply is a settled account of what small sellers are owed, evidence on how sellers and platforms behave once a rule is in place, or measurement of the outcomes that matter to a merchant.

Progress will depend less on new re-ranking algorithms than on connecting three bodies of work that

have grown separately: the technical study of exposure, the economics of multi-sided platforms, and the law on ranking transparency and self-preferencing. Marketplaces that open data and experiments to researchers, under suitable safeguards, would allow that connection to be made on evidence.

Conflict of interest

To be declared by the author(s) before publication.

Funding

To be declared by the author(s) before publication.

References

- Abdollahpouri, H., Adomavicius, G., Burke, R., Guy, I., Jannach, D., Kamishima, T., Krasnodebski, J. and Pizzato, L. (2020). [Multistakeholder recommendation: Survey and research directions](#). *User Modeling and User-Adapted Interaction*, 30(1): 127-158. <https://doi.org/10.1007/s11257-019-09256-1>
- Abdollahpouri, H., Burke, R. and Mobasher, B. (2017). [Controlling popularity bias in learning-to-rank recommendation](#). *Proceedings of the Eleventh ACM Conference on Recommender Systems*, 42-46. <https://doi.org/10.1145/3109859.3109912>
- Biega, A.J., Gummadi, K.P. and Weikum, G. (2018). [Equity of attention: Amortizing individual fairness in rankings](#). *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval*, 405-414. <https://doi.org/10.1145/3209978.3210063>
- Brynjolfsson, E., Hu, Y. and Smith, M.D. (2003). [Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers](#). *Management Science*, 49(11): 1580-1596. <https://doi.org/10.1287/mnsc.49.11.1580.20580>
- Burke, R. (2017). [Multisided fairness for recommendation](#). *arXiv preprint*, arXiv:1707.00093.
- Chaney, A.J.B., Stewart, B.M. and Engelhardt, B.E. (2018). [How algorithmic confounding in recommendation systems increases homogeneity and decreases utility](#). *Proceedings of the 12th ACM Conference on Recommender Systems*, 224-232. <https://doi.org/10.1145/3240323.3240370>
- Craswell, N., Zoeter, O., Taylor, M. and Ramsey, B. (2008). [An experimental comparison of click position-bias models](#). *Proceedings of the 2008 International Conference on Web Search and Data Mining*, 87-94. <https://doi.org/10.1145/1341531.1341545>
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O. and Zemel, R. (2012). [Fairness through awareness](#). *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 214-226. <https://doi.org/10.1145/2090236.2090255>
- Ekstrand, M.D., Das, A., Burke, R. and Diaz, F. (2022). [Fairness in information access systems](#). *Foundations and Trends in Information Retrieval*, 16(1-2): 1-177. <https://doi.org/10.1561/15000000079>
- European Union (2019). [Regulation \(EU\) 2019/1150 of the European Parliament and of the Council of 20 June 2019 on promoting fairness and transparency for business users of online intermediation services](#). *Official Journal of the European Union*, L 186: 57-79.
- European Union (2022). [Regulation \(EU\) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector \(Digital Markets Act\)](#). *Official Journal of the European Union*, L 265: 1-66.
- Fleder, D. and Hosanagar, K. (2009). [Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity](#). *Management Science*, 55(5): 697-712. <https://doi.org/10.1287/mnsc.1080.0974>
- Hagiu, A. and Wright, J. (2015). [Multi-sided platforms](#). *International Journal of Industrial Organization*, 43: 162-174. <https://doi.org/10.1016/j.ijindorg.2015.03.003>
- Joachims, T., Swaminathan, A. and Schnabel, T. (2017). [Unbiased learning-to-rank with biased feedback](#). *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 781-789.

<https://doi.org/10.1145/3018661.3018699>

- Khan, L.M. (2017). *Amazon's antitrust paradox*. *Yale Law Journal*, 126(3): 710-805.
- Mehrotra, R., McInerney, J., Bouchard, H., Lalmas, M. and Diaz, F. (2018). *Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems*. *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, 2243-2251. <https://doi.org/10.1145/3269206.3272027>
- Morik, M., Singh, A., Hong, J. and Joachims, T. (2020). *Controlling fairness and bias in dynamic learning-to-rank*. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 429-438. <https://doi.org/10.1145/3397271.3401100>
- Patro, G.K., Biswas, A., Ganguly, N., Gummadi, K.P. and Chakraborty, A. (2020). *FairRec: Two-sided fairness for personalized recommendations in two-sided platforms*. *Proceedings of The Web Conference 2020*, 1194-1204. <https://doi.org/10.1145/3366423.3380196>
- Rochet, J.-C. and Tirole, J. (2003). *Platform competition in two-sided markets*. *Journal of the European Economic Association*, 1(4): 990-1029. <https://doi.org/10.1162/154247603322493212>
- Schein, A.I., Popescul, A., Ungar, L.H. and Pennock, D.M. (2002). *Methods and metrics for cold-start recommendations*. *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 253-260. <https://doi.org/10.1145/564376.564421>
- Schnabel, T., Swaminathan, A., Singh, A., Chandak, N. and Joachims, T. (2016). *Recommendations as treatments: Debiasing learning and evaluation*. *Proceedings of the 33rd International Conference on Machine Learning*, PMLR 48: 1670-1679.
- Singh, A. and Joachims, T. (2018). *Fairness of exposure in rankings*. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2219-2228. <https://doi.org/10.1145/3219819.3220088>
- Sürer, Ö., Burke, R. and Malthouse, E.C. (2018). *Multistakeholder recommendation with provider constraints*. *Proceedings of the 12th ACM Conference on Recommender Systems*, 54-62.
- Zehlike, M., Bonchi, F., Castillo, C., Hajian, S., Megahed, M. and Baeza-Yates, R. (2017). *FA*IR: A fair top-k ranking algorithm*. *Proceedings of the 2017 ACM Conference on Information and Knowledge Management*, 1569-1578. <https://doi.org/10.1145/3132847.3132938>
- Zehlike, M., Yang, K. and Stoyanovich, J. (2022). *Fairness in ranking, Part I: Score-based ranking*. *ACM Computing Surveys*, 55(6): Article 118. <https://doi.org/10.1145/3533379>